

Coordination in Dynamic Interactions by Converging on Tacitly Agreed Joint Plans

Arthur Le Pargneux (arthur.le-pargneux@warwick.ac.uk), Hossam Zeitoun (hossam.zeitoun@wbs.ac.uk),
Emmanouil Konstantinidis (e.konstantinidis@warwick.ac.uk), Nick Chater (nick.chater@wbs.ac.uk)
University of Warwick, Coventry, CV4 7AL, United Kingdom

Abstract

People solve a myriad of coordination problems without explicit communication every day. A recent theoretical account, virtual bargaining, proposes that, to coordinate, we often simulate a negotiation process, and act according to what we would be most likely to agree to do if we were to bargain. But very often several equivalent tacit agreements — or virtual bargains — are available, which poses the challenge of figuring out which one to follow. Here we take inspiration from virtual bargaining to develop a cognitive modeling framework for dynamic coordination problems. We assume that players recognize their common goal, identify one or more possible tacit agreements based on situational features, observe the history of their partner's choices to infer the most likely tacit agreement, and play their role in the joint plan. We test this approach in two experiments ($n = 125$ and $n = 133$) based on a dynamic coordination game designed to elicit agreement-based behavior. We fit our model at the individual level and compare its performance against alternative models. Across four different conditions, our model performs best among the set of models considered. Behavioral results are also consistent with players sustaining coordination and cooperation in the task by converging on tacitly agreed strategies or “virtual bargains”.

Keywords: coordination; cooperation; agreement; virtual bargaining; computational modeling; social cognition

Introduction

Humans constantly face coordination problems, from navigating crowded corridors and passing the ball to a teammate to playing improvised music and lifting heavy furniture. How do we successfully interact with one another to ensure and sustain coordination? Virtual bargaining, a recent theory of social cognition (Misyak et al., 2014), proposes that instead of engaging in — sometimes costly and inefficient — explicit negotiation, we instead often simulate a negotiation process between the involved parties and act according to “what we would agree to do” (Chater et al., 2022). And when there is a unique “best” solution that “we” would agree, people will coordinate successfully. Very often, however, several equivalent tacit agreements are available. For example, two cars facing one another, both wanting to turn right, can avoid collision equally well if either of them “stops” and the other “goes”. How do people arbitrate between equivalent tacit agreements when communication is impossible? Building on virtual bargaining, we propose a psychological model in which people contemplate several tacit agreements at the same time, keep track of their partner's behavior to figure out which one is most likely, and then play their role in this tacitly agreed joint plan. Thus, observing the other car suddenly slowing down makes the tacit agreement (I go, you stop) more likely than (I stop, you go), which could be sufficient to break the symmetry. We sketch the contours

of a framework to model coordination based on these ideas and test it on a modified version of a coordination task — the “dice game” — developed to elicit agreement-based behavior (Le Pargneux et al., 2023). In this task, players can tacitly agree to use a colored dice as a coordination device (playing a role analogous to a traffic light). This seems to require players to engage in shared intentionality (Tomasello & Carpenter, 2007) — i.e., to identify a shared goal and specific roles to achieve that goal — and some form of collective (Bratman, 1992; Gilbert, 2006; Searle, 1990) or team-reasoning (Bacharach, 1999; Colman & Gold, 2018; Sugden, 2003) — i.e., reasoning about what “we as a team” should do to successfully coordinate. Importantly, this is a task of “impure” coordination in which players must infer opposite moves from the same random signal, a relatively sophisticated inference that goes beyond (pure) coordination based on mere visual salience. We modify the coordination task — originally static and asynchronous — to study real-time consecutive interactions with feedback. This is to investigate if and how tacit agreements can be used to initiate and sustain cooperation over time in a rich dynamic setting.

Past research, primarily in economics, has extensively investigated pure coordination games (Schelling, 1980) in which players use focal points to successfully coordinate by making the same move (e.g., selecting the same date or name) (Bardsley et al., 2010; Mehta et al., 1994). There is also existing work on repeated mixed-motive games (Guyer & Rapoport, 1969), turn-taking in variants of the prisoner's dilemma (Sibly & Tisdell, 2018), and the evolution of social norms in economic games (Bicchieri, 2005) — but these do not involve coordination devices with randomly determined signals. While this rich literature demonstrates people's impressive abilities for tacit coordination, it usually puts relatively little emphasis on the role of tacit *agreements* and virtual bargaining processes in explaining the psychological and cognitive foundations of coordination. Our *psychological* approach departs from such *game-theoretic* accounts (for a game-theoretic analysis of dynamic interactions based on virtual bargaining see Melkonyan et al. (2022)).

Many other models of coordination and social cognition have been proposed over the years (e.g., Frank & Goodman, 2012; Shafto et al., 2014; Tamir & Thornton, 2018). For such “individualistic” models, explaining shared intentionality is often problematic because of issues posed by circularity, infinite regress, and multiple equilibria (Chater et al., 2022). By contrast, our work complements previous models which explicitly build on joint-planning, we-reasoning, and team-reasoning (Ho et al., 2016; Kleiman-Weiner et al., 2016) by viewing the problems not as choosing individual actions but coordinating by converging on a tacit agreement.

Task

Impure Coordination In the modified “dice game” (Le Pargneux et al., 2023) two players are matched at random and play 30 rounds of the game together. They cannot communicate. At each round, they simultaneously and independently make a binary decision: press the button or not. The button works as follows: by pressing the button, each participant steals points (e.g., 2) from the other player. If they both press, they both steal points from each other, to their mutual disadvantage. If only one of them presses, then that player steals points from the other player, but crucially both of them also receive a large bonus (e.g., 5 points) — so that the overall outcome is unequal but mutually advantageous. If neither player presses, then both receive a small bonus (e.g., 1 point). As such, the worst outcome, in which both players steal from each other, is achieved where both players press. By contrast, the best (but unequal) outcome is achieved where players manage to successfully coordinate by having only one of them press (and hence steal from the other), so that both benefit from the large bonus. Thus the formal structure of the game corresponds to that of a finitely repeated “Leader” game, a (mixed motive) game of impure coordination (Guyer & Rapoport, 1969). While coordination is mutually beneficial, each player is better off if they are the one to press, making coordination difficult to achieve.

Coordination Device At the start of the game, each player is assigned a color at random for the whole game: one player is red, the other is blue. At the start of each round, players witness a virtual colored dice roll. The dice takes different colors in different versions of the experiment: e.g., in Experiment 1: 2 red faces, 2 blue faces, and 2 yellow faces; in Experiment 2: 3 red faces and 3 blue faces (“red-blue” condition); 3 yellow faces and 3 green faces (“green-yellow” condition); 6 yellow faces (“yellow only” condition). Thus, to efficiently cooperate, players can tacitly agree to use the dice as a coordination device (similar to a traffic light). In game theory, such combinations of strategies relate to the concept of correlated equilibrium (Aumann, 1987).

Multiple Equilibria The one-shot version of the game has two pure strategy Nash equilibria (NE), (Press, Not press) and (Not press, Press) (which lead to (7, 3) or (3, 7) “on average” in Experiment 1 (variable payoff matrix) and in Experiment 2 (fixed payoff matrix)), and one NE in mixed strategies where players press with probability p (2/3 in Experiment 1 (expected payoff: 2.33) and 3/5 in Experiment 2 (expected payoff: 2.2)). Thus, there are many subgame perfect Nash equilibria in the finitely repeated game, and it is therefore unclear how a “rational” player would play. For both players (Press, Press) leads to 0 (Experiment 1) or -1 (Experiment 2), while (Not Press, Not Press) leads to +1 (small bonus).

Tacit Agreements Several tacit agreements are available to the players to coordinate. A relatively obvious one is to “press only when the dice matches my color” (i.e., press only in “same” rounds). But the dice is an ambiguous signal: a

perfectly equivalent but perhaps less obvious tacit agreement is to “press only when the dice matches my opponent’s color” (i.e., press only in “opposite” rounds). Cooperative players can also ignore the dice entirely and adopt another type of tacit agreement in which they alternate pressing and not pressing. Again, there are two equivalent options: “one of us only presses on odd (even) rounds while the other presses on even (odd) rounds”. More sophisticated alternation patterns, such as “one player presses for n rounds, then the other player presses for n rounds, and so on and so forth”, are also possible, though they are very difficult to implement. Finally, competitive players can try to enforce a mutually beneficial but unfair tacit agreement in which “one player always presses while the other never presses”.

Cognitive Model

General Modeling Approach

First, we assume that when interacting people are able to easily infer a common goal through shared intentionality (Tomasello & Carpenter, 2007). For two cars at an intersection, the common goal could be for “both to turn right without colliding”. In the dice game it is to “get the large bonus”. Second, they identify one or more possible tacit agreements that they can use to achieve that goal from the features of the environment: e.g., “if traffic light is green (red) go, if red (green) stop”, “if the dice matches my (your) color, I press, if it matches yours (mine), you press”. Third, they are uncertain about which tacit agreement their partner will follow, if any. But they can observe their partner’s behavior (e.g., sudden deceleration, past choices) to figure out which tacit agreement seems most likely to be followed. Fourth, they play their respective role (“go (stop)”, “press (not press)”) in the most likely tacit agreement. Fifth, after the interaction they “update” which tacit agreement is most likely to be followed: e.g., “the light was green for me but the other car did not stop, perhaps green means “stop” in this country” or “the dice rolled my color but my partner pressed, perhaps I should press on opposite rounds instead”.

In dynamic interactions, if both players do this, they can converge over time on a tacitly agreed strategy by trial-and-error and approximate the outcome of explicit negotiation. This general approach can be summed up as “attempting to figure out one’s role in the tacitly agreed joint plan”.

Task-Specific Implementation

We implement this modeling approach in our task. We assume that players — who know the rules of the game — infer that they have a common goal (coordination) and mutual interests (the large bonus). Based on their knowledge of the features of the game (e.g., payoff matrix, multiple rounds, dice roll, colors), they identify one or more possible tacit agreements (e.g., “play based on dice roll”, “alternate at each round”) to achieve that goal. They observe their partner’s choices on types of rounds that correspond to identified tacit agreements (red, blue, yellow, green rounds; odd and even rounds etc.) and, based on the features of the current round,

they (probabilistically) adopt their “role” (press or not press) in the followed joint plan.

Our model (henceforth “Agreement”) separates the decision-making process in two steps. First, the decision maker forms a representation of the “utility” $U_i(\text{Press})$ (i.e., how attractive pressing is in round i) associated with pressing the button based on observed past plays of their partner:

$$(1) U_i(\text{Press}) = -\omega_{\text{odd-even}} \cdot \begin{cases} v_{\text{odd}} & \text{if round is odd} \\ v_{\text{even}} & \text{if round is even} \end{cases} - \omega_{\text{colors}} \cdot \begin{cases} v_{\text{red}} & \text{if round is red} \\ v_{\text{blue}} & \text{if round is blue} \\ v_{\text{yellow}} & \text{if round is yellow} \\ v_{\text{green}} & \text{if round is green} \end{cases}$$

v_x are cumulative sums that go up (+1) or down (-1) every time one’s partner presses/does not press on the relevant type of round e.g., if my partner tends to press (not press) on odd rounds, it becomes less (more) attractive for me to press on odd rounds (to achieve the common goal of having only one player press), thus *decreasing* (increasing) the utility I associate with pressing on odd rounds. At each round, the cumulative sums corresponding to the round’s number (odd-even) and color (dice) are used: e.g., if this is round 3 and the dice rolled blue, only the values v_{odd} and v_{blue} are used. This utility depends on two free parameters $\omega_{\text{odd-even}}$ and ω_{colors} which are individual-specific weights attached to each type of tacit agreement. They reflect both a participant’s preference or bias for one type of tacit agreement over another and the importance attached to the partner’s past plays on types of rounds corresponding to each type of tacit agreement (odd and even rounds and colors respectively). They can be equal to 0 — for example if one type of tacit agreement is not recognized as available — or negative — to reflect “misleading” information (e.g., if a player presses on odd rounds only, they will sometimes press on red rounds too (i.e., an odd-red round) but not because the dice is red).

Second, this utility is translated into a probability of pressing the button using a standard logit choice rule:

$$(2) P(\text{Press}) = \frac{1}{1 + e^{-U(\text{Press})}}$$

This specific implementation is flexible in several ways. First, it allows for different types of tacit agreements (based on odd-even rounds and on colors) and different (and incompatible) versions of each tacit agreement (“press on same (opposite) rounds only”, “press on odd (even) rounds only”). Second, it “learns” and “adapts” based on the behavior of the player’s partner and can (to an extent) prescribe switching strategies over time if necessary (e.g., when the partner switches from one strategy to another). Third, it can also handle unfair play (prescribe to never press (or always press) if the partner always (never) presses) and non-agreement-based behavior by the partner (probabilistic play when the partner plays at random or probabilistically). However, the model may be seen as a simplified abstraction of the assumed cognitive and choice processes. With only two free parameters (one for each type of tacit agreement) its

adaptive nature is limited. It could also be improved with the implementation of additional components to account for recency, decay, and/or memory effects for example. While this implementation is specific to the dice game, the overall modeling approach is generalizable to other tasks and naturally lends itself to a Bayesian framework in which players update at each round the subjective probabilities they associate with each tacit agreement given observation of their partner’s behavior. It is also psychologically plausible: players identify possible strategies that can be implemented to successfully coordinate, are uncertain about which of these (if any) their partner will follow, and try to converge on a tacitly agreed strategy by observing their partner’s behavior.

Experiment 1

Methods

The sample size, methods, and exclusion criteria for [study 1](#) and [study 2](#) were preregistered on OSF but analyses should be considered *exploratory*. Data, oTree code, and analysis scripts are available on [OSF](#) for both studies.

Participants 125 participants (50.4% female; age range: 20-76 years, mean = 41.2) were recruited from Prolific. Participants spent 15.4 minutes on the task on average. They were paid a flat fee of £1.50 for participating and could earn up to £2.00 in additional bonus payments (mean = £0.97). 7 participants were excluded from analyses following pre-registered criteria, leading to 118 complete submissions.

Procedure The game was developed and administered on oTree (Chen et al., 2016). After reading the instructions and the rules of the game, players read the description and payoff matrix of a trial round and had 2 attempts to answer 6 comprehension questions. Then they were matched with a partner, randomly assigned a color (red or blue), and played the game with the same partner for 30 rounds. Each participant started with 3 points. Each point was worth £0.01, and, at the end of the game, each participant received a bonus payment corresponding to the number of points they had accumulated. At the start of each round, participants were informed of the characteristics of the game on that round, which were chosen randomly: the outcome of the dice roll (red, blue, or yellow, 2 faces of each color), the number of points stolen by pressing (1, 2 or 3), and the size of the large bonus (4, 5 or 6). This was to test the extent to which players’ decisions were influenced by variations in the payoff matrix. Once both players had (privately) made their decision, each player was informed of their partner’s decision, the number of points that were won and stolen by each player in this round, their current number of points, and the number of rounds remaining. Then a new round started. Finally, participants answered an attention check, and provided their age, gender, and any comments they had.

Tacit Agreements In Experiment 1, there were several potential tacit agreements or “virtual bargains” available to

the players about what to do. A relatively “obvious” one was for each player to press on rounds in which the dice matched their own color (“same” rounds) and to let the other player press on rounds in which the dice matched their opponent’s color (“opposite” rounds). Crucially, this strategy could not be adopted on rounds in which the dice was yellow (“neutral” rounds). Consistent with an agreement-based account, we thus predicted that players would be more likely to press in “same” rounds than in “neutral” rounds, and in neutral rounds than in “opposite” rounds. Such behavior would suggest that the players had engaged in some form of shared attention and intentionality (Tomasello & Carpenter, 2007) or joint reasoning (Chater et al., 2022) by recognizing that they could tacitly agree to use the dice as a coordination device to coordinate without explicit communication (Le Pargneux et al., 2023). As previously discussed (see Task), another type of tacit agreement involved alternating pressing and not pressing at each round (odd-even rounds).

Modeling Results

Model Comparison We compared the Agreement model against various alternatives. One model corresponded to playing (in each round) according to the NE in mixed strategies of the “average” one-shot game (i.e., selecting Press with probability $p = 2/3$) in which the large bonus is equal to 5 points and the number of points stolen is equal to 2 points. Another was a model-free reinforcement learning model where actions are reinforced when they lead to positive rewards and become less likely when they lead to negative rewards. This model (“RL”) learns a Q -value for each action (Press/Not press) (Watkins & Dayan, 1992) and has one free parameter, the learning rate (adding an inverse temperature parameter leads to worse fit (not shown here)). Other models represented simple heuristics that participants could adopt in the task: “repeat” (press with probability .5 in the first round, then model predicts the same choice as in the previous round), “always” (press with probability 1), “random” (press with probability .5), “never” (press with probability 0), “alternate” (press with probability .5 in the first round, then model predicts a different choice than in the previous round).

Model Performance We estimated individual-level parameters for the Agreement and RL models (there are no parameters to estimate for the other models) for each participant via maximum likelihood estimation (MLE) using the “nlminb” optimization procedure (R Core Team, 2023). Table 1 contains the fit measures for all models. As a first evaluation of the performance of our models, we used the following procedure. Using the individual-level parameter estimates, we simulated 1000 series of 30 decisions for each participant. Then, we computed the proportion of correct (i.e., simulated decision is identical to observed decision) simulated decisions for each series. Table 1 presents the averages of these proportions. Based on this metric, the Agreement model performed better than the other models considered: it was correct in 67.0% of cases when the second-best model (RL) was correct in 61.5%. To assess the fit of the

models we computed the AIC and BIC. We report the average individual-level score in Table 1. Lower values indicate a better fit. Our model performed better than any other model considered on both measures. Finally, we report the number of participants best fitted by each model according to the AIC and the BIC. The Agreement model best fitted a substantial subset of participants: 36 based on the AIC and 24 based on the BIC (which penalizes more for additional free parameters). Importantly, other models best fitted other subgroups of participants, suggesting that specific subgroups played the game in different ways. The RL strategy best fitted 30 and 30 participants respectively. The mixed NE strategy best fitted 17 and 22 participants respectively. Also note that a substantial proportion of participants seemed to play “at random” presumably due to the complexity of the task and the difficulty to coordinate (see below). Other participants seemed to follow clear strategies consistently for most of the game (e.g., “alternate”, “always”, “never”).

Table 1: Modeling results for Experiment 1.

Model	Correct predictions	AIC	$n(\text{AIC})$	BIC	$n(\text{BIC})$
Agreement	67.0%	32.7 (12.5)	36	35.5 (12.5)	24
RL	61.5%	35.4 (11.9)	30	36.8 (11.9)	30
Repeat	60.2%	159.5 (103.4)	0	159.5 (103.4)	0
Always	56.6%	180.0 (103.5)	4	180.0 (103.5)	4
Mixed	52.2%	42.4 (10.4)	17	42.4 (10.4)	22
Random	50.0%	41.6 (0.00)	21	41.6 (0.00)	26
Never	43.4%	234.5 (103.5)	4	234.5 (103.5)	5
Alternate	39.8%	244.0 (103.4)	6	244.0 (103.4)	7

Note. Correct predictions: average proportion of correct predictions over all participants (1000 simulated series of 30 decisions per participant). AIC and BIC values are averages and the standard deviation is given in parentheses. $n(\text{AIC})$ and $n(\text{BIC})$ are the total number of participants best fitted by the corresponding model according to each criterion.




Figure 1: Probability to press (A) and coordination (B) per round in Experiment 1.

Behavioral Results

The main statistical model was a mixed-effects logistic regression with random intercepts per participant and pair ($\text{press} \sim \text{type} + \text{stolen} + \text{large} + (1 | \text{participant}) + (1 | \text{pair})$) where *type* denotes the type of round (“same”, “neutral” or “opposite”), *stolen* denotes the number of points (1, 2 or 3) stolen by pressing, and *large* denotes the size of the large bonus (4, 5 or 6 points) that each participant receives if only one of them presses. As predicted, participants were more likely to press in same rounds ($b = 0.63$; $p < .001$; estimated marginal means: 71.8% [95% CI: 65.5%, 77.3%]) than in neutral rounds (57.6% [50.4%, 64.4%]). Participants were also more likely to press in neutral rounds than in opposite rounds ($b = -0.51$; $p < .001$; 44.9% [37.9%, 52.2%]), see Figure 1. These results are consistent with some participants recognizing that they could tacitly agree to use the dice as a coordination device. In addition, participants were slightly more likely to press where pressing led to 2 ($b = 0.25$; $p = .012$; 60.2% [53.1%, 66.9%]) or 3 points ($b = 0.29$; $p = .003$; 61.2% [54.1%, 67.8%]) being stolen as opposed to 1 (54.2% [46.9%, 61.2%]). However, participants were as likely to press whether the large bonus was 4 (57.2% [49.9%, 64.1%]), 5 ($b = 0.08$; $p = .427$; 59.1% [51.9%, 65.9%]) or 6 points ($b = 0.09$; $p = .345$; 59.4% [52.3%, 66.1%]). Finally, exploratory analyses showed that coordination — as measured by the proportion of pairs obtaining the large bonus in a given round — improved over time (going from 47.5% in the first round to 67.8% in the last round; see Figure 1). Thus, partners became better at coordinating as the game progressed even without communication.

While coordination is achieved well above chance we suspect that coordination is made more difficult by the complexity of the task. First, several elements change at each round: the dice’s color, the number of points to steal, and the size of the large bonus. Second, because there are multiple viable tacit agreements, two cooperative players could fail to coordinate for several rounds because they both switch strategies simultaneously. Third, the color-based strategies cannot be adopted in neutral rounds. Fourth, the dice roll being random, it is not rare for participants to obtain the same color multiple times in a row, which could induce strategy switching. Fifth, while not optimal, participants can often do relatively well by pressing “at random” — a strategy that is cheap in terms of attention and effort — which may explain why a substantial proportion of participants are best fitted by the mixed and random strategies. Sixth, some participants spontaneously report being unsure whether they were playing against a bot. To alleviate some of these issues, in Experiment 2, we therefore simplify our experimental design: in the main conditions (see below), the dice can only take 1 of 2 colors and the payoff matrix remains identical in all rounds.

Experiment 2

Participants 133 participants (49.6% female; age range: 20–71 years, mean = 38.3) were recruited from Prolific. Participants spent 11.2 minutes on the task on average. They

were paid a flat fee of £1.50 for participating and could earn up to £2.00 in additional bonus payments (mean = £1.04). 7 participants were excluded from analyses following pre-registered criteria, leading to 126 complete submissions.

Conditions We recruited 2 groups of participants in two separate sessions separated by a few days. In the first session (preregistered) participants were randomly allocated to 1 of 2 conditions: “red-blue” and “green-yellow”. We subsequently decided to recruit an additional condition for exploratory purposes and the second session was not preregistered. In this session all participants were allocated to the “yellow only” condition. In the “red-blue” condition the dice was red and blue (3 faces of each color). In the “green-yellow” condition the dice was green and yellow (3 faces of each color). In the “yellow” condition all faces of the dice were yellow. In all conditions participants were assigned one color at random for the whole game: red or blue. We reasoned that in the “red-blue” condition participants could use their assigned color and the dice to coordinate. By contrast, in the “green-yellow” condition, initially assigned colors (red or blue) could not be used and each player would need to “pick” different colors over time to coordinate. In the “yellow” condition, the dice always rolled yellow and as such it could not be used as a coordination device. “Fair” coordination could only be achieved by alternating (odd-even rounds).

Procedure The procedure was nearly identical to Experiment 1. Each participant started the game with 3 points. The button worked as follows: by pressing the button, each participant stole 2 points from the other player. If they both pressed, they each stole points from each other, and each participant also received a 1-point penalty (players were informed that a negative number of points at the end of the experiment would lead to a bonus payment of £0.00). If only one of them pressed, then the player who pressed stole 2 points from the player who did not press, but crucially both of them also received a large bonus of 5 points. If neither of them pressed, then both received a small bonus of 1 point. The payoff matrix was thus identical in all rounds.

Table 2: Modeling results for Experiment 2.

Model	Correct predictions	AIC	$n(\text{AIC})$	BIC	$n(\text{BIC})$
Agreement	73.6%	26.4 (13.8)	54	28.8 (13.8)	44
RL	59.7%	36.1 (12.4)	19	37.5 (12.4)	18
Repeat	54.3%	184.0 (120.2)	0	184.0 (120.2)	0
Always	53.9%	191.1 (98.6)	7	191.1 (98.6)	8
Mixed	50.8%	41.87 (5.79)	14	41.87 (5.79)	17
Random	50.0%	41.6 (0.00)	13	41.6 (0.00)	19
Never	46.1%	223.4 (98.6)	3	223.4 (98.6)	3
Alternate	45.7%	219.5 (120.2)	7	219.5 (120.2)	7
<i>Condition-specific strategies</i>					
Green	51.5%	201.0 (110.5)	2	201.0 (110.5)	3
Red	51.3%	201.7 (100.9)	1	201.7 (100.9)	1
Blue	48.7%	212.8 (100.9)	3	212.8 (100.9)	3
Yellow	48.5%	213.5 (110.5)	3	213.5 (110.5)	3

Note. See Table 1 note.

Modeling Results

We performed the same model comparison as in Experiment 1, see Table 2. We also included condition-specific strategies for only pressing on “green”, “blue”, “yellow”, and “red” rounds, respectively. The NE in mixed strategies of the one-shot game is to Press with probability $p = 3/5$.

The Agreement model performed best in all three conditions. It correctly predicted choices in 73.6% of simulations, on average, and best fitted the data according to the AIC and the BIC. It also best fitted 54 and 44 participants according to these two measures, respectively. Other models best fitted other smaller subgroups of participants but their overall performance was substantially worse, see Table 2.

Behavioural Results

The main statistical model (red-blue vs yellow-green comparison) was a mixed-effects logistic regression with random intercepts per participant and pair (press ~ type + (1 | participant) + (1 | pair)) where *type* denotes the type of round (“same”, “neutral” or “opposite”). Participants in the red-blue condition were more likely to press in same rounds ($b = 0.45$; $p < .001$; estimated marginal means: 55.2% [95% CI: 45.4%, 64.6%]) than in opposite rounds (44.0% [34.7%, 53.8%]), consistent with some players using the dice as a coordination device. The probability to press was also higher in the green-yellow condition (in which all rounds were neutral: $b = 0.53$; $p = .049021$; 57.3% [48.4%, 65.7%]) than in opposite rounds of the red-blue condition. We computed a cooperation index (Guyer & Rapoport, 1969) to measure the extent of “fair” coordination for each pair:

$$C = \% ((P, NP) + (NP, P)) - | \% (P, NP) - \% (NP, P) |$$

Where (P, NP) is read as “the left player selects P (press) while the right player selects NP (not press)”. The index increases with the number of rounds in which only one player presses (left hand side of the subtraction — henceforth *C1*) and decreases with increasing non-parity between the players in “who gets to press” (right hand side — henceforth *C2*). 100 represents full cooperation while 0 represents complete absence of cooperative interaction. Mean cooperation was remarkably similar across conditions, but non-parity was higher in the “yellow only” ($M = 0.39$, $MC_1 = 0.73$, $MC_2 = 0.34$) than in the “red-blue” ($M = 0.43$, $MC_1 = 0.68$, $MC_2 = 0.25$) and “green-yellow” ($M = 0.38$, $MC_1 = 0.65$, $MC_2 = 0.27$) conditions. Mean cooperation was also higher in each condition of Experiment 2 than in Experiment 1 ($M = 0.28$, $MC_1 = 0.59$, $MC_2 = 0.31$). Coordination (proportion of pairs with only one player pressing, for each round) went up over time in all 3 conditions, see Figure 2. This is remarkable as different tacit agreements were available in each condition. We count pairs for which players played complementary strategies (i.e., one plays red the other blue; or yellow/green; or always/never; or both alternate) for most of the game (each member’s choices are consistent for at least 24 rounds). Many pairs coordinated by using colors (red-blue: 5; green-yellow: 4), alternating (red-blue: 2; green-yellow: 1; yellow only: 6),

or having one player always/never press (red-blue: 3; green-yellow: 3; yellow only: 5). These results are consistent with players flexibly adopting whichever type of tacit agreement is most convenient depending on the features of the environment and their partner. Thus, the simpler design of Experiment 2 produced clearer computational and behavioral evidence in support of players successfully coordinating by progressively converging on tacitly agreed joint plans.




Figure 2: Coordination per round in Experiment 2.

Discussion

In this paper we developed a modeling framework for coordination tasks inspired by shared intentionality, we-reasoning, and virtual bargaining, in which players contemplate several possible tacit agreements available to achieve a common goal and play their role in the most likely agreement given their partner’s past choices. We applied our framework to a task of impure coordination (requiring opposite moves) designed to elicit agreement-based behavior and involving dynamic interactions. Our model makes quantitative predictions and performs well in comparison to simple game-theoretic (mixed NE) and reinforcement learning (Q-value) models, and a number of simple heuristics that could be adopted by players. Behavioral results are also consistent with tacit agreements and virtual bargaining processes playing an important role in this task.

This work can be extended in a number of ways. First, in future work, we plan to enrich our model comparison with more sophisticated models from the behavioral game theory (e.g., level- k and cognitive hierarchy (Camerer et al., 2004), experience weighted attraction (Ho et al., 2007)), reinforcement learning, and social cognition literature (see Table 2 in Chater et al. (2022) for a recent review). Second, we plan to apply our approach to other coordination games, and given previously mentioned difficulties with the complexity of the current task, to test our model on simpler versions of the game. Third, we plan to build on our approach to develop a formal generalizable model in a Bayesian framework. Our work is a first step towards a better understanding of the computational foundations of shared intentionality and virtual bargaining. Formalizing agreement-based reasoning processes could shed light on the cognitive underpinnings of coordination but also on other challenges in cognitive science, including models of “contractualist” moral cognition (Levine et al., 2023) and progress towards developing genuinely social AI systems (Chater, 2023).

Acknowledgments

We thank members of the Behavioural Science Group and Computational Behavioural Science Lab at Warwick, the Moral Psychology Research Lab at Harvard, the Computational Cognitive Science Lab at MIT, and the Morality Lab at Boston College for helpful feedback and discussions. A.L.P is supported by the Economic and Social Research Council [ES/P000711/1 project 244367]. N.C. gratefully acknowledges support from the European Union/UKRI under Horizon Europe Programme Grant Agreement no. 101120763 – TANGO. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HaDEA). Neither the European Union nor the granting authority can be held responsible for them.

References

- Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 1-18.
- Bacharach, M. (1999). Interactive team reasoning: A contribution to the theory of co-operation. *Research in Economics*, 53(2), 117-147.
- Bardsley, N., Mehta, J., Starmer, C., & Sugden, R. (2010). Explaining focal points: Cognitive hierarchy theory versus team reasoning. *The Economic Journal*, 120(543), 40-79.
- Bicchieri, C. (2005). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press.
- Bratman, M. E. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2), 327-341.
- Camerer, C. F., Ho, T.-H., & Chong, J.-K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3), 861-898.
- Chater, N. (2023). How could we make a social robot? A virtual bargaining approach. *Philosophical Transactions of the Royal Society A*, 381(2251), 20220040.
- Chater, N., Zeitoun, H., & Melkonyan, T. (2022). The paradox of social interaction: Shared intentionality, we-reasoning, and virtual bargaining. *Psychological Review*, 129(3), 415.
- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88-97.
- Colman, A. M., & Gold, N. (2018). Team reasoning: Solving the puzzle of coordination. *Psychonomic Bulletin & Review*, 25, 1770-1783.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998-998.
- Gilbert, M. (2006). Rationality in collective action. *Philosophy of the Social Sciences*, 36(1), 3-17.
- Guyer, M., & Rapoport, A. (1969). Information effects in two mixed-motive games. *Behavioral Science*, 14(6), 467-482.
- Ho, M. K., MacGlashan, J., Greenwald, A., Littman, M. L., Hilliard, E., Trimbach, C., Brawner, S., Tenenbaum, J., Kleiman-Weiner, M., & Austerweil, J. L. (2016). Feature-based joint planning and norm learning in collaborative games. *Proceedings of the Annual Meeting of the Cognitive Science Society 2016*.
- Ho, T. H., Camerer, C. F., & Chong, J.-K. (2007). Self-tuning experience weighted attraction learning in games. *Journal of Economic Theory*, 133(1), 177-198.
- Kleiman-Weiner, M., Ho, M. K., Austerweil, J. L., Littman, M. L., & Tenenbaum, J. B. (2016). Coordinate to cooperate or compete: Abstract goals and joint intentions in social interaction. *Proceedings of the Annual Meeting of the Cognitive Science Society 2016*.
- Le Pargneux, A., Chater, N., & Zeitoun, H. (2024, March 6). Contractualist tendencies and reasoning in moral judgment and decision making. <https://doi.org/10.31234/osf.io/p4cyx>
- Levine, S., Chater, N., Tenenbaum, J., & Cushman, F. A. (2023, May 7). Resource-rational contractualism: A triple theory of moral cognition. <https://doi.org/10.31234/osf.io/p48t7>
- Mehta, J., Starmer, C., & Sugden, R. (1994). The nature of salience: An experimental investigation of pure coordination games. *The American Economic Review*, 84(3), 658-673.
- Melkonyan, T., Zeitoun, H., & Chater, N. (2022). The cognitive foundations of tacit commitments: A virtual bargaining model of dynamic interactions. *Journal of Mathematical Psychology*, 108, 102640.
- Misyak, J. B., Melkonyan, T., Zeitoun, H., & Chater, N. (2014). Unwritten rules: Virtual bargaining underpins social interaction, culture, and society. *Trends in Cognitive Sciences*, 18(10), 512-519.
- Schelling, T. C. (1980). *The Strategy of Conflict*. Harvard University Press.
- Searle, J. R. (1990). Collective intentions and actions. *Intentions in Communication*, 401(4), 401.
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55-89.
- Sibly, H., & Tisdell, J. (2018). Cooperation and turn taking in finitely-repeated prisoners' dilemmas: An experimental analysis. *Journal of Economic Psychology*, 64, 49-56.
- Sugden, R. (2003). The logic of team reasoning. *Philosophical Explorations*, 6(3), 165-181.
- Tamir, D. I., & Thornton, M. A. (2018). Modeling the predictive social mind. *Trends in Cognitive Sciences*, 22(3), 201-212.
- R Core Team (2023). *R: A language and environment for statistical computing* (Version 4.3.0) [Computer

- software]. Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- Tomasello, M., & Carpenter, M. (2007). Shared intentionality. *Developmental Science*, 10(1), 121-125.
- Watkins, C. J., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8, 279-292.