(Re-) Inventing the Traffic Light: Designing Recommendation Devices for Play of Strategic Games^{*}

Mikhail Anufriev^a

John Duffy^b Valentyn Panchenko^c Benjamin Young^a

August, 2024

^a Department of Economics, University of Technology Sydney
 ^b Department of Economics, University of California, Irvine
 ^c Economics, UNSW Business School, University of New South Wales

Abstract

We present the results of a novel experiment investigating individuals' ability to design incentive-compatible recommendation devices for five canonical 2×2 games played by rational robot players. Human designers were incentivized to achieve Pareto efficiency and fairness. Most subjects succeeded in Matching Pennies and Battle of the Sexes. However, only a minority did so in the two Chicken games, though many did design an incentive-compatible device. In the Prisoner's Dilemma, the vast majority failed to design such a device, not recognizing the conflict between strategic incentives and social efficiency. The study's task requires a holistic approach to equilibrium reasoning, and our findings suggest it is most challenging when strategic incentives are not aligned with cooperative outcomes.

Keywords: Correlated Equilibrium, Recommendation Device, Nash Equilibrium, Prisoner's Dilemma, Experimental Economics.

JEL Codes: C72, C91, D82, D91.

^{*}We thank Timothy Cason, Nick Feltovich, Gigi Foster, Dan Friedman, Michael McBride, Jean Paul Rabanal, Tridib Sharma, John Wooders, and seminar participants at Chapman University, UCI, and NYU Abu-Dhabi for their helpful comments and suggestions. The experiment was preregistered on the AEA RCT Registry (AEARCTR-0009706), and the experimental protocol was approved by the Irvine Institutional Review Board. The authors acknowledge support from the Australian Research Council's Discovery Project funding scheme (project DP200101438).

1 Introduction

Nash equilibrium (Nash Jr, 1950) and its more general counterpart, correlated equilibrium (Aumann, 1974, 1987), are important solution concepts in game theory that form the basis of a multitude of predictions and rationalizations of behavior in applied economic theory. In one interpretation of correlated equilibrium due to Myerson (1991), a mediator designs a device that makes private but correlated recommendations to each player in the game.¹ The recommendation device constitutes a correlated equilibrium if every player is always willing to follow the device's recommendation. To shed light on human cognition, we design a novel experiment where individuals act as the *designers* of such recommendation devices. Specifically, we examine subjects' ability to coordinate the behavior of players in games towards desirable outcomes, while considering the strategic incentives of the players to follow recommendations. We find that subjects are most successful in this task when strategic incentives align with socially desirable outcomes.

Consider the following metaphorical example. Two cars are approaching an obscured intersection along perpendicular roads. If one car waits, the other prefers to go rather than wait. The worst outcome is both cars proceeding simultaneously, leading to a crash. Without seeing or communicating with each other, coordination is impossible. Now, suppose a traffic light is installed, showing green for one car and red for the other half the time. With the norm that green means 'Go' and red means 'Stop', both cars will *always* follow the traffic light's instructions. In this scenario, the traffic light coordinates the cars' behavior to achieve a fair and efficient outcome. In other words, the traffic light serves as a recommendation device that implements a particularly desirable correlated equilibrium.

Analogies for coordinating devices like traffic lights are plentiful, including mediation in conflict resolution, central bankers' announcements guiding market expecta-

¹Myerson (1991), p. 250 refers to the mediator as "a person or machine that can help the players communicate and share information." Similarly, Gintis (2014), p. 142 refers to such a designer as "the choreographer," – a "new player" in the "augmented" game who issues a "directive".

tions, the design of data exchange protocols between servers, or even the role of a conductor leading an orchestra to ensure harmony. But how effective are people at crafting such devices or recommendations for others? Do they recognize the incentives of the other parties to follow the recommendations; i.e., do their recommendations constitute a correlated equilibria? How do features of the environment, such as prevailing incentives, influence the recommendations they provide? Are there some situations in which crafting effective recommendations is relatively straightforward and others where it is particularly difficult?

To date, research on the empirical relevance of correlated equilibrium has focused on placing subjects in the role of a player in the game to investigate whether they are willing to follow, or learn to follow, the recommendations produced by coordinating devices. However, numerous factors may influence an individual's propensity to follow recommendations. For instance, those deviating from recommendations may struggle to effectively reason about their own strategy or the strategies of others, lack trust in their opponents (i.e., face strategic uncertainty), or exhibit social preferences (such as fairness considerations) beyond the experimenter's control. Thus, asking someone to simply play games does not necessarily identify their degree of strategic sophistication. In contrast, our research is the first to investigate whether experimental subjects can *design* a recommendation device to coordinate the behavior of players within a given game. To focus on the behavior of the designer and exclude the multitude of factors that may shape behavior of the players, we automate the players in the game as *rational robots*. Using the above example, while prior research places subjects in the role of drivers, we task our subjects to design a traffic light.

Our experimental setting is of interest for several reasons. First, designing a recommendation device presumes an understanding of how subjects themselves would react to these recommendations. Therefore, evidence of this designing ability provides a more comprehensive evaluation of the relevance of correlated equilibrium.²

²Our approach also offers insights into the empirical relevance of correlated versus Nash equilibria. While every Nash equilibrium is a correlated equilibrium, the reverse is not true. Thus, it is of interest to understand the extent to which individuals recognize and use the broader concept of

Second, previous tests of recommendation following a conditioned signal structure are certainly interesting but ultimately incomplete. They rely on an exogenous external device, but it is unclear how such a device might have arisen in the first place.³ Thus, our approach can be viewed more broadly as placing subjects in the role of 'recommendation designers', rather than simply observing their responses to various recommendations. Third, by using rational robots for the players who consider the recommendations of the human designers, we eliminate additional noise from individual player behavior. The consistent behavior of rational robot players helps isolate the comprehension of the design problem by the human equilibrium designer. Indeed, designing a correlation device requires a holistic approach that integrates strategic incentives, randomization, and correlation, but it also depends on how designers expect players to engage with the game. Our design removes this latter variable. Finally, advancements in AI technology require a better understanding of human-robot interactions.⁴ Recent literature primarily focused on games where humans play against robots (e.g., Bayer and Renou, 2024) or robo-advice to humans (e.g., Bianchi and Brière, 2024). Instead, we contribute to how humans design systems to coordinate robots.⁵

Our experimental design is as follows. Participants design recommendation devices for five different games under a common objective. They have five attempts to design a device for each game. Each device provides probabilistic, private, potentially correlated recommendations to pairs of *robot players* playing each game. These robot players are Bayesian expected-utility maximizers who determine whether following the recommendations from the device aligns with their self-interest. Thus, a device for which all recommendations are followed implements a correlated equilibrium.

correlated equilibrium.

³For instance, traffic lights are designed and continually adjusted by human engineers.

⁴Consider the importance of crafting effective prompts to fully utilize the large language models like ChatGPT or determining appropriate tasks for robots.

⁵This is increasingly important as more robotic autonomous devices interact without direct human involvement. For example, autonomous vehicles are already navigating our roads and sometimes face difficult moral trade-offs, such as deciding whether to save their own passengers or those in another autonomous vehicle when a collision is unavoidable. Humans need to establish rules for these situations aligning them with our moral and norms, see, Bonnefon *et al.* (2024).

We are interested in subjects' ability to design a device that achieves a correlated equilibrium, and, conditional on this, results in a fair and efficient outcome. This is accomplished as follows. First, subjects earn a positive payoff from their device only if all of its recommendations are followed by the robot players. Second, the payoff earned from a device for which all recommendations are followed increases with the minimum expected payoff across the robot players.⁶ As all the games we consider are symmetric, the *desirable correlated equilibrium* (DCE), that maximizes a subject's payoff is one that is both Pareto efficient and fair. We focus on this selection rule from the set of correlated equilibria because it is naturally applicable in practice. For example, the traffic light achieves this social objective by being Pareto efficient, – allowing cars from one direction only to pass through the intersection at a time, – and ensuring fairness, as cars from different directions have an opportunity to cross the intersection.

We ask subjects to design recommendation devices for five games: Prisoner's Dilemma, Matching Pennies, Battle of the Sexes, and two versions of Chicken that differ in both their set of correlated equilibria and their DCE. These games are selected because the structure of their DCE differs in important aspects of equilibrium reasoning, such as (1) the necessity of randomization, (2) the degree of correlation required, and (3) the alignment of strategic incentives with the social objective of fairness and efficiency. For example, the DCE in the Prisoner's Dilemma is the unique Nash equilibrium in pure strategies, which does not require randomization or correlation and has payoff incentives that conflict with efficient outcomes. In contrast, the DCE in the Battle of the Sexes involves perfectly correlated randomization between its two pure-strategy Nash equilibria and has payoff incentives that are perfectly aligned with the objective of fairness and efficiency. These differences allow us to explore whether these distinct aspects of equilibrium reasoning impact subjects' ability to successfully

⁶In this sense, subjects are induced to apply the *Rawlsian criterion* (see Rawls, 1971) – to make the worst-off player as well off as possible, – as a secondary objective in selecting a correlated equilibrium. Subjects aim to satisfy this criterion subject to the constraint that the device is incentive compatible (i.e., it constitutes a correlated equilibrium).

implement the DCE through their recommendation device.

We find significant heterogeneity in subjects' proficiency at crafting recommendation devices to achieve a correlated equilibrium and the DCE across different games. Specifically, a significant portion of participants successfully designed the DCE in Battle of the Sexes and Matching Pennies, with success rates of 79% and 73%, respectively. In contrast, fewer subjects achieved this in the Prisoner's Dilemma and two versions of Chicken, with success rates of 37%, 18%, and 30%, respectively. In the Chicken games, however, most subjects were able to implement some correlated equilibrium (93% and 96%). An analysis using a distributional distance metric reveals that subjects' designs closely align with the ideal outcomes in all games except for the Prisoner's Dilemma.

Our results reveal that subjects' success in designing recommendation devices correlates positively with the alignment between strategic incentives and socially desirable outcomes. Specifically, success rates are highest in games where a fair and Pareto efficient outcome coincides with a correlated equilibrium, as seen in Battle of the Sexes and Matching Pennies. Subjects are adept at managing randomization and correlation levels tailored to these games, despite differing requirements (perfect correlation in Battle of the Sexes versus none in Matching Pennies). Conversely, success rates decline when strategic incentives conflict with social desirability, as seen in the Prisoner's Dilemma and Chicken games. The difficulty appears to arise from a reluctance to adopt a non-cooperative mindset, which is crucial for responding to robot players' strategies. This is particularly stark in Prisoner's Dilemma, where subjects fail to recognize the strong strategic incentives of each robot player, governed by the presence of a strictly dominant strategy. Thus, we find that individuals *can* effectively design recommendation devices with desirable characteristics but struggle when strategic and social goals are misaligned.

2 Related Literature

There are a number of papers that experimentally investigate correlated behavior in games. These papers find that individuals tend to correlate their behavior when pre-play communication is permitted (Moreno and Wooders, 1998), follow correlatedequilibrium recommendations when those improve upon Nash equilibrium outcomes (Duffy and Feltovich, 2010), and may learn to correlate their behavior in the absence of coordinating devices (Friedman *et al.*, 2022). The literature also finds, however, that the propensity of subjects to follow recommendations that constitute a correlated equilibrium is mediated by factors such as strategic uncertainty (Cason and Sharma, 2007), the directness of recommendations (Duffy *et al.*, 2017), payoff asymmetries (Anbarci *et al.*, 2018), whether recommendations are private or public (Bone *et al.*, 2013), and whether subjects must commit to the coordinating device as a whole (Georgalos *et al.*, 2020). These papers focus mostly on the willingness of individuals to adhere to the recommendations made by correlation devices. By contrast in this paper, we ask whether individuals are able to design the correlation device themselves.

An exception to studies that investigate subjects' willingness to follow correlatedequilibrium recommendations is the work of Cason *et al.* (2022). In their paper, experimental subjects are asked to predict the distribution over actions that will be generated by other individuals playing that game. They find that subjects tend to believe that play will be correlated even though the people playing the games do not have access to any correlating device. In contrast, instead of asking subjects to predict what *will* happen in the game, we ask subjects to design a coordination device that implements what they believe *should* happen in the game. Moreover, our subjects design devices for rational robot players to avoid the strategic uncertainties created by human players.

There is also a related literature on mechanism design experiments, surveyed in Chen and Ledyard (2010). These experiments generally involve evaluation of subjects' play under various imposed mechanisms with most applications concerning public good provision, auctions, contract theory, matching markets, and prediction markets, among others. In this literature, it is typically *not* the case that subjects are placed in the designer role as in our study. Still, we think that for simple games, such as those that we study here, placing subjects in the designer role is a reasonable approach. That is because it reveals the design preferences of non-expert designers, which may, in turn, hold greater appeal to those operating under such designs.

Our work is also related to the literature on experiments regarding information design and Bayesian persuasion (Kamenica and Gentzkow, 2011). These papers ask whether subjects can use information structures that successfully correlate behavior with some payoff-relevant state of the world (Aristidou *et al.*, 2019; Au *et al.*, 2023; Fréchette *et al.*, 2022; Wu and Ye, 2023; Ziegler, 2022). In our paper, we focus on complete information games (i.e., a single payoff-relevant state) and subjects seek to create signal structures or recommendations that correlate the behavior of players in that complete information game (i.e., implement correlated equilibria).

Finally, our work is related to studies exploring the ability of individuals to draw inferences from correlated signals and the notion of correlation neglect (Enke and Zimmermann, 2019; Hossain and Okui, 2021). These issues are more likely to affect the players in the game who must decide whether to follow correlated recommendations made privately. We, instead, ask the reverse question: can subject actually create the correlated signals from which inferences must be drawn? As such, our subjects need to think about signal correlation *hypothetically* rather than drawing inferences from *actual* signal realizations.⁷

3 Theoretical Framework

A finite game of complete information can be described by the tuple

$$G = (N, (A_i)_{i \in N}, (u_i)_{i \in N}),$$

⁷This is an important distinction as there is evidence that thinking about hypothetical events is fundamentally different from extracting information from actual events (Esponda and Vespa, 2014).

where N is a set of players, A_i is a finite action set for player $i \in N$, and $u_i : A \to \mathbb{R}$ is the utility of player *i* defined over the set of action profiles $A \equiv \prod_{i \in N} A_i$. Let $\Delta(A)$ be the set of probability distributions over A. For a game G, we can define a distribution over action profiles $\mu \in \Delta(A)$ to be a correlated equilibrium as follows.

Definition. Distribution μ is a correlated equilibrium (CE) for game G if

$$\sum_{a_{-i}} \mu(a_i, a_{-i}) [u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i})] \ge 0$$

for all $i \in N$, $a_i, a'_i \in A_i$.

In one interpretation of a CE, μ serves as a statistical summary of a 'recommendation device'. In this framework, the action profile a is drawn with probability $\mu(a)$, and each player i receives only their own recommended action a_i from profile a. The device implements a CE if, for any drawn profile a, each player i finds it optimal to follow their recommendation, given that all other players also follow their respective recommendations.⁸ Note that a Nash equilibrium is a special case of a CE where μ is the product of its marginal distributions across players; that is, action recommendations are made *independently* to the players in the game.

In our experiment, subjects will design recommendation devices in an attempt to implement CE. Since some games have multiple CE, we specifically focus on the CE that solves the following constrained optimization problem:

$$\max_{\mu \in \Delta(A)} \min_{i \in N} \sum_{a \in A} \mu(a) u_i(a), \quad \text{s.t. } \mu \text{ is a correlated equilibrium.}$$
(1)

In other words, we are interested in the CE that maximizes the minimum average payoff across players, following the Rawlsian criterion (Rawls, 1971). Each game for which our subjects design recommendation devices is symmetric, making the problem in (1) equivalent to finding the CE that is both efficient (i.e., surplus maximizing) and

⁸Alternative theories explain why observed behavior in games might constitute a CE without requiring the interpretation of a recommendation device. For example, CE can arise as Bayes' Nash equilibria when players condition on extraneous signals, or from simple adaptive procedures with repeated play (Hart and Mas-Colell, 2000). We focus on the recommendation device interpretation because it provides a natural framework to investigate individuals' ability to implement CE.

fair (i.e., equally distributing surplus among players). We refer to this distribution as the **desirable correlated equilibrium** (DCE) and denote it by μ^* .

We focus on the CE that is fair and Pareto efficient as these are natural objectives in practice, as the traffic light example from the Introduction illustrates. Moreover, experimental evidence suggests that individuals are more likely to follow CE recommendations when they offer efficiency gains compared to Nash equilibria (Duffy and Feltovich, 2010) and minimize payoff asymmetries (Anbarci *et al.*, 2018). Consequently, recommendation devices that solve (1) should be more effective at coordinating real-world behavior.

Throughout the paper, we will also refer to the **socially desirable outcome** (SDO) which is the *unconstrained* solution to problem (1): the probability distribution over action profiles that is fair and efficient but not necessarily a CE. We denote this distribution by μ^{SD} . The SDO differs from the DCE when μ^{SD} is not incentive compatible. In such cases, we will say that the strategic incentives of the game are not perfectly aligned with the SDO.

4 Experimental Design

Our experiment employed a within-subjects design and was coded using oTree (Chen *et al.*, 2016). Each subject designed recommendation devices for five different games, presented in random order. For each game, subjects were given five consecutive attempts (referred to as rounds) to design their device.⁹

Our subjects are 100 undergraduate students from the University of California, Irvine, pursuing various programs of study. In total, we collected 2,500 designed recommendation devices (100 participants×5 games×5 rounds per game). The sample was gender balanced with 54 females and 46 males. Subjects had no prior experience

⁹We selected five games to explore how variations in game structure impact subjects' understanding of strategic incentives and the appropriate use of correlation and randomization. The decision to include five rounds per game was based on a pilot and aimed to balance the opportunity for subjects to learn while maintaining high stakes for each attempt.

with the experiment and participated in a single session with five games. Payments were determined by converting points earned from the usage of their recommendation devices into USD (more details are provided in Section 4.2). In addition, subjects received a 10 USD show-up payment. The average total earnings for subjects were 27.64 USD for an experiment lasting, on average, 90 minutes (including approximately 25 minutes spent reading instructions and 20 minutes for a comprehension quiz).

Given the complexity of the experimental task, we implemented a simple visual interface for designing recommendation devices and took extensive measures to ensure that subjects understood the task. Written instructions (see Online Appendix A) were provided and read aloud at the start of the experiment, which included detailed explanations and examples for a specific game different from those used in the experiment. The instructions explained how to use the interface to design recommendation devices, how robot players would analyze the game and respond to those devices, and how payments would be determined.

Subjects were then required to complete a nine-question comprehension quiz to assess their understanding of the experiment's key aspects, which included explicitly designing devices. They could not begin the actual experiment until all questions were answered correctly, with detailed feedback provided after each response. Additionally, three practice rounds were conducted using the same game from the instructions.

At the end of the experiment, we collected additional information to investigate potential correlates with performance in the design task. This included standard demographic information, field and year of study, performance on a five-question version of the cognitive reflection test (CRT) (Frederick, 2005), a measure of strategic reasoning related to the hypothetical play of the 11-20 game (Arad and Rubinstein, 2012), and the self-reported strategies utilized by subjects.¹⁰

¹⁰We used CRT questions because they evaluate inhibitory control and analytical thinking and are often found to correlate with performance in logical tasks (Campitelli and Gerrans, 2014). The 11-20 game (Arad and Rubinstein, 2012) was included to measure players' levels of reasoning, as our design task requires high levels of reasoning. For the precise wording of the CRT questions and the 11-20 task, see Online Appendix B.

P1/P2	Red (R)	Blue (B)	P1/P2	$\ \operatorname{Red} (R)$	Blue (B)
Red (R)	$\left \begin{array}{c} (u_1^{RR}, u_2^{RR}) \end{array} \right $	(u_1^{RB}, u_2^{RB})	Red (R)	μ_{RR}	μ_{RB}
Blue (B)	$\left \left(u_1^{\boldsymbol{BR}}, u_2^{\boldsymbol{BR}} \right) \right.$	(u_1^{BB}, u_2^{BB})	Blue (B)	$\ \mu_{BR}$	μ_{BB}

Figure 1: Left: Normal form for a 2×2 game, where both players choose between actions Red and Blue. Right: Table representation of a recommendation device.

4.1 The Design Problem

Subjects designed recommendation devices for five 2×2 games which take the general normal form shown in the left panel of Fig. 1. Each player in the game has two actions, either Red (denoted R) or Blue (denoted B). We chose to use colors to provide a clear visual representation of the recommendation device that subjects needed to design. Subjects were told that their goal was to make recommendations to 'intelligent' robot players that would follow these recommendations *only if* it is in their best interest to do so (i.e., they strictly prefer to follow the recommendations or are indifferent).¹¹

In theory, any recommendation device for a 2×2 game can be summarized using a table, as shown in the right panel of Fig. 1. In this table, the entries $\mu_{a_1a_2}$ represent the probabilities that the device jointly recommends actions $a_1 \in \{R, B\}$ and $a_2 \in \{R, B\}$ to player 1 and player 2, respectively.

We implemented the task of constructing such a table as follows. Subjects were provided on the screen with a bucket containing 24 balls, each labeled with '1' on one half and '2' on the other half. To design the device, subjects colored each half of these balls with either red or blue. They were instructed that for each pair of robot players, a random ball would be selected from the bucket and split in half. The half labeled '1' would be given to robot-player 1, and the half labeled '2' would be given to robot-player 2. The color of the half given to each robot-player would represent the recommended action for that player.

¹¹Essentially we explained in layman's terms that the robot players are Bayesian expected-utility maximizers. See page 2 of the instructions in the Online Appendix A for the precise explanation. Pages 6-8 of the instructions provide examples of the robots' reasoning.



Figure 2: Example of the decision screen used by subjects to input a recommendation device. The game shown here was used in the instructions and for practice rounds.

The balls were colored by the subjects using sliders. First, subjects used a slider to determine how many of the 24 balls' halves labeled '1' should be colored red, with the remainder colored blue. Next, subjects used two additional sliders to color the half labeled '2': one slider for balls where the half labeled '1' was red and another one for balls where the half labeled '1' was blue. Once all balls were colored, the design of the recommendation device was complete and subjects were provided with a summary table of the recommendations their device would make; see an example of the screen with a designed device in Fig. 2.

Subjects were then asked whether they wanted to submit their device for possible use by the robot players. Before confirming their submission, subjects could continue to make any alterations to their device without any time limits. Both the device and the summary table were updated anytime as subjects adjusted their color choices for the two halves of the 24 balls.

While the finite number of balls imposes some limitations, this device still enables subjects to implement a large number of probability distributions. Specifically, subjects could feasibly design 2925 distinct probability distributions, including the DCE distribution of each game. Therefore, it was unlikely that subjects would find the DCE distribution *randomly* within the five attempts provided for each game.

4.2 Incentivization and Feedback Between Rounds

For each game, subjects' payoffs were aligned with the objective in (1): to design a recommendation device that implements the DCE, i.e., a CE where the minimum expected payoff across the two players is as large as possible. We achieved this by structuring the point system as follows.

Subjects were informed that a large number of pairs of robot players would play the game against each other, with each pair receiving a random recommendation made by their designed device.¹² If even a single instance occurred where a robot player chose not to follow a recommendation made by the device, no points would be awarded for that device. Conversely, if recommendations from a device were *always* followed, the subject would earn points equal to the minimum of the average payoff across *all* robot players 1 and the average payoff across *all* robot players 2.

With this payment structure, subjects were motivated to target the set of CE as a whole, since a *positive* payoff was only awarded for a device for which all recommendations were followed. However, the maximum number of points was awarded for the device that implemented the solution of (1), i.e., the DCE μ^* .

At the end of the experiment, a random game was selected, and subjects were paid for the device that earned the *highest* number of points for that game across five rounds. This approach was intended to motivate subjects who had found a CE to keep experimenting to identify the device that implements the DCE.¹³ Subjects were

¹²Although this explanation may suggest a dynamic context with repeated interactions, we emphasized that robots would treat each round as a one-shot game when deciding whether to follow the recommendations. See page 2 of the instructions in the Online Appendix A.

¹³Choosing a random round within a random game might have led subjects to stop experimenting once they found *any* CE (not necessarily the DCE), in order to maximize their chances of receiving some payment.

paid at a rate of 6 USD per point earned, in addition to a show-up fee of 10 USD.

We provided feedback to subjects between rounds. They were informed whether all recommendations from their device were followed. If so, we shared the points earned and whether the device could be further improved to earn more (see Figs. C1 and C2 in Online Appendix C). This feedback essentially indicated whether they had found the DCE, only a CE, or neither. Along with the rule that only the best device across the five rounds in a game was considered for payment, this incentivized subjects to experiment with their designs.

If some recommendations were not followed (i.e., a CE was not achieved), we presented a ball from their device where one robot player did not follow the recommendation. This ball was selected randomly based on the device's probability distribution. We explained which player did not follow the recommendation and why, emphasizing that subjects must consider their device as a whole to ensure incentive compatibility.¹⁴ To avoid overwhelming subjects, we revealed only one ball at a time and did not show the full calculations. This approach encouraged deeper reasoning without revealing the correct answer. Subjects could refer to the instructions for examples of reasoning. Evidence provided in the Online Appendix E.1 suggests that subjects responded to this feedback.¹⁵

4.3 Games Considered in the Experiment

Subjects designed recommendation devices for five games: Prisoner's Dilemma, Matching Pennies, Battle of the Sexes, and two versions of the Chicken game. The payoffs of each game are presented in the left column of Table 1, together with the volume of the set of all CE as a fraction of the volume of the set of probability distributions over action profiles and the payoff in the DCE (the distribution that solves problem

¹⁴See Fig. C3 in Online Appendix C, where the reasoning was "Given what P2 perceived about what P1 was recommended, which is based on the composition of all the balls in the container, P2 decided that it is in its interest **not to follow** the recommendation."

¹⁵In particular, when informed that a recommendation from a particular ball was not followed, subjects tended to decrease the number of this ball in their subsequent device.

(1)). We list below the DCE for each game.¹⁶

- (i) **Prisoner's Dilemma:** The DCE is $\mu_{RR}^* = 1$ and $\mu_{RB}^* = \mu_{BR}^* = \mu_{BB}^* = 0$.
- (ii) Matching Pennies: The DCE is $\mu_{RR}^* = \mu_{RB}^* = \mu_{BR}^* = \mu_{BB}^* = 1/4$.
- (iii) Battle of the Sexes: The DCE is $\mu_{RB}^* = \mu_{BR}^* = 1/2$, and $\mu_{RR}^* = \mu_{BB}^* = 0$.
- (iv) Chicken, version 1: The DCE is $\mu_{RB}^* = \mu_{BR}^* = \mu_{BB}^* = 1/3$, and $\mu_{RR}^* = 0$.
- (v) Chicken, version 2: The DCE is $\mu_{BB}^* = 1/2$, $\mu_{RB}^* = \mu_{BR}^* = 1/4$, and $\mu_{RR}^* = 0$.

Table 1 also provides a graphical representation of each game in both the joint distribution space (middle panels) and the payoff space (right panels). In these plots, white discs correspond to the Nash equilibria of the game, and the black dot shows the DCE. The distribution space is a three-dimensional simplex, represented in coordinates ($\mu_{RB}, \mu_{BR}, \mu_{BB}$). Here, the origin corresponds to giving full weight to the (R, R) recommendation, which is the Nash equilibrium for the Prisoner's Dilemma, for example. In the first two games, there is a unique CE, so their sets are singletons with zero volume. For the last three games, the set of CEs forms a polytope; its edges are shown in black and the faces are colored. Projections of the DCE onto the axes are indicated by dashed lines.

The space of expected payoffs for two robot players is two-dimensional, offering a clearer illustration. However, it should be emphasized that these are projections from the distribution space. For each game, we display (i) the set of feasible payoffs (the light-grey shaded area), (ii) the set of CEs (the light-blue shaded area with a thick solid boundary), and (iii) the convex hull of Nash equilibrium payoffs (the dark-blue shaded area). In both Chicken games, the convex hull is a proper subset of the set of CEs, but in the Battle of the Sexes, these two sets coincide. For each game, the DCE corresponds to the CE that is both Pareto efficient and fair, meaning that it lies on the upper frontier of CE payoffs and the expected payoffs of both robot players are equal.

¹⁶Online Appendix D shows derivations and how our interface implements them.



Table 1: Games used in the experiment.

4.4 Research Questions

With this experimental design, we address several research questions for each game:

- 1. Do subjects design devices that implement the DCE? If not, how close are they?
- 2. What devices do subjects tend to design? How do these evolve with experience?
- 3. Do subjects' performance and their individual characteristics correlate?

As discussed in the Introduction, evaluating the ability to design recommendations that support a CE offers a more comprehensive assessment of equilibrium thinking, revealing aspects of individuals' strategic reasoning that previous experiments could not address due to various confounds. Generally, the implementation of CEs in strategic interactions requires understanding such concepts as individual incentives, randomization, and correlation. Therefore, we have chosen five classic games with varying CE and DCE structures to investigate how these structural differences impact answers to our research questions. These differences are summarized in Table 2.

All games require randomization except for Prisoner's Dilemma where the unique CE is the unique Nash equilibrium in pure strategies. Neither Prisoner's Dilemma nor Matching Pennies require correlation, as their DCEs are also Nash equilibria. In contrast, Battle of the Sexes requires perfect correlation (i.e., it can be implemented using a public randomization device), while the two Chicken games require more nuanced, partial correlation (i.e., private recommendations are required).

	Randomization	Correlation	Alignment of Incentives
PD	No	None	Strongly Unaligned
MP	Yes	None	Aligned
BoS	Yes	Perfect	Aligned
C1	Yes	Partial	Weakly Aligned
C2	Yes	Partial	Weakly Aligned

Table 2: Differences in the types of reasoning required to solve for the DCE across games.

We consider individual incentives in the games from the perspective of their alignment with the socially desirable outcome (SDO). In Matching Pennies and Battle of the Sexes, the DCEs are socially desirable (i.e., the constrained and unconstrained versions of (1) have the same solution), indicating perfect alignment. In Prisoner's Dilemma and both versions of Chicken, however, the SDO is $\mu_{BB}^{SD} = 1$, which is not a CE. We describe the alignment as weak in the Chicken games, where the DCE assigns positive weight on the SDO ($\mu_{BB}^{SD} = 1/3$ in C1 and $\mu_{BB}^{SD} = 1/2$ in C2). In contrast, we consider the incentives strongly unaligned with the SDO in the Prisoner's Dilemma, where $\mu_{BB}^{SD} = 0.17$

5 Results

In Sections 5.1 to 5.3, we address in turn the three research questions posed above for each game. Then, in Section 5.4, we compare the results across games and discuss the insights gained from this comparison, emphasizing four key findings.

5.1 How Well Did Subjects Perform?

Recall that the subjects' objective in each game was to design a device corresponding to the DCE, which would maximize their points. In three games, where other CEs are also possible, designing a device corresponding to any of them would still earn some points. Additionally, when the DCE is not found, it is interesting how far the subjects were from designing the optimal device.

These considerations lead us to focus on three performance measures.¹⁸

TopScore: An indicator equal to 1 if a subject designed the DCE, and 0 otherwise.

Score: The number of points a subject earned, relative to the maximum points pos-

 $^{^{17}}$ A quantitative measure of the alignment between strategic incentives and the SDO is provided in Section 5.4.1.

 $^{^{18}}$ Online Appendix E.2 presents two additional measures: an indicator whether any CE was found and Selten measure, which penalizes the CE indicator by the volume of the CE-set.

sible. Score is 1 if the DCE is implemented, and positive if any CE is implemented.¹⁹ **N-RMSE**: A normalized Root Mean Square Error (RMSE) measure capturing how close a subject's designed distribution μ is to the DCE μ^* . The RMSE between two arbitrary distributions μ and ν is defined as:

$$\operatorname{RMSE}(\mu,\nu) \equiv \sqrt{\sum_{a \in \{\boldsymbol{R},\boldsymbol{B}\}^2} (\mu(a) - \nu(a))^2}.$$

N-RMSE scales the measure to the [0, 1] interval and increases as μ approaches μ^* :

$$N-RMSE(\mu;\mu^*) \equiv 1 - \frac{RMSE(\mu,\mu^*)}{\max_{\mu'}RMSE(\mu',\mu^*)}.$$
(2)

This measure complements the two point-based measures, as a subject's device may implement a distribution close to the DCE but still earn zero points if it is not a CE. Note that N-RMSE equals 1 if and only if $\mu = \mu^*$, and 0 when μ is maximally distant from the DCE μ^* in terms of RMSE.

All three measures are defined for a given game and round. However, we focus on participants' *best attempt* in each game. Therefore, for each participant, we calculate all three performance measures for each of the five rounds. Then, for each measure, we take the participant's maximum across the five rounds and average these across the entire sample. Fig. 3 provides a graphical comparison of the three aggregate performance measures, computed in this way, across the five games in our study.

We begin by examining whether subjects designed devices that implement the DCE in each game. TopScore shows that performance is highest in Battle of the Sexes (nearly 80% found the DCE) and Matching Pennies (nearly 75% found the DCE), with no statistically significant difference between the two.²⁰ In contrast, a minority of subjects implemented the DCE in the other games: approximately 40% in Prisoner's Dilemma, 30% in Chicken 2, and 20% in Chicken 1.

¹⁹Normalization is useful because the maximum points varied across games. For example, the second version of Chicken has a maximum payoff of 4.5 points, while all other games have a maximum of 4 points.

 $^{^{20}}$ This is based on a two-sided *t*-test for the difference of two averages at the 5% significance level, which we use hereafter unless specified otherwise.



Figure 3: The average (across all subjects) of the highest value in five rounds for each of the three performance measures in each game. Error bars represent 95% confidence intervals.

Next, we assess how close the designed devices are to implementing the DCE by examining both Score and N-RMSE. N-RMSE is high for Battle of the Sexes and Matching Pennies, indicating that even subjects who did not find the DCE were close to it. Despite this similarity, the Score in Battle of the Sexes is significantly higher than in Matching Pennies. This difference arises because Matching Pennies has a unique CE, while Battle of the Sexes has multiple CEs. While subjects struggled to implement the DCE in the Chicken games, they were near the objective, as reflected by a high Score. That is, subjects consistently found a CE that was reasonably close to the optimal one in terms of Pareto efficiency and fairness. Although N-RMSE is significantly lower in the Chicken games compared to Battle of the Sexes and Matching Pennies, it is significantly higher than in Prisoner's Dilemma. In Prisoner's Dilemma, performance was the lowest: subjects struggled to find the DCE and, overall, designed devices that were maximally distant from it in terms of N-RMSE.

In summary, overall performance was highest in Battle of the Sexes and Match-



Figure 4: Average designed device for each game across rounds (from round 1 to round 5, reading left to right within each game bin). The final bar in each bin, marked with a star \star , represents the DCE for the game. Black circles indicate the average Score per round.

ing Pennies, intermediate in the Chicken games, and lowest in Prisoner's Dilemma. Section 5.4 explores how the structures of these games may underlie this finding.

5.2 What Devices Did Subjects Design?

We now examine the recommendation devices designed by subjects in each game. To do this, we adjust our raw dataset by imputing the DCE device for all rounds following the first round in which a subject achieved the DCE.²¹

We begin by illustrating in Fig. 4 the *average* designed device for each round in each game. This average is calculated by taking the mean of the designed distribu-

²¹Subjects completed all five rounds for a game, even if they achieved the DCE earlier. In 687 out of 711 cases, subjects did not change their device once the DCE was achieved. However, in 24 cases (3.38%), subjects continued to experiment with other devices despite already having designed the DCE. Although this experimentation had no impact on payoffs (since they were paid for the best round), these instances could distort our performance statistics if not corrected. Therefore, we imputed the use of the DCE for all subsequent rounds after it was first achieved.

tions across all 100 subjects for each round.²² Each bin shows the fractions of the corresponding action profiles (reflecting the four types of colored balls placed by subjects in their containers) for each of the five rounds, with the final bin for each game displaying the DCE distribution. We also superimpose the evolution of the Score measure as a black line.²³

Starting with Prisoner's Dilemma, recall that the unique CE in this game is also the unique Nash equilibrium in pure strategies, i.e., where $\mu_{RR}^* = 1$. Although subjects increased the weight on action profile (R, R) round-by-round, convergence to the DCE was slow: by round 5, the average device allocated only approximately 50% weight to (R, R). Subjects did not fully abandon the socially desirable profile (B, B), as the average weight on this profile remained relatively constant from rounds 2 to 5. This further corroborates the finding that performance in Prisoner's Dilemma showed the greatest distance from the DCE (recall Fig. 3).

In contrast, for Matching Pennies and Battle of the Sexes, the average device by round 5 is very close to the corresponding DCE. This is in line with Score and other performance measures as shown in Fig. 3. Note that learning in these games is relatively quick, with significant improvements occurring already after round 1.

For the two versions of the Chicken game, subjects, on average, successfully eliminated the extremely undesirable action profile (R, R), but did not consistently implement the precise DCE distribution. By round 5, subjects appear to *underestimate* the extent to which they can recommend the socially desirable profile (B, B).

While inspecting the aggregate data provides insight into the average recommendation behavior of participants, investigating the distribution of recommendation devices at the individual level offers a deeper understanding of how subjects attempted to solve each game. For each game, we compare the distribution of recommendation devices designed by subjects in round 1 (their initial attempt) with the distribution

²²Note that for a device μ_i designed by subject *i*, the average over *N* subjects, $\overline{\mu} \equiv \sum_{i=1}^{N} \mu_i / N$, is a feasible distribution over action profiles $\{\underline{R}, B\}^2$; that is, $\overline{\mu} \in \Delta(\{\underline{R}, B\}^2)$.

²³Online Appendix E.3 shows the dynamics over the five rounds for other performance measures.



(c) Distribution Space: Round 1 (d) Distribution Space: Round 5

Figure 5: Bubble plots displaying, in both the payoff-space and distribution-space, the distribution of designed devices for **Prisoner's Dilemma** in round 1 vs. round 5.

in round 5 (their final attempt or the DCE, if achieved earlier). The results for the Prisoner's Dilemma and the two Chicken games – identified as the most challenging – are discussed here, while Online Appendix E.4 discusses the remaining two games. Figs. 5 to 7 illustrate the results both in the space of players' expected payoffs (top panels of each figure) and in the distribution space with coordinates ($\mu_{RB}, \mu_{BR}, \mu_{BB}$) (bottom panels of each figure), cf. Table 1.

Prisoner's Dilemma. Fig. 5c reveals that, in Round 1 of the Prisoner's Dilemma, several salient distributions are attempted by subjects, including the DCE, the SDO, and a 50-50 mix between these two profiles. In general, subjects tend to place too much weight on the (B, B) profile, as indicated by the distinct number of devices implying expected payoffs in the upper-right quadrant relative to the DCE in Fig. 5a. By round 5, as shown in Fig. 5d, the modal response is the DCE. However, a majority of recommendation devices remain suboptimal, appearing almost randomly dispersed across the distribution space. Thus, no consistent pattern emerges among subjects who failed to reach the DCE by round 5. Nonetheless, there remains a tendency for devices in round 5 to place substantial weight on the socially desirable (B, B) profile, as demonstrated by the significant number of devices implying expected payoffs that lie along the 45-degree line above the DCE (Fig. 5b).

Chicken, Version 1. Recall that in the Chicken 1 game, the DCE places zero probability on profile (R, R) and equal probability on the remaining profiles, including the socially desirable profile, (B, B). As seen in Fig. 6c, subjects begin optimistically: they assign more weight to the socially desirable profile than strategic incentives justify. This is strongly evidenced by the modal recommendation device, which coincides with the SDO, and by many devices implying expected payoffs that lie to the upperright of the DCE payoffs (Fig. 6a). By the final round, however, subjects seem to have learned that they cannot recommend (B, B) too often. Indeed, the majority of devices now fall within the set of CE (Fig. 6d).²⁴ The two most salient devices in round 5 are the DCE (now the modal device) and a device that mixes 50-50 between the two asymmetric Nash equilibria, (R, B) and (B, R). Additionally, some devices (Fig. 6b). This suggests that subjects recognized that partial correlation, by mixing some (B, B) with the asymmetric action profiles, was desirable given the objective. Very few subjects continue to design devices that over-weight the socially desirable

 $^{^{24}{\}rm The}$ percentage of subjects who designed a device implementing any CE increased from 37% in round 1 to 93% in round 5.



(c) Distribution Space: Round 1

(d) Distribution Space: Round 5

Figure 6: Bubble plots displaying, in both the payoff-space and distribution-space, the distribution of designed devices for **Chicken 1** in round 1 vs. round 5.

profile. In fact, most appear to underestimate the extent to which (B, B) can be recommended. This aligns with the average device for Chicken 1 shown in Fig. 4.

Chicken, Version 2. This game is similar to Chicken 1 in that the DCE assigns zero weight to the inefficient action profile (R, R). The main difference is that subjects can place more weight on the socially desirable profile in this version of Chicken $(\mu_{BB}^* = 1/2 \text{ rather than } 1/3 \text{ as in Chicken } 1)$. Examining Fig. 7, we observe that



(c) Distribution Space: Round 1

(d) Distribution Space: Round 5

Figure 7: Bubble plots displaying, in both the payoff-space and distribution-space, the distribution of designed devices for **Chicken 2** in round 1 vs. round 5.

the distribution of designed devices for Chicken 2 resembles the distribution found in Chicken 1. In round 1, subjects overestimate the extent to which they can recommend the socially desirable action profile: the modal device implements the SDO, μ^{SD} . However, as with Chicken 1, by round 5, subjects appear to have understood qualitatively that they cannot recommend (B, B) too frequently. Instead, most recommendation devices now constitute CEs.²⁵ The modal device implements the DCE,

 $^{^{25}{\}rm The}$ percentage of subjects who designed a device implementing any CE increased from 54% in round 1 to 96% in round 5.

Variable	Average	Description
StudyYear	2.96	year of study (from 1-freshman to 6-PhD Student)
CRTNcorr	2.44	number of correct answers on CRT (from 0 to 5)
CRTIntuit	1.68	number of intuitive answers on CRT (from 0 to 5)
StratSoph	0.60	level of reasoning (from 0 to 2)
GenderF	0.54	1 if gender is Female, $= 0$ if gender is Male
GameOrder	3.00	order in which the game was played (1-first : 5-last)
Study major	rs:	Arts, BiologicalSc, Business, Design, Economics, Engineering,
		HealthMed, Inform&CS, SocialSc, PhysicalSc

Table 3: List of all potential correlates. The second column reports their average values.

yet the device that mixes 50-50 between the two asymmetric Nash equilibria, (\mathbf{R}, \mathbf{B}) and (\mathbf{B}, \mathbf{R}) , is also salient (Fig. 7d). Fig. 7b further shows that several devices implement expected payoffs along the 45-degree line between these two salient devices, again suggesting that subjects understood that partial correlation of the actions of the robot players was necessary to some extent. Consequently, similar to Chicken 1, the average designed device assigns an insufficient probability to the socially desirable profile, as observed in Fig. 4.

5.3 Do Subject Characteristics Correlate with Performance?

At the end of the experiment, we collected demographic information from the subjects. We also presented them with a five-question CRT test and the 11-20 game task designed to measure their strategic reasoning. We now investigate whether any of this information sheds light on the sources of individual heterogeneity in performance. Since the game order was randomized across subjects, we also check whether this had any impact on performance. Table 3 provides a list of all the correlates we consider.²⁶

As the Score is directly linked to the subjects' payoffs, we focus on analyzing this measure here. The results for the TopScore performance measure are presented

²⁶See Online Appendix B for the correct and intuitive answers to CRT questions. For the 11-20 game task, we followed Arad and Rubinstein (2012), estimating the highest level as 2 and categorizing subjects who answered 18 or 19 as levels 2 and 1, respectively. All other subjects were categorized as level 0. For the proportions of study majors, see Table E2 in Online Appendix E.5. This table also reports all pairwise correlations among the potential explanatory variables.

	PD	MP	BoS	C1	$\mathbf{C2}$
Com o Ondon	-0.0627^{*}	0.0346	-0.0244	0.0235	0.0075
GameOrder	(0.0317)	(0.0310)	(0.0165)	(0.0211)	(0.0184)
CRTNcorr	-0.0159	0.0618^{**}	0.0349^{**}	-0.0025	0.0311^{*}
	(0.0299)	(0.0309)	(0.0146)	(0.0187)	(0.0162)
Qtara t Q a ra la	0.0809	0.0746	0.0144	0.0728^{*}	0.0528
StratSoph	(0.0623)	(0.0641)	(0.0301)	(0.0388)	(0.0336)
CondorF	-0.0392	0.0834	-0.0221	-0.0805	0.0368
Genderr	(0.1029)	(0.1055)	(0.0509)	(0.0644)	(0.0573)
$\mathbf{C}_{\mathbf{L}} = \mathbf{L}_{\mathbf{V}_{\mathbf{L}}}$	0.1184***	-0.0247	0.0096	0.0020	0.0033
Study rear	(0.0446)	(0.0457)	(0.0216)	(0.0277)	(0.0241)
Diagnostic: R^2	0.3326	0.1689	0.2865	0.1869	0.1402
* 10%, ** 5%, *** 1% levels of significance (two-sided t -test)					

Table 4: Regressions of the highest Score (over five rounds) on the set of correlates. Study major dummies were included as additional controls. These and the intercept are not shown here (see Online Appendix E.5 for complete results). Each regression uses 100 observations. Standard errors are reported in parentheses.

in Online Appendix E.5. For each game, we regress the highest value (over the five rounds) of the Score for each subject on the set of correlates. The results are presented in Table 4.

First, note that performance on the CRT significantly increases the Score in Matching Pennies and Battle of the Sexes at the 5% level. This can be attributed to the fact that our experimental task required substantial contemplative effort rather than intuitive reasoning. However, these two games were among the easiest for the subjects, suggesting that high CRT performance alone was not sufficient for the remaining three games.²⁷ Second, the year of study is a significant determinant of performance in the Prisoner's Dilemma.²⁸

It is also informative to examine potential correlates that do not exhibit a signif-

 $^{^{27}}$ The CRTNcorr is also significant in Chicken 2 but only at the 10% level. For the TopScore regressions, it is *not* significant for both Chicken games, but it is significant for the Battle of the Sexes.

²⁸From more general regressions in Online Appendix E.5, subjects who major in Economics, Engineering, and Physical Sciences perform significantly better in the Prisoner's Dilemma compared to those who major in other fields. This likely reflects a higher familiarity with the Prisoner's Dilemma from their study among these experimental subjects.

icant relationship. First, there is no effect of strategic sophistication on the Score, except for marginal significance at the 10% level in Chicken 1. This may be because our measure is too crude to capture the necessary elements for our task or because it is positively correlated with CRTNcorr. Second, no order effects are observed, indicating that performance in a particular game was not influenced by when subjects engaged with that game.²⁹ Thus improvements in performance primarily occur with experience *within* a game across rounds, as discussed in Section 5.2.

These findings suggest that individual characteristics (at least those we collected) may explain less of the variation in performance compared to the characteristics of the games themselves. We explore this further in the following section.

5.4 Discussion

Our analysis shows significant performance differences across the five games (Fig. 3). Subjects successfully found the DCE in Battle of the Sexes and Matching Pennies but struggled in Prisoner's Dilemma, Chicken 1, and Chicken 2 (though many did find some CE in the latter two). Learning was also slow in these latter three games. By the fifth round, the DCE was the modal response, but the average device deviated significantly from the DCE, particularly in Prisoner's Dilemma (Fig. 4). This considerable variation in performance cannot be solely attributed to subjects' characteristics.

Implementing the DCE in each game involves three key aspects of equilibrium reasoning: (a) aligning strategic incentives with socially desirable outcomes, (b) the necessity of randomization, and (c) the degree of correlation required, as detailed in Table 2. We will now examine our results through these lenses.

 $^{^{29}{\}rm The}$ only exception is the Prisoner's Dilemma, where, at the 10% significance level, performance decreased as the game was played later.

	$\mu^{\rm SD}$	μ^*	$ 1 - \text{N-RMSE}(\mu^{\text{SD}}; \mu^*)$
PD	(0,0,0,1)	(1, 0, 0, 0)	1
MP	$\left(\frac{1}{4},\frac{1}{4},\frac{1}{4},\frac{1}{4},\frac{1}{4}\right)$	$\left(\frac{1}{4},\frac{1}{4},\frac{1}{4},\frac{1}{4},\frac{1}{4}\right)$	0
BoS	$\left(0,\frac{1}{2},\frac{1}{2},0\right)$	$\left(0,\frac{1}{2},\frac{1}{2},0\right)$	0
C1	(0, 0, 0, 1)	$\left(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$	0.707
C2	(0,0,0,1)	$\left(0, \frac{1}{4}, \frac{1}{4}, \frac{1}{2}\right)$	0.522

Table 5: The SDO, μ^{SD} and the DCE, μ^* , in coordinates $(\mu_{RR}^*, \mu_{RB}^*, \mu_{BR}^*, \mu_{BB}^*)$; and the N-RMSE between them for each game.

5.4.1 The Alignment of Strategic Incentives

We first examine how well the strategic incentives in each game align with socially desirable outcomes. Subjects aimed to design a recommendation device that generates a fair and Pareto-efficient outcome for rational players, targeting the DCE, μ^* , which solves the constrained optimization problem (1). However, the objectives of fairness and efficiency might also lead towards the socially desirable outcome, μ^{SD} , the unconstrained solution of (1) that ignores the strategic incentives of rational players. As discussed in Section 4.4, the alignment between the DCE and the SDO varies significantly across five games, which may explain the differences in design success.

We quantify this alignment, previously shown qualitatively in Table 2, by computing the normalized distance between the DCE and the SDO, see Table 5.³⁰ Strategic incentives align perfectly with the SDO in Matching Pennies and Battle of the Sexes, where μ^* and μ^{SD} coincide. In contrast, Prisoner's Dilemma shows maximal misalignment, as μ^* is far from μ^{SD} . Both Chicken games exhibit partial alignment, with Chicken 2 aligning more closely with μ^{SD} than Chicken 1. The alignment order is:

$$\mathbf{MP} \sim \mathbf{BoS} \succ \mathbf{C2} \succ \mathbf{C1} \succ \mathbf{PD}.$$
 (3)

We observe that this order relates to performance in the experiment and formulate:

 $^{^{30}}$ Our performance measure, N-RMSE, defined in (2), is the normalized RMSE transformed to decrease with distance. Here, we use 1 - N-RMSE.

Finding 1. Performance in the experiment is inversely related to how well strategic incentives align with the social objective.

Support: Subjects performed best in Matching Pennies and Battle of the Sexes, where incentives fully align with the social objective. They performed less well in the Chicken games, where incentives are only weakly aligned but still managed to find CE close to the optimum (evidenced by high Score and N-RMSE in Fig. 3). Moreover, in their final attempts, subjects tended to *underweight* the socially optimal action profile (B, B), as discussed in Section 5.2. This suggests that subjects, at least qualitatively, understood the role of strategic incentives in the Chicken games, which may be due to the fact that these incentives are weakly aligned with the SDO. Furthermore, according to (3), strategic incentives are more closely aligned with the SDO in Chicken 2 compared to Chicken 1. As shown in Fig. 3, subjects were more likely to find the DCE in Chicken 2 (approximately 10% more subjects) and achieved a higher Score, although these differences are not statistically significant. The poorest performance was in Prisoner's Dilemma, where strategic incentives are completely misaligned with the social objective: most subjects did not find the DCE, and their devices were relatively distant from the optimum (Fig. 3). Subjects consistently placed a positive weight on the socially optimal action profile (B, B) throughout the experiment (Figs. 5b and 5d), and the average weight on this action profile did not decrease significantly with experience (Fig. 4).

The analysis in Section 5.2 provides insight into how the relationship in Finding 1 arises. Initially, subjects tend to implement the unconstrained optimal device for each game. The degree of alignment between the DCE and the SDO affects their success. Subjects then adjust their devices towards the DCE but do this slowly. We formulate

Finding 2. Subjects initially design a device for each game that implements the socially desirable outcome, μ^{SD} . By their final attempt, the modal device is the DCE, μ^* ; however, adjustments to this device are slower the further μ^* is from μ^{SD} .

Support: In Round 1, the unconstrained optimal device, μ^{SD} , is the modal choice in

Matching Pennies, Battle of the Sexes, and the two Chicken games, and is a salient choice in Prisoner's Dilemma. This device works well in Matching Pennies and Battle of the Sexes, where the SDO aligns with the DCE. In other games, subjects need to adjust their devices to account for incentive constraints. By Round 5, the DCE becomes the modal choice across all games. While subjects do not, in the aggregate, successfully implement the DCE in either version of Chicken, μ^* is close enough to μ^{SD} that subjects are able to find a CE with positive payoffs. In contrast, since the DCE is maximally distant from the SDO in Prisoner's Dilemma, adjustments are insufficient, resulting in the lowest final performance in Round 5 across all games.

5.4.2 The Necessity of Randomization

Our performance results suggest that subjects had no difficulty designing devices that generate recommendations at random. In particular, they successfully designed the DCE when randomization (at least in the symmetric form) was required. We have:

Finding 3. Subjects are capable of implementing randomization of recommendations with their devices.

<u>Support</u>: Performance is highest in Matching Pennies and Battle of the Sexes, both of which necessitate randomization. In contrast, the DCE for Prisoner's Dilemma does not require randomization, and this is the game where aggregate performance was lowest.

One might suspect that our experimental design encouraged subjects to *overuse* randomization. However, we found no evidence of aversion to design devices offering pure-strategy recommendations. Many subjects designed pure-strategy devices, including those implementing the SDO in the Prisoner's Dilemma and Chicken games (Figs. 5a, 6a and 7a). Subjects also designed devices implementing the asymmetric pure-strategy Nash equilibria in the Battle of the Sexes and Chicken games (Figs. 6, 7 and E5). Furthermore, we provided examples of devices recommending only a single action profile in the instructions to demonstrate that such devices could be optimal

given the objective.³¹

5.4.3 The Degree of Correlation

Focusing finally on the degree of correlation required in the DCE, we formulate:

Finding 4. Subjects are capable of designing devices with varying degrees of correlation. However, they tend to struggle more with implementing partial correlation compared to perfect correlation.

<u>Support</u>: Subjects achieved the highest TopScore in games requiring perfect correlation (Battle of the Sexes) and no correlation (Matching Pennies). In contrast, the TopScore was lower in each version of the Chicken game, where partial correlation was required. However, this difficulty seems to be quantitative rather than qualitative. The relatively high Score in the Chicken games indicates that subjects were able to achieve at least partial correlation. Specifically, they successfully eliminated the inefficient profile (R, R) and assigned some weight to the unconstrained optimal profile (B, B).

While subjects had the least success in the Prisoner's Dilemma, a game that requires no correlation, they performed well in Matching Pennies, which also does not require correlated behavior. Thus, the lack of correlation alone does not account for the low performance in the Prisoner's Dilemma.

6 Conclusion

Our main contribution in this paper is the development of a methodology for examining strategic sophistication while mitigating common issues such as strategic

³¹Romero and Rosokha (2023) found that individuals can randomize effectively in the indefinitely repeated Prisoner's Dilemma when given an appropriate interface. In contrast, our experiment involves one-shot Prisoner's Dilemma, where randomization does not yield a payoff to the subject. Nevertheless, subjects still tend to use randomization strategies.

uncertainty and social preferences. We achieve this by having subjects design recommendation devices for the play of games for rational robot players, that only follow their own recommendations if it is in their best interests given the distribution of recommendations made to other players implied by their device.

Overall, subjects successfully designed recommendation devices for various games, aiming to implement correlated equilibria that are Pareto efficient and fair. Their performance significantly improved across rounds, with the modal device designed by subjects in every game effectively implementing the desirable correlated equilibrium.

Performance, however, varied across games. Subjects excelled in designing devices for Battle of the Sexes and Matching Pennies, where most successfully identified the desirable correlated equilibrium, and those who did not still created devices that were nearly optimal. While subjects found it more challenging to design the best device for the two versions of the Chicken game, they were still able to implement some correlated equilibria. In contrast, in Prisoner's Dilemma, subjects struggled significantly with the qualitative aspects of equilibrium reasoning driven by the game's strategic incentives. Most subjects did not solve the Prisoner's Dilemma problem, as they could not let go of the idea of recommending a socially desirable profile with a strictly positive probability. This suggests that individuals face difficulties reconciling strategic incentives with socially desirable outcomes.

Our findings imply that people approach equilibrium reasoning with a cooperative mindset, even when dealing with non-cooperative incentives, particularly in scenarios aimed at achieving socially desirable outcomes. It appears challenging for subjects to accept that strategic incentives may conflict with desirable outcomes, and they often do not realize that no device can utilize the socially efficient profile in the Prisoner's Dilemma. This tendency holds despite the strong strategic incentives present, where each rational robot player has a strictly dominant strategy. That is, humans seem to project their cooperative mindset onto robots.

The experimental framework proposed in this paper is flexible and can be adapted

for investigating correlated equilibria in any two-player game. While we focused on symmetric games, future research could explore asymmetric games and coordination among more than two players. Moreover, although our experimental design with robot players was intended to isolate individuals' reasoning about equilibrium, it would be valuable to study whether recommendations differ when made for humans instead of robots. This would help distinguish between individuals' strategic reasoning capabilities and their perceptions of others' abilities in strategic contexts. This promising direction will be explored in future research.

References

- Anbarcı, N., Feltovich, N. and Gürdal, M. Y. (2018) Payoff inequity reduces the effectiveness of correlated-equilibrium recommendations, *European Economic Re*view, 108, 172–190.
- Arad, A. and Rubinstein, A. (2012) The 11–20 money request game: A level-k reasoning study, American Economic Review, 102, 3561–3573.
- Aristidou, A., Coricelli, G. and Vostroknutov, A. (2019) Incentives or persuasion? An experimental investigation, Working paper.
- Au, P. H., Kwon, O. and Li, K. K. (2023) A simple experiment on simple Bayesian persuasion, Working Paper.
- Aumann, R. J. (1974) Subjectivity and correlation in randomized strategies, Journal of Mathematical Economics, 1, 67–96.
- Aumann, R. J. (1987) Correlated equilibrium as an expression of Bayesian rationality, *Econometrica*, 55, 1–18.
- Bayer, R.-C. and Renou, L. (2024) Interacting with man or machine: When do humans reason better?, *Management Science*, forthcoming.
- Bianchi, M. and Brière, M. (2024) Human-robot interactions in investment decisions, Management Science, forthcoming.
- Bone, J., Drouvelis, M. and Ray, I. (2013) Coordination in 2 x 2 games by following recommendations from correlated equilibria, Working Paper.
- Bonnefon, J.-F., Rahwan, I. and Shariff, A. (2024) The moral psychology of artificial intelligence, Annual review of psychology, 75, 653–675.
- Campitelli, G. and Gerrans, P. (2014) Does the cognitive reflection test measure cognitive reflection? a mathematical modeling approach, *Memory & cognition*, 42, 434–447.
- Cason, T. N. and Sharma, T. (2007) Recommended play and correlated equilibria: An experimental study, *Economic Theory*, **33**, 11–27.
- Cason, T. N., Sharma, T. and Vadovic, R. (2022) Correlated equilibria and forecasts based on naïve play in hawk-dove games, Working Paper, SSRN-4174461.
- Chen, D. L., Schonger, M. and Wickens, C. (2016) oTree An open-source platform for laboratory, online, and field experiments, *Journal of Behavioral and Experimental Finance*, 9, 88–97.
- Chen, Y. and Ledyard, J. O. (2010) Mechanism design experiments, in *Behavioural and Experimental Economics* (Eds.) S. N. Durlauf and L. E. Blume, Palgrave Macmillan, pp. 191–205.
- Duffy, J. and Feltovich, N. (2010) Correlated equilibria, good and bad: An experimental study, *International Economic Review*, 51, 701–721.
- Duffy, J., Lai, E. K. and Lim, W. (2017) Coordination via correlation: An experimental study, *Economic Theory*, 64, 265–304.
- Enke, B. and Zimmermann, F. (2019) Correlation neglect in belief formation, The Review of Economic Studies, 86, 313–332.
- Esponda, I. and Vespa, E. (2014) Hypothetical thinking and information extraction in the laboratory, American Economic Journal: Microeconomics, 6, 180–202.
- Fréchette, G. R., Lizzeri, A. and Perego, J. (2022) Rules and commitment in communication: An experimental analysis, *Econometrica*, **90**, 2283–2318.
- Frederick, S. (2005) Cognitive reflection and decision making, Journal of Economic

Perspectives, **19**, 25–42.

- Friedman, D., Rabanal, J. P., Rud, O. A. and Zhao, S. (2022) On the empirical relevance of correlated equilibrium, *Journal of Economic Theory*, 205, 105531.
- Georgalos, K., Ray, I. and SenGupta, S. (2020) Nash versus coarse correlation, *Experimental Economics*, 23, 1178–1204.
- Gintis, H. (2014) The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences-Revised Edition, Princeton University Press.
- Hart, S. and Mas-Colell, A. (2000) A simple adaptive procedure leading to correlated equilibrium, *Econometrica*, 68, 1127–1150.
- Hossain, T. and Okui, R. (2021) Belief formation under signal correlation, Working Paper, SSRN 3218152.
- Kamenica, E. and Gentzkow, M. (2011) Bayesian persuasion, American Economic Review, 101, 2590–2615.
- Moreno, D. and Wooders, J. (1998) An experimental study of communication and coordination in noncooperative games, *Games and Economic Behavior*, 24, 47–76.
- Myerson, R. B. (1991) Game Theory: Analysis of Conflict, Harvard University Press.
- Nash Jr, J. F. (1950) Equilibrium points in N-person games, Proceedings of the National Academy of Sciences, 36, 48–49.
- Rawls, J. (1971) A Theory of Justice: Original Edition, Harvard University Press.
- Romero, J. and Rosokha, Y. (2023) Mixed strategies in the indefinitely repeated prisoner's dilemma, *Econometrica*, **91**, 2295–2331.
- Selten, R. and Krischker, W. (1982) Comparison of two theories for characteristic function experiments, Working paper, Center for Mathematical Economics, Bielefeld University.
- Wu, W. and Ye, B. (2023) Competition in persuasion: An experiment, Games and Economic Behavior, 138, 72–89.
- Ziegler, A. (2022) Persuading an audience: Testing information design in the laboratory, Working Paper.

Appendix For Online Publication Only

A Experimental Instructions

Welcome

Welcome to this experiment in economic decision-making. Please read these instructions carefully as they explain how you earn money from the decisions you make. For completing this experiment you are guaranteed the \$10 show-up fee. Depending on your actions, you can earn an additional amount as explained below. Your earnings during the experiment will be computed in *points* that will be converted to the real dollars at the end of the experiment at the rate of 1 point = \$6, so more points means more money.

Please do not talk with others for the duration of this experiment and silence all mobile devices. If you have any questions, please raise your hand.

Overview

Your job will be to design recommendations to two intelligent robot players on how they should play a game with one another. In this game, the two robot players have to simultaneously choose between two actions, labeled **Red** and **Blue**. After making their choices, each robot player receives a payoff in points that depends: (1) on the action color they chose and (2) on the action color that the other player chose.

The payoffs of robot Player 1 and of robot Player 2 are shown in the payoff table of the game. Each cell of the table corresponds to the chosen actions of both players. The payoffs of robot Player 1 are shown in the lower left part of the cell, and the payoffs of robot Player 2 are shown in the upper right part of the cell.

The figure below provides an **example** of the payoff table for a game. In this case, if, say, Player 1 (P1) plays Blue and Player 2 (P2) plays Red, then the payoff

to Player 1 is 2 points and the payoff to Player 2 is 5 points.

P2	Red		Blue	
P1				
	\searrow	1	$\overline{\ }$	3
Red	0	$\overline{\ }$	3	\searrow
	\sim	5	$\overline{\ }$	1
Blue	2	$\overline{\ }$	3	\searrow

Games will differ in their payoffs to each player. The specific payoffs will be known to both the robot players and you.

The two robot players are not able to communicate with one another. Further, each game will be played for multiple periods but the robot players in each period will always be different (so that their actions will not depend on what was played in the past).

Your task is to design a device that will make recommendations to both players about which color action they should choose in each period. These recommendations are separate, that is, each robot player knows only their own recommendation from your device. However, the robot players know the design of your device and use this knowledge to infer as much as possible about the recommendation the other robot player received.

The robot players are not easily influenced and **do not follow recommendations blindly**. They are self-interested, intelligent and take into account that the other robot player is self-interested and intelligent as well. When a particular color action is recommended to a robot, it will investigate whether it is truly worthwhile to follow this recommendation, using all the information from the design of your recommendation device. The robot player will do this by answering two questions: (a) given the device and the recommendation received, what is the percentage chance that the other robot player received each of their possible recommendations?; and (b) given that the other robot player follows their recommendation, do I (the robot) earn a higher payoff by following the recommendation given to me or by choosing the other color action? You will design this recommendation device for **five** games that will differ in their payoffs to the two players. You will design a separate device for each game. This will happen sequentially. This means that you will first design the device for one game, then, once completed, for the next game, and so on. For each game, there will be **five** rounds during which you can update your recommendation device to try to increase the amount of money that you take home with you today.

After reading these instructions, you will be asked several control questions. Next, you will play three practice rounds. Right after the practice rounds, the main part of the experiment (5 games with 5 rounds each) will start. After the main part, you will be asked to answer several quiz questions and fill out a short survey. Following completion of the survey, your participation in the study is over and you will be paid your earnings.

Designing the recommendation device

For each game, your task is to design a device that makes separate recommendations to Player 1 and to Player 2 as to which color action to choose, either Red or Blue. With your device, the game will be played multiple periods. In each period, two intelligent robot players will participate in the game, the row robot "Player 1" and the column robot "Player 2".

Your recommendation device will make use of a container filled with 24 balls. Each ball can be split into two halves: one half is labeled **1** and the other half is labeled **2**. Your task is **to select a color for each half of each ball**, Red or Blue. In other words, you will fill the container with 24 balls of four types, RR, RB, BR, and BB, where the first letter corresponds to the half labeled **1** and the second letter to the half labeled **2**.

The picture below shows an example of the screen where you design a device. You first choose the number of balls (out of 24) for which you want to color the half labeled **1** Red or Blue. This choice is in the lower left part of your decision screen, labeled **Recommendation to P1**. Then, of those balls for which you colored the half labeled **1** Red, you will chose the number of balls for which you want to color the half labeled **2** Red or Blue. Similarly, out of those balls for which the half labeled **1** were colored Blue, you will chose the number of balls for which you want to color the half labeled **2** Red or Blue. These two choices are in the upper right part of your decision screen called **Recommendation to P2**. You can color the balls using either a slider or by entering the exact number of Red balls in the box below the slider.



After you have made all selections, your device with the colored balls will be complete. Below the device, the resulting **Summary** table will appear showing the total number of balls of each type. You may wish to consult this summary table to better comprehend how many balls of each type your device contains.

<u>NOTE</u>: you can achieve any desired device! For this, you do the following:

- Count the number of balls in your desired device where the half labeled 1 is Red. Insert this number as the Recommendation to P1.
- 2. Count the number of **RR** balls in the desired device, and insert this number in the upper right part of the screen, under **Recommendation to P2**.
- 3. Count the number of BR balls in the desired device, and insert this number in the middle right part of the screen, under **Recommendation to P2**.

In the previous example, the device has 15 balls where the left part is Red: you can count them inside the container, and you also can find it as a sum of 10 RR and 5 RB balls in the second row of the Summary table. This number 15 is entered in the lower left part of the screen as the **Recommendation to P1**. There are 10 RR balls, and 10 is entered in the upper right as the **Recommendation to P2**. Finally, there are 2 BR balls, and 2 is entered in the middle right part as the **Recommendation to P2**.

You can adjust your device as many times as you want, until you click on the **Next** button, at which point your device is finalized for the round.

After you design the device, it will be used to provide recommendations for many periods of the game. Each period, the game will be played by a new pair of robot players. One ball will be drawn **at random** and will be split in two halves. The half labeled **1** will be given to Player 1 and the half labeled **2** will be given to Player 2. The color of the piece of the ball given to each player represents the recommended color action for each player.

The robot players **know the design of your device**, that is, the proportion of colored balls of each type in the container. They use this knowledge in deciding whether or not to follow the recommendation they received. If they can do better by not following the recommendation of your device, then they will not follow the recommendation of your device. Otherwise, they will follow it.

Examples of Recommendation Devices and Robots' Reasoning

We discuss three examples for the game we introduced earlier:



Example 1. Suppose you color all balls **RR** (i.e., *both halves* of all 24 balls are **Red**), so that your resulting device looks as follows.





This device always recommends that Player 1 chooses Red and that Player 2 also chooses Red. Even though Player 1 does not see the recommendation given to Player 2, by the design of the device, Player 1 knows that Player 2 is recommended to choose Red. As such, Player 1 compares following the recommendation and receiving a payoff of 0 points (Player 1's payoff when both players play Red) to not following the recommendation, choosing Blue instead and receiving a payoff of 2 points (Player 1 chooses Blue and Player 2 chooses Red). Since a payoff of 0 from playing Red is smaller than 2 from playing Blue, Player 1 does not follow the recommendation of Red. (Player 2 does not follow the recommendation as well. Player 2 infers that Player 1 was recommended Red, and prefers to play Blue.)

In Example 1, Player 1 knew exactly what is recommended to Player 2 and Player 2 knew exactly what is recommended to Player 1. More generally, this will not be the case.

Example 2. Suppose that you color 12 balls as **RB** and 12 balls as **BB** so that your device is as shown below:



Essentially, Player 2 is always recommended Blue, and Player 1 is recommended Red or Blue with a chance of 50% each. While Player 1 knows that Player 2 is always recommended Blue, Player 2 does not know the recommendation given to Player 1. Player 2 only knows that Player 1 gets a recommendation of Red with a chance of 50% and gets a recommendation of Blue with a chance of 50%.

With this knowledge, Player 2 <u>does not follow</u> the recommendation of Blue. This is because, if Player 2 chooses Blue, then Player 2 receives 3 points when Player 1 chooses Red (which occurs 50% of the time) and 1 point when Player 1 chooses Blue (which occurs 50% of the time). This gives on average $0.5 \times 3 + 0.5 \times 1 = 2$ points. If Player 2 instead chooses Red, Player 2 receives 1 point when Player 1 chooses Red and 5 points when Player 1 chooses Blue, which is on average $0.5 \times 1 + 0.5 \times 5 = 3$ points. Hence, the average payoff from not following the recommendation and choosing Red for Player 2 is higher than following the recommendation of Blue.

This illustrates the calculations that the self-interested, intelligent robots make in deciding whether or not to follow the recommendations of your device. Note that the ultimate reason Player 2 does not follow the recommendation of Blue in this example, is because Player 1 is recommended to choose Blue too often (thereby incentivizing Player 2 to play Red and grab the high payoff of 5). Just because Player 2 does not follow the Blue recommendation does not necessarily mean Player 2 should be

recommended to play Red!

Example 3. Consider the following device, where 18 balls are colored as RB and 6 balls are colored as BB.



Compared with Example 2, Player 2 is again always recommended Blue, but Player 1 is now more often recommended to play Red than Blue.

Consider the game from the perspective of Player 2. If Player 2 follows the recommendation to play Blue, and Player 1 follows its recommendations, then Player 2 will receive a payoff of 3 in 75% of all periods (that is when Player 1 is recommended to play Red since 18/24 = .75 balls have the half labeled 1 colored Red) and a payoff of 1 in 25% of all periods (that is when Player 1 is recommended to play Blue since 6/24=.25 balls have the half labeled 1 colored Blue). Player 2 thus expects to get 2.5 on average by following the recommendation (as $0.75 \times 3 + 0.25 \times 1 = 2.5$). Not following the recommendation and playing Red, will bring a payoff of 1 in 75% of all periods and a payoff of 5 in 25% all periods i.e., 2 on average (as $0.75 \times 1 + 0.25 \times 5 = 2$). It is now better for Player 2 to follow the recommendation to play Blue.

We can also consider the game from the perspective of Player 1 who knows that Player 2 was recommended Blue. Player 1 is indifferent between playing Red and Blue, as the payoff will be 3 in either case (that is, when Player 2 plays Blue and Player 1 plays either Red or Blue). Thus, Player 1 will follow the recommendation Red, if received, and the recommendation Blue, if received. In this example, <u>all recommendations</u> from this device will always be followed.

Your points for a round

For each game, you will participate in 5 rounds. At the beginning of the first round, you will have to design the device. At the end of each round, you will receive feedback on how well your device performed. At the start of each subsequent round, you will have the opportunity to re-design the device.

Each round will give you a certain number of points, which convert into money earnings at a fixed rate. We now describe how this amount of points is determined.

As stated, each round consists of robots playing the game for a **very large** number of periods. In each period, a ball is drawn at random, which represents the recommendation made to each robot player. The robot players decide whether or not to follow their recommendations and play the game once. Then the ball is placed back into the container, and a new period begins. A ball is drawn again for two new robot players, and so on. The number of periods is so large that any particular ball in the container will eventually be drawn at some point.

If in *at least in one period*, one of the robot players did not follow your recommendation, then you will not receive any points in the round. That is, for you to receive points, the robot players must *always* (that is, in all periods) follow the recommendations that your device makes to them. In this case, your **point earnings** are equal to **the smallest of the two average payoffs**, the average payoff earned by player 1 and the average payoff earned by player 2.

We return to our game and provide some examples of payoff calculations.



Example 1. If you color all balls **RR**, then neither Player 1 nor Player 2 would follow the recommendations from this device. You earn **no points**.

Example 2. If you color 12 balls as **RB** and 12 balls as **BB**, then Player 2 would not follow the recommendations, as previously discussed. You earn **no points**.

Example 3. If you color 18 balls as **RB** and 6 balls as **BB**, then players would follow your recommendations. Player 1 gets 3 points always with this device, and player 2 gets 2.5 points on average, as we calculated. The minimum of 3 and 2.5 is 2.5. Thus, you earn 2.5 points.

Example 4. If you color all balls **BR**, you get this device.



Player 1 knows that Player 2 is recommended Red and follows the recommendation Blue (as the payoff of 2 is not less than the payoff of 0 from not following this recommendation). Player 2 knows that Player 1 is recommended Blue and follows the recommendation Red (as the payoff of 5 is not less than the payoff of 1 from not following this recommendation). Hence, you earn additional points with this device. The minimum average payoff of 2 and 5 is 2. Thus, **you earn 2 points**.

Feedback and updating the device

Between rounds, you will have a chance to update the design of your device (that is, re-color the balls in the container of 24 balls).

After each round, you will receive *feedback* consisting of two parts. Above all, you will learn whether the recommendations produced by your device were always followed by the robot players or not.

Then, in the case the robot players always followed the recommendations of your device, you will also learn whether it is possible to improve your device to receive more points or, instead, the maximum possible payoff has been achieved. In the case where the robot players did not always follow the recommendations, you will see an example of a one-period realization (i.e., one ball drawn from your device) when the recommendations were not followed and why one player chose not to follow the recommendation.

Your total payoff for the experiment

At the end of the experiment, the computer will choose one of the five games at random and you will receive the number of points you earned in **the round for that game with the highest payoff**. This amount will be converted into dollars at the rate of 1 point equal to \$6. You will be paid only if you complete all rounds of the experiment, the quiz, and the short survey.

This means that you will definitely earn the show-up fee of \$10 in this experiment. If, in addition, your maximum round earnings in the randomly selected game is, for instance, 3 points, then you will receive a payment of $3 \times \$6 = \18 (in addition to your guaranteed earnings of \$10).

To summarize, the payoff structure of this experiment implies:

- you will definitely earn \$10 in this experiment
- you want to do your best in every game to earn an additional amount of points/-

money.

- you want to design a recommendation device where both players always follow the recommendations
- you can improve your payoff by improving your recommendation device in a given round (if the device is not yet giving you the highest possible payoff)
- you are encouraged to experiment between rounds by changing the color of the balls in order to try to improve your payoff.

B Experiment Questions

Pre-Experiment Quiz

Question 1: Suppose that the game that the robot players participate in has the following specific payoffs:



If Player 1 chooses Blue and Player 2 chooses Blue, then:

- (a) The payoff to Player 1 is 3 and the payoff to Player 2 is 3.
- (b) The payoff to Player 1 is 5 and the payoff to Player 2 is 2.
- (c) The payoff to Player 1 is 3 and the payoff to Player 2 is 1. [Correct]
- (d) The payoff to Player 1 is 1 and the payoff to Player 2 is 3.

<u>Feedback:</u> "The payoffs of two players are in the cell corresponding to their actions (Red or Blue). The payoff of player 1 is in the lower left part of the cell and the payoff of player 2 is in the upper right part of the cell."

Question 2: Consider the same game



Suppose that Player 2 thinks that Player 1 is going to choose Blue. Then,

- (a) It is in Player 2's best interest to choose Blue.
- (b) It is in Player 2's best interest to choose Red. [Correct]
- (c) It doesn't matter whether Player 2 chooses Blue or Red.
- (d) None of the above are true.

<u>Feedback:</u> "Player 2 thinks that player 1 plays Blue and thus look at the bottom row of the table. Then, player 2 compares his/her payoff from playing Red (that is 5) with his/her payoff from playing Blue (that is 1). Player 2 will choose the action leading to the maximum of these two payoffs."

Question 3: Consider the same game



Suppose that Player 2 gets a recommendation to play Red. Then,

(a) Player 2 will follow this recommendation blindly.

- (b) Player 2 will not follow this recommendation and will play Blue in any case.
- (c) Player 2 will ask Player 1 what was recommended to Player 1 and will choose an action depending on the response.
- (d) None of the above are true. [Correct]

<u>Feedback</u>: "Players cannot communicate. Also they do not follow their recommendations blindly but also do not reject all recommendations. Instead, players investigate whether it is truly worthwhile for them to follow the recommendation using all the information from the design of the recommendation device. They need information about the recommendation device."

Question 4: Your task in this experiment:

- (a) To play four games, several times each, against humans.
- (b) To play four games, several times each, against robots.
- (c) To observe and describe how others play the games.
- (d) To design recommending devices for given games that will be played by robots.[Correct]

<u>Feedback:</u> "As instructions explain, your task is **to design recommendations** to two intelligent robot players on how they should play a game with one another. You will go through a sequence of games and have several chances to improve the device."

Question 5: The following picture displays a recommendation device.



Suppose that one ball is drawn out of this container at random. What is the chance that this ball recommends to player 1 to choose Blue and, at the same time, it recommends to player 2 to choose Red.

- (a) 2 out of 24.
- (b) 5 out of 24.
- (c) 7 out of 24. [Correct]
- (d) 10 out of 24.

<u>Feedback:</u> "To answer this question, you need to count the number of balls with the left side (1) being Blue and the right side (2) being Red."

Question 6: Suppose that the game that the robot players participate in has the following specific payoffs:

P2	Red		Blue	
P1				
	\smallsetminus	1	$\overline{\ }$	3
Red	0	$\overline{\ }$	3	\searrow
	\smallsetminus	5	$\overline{\ }$	1
Blue	2	$\overline{\ }$	3	\searrow

Moreover, suppose that you design the following device:



Which of the following statements is true:

- (a) Neither player 1 nor player 2 are willing to follow the recommendation made by the device and you will earn no points in this round.
- (b) Player 1 and player 2 are both willing to follow the recommendation made by the device and you will earn 3.5 points in this round.
- (c) Player 1 and player 2 are both willing to follow the recommendation made by the device and you will earn 2 points in this round.[Correct]
- (d) Player 1 and player 2 are both willing to follow the recommendation made by the device and you will earn 5 points in this round.

<u>Feedback</u>: "This device always recommends player 1 to play Blue and player 2 to play Red. With this device, each player will have a certainty about the recommendation to another player. Checking the table of payoffs, we can see that player 1 prefers to follow the recommendation (when player 2 plays Red, it is beneficial to player 1 to play Blue). Similarly, player 2 prefers to follow the recommendation.

According to the instructions, for this round, you would get points based on the smallest of the two average payoffs of players 1 and 2. In this case, player 1 will always (for any of the balls drawn) get 2 and player 2 will always get 5. The smallest of 2 and 5 is 2."

Question 7: Suppose that in one of four games you are presented in the experiment, in the first three rounds player 2 never followed the recommendation of the devices you made.

Which of the following statements is true:

- (a) You have a chance to get additional points for this game, only in the case when you design a new device in the fourth round. [Correct]
- (b) You have no chance to get additional points for this experiment.
- (c) You still can earn additional points in the experiment, but not for this game.
- (d) If player 2 will follow the recommendation of the device you redesigned for the 4th round in all periods of that round, you will get additional points in that round.

<u>Feedback</u>: "As the instructions explain, for a given game, you receive the number of points you earned in the round for that game with the highest amount of points. This means that if in the last round you will design such device that the recommendations are always followed (by both players 1s and player 2s), you will get points from that round and for this game."

Question 8: Once you design your recommendation device in a round for a given game, the game is played by robot players for multiple periods. Can a robot player choose its action depending on how the game was played in previous periods?

- (a) Yes, a robot player observes the actions played in the previous periods and can react to them.
- (b) Yes, if it allows a robot player to maximize a total payoff over all periods.
- (c) No, in each period, two new robots play the game and they do not have information about the previous periods. [Correct]
- (d) None of the above is true.

<u>Feedback</u>: "Each game will be played for multiple periods but the robot players in each round will always be different (so that their actions will not depend on what was played in the past). In each period, when a particular color action is recommended to them, they will investigate whether it is truly worthwhile for them to follow this recommendation using all the information from the design of your recommendation device."

Question 9: Please follow the instructions shown on the screen below. [By moving the sliders, subjects had to implement the distribution $\mu_{RR} = 11/24$, $\mu_{RB} = 7/24$, $\mu_{BR} = 2/24$ and $\mu_{BR}^* = 4/24$.]



Post-Experiment Cognitive Reflection Test (CRT)

- Question 1: Together, a mobile phone and laptop cost \$1100. The laptop costs 1000 more dollars than the mobile phone. How much does the mobile phone cost in dollars? [Correct: 50, Intuitive: 100]
- Question 2: It takes 5 bakers 5 hours to bake 5 cakes. How many hours would it take 100 bakers to bake 100 cakes? [*Correct: 5, Intuitive: 100*]
- Question 3: In a field there is a group of rabbits. The population of the rabbits doubles in size every day. If it would take 50 days for the rabbit population to

completely cover the field, how many days would it take for the rabbits to cover half of the field? [Correct: 49, Intuitive: 25]

- Question 4: If Charlie drinks one gallon of milk in 3 days, and Emerson drinks one gallon of milk in 6 days, how many days would it take them to drink one gallon of milk together? [*Correct: 2, Intuitive: 4.5*]
- Question 5: Simon decided to invest \$2,000 in the stock market one day early in 2008. Six months after he invested, on July 17, the stocks he had purchased were down 50%. Fortunately for Simon, from July 17 to October 17, the stocks he had purchased went up 75%. As of October 17, Simon has: [Answers: Lost money in the stock market; Broken even in the stock market; Made money in the stock market] [Correct: Lost money in the stock market, Intuitive: Made money in the stock market]

Post-Experiment 11-20 Game Question

Two highly intelligent and competitive individuals are playing the following game. Each player must simultaneously announce an integer number between 11 and 20. Then, the payoffs are determined as follows. Each player receives an amount of dollars equal to the number the player announced. Moreover, if one of the players announces a number that is exactly one less than their opponent, then this player receives an additional reward of \$20. Finally, if both players announce the same number, each of them receives an extra \$10.

For example, if player 1 says 17 and player 2 says 19, then player 1 receives 17 and player 2 receives 19. Or, if player 1 says 12 and player 2 says 13, then player 1 receives 32 and player 2 receives 13.

What do you think would be the number that one of these players would announce? Or, in other words, if you were one of these two players which number would you announce?

C Feedback Screens



Figs. C1 to C3 display feedback provided to subjects.

Figure C1: Feedback for a device that implements a sub-optimal correlated equilibrium.







Figure C3: Feedback for a device that is not a correlated equilibrium.

D Correlated Equilibria of Tested Games

All 5 games we tested, are presented in the left part of Table 1 in the paper. In this appendix, for each of these games, we derive the set of correlated equilibria, that is distributions over action profiles in $\{Red, Blue\} \times \{Red, Blue\}$ that satisfy Section 3. We also solve problem (1) and find the desirable correlated equilibrium μ^* . Finally, we show the recommendation device that implements this DCE, i.e., the device that gives subjects the maximal payoff in each game.

D.1 Prisoner's Dilemma

In Prisoner's Dilemma, each player has a strictly dominant strategy to choose the action Red. As a consequence, in the unique correlated equilibrium (which is, indeed, the unique Nash equilibrium) $\mu_{RR} = 1$ and $\mu_{RB} = \mu_{BR} = \mu_{BB} = 0$. Consequently, this is also the DCE. Thus, the maximum achievable payoff that can be earned in this game is 4. Fig. D1 shows the device that earns this maximal payoff.



Figure D1: Optimal Device in Prisoner's Dilemma

D.2 Matching Pennies

Using Section 3, distribution μ is a correlated equilibrium in the Matching Pennies game if it satisfies the following inequalities:

$$5\mu_{RR} + 3\mu_{RB} \ge 3\mu_{RR} + 5\mu_{RB} , \qquad 3\mu_{BR} + 5\mu_{BB} \ge 5\mu_{BR} + 3\mu_{BB} ,$$
$$3\mu_{RR} + 5\mu_{BR} \ge 5\mu_{RR} + 3\mu_{BR} , \qquad 5\mu_{RB} + 3\mu_{BB} \ge 3\mu_{RB} + 5\mu_{BB} ,$$

which can be summarized as

$$\mu_{RR} \ge \mu_{RB} \ge \mu_{BB} \ge \mu_{BR} \ge \mu_{RR}.$$

We conclude that there is a unique correlated equilibrium of this game with $\mu(a) = 1/4$ for all $a \in \{R, B\} \times \{R, B\}$. This is also precisely the unique Nash equilibrium in mixed strategies, where each player independently randomizes over their available actions. Consequently, this is also the DCE and the maximum achievable payoff for this game is $(1/2) \times 5 + (1/2) \times 3 = 4$. Fig. D2 shows the device that would need to be designed in the experiment to earn this maximal payoff.



Figure D2: Optimal Device in Matching Pennies

D.3 Battle of the Sexes

Using Section 3, distribution μ is a correlated equilibrium if it satisfies the following four inequalities:

$$6\mu_{\mathbf{RB}} \ge 2\mu_{\mathbf{RR}}, \quad 2\mu_{\mathbf{BR}} \ge 6\mu_{\mathbf{BB}}, \quad 6\mu_{\mathbf{BR}} \ge 2\mu_{\mathbf{RR}}, \quad 2\mu_{\mathbf{RB}} \ge 6\mu_{\mathbf{BB}},$$

which is equivalent to the condition that

$$\min\{\mu_{RB}, \mu_{BR}\} \ge \max\left\{3\mu_{BB}, \frac{1}{3}\mu_{RR}\right\}.$$

The DCE solves problem (1). Clearly, at the optimum, $\mu_{BB} = \mu_{RR} = 0$, as the strategies (B, B) and (R, R) result in 0 payoff for both robot players. Finally, to maximize the minimum expected payoff of both robot players, we should set $\mu_{RB} = \mu_{BR} = 1/2$. This achieves a maximum payoff of $(1/2) \times 6 + (1/2) \times 2 = 4$ in this game. Figure D3 displays the device that would need to be designed in the experiment to earn this maximal payoff.



Figure D3: Optimal Device in Battle of the Sexes

D.4 Chicken Version 1

Using Section 3, distribution μ is a correlated equilibrium if and only if it satisfies the following four inequalities:

$6\mu_{RB} \ge \mu_{RR} + 5\mu_{RB} ,$	$\mu_{BR} + 5\mu_{BB} \ge 6\mu_{BB} ,$
$6\mu_{BR} \ge \mu_{RR} + 5\mu_{BR} ,$	$\mu_{RB} + 5\mu_{BB} \ge 6\mu_{BB},$

which is equivalent to the condition

 $\min\{\mu_{RB}, \mu_{BR}\} \ge \max\{\mu_{RR}, \mu_{BB}\}.$

Given the payoff matrix, the objective function in (1) is maximized, when as much weight as possible is put on μ_{BB} , as little weight as possible is put on μ_{RR} , and there is a symmetry in the remaining two weights. Thus, the DCE in this case is

$$\mu_{RB} = \mu_{BR} = \mu_{BB} = 1/3,$$

which earns a maximal payoff of $(1/3) \times 1 + (1/3) \times 6 + (1/3) \times 5 = 4$. Fig. D4 displays the recommendation device that yields this maximal payoff.



Figure D4: Optimal Device in Version 1 of Chicken

D.5 Chicken Version 2

Using Section 3, distribution μ constitutes a correlated equilibrium in this case if and only if the following four constraints are satisfied:

$$6\mu_{RB} \ge 2\mu_{RR} + 5\mu_{RB} , \qquad \qquad 2\mu_{BR} + 5\mu_{BB} \ge 6\mu_{BB} ,$$

$$6\mu_{BR} \ge 2\mu_{RR} + 5\mu_{BR} , \qquad \qquad 2\mu_{RB} + 5\mu_{BB} \ge 6\mu_{BB} ,$$

which is equivalent to the condition

$$\min\{\mu_{RB}, \mu_{BR}\} \ge \left\{2\mu_{RR}, \frac{1}{2}\mu_{BB}\right\}.$$

As in the first Chicken game, the objective function in (1) is maximized, when as much weight as possible is put on μ_{BB} , as little weight as possible is put on μ_{RR} , and there is a symmetry in the remaining two weights. The solution has

$$\mu_{RB} = \mu_{BR} = \frac{1}{2}\mu_{BB},$$

so that $\mu_{RB} = \mu_{BR} = 1/4$ and $\mu_{BB} = 1/2$. This yields a maximal payoff equal to $(1/4) \times 2 + (1/4) \times 6 + (1/2) \times 5 = 4.5$. Fig. D5 displays the device that yields this maximal payoff.



Figure D5: Optimal Device in Version 2 of Chicken

E Additional Data Analysis

E.1 Feedback

When subjects failed to construct a device consistent with any CE in a given round, they were shown a ball from their device where one robot player did not follow the recommendation, along with an explanation for why this occurred (see Fig. C3). The ball was selected randomly based on the device's probability distribution. Here, we discuss how subjects responded to this feedback.

Specifically, we are interested in how subjects adjusted their devices in response to this feedback. A simple (though not necessarily optimal) heuristic is to reduce the probability mass on the action profile corresponding to the shown ball. Table E1 reports the proportion of times subjects adjusted (either decreased, made no change, or increased) the probability mass on the shown ball relative to the total number of feedback instances for a given game. The largest possible number of feedback instances for each game is 4×100 , as subjects could only respond to feedback during the first four rounds of each game. However, if a CE was found in a round, no feedback was provided, resulting in fewer actual feedback instances, N. (We also excluded any

Game	N	Decrease	No Change	Increase
PD	302	0.712	0.086	0.202
MP	170	0.753	0.118	0.129
BoS	51	0.686	0.098	0.216
C1	170	0.753	0.100	0.147
C2	116	0.810	0.060	0.129

Table E1: The proportion of times subjects adjusted the probability mass on the ball shown during feedback, based on N feedback instances provided during the first four rounds of each game, before the DCE was reached.

reactions to feedback after subjects had found the DCE.) We observe that across all games, subjects decreased the probability mass on the shown ball in more than 50% of cases (significant at the 1% level). This suggests that subjects mostly reacted in a systematic manner.

E.2 Additional Performance Measures

In the main text, we discuss performance using three performance measures: Top-Score, Score, and N-RMSE. Here, we consider two additional performance measures: CE and Selten.

CE is an indicator variable equal to 1 if a subject designed a CE and 0 otherwise. This measure shows the fraction of subjects who succeeded in finding at least one CE. Thus, it is equal to TopScore for games with a unique CE, and is different from Score in all other games, as it does not consider a payoff that subjects get in different CEs.

Selten is defined in Selten and Krischker (1982) as $\mathbf{CE} - V$, where V is the volume of the CE set relative to the space of all possible strategies. This volume is reported in Table 1. This measure essentially penalizes the CE measure when finding a CE randomly is easier.

Fig. E1 displays the three performance measures discussed in the paper, along with these two additional measures. While the CE measure indicates that almost all



Figure E1: The average (across all subjects) of the highest value in five rounds for each of all performance measures in each game. Error bars represent 95% confidence intervals.

subjects found at least one CE in the Battle of the Sexes and both Chicken games, the Selten measure, after correcting for the volume, reduces the CE and aligns more closely with the Score measure used in the main text.

E.3 Performance Dynamics

We describe how the aggregate performance measures evolved over the five rounds within each of the five games.³²

Fig. E2 displays the dynamics across rounds for TopScore and Score measures. We observe that average performance improves across rounds. In fact, performance under both measures is significantly better in round 5 (R5) compared to round 1 (R1) for all games. However, the rate of improvement diminishes over the five rounds: the difference in either TopScore or Score between round 5 (R5) and round 3 (R3) is not

 $^{^{32}}$ Recall from footnote 21 that the raw data for these statistics had to be adjusted. Once a subject achieved the DCE in a game, the devices submitted in all subsequent rounds were replaced with the DCE device. We refer to the new data as an 'amended' data set.



(b) Dynamics of Score.

Figure E2: Evolution of the average of points-based performance measures, TopScore and Score, across rounds for each game. Error bars represent 95% confidence intervals.



Figure E3: Evolution of the average of distance-based performance measure, N-RMSE, across rounds for each game. Error bars represent 95% confidence intervals.

significant for any game.

Given our amended dataset, we interpret TopScore in a given round as the fraction of subjects who found the DCE in that round or earlier. According to Fig. E2a, most subjects found the DCE quickly in both Battle of the Sexes and Matching Pennies: on their first attempt in Battle of the Sexes and by their second attempt in Matching Pennies. Furthermore, the rate at which subjects discover the DCE for the first time decreases in subsequent rounds, suggesting that these games were relatively easy to solve. In contrast, for Prisoner's Dilemma and the two Chicken games, the rate of finding the DCE remains relatively constant across rounds. This indicates that solving these games requires more experimentation and deliberation, providing evidence that these games are more complex to solve.

We perform the same analysis using our distance-based performance measure, N-RMSE, as shown in Fig. E3. We observe a pattern similar to the points-based measures in Fig. E2: performance improves significantly from round 1, but this improvement tapers off relatively quickly. The difference between N-RMSE in rounds 2 and 5 is insignificant for any game.

Overall, we conclude that the experimental task allowed subjects to learn and improve both through experience and by utilizing the feedback provided between rounds. However, there are diminishing returns to performance improvement, suggesting that full convergence to the DCE may require a larger number of rounds, particularly for games where only a minority of subjects found the DCE by Round 5 (Prisoner's Dilemma and both versions of Chicken).

E.4 Distribution of Recommendation Devices

Figs. E4 and E5 illustrate the evolution of the distributions of the recommendation devices for the Matching Pennies and the Battles of the Sexes games both in the space of expected payoffs (the top panels) and in the distribution space (the bottom panels). These figures are analogous to Figs. 5 to 7 for three other games, discussed in the main text.

Matching Pennies. Recall that overall performance in Matching Pennies was high according to our performance measures (Fig. 3). This suggests that individual subjects were successful in their design of recommendation devices to meet the DCE objective in this game. Fig. E4d illustrates this point: the DCE (which coincides with the unique Nash equilibrium in mixed strategies) is already the modal device by round 1 and is the clear modal device by round 5. Moreover, recommendation devices in round 5 that do not constitute correlated equilibria appear to be distributed somewhat randomly around the unique Nash equilibrium, with all of these implying expected payoffs close to those implied by the DCE (Fig. E4b). Thus, subjects appear to recognize the zero-sum nature of Matching Pennies and no additional recommendation devices emerge as focal by the final round.



(c) Distribution Space: Round 1

(d) Distribution Space: Round 5

Figure E4: Bubble plots displaying, in both the payoff-space and distribution-space, the distribution of designed devices for **Matching Pennies** in round 1 vs. round 5.



(c) Distribution Space: Round 1

(d) Distribution Space: Round 5



Battle of the Sexes. Battle of the Sexes was another game in which subjects performed well in terms of finding the DCE (Fig. 3). From Figs. E5a and E5c, we see that the DCE is the modal response already in round 1, although some subjects implement one of the two asymmetric Nash equilibria with their device. Moreover, already in round 1, there are very few recommendation devices that yield payoffs outside of the set of correlated-equilibrium payoffs. Experience serves to consolidate this finding: by round 5 designed distributions are almost entirely in the set of correlated-

related equilibria (Fig. E5d) or the set of correlated equilibrium payoffs (Fig. E5b).³³ The vast majority of these distributions fall on the DCE, as expected from the high aggregate performance in this game.

E.5 Extended Correlate Analysis

Table E2 reports the average values and proportions (for categorical variables) of the considered individual correlates including performance in the CRT test (CTRNcorr and CTRNintuit) and the 11-20 game (StratSoph). To assess potential collinearity among the correlates, we also report their pairwise correlations.

	Ave/prop	CRTNcorr	CRTNintuit	StratSoph	GenderF	StudyYear
CRTNcorr	2.44		-0.829	0.378	-0.439	0.347
CRTNintuit	1.68	-0.829		-0.288	0.316	-0.344
StratSoph	0.60	0.378	-0.288		-0.195	0.198
GenderF	0.54	-0.439	0.316	-0.195		-0.188
StudyYear	2.96	0.347	-0.344	0.198	-0.188	
SocialSc	0.46	-0.240	0.140	-0.042	0.127	-0.315
Arts	0.03	0.053	-0.078	0.015	0.045	0.006
BiologicalSc	0.06	-0.128	0.107	0.022	0.064	0.118
Business	0.06	0.212	-0.221	-0.033	-0.189	0.191
Design	0.01	-0.132	0.151	-0.079	0.093	0.004
Economics	0.12	-0.038	0.117	-0.089	0.094	0.199
Engineering	0.11	0.192	-0.197	0.101	-0.060	0.206
HealthMed	0.02	-0.034	0.030	-0.019	0.132	0.005
Inform&CS	0.10	0.118	0.004	0.044	-0.294	-0.161
PhysicalSc	0.03	0.147	-0.154	0.092	0.045	0.006

Table E2: Average values or proportions (for categorical variables) and pairwise correlations among the potential explanatory variables, with color coding ranging from dark blue (strong negative) to intense red (strong positive). There are 100 individual observations.

The number of correct answers on the CRT test (CTRNcorr) is strongly negatively correlated with the number of intuitive answers (CTRNIntuit). Consequently, only CRTNcorr is included in the regression analysis.

 $^{^{33}{\}rm The}$ percentage of subjects who designed the device implementing any correlated equilibrium changed from 78% in round 1 to 97% in round 5.
Measure	Score=Te	op Score	Score			Top Score		
	PD	MP	BoS	C1	C2	BoS	C1	C2
Intercept	0.1326	0.4631**	0.8958***	0.6969***	0.6798***	0.7787***	-0.2664	0.2179
	(0.1703)	(0.1785)	(0.1010)	(0.1228)	(0.0990)	(0.1915)	(0.1685)	(0.1816)
GameOrder	-0.0627^{*}	0.0346	-0.0244	0.0235	0.0075	-0.0564^{*}	0.0725^{**}	0.0162
	(0.0317)	(0.0310)	(0.0165)	(0.0211)	(0.0184)	(0.0314)	(0.0290)	(0.0337)
CRTNcorr	-0.0159	0.0618^{**}	0.0349**	-0.0025	0.0311^{*}	0.0581**	0.0101	0.0255
	(0.0299)	(0.0309)	(0.0146)	(0.0187)	(0.0162)	(0.0278)	(0.0257)	(0.0297)
StratSoph	0.0809	0.0746	0.0144	0.0728^{*}	0.0528	0.0666	0.0742	-0.0293
	(0.0623)	(0.0641)	(0.0301)	(0.0388)	(0.0336)	(0.0572)	(0.0532)	(0.0616)
GenderF	-0.0392	0.0834	-0.0221	-0.0805	0.0368	-0.0426	0.0856	-0.2222^{**}
	(0.1029)	(0.1055)	(0.0509)	(0.0644)	(0.0573)	(0.0965)	(0.0883)	(0.1052)
StudyYear	0.1184^{***}	-0.0247	0.0096	0.0020	0.0033	0.0035	0.0151	0.0213
	(0.0446)	(0.0457)	(0.0216)	(0.0277)	(0.0241)	(0.0410)	(0.0380)	(0.0443)
Arts	-0.2799	-0.0939	-0.1339	0.0043	0.0188	-0.1814	-0.1815	0.0834
	(0.2573)	(0.2652)	(0.1250)	(0.1607)	(0.1396)	(0.2370)	(0.2204)	(0.2561)
BiologicalSc	-0.0296	-0.3279	-0.2123^{**}	-0.0226	-0.0448	-0.0020	0.0410	-0.0368
	(0.1917)	(0.1974)	(0.0944)	(0.1199)	(0.1036)	(0.1791)	(0.1645)	(0.1900)
Business	-0.1185	0.0799	-0.0783	0.0175	-0.0346	-0.0954	0.5297^{***}	0.4430**
	(0.2053)	(0.2068)	(0.0975)	(0.1245)	(0.1078)	(0.1849)	(0.1709)	(0.1976)
Design	-0.1354	0.4238	0.1463	0.1586	0.0994	0.3661	-0.1546	-0.0759
	(0.4439)	(0.4493)	(0.2128)	(0.2736)	(0.2395)	(0.4036)	(0.3754)	(0.4393)
Economics	0.2895^{*}	0.0141	0.0359	0.1044	0.0344	0.1502	0.1401	-0.1565
	(0.1472)	(0.1520)	(0.0711)	(0.0920)	(0.0793)	(0.1349)	(0.1262)	(0.1455)
Engineering	0.3360^{**}	0.0349	-0.0230	0.0810	0.0526	0.0559	0.1256	0.2326
	(0.1525)	(0.1570)	(0.0739)	(0.0952)	(0.0833)	(0.1402)	(0.1307)	(0.1529)
HealthMed	-0.2379	-0.2026	-0.3818^{**}	-0.4666^{**}	0.0482	-0.1707	-0.2118	0.3798
	(0.3135)	(0.3225)	(0.1529)	(0.1956)	(0.1715)	(0.2899)	(0.2685)	(0.3145)
Inform&CS	0.0491	0.0437	0.0005	-0.0997	0.0338	0.0876	0.0064	0.1510
	(0.1587)	(0.1630)	(0.0770)	(0.0992)	(0.0858)	(0.1460)	(0.1361)	(0.1573)
PhysicalSc	0.7509^{***}	0.1067	-0.1574	-0.1185	-0.1755	-0.2242	-0.0472	-0.2602
	(0.2633)	(0.2690)	(0.1269)	(0.1686)	(0.1416)	(0.2407)	(0.2313)	(0.2597)
\mathbb{R}^2	0.3326	0.1689	0.2865	0.1869	0.1402	0.2075	0.2272	0.2730
* 10%, ** 5%, *** 1% levels of significance (two-sided t -test)								

Table E3 reports the complete results of the regression analysis for the Score and Top Score measures.

Table E3: Regressions of the highest (over five rounds) Score and Top Score measures on the set of correlates. The dummy variable for the Social science major is the base category. Each regression uses 100 observations.