Journal of Experimental Psychology: General
2015, Vol. 144, No. 5, 967–981

0096-3445/15/$12.00 http://dx.doi.org/10.1037/xge0000096

Compete, Coordinate, and Cooperate: How to Exploit Uncertain Environments With Social Interaction

Christin Schulze and Ben R. Newell
University of New South Wales

Countless decisions, from the trivial to the crucial, are made in complex social contexts while facing uncertain consequences. Yet a large portion of decision making research focuses on either the effects of social interaction or the effects of environmental uncertainty by examining strategic games against others or individual games against nature. Drawing a connection between these approaches, the authors extend a standard individual choice paradigm to include social interaction with 1 other person. In this paradigm, 2 competing decision makers repeatedly select among 2 options, each offering a particular probability of a fixed payoff. When both players choose the same, correct option, the payoff is evenly split; when they choose different options, the player choosing the correct option receives the full payoff. The addition of this social dimension gives players an opportunity to fully exploit an uncertain environment via cooperation: By consistently choosing opposite options, two players can exploit the uncertain environment more effectively than a single player could. We present 2 experiments that manipulate environmental (Experiment 1) and social (Experiment 2) aspects of the paradigm. In Experiment 1, the outcome probabilities were either known or unknown to participants; in Experiment 2, participants' attention was drawn to individual or group gains by introducing either within- or between-group competition. Efficient cooperation did not emerge spontaneously in Experiment 1. Instead, most people probability maximized, mirroring the behavior observed in individual choice. By contrast, between-group competition in Experiment 2 facilitated efficient—but not always equitable—exploitation of uncertain environments. This work links the concepts of individual risky choice and strategic decision making under both environmental and social uncertainty.

Keywords: decisions under uncertainty, cooperation, competition, risky choice

What are the chances of a Formula 1 driver successfully overtaking an opponent who is leading the race? Setting the expertise, skill level, and personal characteristics of both drivers aside, whether or not the move succeeds depends on the reaction of the leading driver—the situation is social and competitive. But it also depends on uncertain environmental factors. Heavy rainfall during the race, for instance, may affect the track conditions and impair the pursuing driver's ability to gauge the odds of his maneuver succeeding. Countless other choice situations share these two broad characteristics—they are both socially interactive *and* environmentally uncertain. Yet most decision making research to date has focused on either the effects of social interaction in situations of conflict, coordination, and cooperation *or* the effects of uncertain or risky outcomes on people's choices—the former by means of strategic games, the latter by means of games against nature.

As an example, consider a prominent game against nature used in the probability learning literature (e.g., Estes, 1964): A single decision maker is asked to repeatedly choose between two options with differing probabilities of yielding a reward (e.g., $p = .70$ vs. $1 - p = .30$). Individual decision makers are surprisingly bad at solving this task effectively. Instead of maximizing their reward by choosing only the option that nature is more likely to reward, many people respond proportionally to the probabilities of nature's move—that is, they allocate 70% of their choices to one option and 30% to the other. This choice anomaly is known as probability matching (see Vulkan, 2000, for a review). A likely reason for this failure to adopt a probability (and reward) maximizing strategy in simple repeated choice is that decision makers are unwilling to tolerate nature outmaneuvering them on a certain 30% of choices (see, e.g., Arkes, Dawes, & Christensen, 1986; Fantino & Esfandiari, 2002). Consequently, they invest considerable effort in finding a strategy that might offer a higher, or better yet, perfect success rate—for example, by searching for a predictable pattern in the outcome sequence (Gaissmaier & Schooler, 2008; Peterson & Ulehla, 1965). Probability matching thus may emerge in a misguided effort to improve choice accuracy. To put it another way, in a game against nature, people—naturally—try to *beat* nature but inevitably fail.

In this article, we turn the tables and allow people to fully exploit nature's uncertain move by extending a classic individual choice task to include social interaction with another person. Specifically, two decision makers seek to exploit a monetary

Christin Schulze and Ben R. Newell, School of Psychology, University of New South Wales.

Christin Schulze is now at Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany.

We thank Susannah Goss for editorial assistance. This research was supported by Australian Research Council grants to Ben R. Newell (DP110100797; FT110100151).

Correspondence concerning this article should be addressed to Christin Schulze, Center for Adaptive Rationality, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany. E-mail: cschulze@mpib-berlin.mpg.de

resource that an indifferent nature repeatedly places at one of two choice alternatives with unequal odds. In addition to competing against nature, the two decision makers compete against each other. When both converge on the same, correct option, the payoff is split evenly between them; when their choices diverge, the player choosing the correct option receives the full payoff (for an illustration of the paradigm's general setup, see Figure 1). By simply selecting opposite options on each turn—that is, covering all bases—two people can exploit this uncertain environment more effectively than a single decision maker could. This setup therefore links two prominent approaches to understanding human choice: studies of individual decision makers facing an uncertain nature (see Estes, 1964) and studies of two (or more) decision makers competing in a stable environment (see Camerer, 2003).

The relationship between these two lines of research has also been explored along other avenues. For example, stochastic forms of classic strategic games examine how environmental risk affects social interaction. In social dilemmas, such as the prisoner's dilemma or public goods games, environmental risk has been found to exert a primarily detrimental influence on decision making: Outcome efficiency and the speed of learning diminish under uncertainty. Bereby-Meyer and Roth (2006), for instance, showed that noisy payoffs increase the rate of (dominated) cooperation in one-shot prisoner's dilemma games but decrease (optimal) cooperation in repeated prisoner's dilemma games (relative to equivalent games with deterministic payoffs). Similarly, probabilistic losses reduce the rate of cooperation observed under deterministic returns on investment in public goods games; that is, environmental risk increases the tendency to free ride (Berger & Hershey, 1994). Another way to explore the interplay of social encounters and uncertain environmental profitability is to examine collective human choice in experimental tasks that mimic foraging situations (e.g., Goldstone & Ashpole, 2004; Kraft & Baum, 2001; Madden, Peden, & Yamaguchi, 2002; Sokolowski, Tonneau, & Freixa i Baqué, 1999). Faced with a foraging-type environmental challenge, decision makers typically distribute themselves among resources in proportion to the resources' relative yield, consistent with the predictions of the ideal free distribution model (see Fretwell, 1972). The primary unit of analysis in such foraging analogs, however, is the group itself—its distribution, migration patterns, time frame for efficiency, and so forth—rather than the individual decisions of the members forming that group (Goldstone & Gureckis, 2009).

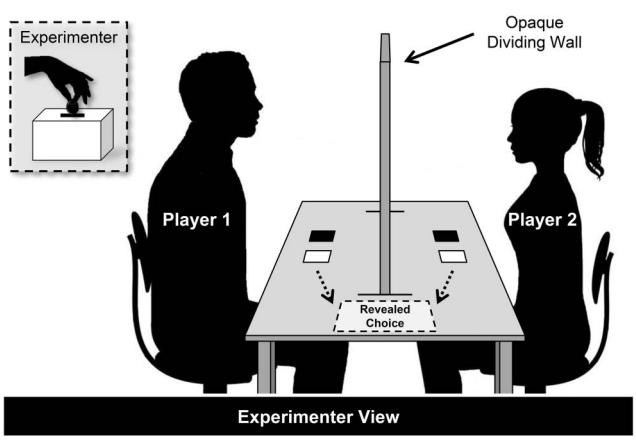

Figure 1. General task setup and experimental procedure (see the Method section of Experiment 1 for details). Two participants were seated at opposite sides of a table; the experimenter sat at the end of the table, facing both participants. An opaque dividing wall prevented communication between the players. Each participant was given a black and a white card to indicate choices. In each round, participants first made their predictions privately and then revealed them simultaneously by sliding the respective card to the edge of the table (in view of both players and the experimenter). Following each choice, the experimenter drew a token out of a box containing 70 black and 30 white tokens (counterbalanced), revealed the outcome and payoffs for that round, and returned the token to the box. This trial structure was repeated for 200 rounds.

Applications of game theoretic models to behavioral ecology—for example, the hawk–dove game—connect these two approaches and have been used to describe and predict the frequency of aggressive encounters when groups of animals compete for shared resources, as in foraging situations (Smith & Price, 1973). Moreover, the stochasticity of the foraging environment itself has been considered in foraging-related behavioral games. Avrahami, Güth, and Kareev (2005), for instance, examined people's choices in a parasite or producer–scrounger game related to the paradigm used here. In the stochastic form of this game, an indifferent nature places a resource at one of two patches with some probability, and both producer and scrounger simultaneously choose to forage at one of the patch locations. In contrast to the present paradigm, however, players are asymmetric: The producer is rewarded if her decision converges with nature's choice but not with that of the scrounger; the scrounger is rewarded if both players' choices coincide with nature's choice; and neither player is rewarded if the producer does not correctly predict nature's choice. Avrahami et al. (2005) found that "efficiency by trust"—a strategy in which the producer always predicts nature's more likely choice, but the scrounger exploits this regularity only half of the time, thus ensuring equal pay between players—did not emerge over 100 rounds of this game.

In this article, we take a different approach to linking environmental uncertainty and social interaction. Instead of asking how uncertainty affects cooperation and conflict (stochastic games), or how groups behave collectively when faced with resources of uncertain value (foraging research), we ask how social interaction can benefit individual choice under uncertainty. There are a number of key differences between our paradigm and the reviewed literature. First, in our approach, decisions have no inherent strategic value (e.g., cooperate vs. defect), but each choice represents a mere preference for one of two environmental resources, as is the case in individual choice under risk and uncertainty. Second, we are concerned with the dynamics and adequacy of individual choice, not the group-level properties that emerge from collective behavior of individuals (as evaluated in group foraging studies). In short, the central goal of this article is to determine if and how people learn to fully exploit nature in a social context. To this end, we introduce a game derived directly from an individual choice task, in which people play against both nature *and* one other person.

To align our paradigm with the individual choice case, we exclude any form of direct verbal or nonverbal communication between the decision makers in the experiments presented here. Decision makers can only signal their intended strategy use indirectly, through each choice they make. This minimal form of

communicating through decisions preserves the constraints of a typical individual game against nature, in which each player single-handedly works out an effective choice strategy. In our tasks too, decision makers need to figure out how to make profitable choices without being able to communicate openly. Do players seize this social opportunity to fully exploit an uncertain environment cooperatively despite the constraints on communication? And if so, how does such efficient behavior emerge? To approach these questions, we discuss predictions derived from the theoretical frameworks of behavioral game theory and human foraging analogs alongside the behavioral outcomes expected under the fully efficient solution.

The payoffs of the two decision makers in our paradigm are maximized if both never attempt to exploit the same option at the same time. This is because if both players choose the same option, any payoff is shared between them, whereas any potential payoff at the unattended option remains unexploited. In the experiments presented here, the two available options have unequal probabilities of yielding a reward (specifically, $p = .70$ and $1 - p = .30$). To guarantee that both players receive equal pay, they thus need to adopt a divergent choice strategy in which each player selects each option 50% of the time—for example, by following an alternating choice pattern. Hence, our two-player extension of a simple individual choice task is essentially a coordination problem that, despite the competitive payoff structure, demands trust and cooperation. Because we prevented direct communication between the players, such coordination can arise only by them communicating through their decisions. In short, two people who seek to exploit nature efficiently in this task face three discrete challenges: working out the optimal policy; signaling and coordinating its implementation (e.g., by following a strict alternation pattern); and upholding that implementation against the temptation of trying to maximize earnings independently. Two players who rise to this challenge can fully exploit nature via coordinated cooperation in the form of what we term *cooperative divergence*, with each obtaining 50% of the total rewards in a game.

Alternatively, one or both decision makers may attempt to maximize gain independently: Because one option is rewarded more often than the other (with $p = .70$), decision makers may be enticed to select this option exclusively. In fact, game theory predicts that both players will adopt this pure strategy of always selecting nature's more probable move, which constitutes a unique Nash equilibrium (Nash, 1951). Specifically, if both players always select the more likely option, all rewards will be shared, and each player will earn 35% of the total revenue in a game. Although smaller than the efficient outcome to be expected from coordinated cooperation, this shared payoff is larger than that associated with nature's less likely move ($1 - p = .30$). Thus, neither player is motivated to shift from the equilibrated outcome, although neither receives the maximal payoff ($p = .70$) and both forgo the long-term benefit of coordinated cooperation.

The ideal free distribution model makes no predictions about the choices of individual decision makers but, at the group level, assumes that the number of individuals exploiting a resource will be proportionate to the abundance of that resource. Situations typically described by this model involve large groups of foragers competing for distributed resources. In these situations, each forager maximizes its resource intake at the optimal model-predicted distribution, and available resources are exploited efficiently. Applied to our two-person paradigm, the ideal free distribution model would imply that, within each dyad, 70% of choices will be allocated to nature's more abundant resource and 30% to its alternative. This interpretation of the model, however, would suggest that the minimal group of two foragers will underexploit available resources when following the optimal solution. This prediction appears to be at odds with the general premise of the ideal free distribution approach. Because of this mismatch between the assumptions of the ideal free distribution model and the situation we study here, we do not consider these predictions further.

In what follows, we present two experiments in which we evaluate the determinants of efficient choice by systematically manipulating environmental (Experiment 1) and social (Experiment 2a and Experiment 2b) aspects of the paradigm.

Experiment 1

In Experiment 1, we examine the influence of two environmental scenarios on spontaneous strategy adoption. In one condition, outcome probabilities are known to participants from the start; in a second condition (between-subjects), outcome probabilities need to be learned over the course of multiple rounds. Pairs of participants in both conditions completed four blocks of 50 rounds each; that is, our experimental design is a mixed-model 2 (probability knowledge) × 4 (block of rounds) design. Revealing outcome probabilities at the start relieves participants of the onus of having to learn about environmental characteristics in addition to figuring out effective choice strategies in a social context. Research on individual binary choice has demonstrated that rates of probability maximizing are typically higher when outcome probabilities are known than when they have to be learned during the task (e.g., Fantino & Esfandiari, 2002). Extending an individual binary choice paradigm to include social interaction renders the relative outcome probabilities (and the knowledge thereof) less relevant because cooperative divergence leads to optimal, efficient outcomes regardless. Determining how to implement a divergence strategy, on the other hand, does require probability knowledge. Specifically, if nature rewards both options with equal odds, coordination and cooperation are simplified because each player can exploit one option exclusively. If nature rewards the options with unequal odds, each player should strive to select each option with the same frequency, which is more difficult to coordinate. Knowing from the start that one option is superior to the other may entice players to favor this option. Such preferential decisions of one option would hinder fair cooperative divergence. Therefore, contrary to findings from the individual choice case, we predicted that outcome probability knowledge would result in less successful cooperative divergence in this social context.

Method

Participants. Sixty (31 female) undergraduate students from the University of New South Wales with a mean age of 19.18 years ($SD = 1.58$ years) participated in this experiment. Sample size was determined prior to recruitment based on the magnitude of previously reported effect sizes and participant numbers in similar human choice experiments (e.g., Newell &

Rakow, 2007; Schulze, van Ravenzwaaij & Newell, 2015). In exchange for participation, students received course credit and could earn a small performance-based monetary payoff; earnings ranged from AUD$2.50 to AUD$3.80 (AUD$1 ≈ USD$0.97 at the time of the experiment).

Material and choice task. As depicted in Figure 1, two participants were seated at opposite sides of a table; the experimenter sat at the end of the table, facing both participants. No attempts were made to establish anonymity. To prevent any form of verbal or nonverbal communication between participants, the table was divided by an opaque screen. Each participant was given one black and one white card (8 × 10 cm) to indicate their decisions. Participant pairs were asked to simultaneously predict the color of a token that was repeatedly drawn (200 draws with replacement) from a box containing 70 black and 30 white tokens (the majority color was counterbalanced across participant pairs). Correct predictions were rewarded competitively: If only one participant predicted the correct color of the token, she received the full payoff (4 cents); if both participants converged on the correct response, the payoff was split equally between them (2 cents each). There was no penalty for incorrect predictions.

Design. A 2 × 4 mixed-model design with outcome probability knowledge (known vs. unknown) as between-subjects factor and trial block (four blocks of 50 rounds each) as within-subjects factor was used. Knowledge of outcome probabilities was manipulated by revealing or concealing the color configuration of the tokens in the box. Participants in the unknown/known probabilities condition, respectively, were told that the box contained 100 tokens and that some/70 of these tokens were black and some/30 were white. The main dependent measure was the proportion of divergent choices—in a two-person binary choice task, pairwise payoffs are maximized if both players choose opposite options on each round, allowing all potential rewards to be collected.

Procedure. Participants were invited to the experiment in pairs. Each pair was randomly allocated to the known-probabilities (n = 30) or unknown-probabilities (n = 30) condition and seated at opposite sides of a table, separated by an opaque screen (see Figure 1). The experimenter sat at the end of the table between the two participants and read the instructions aloud. Participants were asked to cease all communication with the other participant for the duration of the experiment and to attempt to earn as much money as possible during the task. The choice task required participants to make a series of predictions about the color of a token visibly drawn from a box by the experimenter. Correct choices were remunerated competitively—payoffs were split if both participants selected the same color. In the known-probabilities condition, participants were informed about the contents of the box and able to observe the experimenter place all tokens in an empty box before being given an opportunity to shuffle its contents. Instructions emphasized that the contents of the box would remain constant throughout the task and that each token would be drawn at random—no pattern or system would make it possible for participants to correctly predict the outcome of each draw (Shanks, Tunney, & McCarthy, 2002).

In each round, participants made their choices privately and then revealed them simultaneously by moving one colored card to the edge of the table (in view of both participants and the experimenter, see Figure 1) when prompted to do so by the experimenter. The experimenter then drew a token out of the box, revealed the outcome and payoff for that round, and replaced the token to ensure constant outcome probabilities. All choices, outcomes, and associated payoffs were recorded by the experimenter. This trial structure was repeated for 200 rounds, segmented into four blocks of 50 rounds each. After each block, there was a short break, during which participants were informed about both players' earnings in that block.

After the choice task and feedback on total earnings, each participant separately completed a short questionnaire on a computer. First, they were asked to estimate the chances of a black or white token being drawn. Second, they were asked to consider six choice strategies for 10 successive rounds of the interactive choice task: selecting the dominant color for (a) 10, (b) seven, (c) five, or (d) zero of the 10 draws of a token; or trying to select (e) the same or (f) the opposite color as the other player for all 10 draws of a token. Strategies a and b correspond to probability maximizing and matching, respectively; strategy f corresponds to efficiency by (attempted) cooperative divergence. Participants were asked which strategy they had used during the experiment and which strategy they would use for another hypothetical game. Finally, participants completed the Cognitive Reflection Test (CRT; Frederick, 2005) and the Berlin Numeracy Test (BNT; Cokely, Galesic, Schulz, Ghazal, & Garcia-Retamero, 2012).[1]

Results

In addition to conventional null hypothesis significance testing, we conducted Bayesian analyses: the default Bayesian analyses of variance (ANOVAs) outlined by Rouder, Morey, Speckman, and Province (2012), the default Bayesian t tests proposed by Rouder, Speckman, Sun, Morey, and Iverson (2009), and the default Bayesian hypothesis tests for correlations described by Wetzels and Wagenmakers (2012). For all analyses, we report Bayes factors (BF) that quantify the strength of evidence in favor of the presence of an effect.[2]

The proportion of individual participants' dominant color choices in each block of 50 rounds was subjected to a 2 (probability knowledge) × 4 (block of rounds) mixed-model ANOVA and is shown in Figure 2A. The figure also shows the optimal choice proportion resulting from cooperative divergence (recall that payoffs would be evenly maximized between players in a pair if both players always selected opposite colors, but each color 50% of the time). We found a significant main effect of block, F(2.33,

[1] The questionnaire also included measures of self-reported mathematical ability and degree of mathematical education. Because these measures were not related to our dependent variables, however, they are not considered further.

[2] All Bayesian analyses were carried out in R. For Bayesian ANOVAs and t tests, we used the anovaBF and ttestBF functions included in the BayesFactor package (v.0.9.9; Morey & Rouder, 2014) with their respective default settings, with one exception: For the anovaBF function, the number of Monte Carlo samples used to estimate BFs was increased to 50,000. For Bayesian ANOVAs, we computed BFs by contrasting the performance of a model including the relevant effect with one omitting (only) that effect. All models were constructed hierarchically such that the presence of an interaction term always involved the presence of all lower order interactions and main effects involving the components of that higher order interaction. For Bayesian t tests and Bayesian correlations, we report BFs in favor of the alternative hypothesis. A BF of 5, for instance, suggests that the data are five times more likely to have occurred under the model assuming the relevant effect than under a model omitting this effect, whereas a BF of 0.20 indicates that the data are 5 times more likely to have occurred under the model omitting the relevant effect than under a model including this effect.

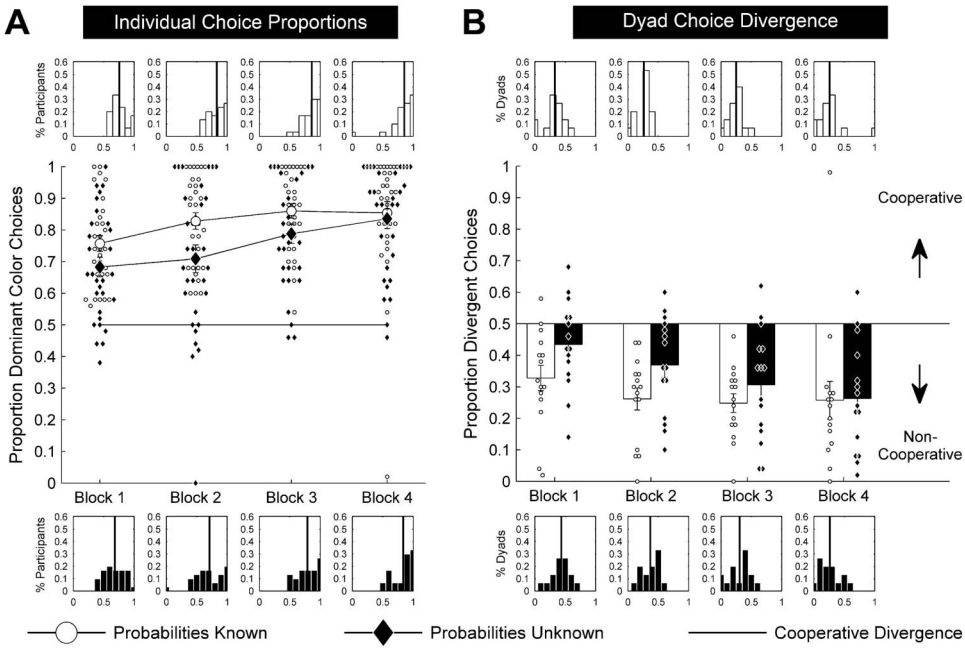

Figure 2. Proportion of individuals' dominant color choices (A) and dyads' divergent choices (B) in Experiment 1. The line graph in panel A plots the mean (± standard error) proportion of choices of the dominant color averaged across participants for all four blocks of 50 rounds. Cooperative divergence is indicated by the black line at .50. The bar graph in panel B plots the mean (± standard error) proportion of divergent choices within participant pairs averaged across dyads for all four blocks of 50 rounds. The expected choice divergence for two randomly responding players marks the baseline at .50; above/below this line, participants within a dyad cooperatively/noncooperatively chose different colors more/less often than not. In both panels, each small white circle/black diamond represents the average dominant color choice proportion of one participant (A) or the average choice divergence of one dyad (B) in the known/unknown probabilities condition during a specific block. The small bar graphs above/below the middle panel show the distribution of dominant color (A) or divergent choices (B) in the known/unknown probabilities condition, respectively, in percentage of participants or dyads for each block.

135.24) = 9.40, $p < .001$, $\eta_p^2 = .139$, $BF = 2139.05$; however, choice proportions showed an upward trajectory *away* from cooperative divergence across rounds.[3] The between-subjects factor of outcome probability knowledge was significant, $F(1, 58) = 5.29$, $p = .025$, $\eta_p^2 = .084$, $BF = 2.10$; participants in the known probabilities condition selected the dominant color more frequently—and thus showed less cooperative divergence—than did participants who were required to learn about outcome contingencies during the task ($M = .82$ and $M = .75$ across all four blocks of 50 rounds for participants in the known/unknown probabilities conditions, respectively). The Probability Knowledge × Block interaction was not statistically significant; in fact, the Bayesian analysis provided evidence for the absence of an effect, $F(2.33, 135.24) = 1.26$, $p = .290$, $\eta_p^2 = .021$, $BF = 0.19$. In addition to group averages, Figure 2A displays individual choice proportions and illustrates that most participants probability maximized by selecting the dominant color in at least 95% of rounds of the final block (33% of participants in both conditions; see Newell & Rakow, 2007). By contrast, only one participant in the known probabilities condition and two participants in the unknown probabilities condition fell within the range of optimal cooperative divergence (.50 ± .05 of dominant color choices) in the final block. In other words, most participants conformed to the equilibrium solution prescribed by game theory rather than attempting to achieve efficiency via cooperative divergence.

The choice proportions of single participants provide only partial information about choice behavior in a two-player task: These individual choice proportions would also average at .50 if two players in a dyad gave responses entirely at random rather than showing cooperative divergence. Hence, Figure 2B shows the mean and distribution of divergent choice proportions within participant pairs for each block of 50 rounds and both experimental conditions. A mixed-model ANOVA revealed a significant main effect of block, with choice divergence *decreasing* noncooperatively across rounds, $F(3, 84) = 4.34$, $p = .007$, $\eta_p^2 = .134$, $BF = 6.78$. Neither the main effect of probability knowledge nor the Probability Knowledge × Block interaction was significant, $F(1, 28) = 3.09$, $p = .090$, $\eta_p^2 = .099$, $BF = 0.92$ for the between-subjects main effect and $F(3, 84) = 0.89$, $p = .451$, $\eta_p^2 = .031$, $BF = 0.22$ for the interaction.

[3] For all conventional ANOVAs, degrees of freedom were corrected using the Greenhouse and Geisser (1959) coefficient when the sphericity assumption was violated.

Turning to the questionnaire data, we found that the probability estimates participants provided after the choice task did not differ between outcome probability knowledge conditions for either the dominant, $t(49.29) = -1.48$, $p = .145$, $BF = 0.66$, or the infrequent event, $t(49.29) = 1.48$, $p = .145$, $BF = 0.66$ ($M = .75$ and $M = .25$ for the dominant/infrequent outcome across probability knowledge conditions, respectively). Participants with no initial knowledge of the outcome probabilities therefore showed adequate probability learning. Prior research has demonstrated that performance on the CRT is positively related to accurate choice in binary prediction paradigms (see, e.g., Koehler & James, 2010; Schulze & Newell, 2015). We found no such correlation between CRT score and the proportion of divergent choices within dyads, $r(58) = .193$, $p = .140$, $BF = 0.30$, or the proportion of individuals' dominant color choices, $r(58) = -.213$, $p = .102$, $BF = 0.38$, across all rounds of the task. In fact, in both cases, the Bayesian analysis provided evidence for the null hypothesis. Neither of these relationships reached statistical significance when we examined the two probability knowledge conditions separately and/or controlled for participants' numeracy as indicated by the BNT score (all ps $\geq .125$ and all BFs ≤ 0.51).

Discussion

Participants did not spontaneously construe our interactive sequential choice problem as a social opportunity to cooperate and coordinate their efforts to fully exploit an uncertain environment. Instead, their responses strongly resembled those observed in the individual choice case: Most participants probability maximized, in accordance with the solution prescribed by the Nash equilibrium. This pattern emerged largely irrespective of participants' a priori knowledge of outcome probabilities. In fact, as predicted, outcome probability knowledge was a slight impediment to successful cooperative divergence. Only a single participant pair diverged fully by the end of the learning process, thus exploiting all obtainable resources. However, this dyad achieved efficiency not by mutual cooperation, but at the expense of one player's payoff—one player exclusively selected the dominant outcome, the other, the infrequent event. Moreover, when presented with a range of choice strategies, including efficiency by cooperative divergence, after the choice task, most participants indicated that they would probability maximize (43%) or match (32%), if they were to participate again. Only 17% of participants endorsed cooperative divergence ("choose oppositely to the other player").

What prevented participants from discovering, implementing, and maintaining the solution offering the highest long-term earnings in this paradigm? The game theoretic notion of equilibrium suggests that this outcome represents the rational solution given the payoff structure at hand. Yet, from a pure learning perspective, another explanation is that the goal of learning—maximization of personal payoffs—may have prevented participants from recognizing that efficient exploitation of the environment *could* be achieved via cooperation. Even if players recognized this possibility, they may reasonably have assumed that their partner would be less insightful and thus decided not to risk forgoing personal gain by signaling (ultimately futile) cooperation. To test this hypothesis, we introduced competitive dynamics that either highlighted or deemphasized individual earning goals when outcome probabilities were initially unknown (Experiment 2a) or known (Experiment 2b).

Experiment 2a

In this experiment, we contrasted the effects of within-group competition with those of between-group competition to manipulate the salience of individual earning goals. In one condition (within-dyad competition), individual players were encouraged to earn more than the other player in their dyad; in the other condition (between-dyad competition), dyads were encouraged to earn more than another pair of players. Dyads in both experimental conditions again completed four blocks of 50 rounds in the choice task. We thus applied a 2 (competitive focus) $\times$ 4 (block of rounds) mixed-model design. Research on coordination games and social dilemmas suggests that collective efficiency and contributions toward a public good are higher under between-group competition than in single-group settings (e.g., Bornstein, 2003; Bornstein, Gneezy, & Nagel, 2002). Accordingly, we predicted that highlighting between-dyad competition would facilitate efficient cooperative divergence by de-emphasizing individual earning goals. By contrast, we predicted that highlighting within-dyad competition would facilitate inefficient probability maximizing by emphasizing individual maximization goals as well as concerns for competitive distrust.

Method

Participants. Sixty (38 female) undergraduate students from the University of New South Wales with a mean age of 19.42 years ($SD = 4.85$ years) participated in this experiment in exchange for course credit. In addition, participants could earn a performance-based monetary payoff; earnings ranged from AUD$1.30 to AUD$8.65 (AUD$1 $\approx$ USD$1.04 at the time of the experiment).

Choice task and design. We used the same interactive sequential choice task as described for Experiment 1, but withheld information about the precise color configuration of the tokens in the box from all participants. That is, participants were informed that the box contained 100 tokens, some black and some white, but that they would not be told how many of the tokens were black/white. Instead, we manipulated the competitive frame of the task, using a 2 $\times$ 4 mixed-model design with competitive focus (within-dyad vs. between-dyad competition) as between-subjects factor and trial block (four blocks of 50 rounds each) as within-subjects factor. Thirty participants (i.e., 15 dyads) were randomly assigned to each competitive focus condition. Those with a within-dyad competitive focus were instructed that they were in competition with the other player in their dyad and that their goal was to earn more than that person. Those with a between-dyad competitive focus were instructed that they were competing against another team of two players and that their goal was to achieve higher combined earnings than the opposing dyad. However, each player could only keep what they earned. That is, payoffs were not split at the end of the experiment, and the experimental procedure ensured that participants in each pair were paid, debriefed, and dismissed separately. In addition, we offered further incentives to enhance performance in each block of rounds: Participants were informed that if they reached their competitive goal (i.e., won the

within- or between-dyad competition) their payoff for that block would be doubled, whereas if they lost it, their payoff for the block would be halved.[4] Upon completion of each block of 50 rounds, this comparison was made publicly and payoffs were adjusted accordingly. Participant pairs in the between-dyad competitive focus condition were told that they were competing against the participant pair who had completed the task before them. In fact, they played against a virtual team to ensure similar competitive pressure for all pairs. The block payoffs of this virtual team were computed separately for each block and participant pair by randomly drawing a team choice divergence from a normal distribution ($\mu = .90$, $\sigma = .05$). As before, the main dependent measure was the proportion of divergent choices.

Results and Discussion

Figure 3A shows the results of a 2 (competitive focus) × 4 (block of rounds) mixed-model ANOVA on the proportion of individual players' dominant color choices. The main effect of competitive focus was significant, $F(1, 58) = 7.40$, $p = .009$, $\eta_p^2 = .113$, $BF = 5.26$; participants competing against the other player in their dyad selected the dominant color more frequently ($M = .74$ across blocks of 50 rounds) than participants competing against another pair of players ($M = .60$ across blocks of 50 rounds). We found a significant main effect of probability learning across blocks, although the Bayesian evidence was somewhat ambiguous, $F(2.67, 154.88) = 3.57$, $p = .019$, $\eta_p^2 = .058$, $BF = 1.38$. Competitive focus significantly interacted with the within-subjects block factor, $F(2.67, 154.88) = 4.26$, $p = .009$, $\eta_p^2 = .068$, $BF = 7.11$; the learning slopes of participants under between-dyad competition remained flatter and closer to optimal cooperative divergence than did those of participants under within-dyad competition, which increased steeply toward probability maximizing (see Figure 3A). Examining individual choice proportions revealed that the between-dyad competition group reached cooperative divergence because individuals either selected both colors with equal rates (see black diamond cluster at .50 in Figure 3A) or exclusively favored one color (see black diamond clusters at 1 and 0 in Figure 3A).

These individual choice proportions were indeed coordinated within pairs of participants; pairs with a between-dyad competitive focus increasingly diverged on different choice options, as shown in Figure 3B. By contrast, pairs competing against each other (within-dyad) diverged decreasingly across blocks, and both the Competitive Focus ×Block interaction, $F(2.18, 61.01) = 15.91$, $p < .001$, $\eta_p^2 = .362$, $BF = 517692.90$, and the competitive focus main effect, $F(1, 28) = 17.00$, $p < .001$, $\eta_p^2 = .378$, $BF = 89.07$, were significant. The main effect of the within-subjects block factor did not reach statistical significance; in fact, the Bayesian analysis provided evidence for the absence of an effect, $F(2.18, 61.01) = 1.18$, $p = .318$, $\eta_p^2 = .040$, $BF = 0.11$.[5]

To characterize the learning and coordination processes of individual dyads under between-dyad competition in more detail, Figure 4 displays—for all 15 dyads separately—the blocked proportion of players' dominant color choices (black circles and white diamonds) together with a 10-round running window of the choice divergence ensuing from these decisions (gray shading). The plotted dyad trajectories are sorted approximately by final choice divergence: the top eight dyads reached full divergence by the last block of rounds (at least 95% divergent choices); the bottom seven fell short of this criterion. We examined whether choice divergence increased gradually as each dyad slowly negotiated a divergence strategy (nonverbally by signaling through decisions) or more rapidly as the proverbial "penny dropped" for one or both players in each dyad. That is, we wanted to know whether differences in choice divergence between dyads stemmed from differences in when cooperative divergence was *recognized* as the efficient solution or differences in how long dyads needed to *coordinate* its implementation. To this end, each subplot indicates rounds on which one player signaled a divergence strategy to the other (gray crosses at the top of each graph). We defined these signals based on 10-round windows of previous choices: If a player selected the infrequent event for 10 successive rounds or followed a perfect alternation pattern for 10 successive rounds, we marked this 10-round signal with a cross at round 10.

A number of conclusions can be drawn from visual inspection of these trajectories. First, only a single dyad implemented a round-by-round alternating pattern (Dyad 1) and followed this pattern for the length of only a single block. Second, five of the top eight dyads diverged at the expense of one player who almost always selected the infrequent event throughout the experiment (Dyads 4–8). Nevertheless, for nine of the 15 dyads (60%), divergent choices were distributed evenly—overall payoff differed less than one dollar between the two players. Third, for all top eight dyads, choice divergence increased rapidly at the first signal (disregarding few dispersed alternations at the start of

[4] Despite these additional incentives to maximize personal gain in the within-dyad competitive focus condition, cooperative divergence still returns higher expected reward than probability maximizing. This is because two probability maximizing players would share all payoffs, tie at the end of each block, and no adjustments (halving or doubling) would apply. Over 200 rounds of the choice task, the expected earnings of both players would average at \$2.80—that is, 200 (number of rounds) × .70 (majority outcome probability) × 0.02 (shared payoff amount). By contrast, if both players cooperatively diverged on all choices on a trial-by-trial basis (i.e., alternate), they would share *no* payoffs but—on average—also tie at the end of each block. The expected payoff from this strategy is \$4.00 for each player (200 × .50 × 0.04). Moreover, if both players cooperatively diverged on a block-by-block basis (i.e., each player chose all black/white for an equal number of blocks), they would also share no payoffs and exploit the payment adjustment to their advantage. The expected payoff from this strategy is \$6.20 for each player—100 × .70 × 0.04 × 2 (100 rounds of probability maximizing with doubled earnings) + 100 × .30 × 0.04/2 (100 rounds of exclusively selecting the infrequent outcome with halved earnings).

[5] Numerous studies report gender effects on competitive and cooperative behavior in strategic games (see Camerer, 2003). Including dyads' gender configuration (both female, both male, or mixed gender) as a factor in the ANOVA on within-dyad divergence revealed no significant gender configuration main effect or interactions with this factor (all $ps \geq .345$ and all $BFs \leq 0.71$), although male-only dyads showed a slight tendency toward higher cooperative divergence across blocks. Similarly, there was a small tendency for male participants to select the dominant color more frequently; however, including gender in the ANOVA on dominant color choices revealed no significant gender main effect or interactions with this factor (all $ps \geq .481$ and all $BFs \leq 0.41$). Moreover, these minor trends may be confounded by gender differences in cognitive thinking styles and mathematical abilities—gender (male coded as 1, female as 0) correlated positively with CRT score, $r(58) = .378$, $p = .003$, $BF = 8.22$, and BNT score, $r(58) = .315$, $p = .014$, $BF = 2.03$.

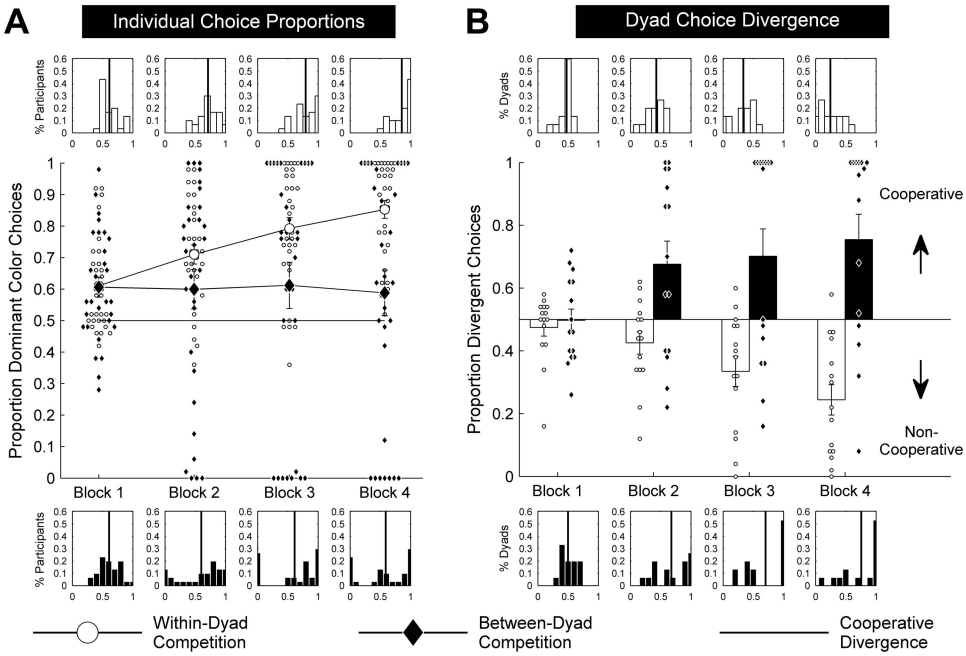

Figure 3. Proportion of individuals' dominant color choices (A) and dyads' divergent choices (B) in Experiment 2a. The line graph in panel A plots the mean (± standard error) proportion of choices of the dominant color averaged across participants for all four blocks of 50 rounds. Cooperative divergence is indicated by the black line at .50. The bar graph in panel B plots the mean (± standard error) proportion of divergent choices within participant pairs averaged across dyads for all four blocks of 50 rounds. The expected choice divergence for two randomly responding players marks the baseline at .50; above/below this line, participants within a dyad cooperatively/noncooperatively chose different colors more/less often than not. In both panels, each small white circle/black diamond represents the average dominant color choice proportion of one participant (A) or the average choice divergence of one dyad (B) in the within-/between-dyad competition condition during a specific block. The small bar graphs above/below the middle panel show the distribution of dominant color (A) or divergent choices (B) in the within-/between-dyad competition condition, respectively, in percentage of participants or dyads for each block.

probability learning). This suggests that successful strategy use was characterized by one player abruptly realizing that full exploitation is possible via choice divergence, followed by the other rapidly coordinating her part in implementing this strategy.[6] And fourth, dyads who did not diverge fully by the final block (bottom seven dyads) also showed abrupt changes in choice divergence, if at least one player recognized the optimal strategy (yet the other may not have caught on). For example, in Dyad 9, Player 2 persistently probability maximized while Player 1 oscillated between attempting to signal a fair divergence strategy (alternation) and surrendering to diverge by selecting the infrequent event exclusively.

Probability estimates elicited after the choice task ($M = .75$ and $M = .26$ for the dominant/infrequent outcome, respectively, across competitive focus conditions) did not differ between competitive focus conditions for the dominant, $t(58) = 1.44$, $p = .154$, $BF = 0.63$, or for the infrequent outcome, $t(58) = -1.64$, $p = .106$, $BF = 0.80$. Therefore, participants in both conditions showed adequate probability learning throughout the choice task. For participants under between-dyad competition, the CRT score correlated positively with the proportion of divergent choices within dyads, $r(28) = .478$, $p = .007$, $BF = 4.89$, but negatively with the proportion of dominant color choices across all rounds of the task, $r(28) = -.373$, $p = .043$, $BF = 1.09$. No such relationships existed among participants who competed against the other player in their dyad; $r(28) = -.099$, $p = .603$, $BF = 0.16$ for proportion of divergent choices and $r(28) = .285$, $p = .127$, $BF = 0.45$ for proportion of dominant color choices.[7]

In sum, between-dyad competition strongly facilitated optimal cooperative divergence and resulted in efficient but not always equitable exploitation of the uncertain environment. Within-dyad competition, on the other hand, amplified inefficient probability maximizing. To

[6] To corroborate this conclusion, we also examined changes in choice divergence between blocks for each dyad. Seven of the top eight dyads showed a single increase in choice divergence between two blocks (mostly the first and second) that accounted for more than 50% of the total increase in choice divergence across all blocks. Only one dyad (Dyad 6) showed equal increases in choice divergence between Blocks 1–2 and Blocks 2–3, indicating a more gradual coordination process.

[7] When we controlled for numeracy as indicated by the BNT score, these relationships between CRT score and the proportion of divergent/dominant color choices remained statistically significant for the between-dyad competition condition—$r(27) = .402$, $p = .031$, $BF = 2.15$ and $r(27) = -.371$, $p = .047$, $BF = 1.54$ for the partial correlation between CRT score and proportion of divergent/dominant color choices, respectively—but did not reach statistical significance for the within-dyad competition condition, although the Bayesian evidence was ambiguous (all ps $\geq .053$ and all BFs ≤ 1.38).

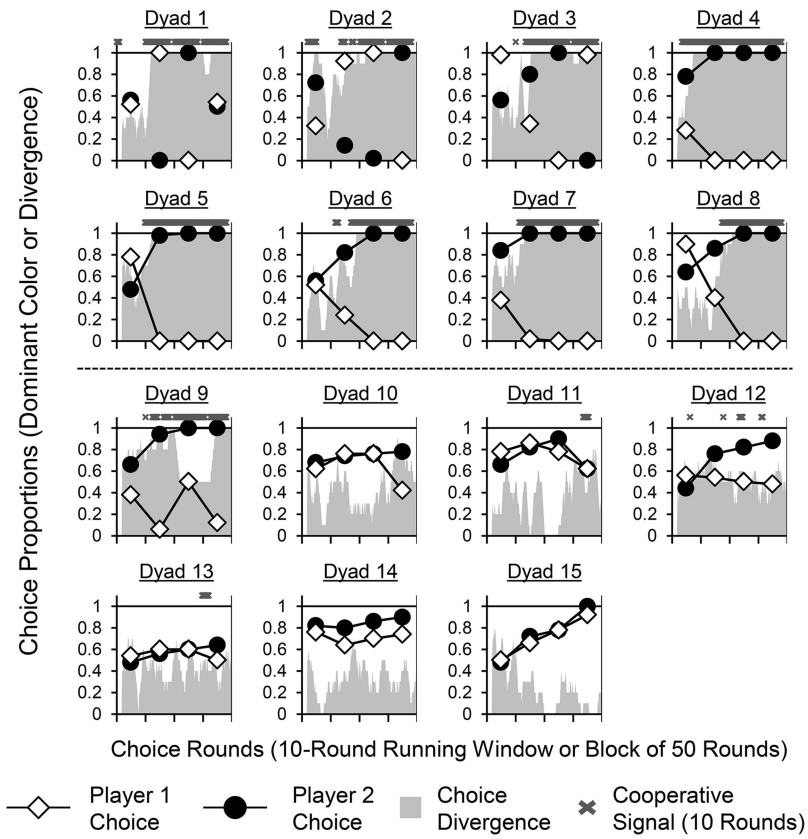

Figure 4. Choice and divergence trajectories for each of 15 participant pairs in the between-dyad competition condition in Experiment 2a. For each dyad, one subplot displays both players' blocked proportions of dominant color choices (black circles and white diamonds) alongside a 10-round running window of the choice divergence ensuing from these choices (gray shading). In addition, at the top of each graph, gray crosses indicate rounds on which one player signaled a divergence strategy to the other (see main text for details). The 15 subplots are organized approximately by final choice divergence; the top eight dyads reached full divergence by the last block (at least 95% divergent choices); the bottom seven fell short of this criterion. Where choice divergence was implemented via block-wise alternation (Dyads 1–3), connecting lines between blocked-choice proportions were omitted for both players.

test the role of outcome probability knowledge on this competitive focus effect, in Experiment 2b we replicated the design and procedure of Experiment 2a but disclosed outcome probabilities to participants at the start. Similar to Experiment 1, we predicted that outcome probability knowledge would hinder cooperative divergence further when individual earning goals are emphasized (within-dyad competition). This is because definite knowledge that one option is more profitable than the other is likely to increase the appeal of that option as a means to maximizing personal gain. When dyad earning goals are emphasized (between-dyad competition), in contrast, outcome knowledge may promote rather than impede (equitable) cooperative divergence. This is because outcome probabilities provide the dyad with the information needed to work out an effective and equitable divergence strategy while reducing overall task difficultly.

Experiment 2b

Method

Participants. Sixty (38 female) undergraduate students from the University of New South Wales with a mean age of 19.05 years (SD = 3.53 years) participated in this experiment in exchange for course credit. In addition, participants could earn a performance-based monetary payoff; earnings ranged from AUD$1.20 to AUD$8.65 (AUD$1 ≈ USD$1.04 at the time of the experiment).

Choice task and design. The choice task, design, and procedure were exact replications of those described for Experiment 2a, except that all participants were informed about the color configuration of the 100 tokens in the box—that is, that it contained 70 black/white and 30 white/black tokens (counterbalanced across participant pairs). Thirty participants (i.e., 15 dyads) were randomly assigned to each competitive focus condition.

Results and Discussion

When outcome probabilities were known at the outset, participant pairs competing against another dyad learned to diverge cooperatively to similar extents as they did under probability learning from experience (Experiment 2a). By contrast, participants competing within their dyad again mostly learned to probability maximize (see Figure 5A). A mixed-model ANOVA on proportion of dominant color choices showed that both the main effect of competitive focus, $F(1, 58) = 24.67$, $p < .001$, $\eta_p^2 = .298$, $BF = 2225.89$, and the Compet-

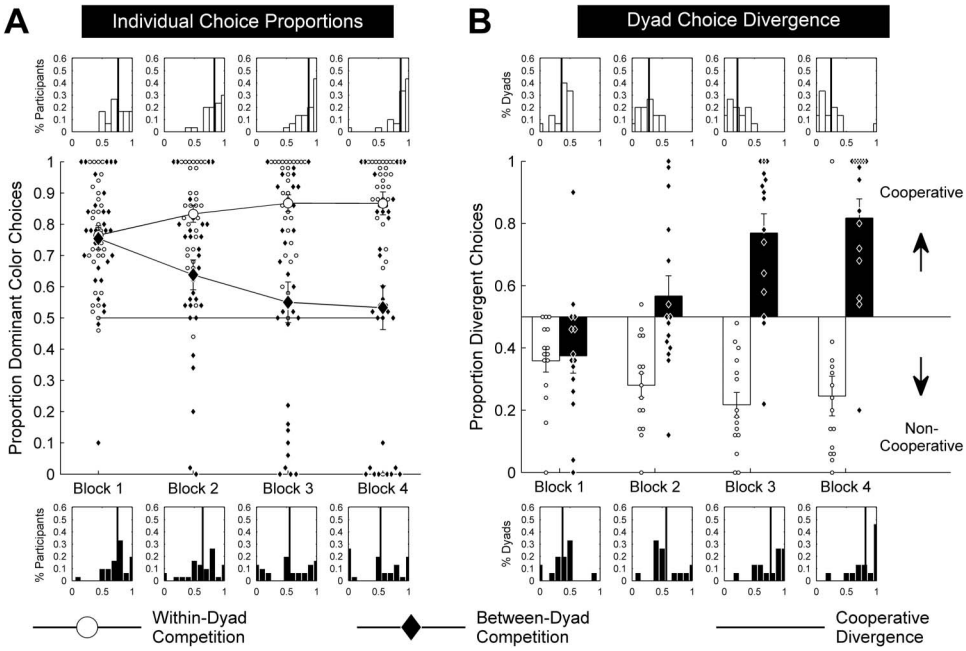

Figure 5. Proportion of individuals' dominant color choices (A) and dyads' divergent choices (B) in Experiment 2b. The line graph in panel A plots the mean (± standard error) proportion of choices of the dominant color averaged across participants for all four blocks of 50 rounds. Cooperative divergence is indicated by the black line at .50. The bar graph in panel B plots the mean (± standard error) proportion of divergent choices within participant pairs averaged across dyads for all four blocks of 50 rounds. The expected choice divergence for two randomly responding players marks the baseline at .50; above/below this line participants within a dyad cooperatively/noncooperatively chose different colors more/less often than not. In both panels, each small white circle/black diamond represents the average dominant color choice proportion of one participant (A) or the average choice divergence of one dyad (B) in the within-/between-dyad competition condition during a specific block. The small bar graphs above/below the middle panel show the distribution of dominant color (A) or divergent choices (B) in the within-/between-dyad competition condition, respectively, in percentage of participants or dyads for each block.

itive Focus × Block interaction, $F(2.37, 137.39) = 7.49$, $p < .001$, $\eta_p^2 = .114$, $BF = 307.04$, were significant. The main effect of the within-subjects block factor was not significant; in fact, the Bayesian analysis provided evidence for the absence of an effect, $F(2.37, 137.39) = 0.99$, $p = .387$, $\eta_p^2 = .017$, $BF = 0.06$. As in Experiment 2a, in the between-dyad competition group, the proportion of divergent choices increased across blocks (see Figure 5B). By contrast, participants in the within-dyad competition group increasingly converged on the same option across blocks. The mixed-model ANOVA on proportion of divergent choices within participant dyads indicated a significant main effect of competitive focus, $F(1, 28) = 44.31$, $p < .001$, $\eta_p^2 = .613$, $BF = 25155.47$, and a significant Competitive Focus × Block interaction, $F(3, 84) = 17.32$, $p < .001$, $\eta_p^2 = .382$, $BF = 2298817.71$. The level of choice divergence among participant pairs in the between-dyad competition group was higher than observed in Experiment 2a, such that the overall level of choice divergence also increased significantly across blocks, $F(3, 84) = 5.43$, $p = .002$, $\eta_p^2 = .162$, $BF = 2.54$ (from .37 to .53 between the first and last block, averaged across competitive focus conditions).[8] Examining individual choice proportions again indicated that the between-dyad competition group reached cooperative divergence because participants either selected both colors with equal rates (see black diamond cluster at .50 in Figure 5A) or exclusively favored one color (black diamond clusters at 1 and 0 in Figure 5A).

Inspection of the coordination trajectories of each dyad in the between-dyad competition group (see Figure 6) again indicated rapidly increasing choice divergence after the first signal in dyads with full final divergence (at least 95% divergent choices in the last block—Dyad 8 narrowly missed this criterion because of coordination error during the first few rounds of the last block). Relative to Experiment 2a, fewer dyads diverged efficiently at the expense of one disadvantaged player who almost exclusively selected the low-probability option (Dyads 5–7), and more dyads implemented a round-by-round alternating pattern for longer lengths of rounds (Dyads 1–2). However, three dyads that did not reach criterion by the final

[8] Again, no significant gender configuration main effect or interactions were found when within-dyad choice divergence was subjected to a 3 (gender configuration) × 2 (competitive focus) × 4 (block of rounds) mixed-model ANOVA (all $ps \geq .066$ and all $BFs \leq 0.94$). However, in contrast to Experiment 2a, when we included gender in the ANOVA on dominant color choices, we found a significant gender main effect, $F(1, 56) = 11.10$, $p = .002$, $\eta_p^2 = .165$, $BF = 11.85$; male participants selected the dominant color *less* frequently than female participants did ($M = .65$ and $M = .77$ across all blocks of 50 rounds for male and female players, respectively). None of the interactions with the gender between-subjects factor reached statistical significance (all $ps \geq .070$ and all $BFs \leq 0.97$) and, again, we found a potentially confounding correlation between gender (male = 1) and CRT scores, $r(58) = .383$, $p = .003$, $BF = 9.50$.

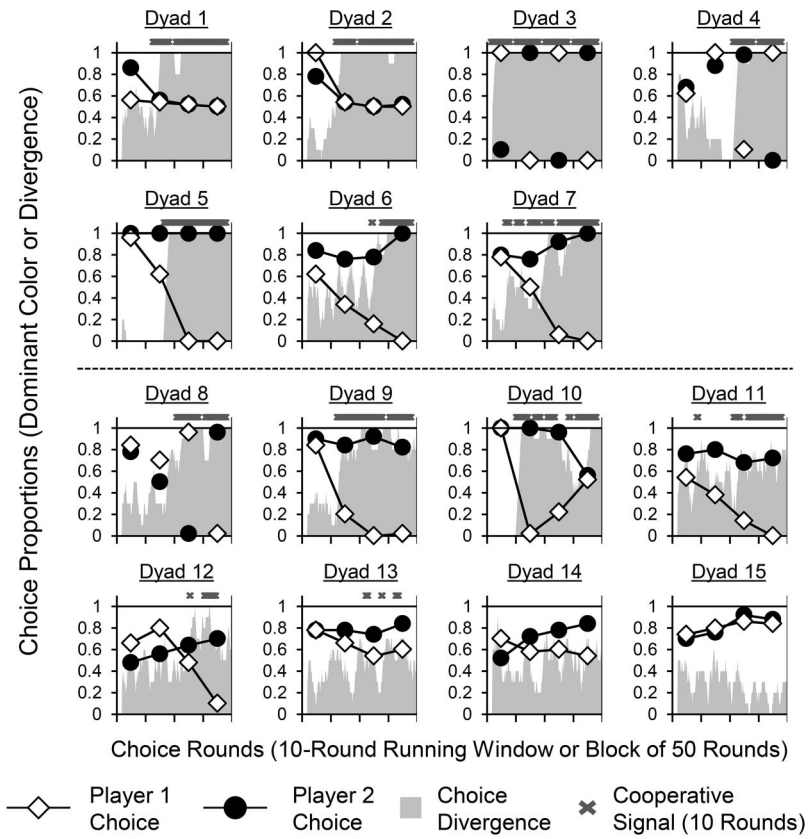

Figure 6. Choice and divergence trajectories for each of 15 participant pairs in the between-dyad competition condition in Experiment 2b. For each dyad, one subplot displays both players' blocked proportion of dominant color choices (black circles and white diamonds) alongside a 10-round running window of the choice divergence ensuing from these choices (gray shading). In addition, at the top of each graph, gray crosses indicate rounds on which one player signaled a divergence strategy to the other (see the Results section of Experiment 2a for details). The 15 subplots are organized approximately by final choice divergence; the top seven dyads reached full divergence by the last block (at least 95% divergent choices); the bottom eight fell short of this criterion. Where choice divergence was implemented via block-wise alternation (Dyads 3–4 and Dyad 8), connecting lines between blocked-choice proportions were omitted for both players.

block (Dyads 9–11) disadvantaged a single player who attempted (with only partial success) to signal divergence in numerous rounds. Thus, overall, knowledge of probabilities did not facilitate fairness in cooperative divergence and, again, only nine of 15 dyads (60%) distributed choices equally among the two choice options, yielding overall payoff differences of less than one dollar.

We again found a positive correlation between CRT scores and the proportion of divergent choices among participants under between-dyad competition, $r(28) = .554$, $p = .001$, $BF = 21.07$. However, the relationship between CRT score and proportion of dominant color choices was only marginally negative for this group, and the Bayesian analysis provided weak evidence for the null hypothesis, $r(28) = -.351$, $p = .057$, $BF = 0.85$. For participants under within-dyad competition, neither of these correlations was statistically significant: $r(28) = -.003$, $p = .988$, $BF = 0.14$ for proportion of divergent choices and $r(28) = -.297$, $p = .111$, $BF = 0.50$ for proportion of dominant color choices.[9]

In sum, choice behavior in Experiment 2b mirrored the findings from Experiment 2a, in which outcome probability information was initially unknown. That is, the effects of competitive focus remained largely unaffected by the provision of outcome probabilities.

General Discussion

Countless everyday decisions oblige people to simultaneously navigate complex social contexts and uncertain environmental consequences. Studying human decision making by examining

[9] Again, when we controlled for numeracy as indicated by the BNT score, these inferences remained largely the same. For participants under between-dyad competition, there was a partial correlation between CRT score and proportion of divergent choices, $r(27) = .437$, $p = .018$, $BF = 3.37$, but not between CRT score and proportion of dominant color choices, $r(27) = -.230$, $p = .231$, $BF = 0.43$. For the within-dyad competition condition, there was no partial correlation between CRT score and proportion of divergent choices, $r(27) = .033$, $p = .866$, $BF = 0.22$, but, surprisingly, the partial correlation between CRT score and proportion of dominant color choices was significant, $r(27) = -.415$, $p = .025$, $BF = 2.66$.

either a single person choosing under environmental uncertainty (e.g., with probability learning paradigms; Estes, 1964) or two people competing, cooperating, or coordinating in an environmental void (e.g., standard behavioral game theory; Camerer, 2003) does not adequately reflect this challenge. Our work connects these two approaches by extending a standard individual choice paradigm to include social interaction with one other person. Although this additional social aspect made the task more complex, it also offered participants a unique opportunity to fully exploit an uncertain environment via cooperation.

Experiment 1 demonstrated that people did not spontaneously seize this opportunity to choose efficiently via coordinated cooperation. Instead, response patterns were mostly congruent with those seen in the individual learning case: participants tended to probability maximize irrespective of outcome probability knowledge. This pattern coincides with the Nash equilibrium prescribed by game theory. It does not, however, constitute an efficient solution within the environment at hand. That is, on each separate round, two probability maximizing players can rightly expect larger (shared) payoffs than any player defecting to the less abundant option could. Yet focusing on each immediate choice means forgoing the opportunity to capitalize on the social dimension of the task. Experiment 2a and Experiment 2b showed that drawing people's attention to group gains rather than individual gains via between-group competition helped players to depart from the inefficient equilibrium by cooperating with potential competitors. That is, emphasis on dyad earning goals facilitated efficient—but not always equitable—exploitation of uncertain choice environments. By contrast, emphasis on individual earning goals (within-group competition) increased inefficient probability maximizing.

How did between-group competition help participants to overcome their initial inertia to cooperate? To answer this question, let us consider the discrete requirements facing two players who seek to exploit nature efficiently in this task. First, one or both players need to work out the optimal policy: Payoffs can only be maximized if all bases are covered on each round—that is, if their choices diverge. Second, the two players need to coordinate the implementation of this optimal strategy. And third, the divergence strategy needs to be upheld against the potential allure of attempting to maximize individually. Our analyses of dyads' coordination trajectories suggest that between-group competition did not facilitate the initial availability of optimal cooperative divergence as a strategy; none of the dyads in Experiment 2a or Experiment 2b diverged efficiently from the start. Rather, participants needed to *experience* at first hand the added incentive to diverge to discover the best strategy (see Newell, Koehler, James, Rakow, & van Ravenzwaaij, 2013; Newell & Rakow, 2007). Once discovered, optimal cooperative divergence was implemented rapidly by most dyads, indicating that the challenge lay in the proverbial penny waiting to drop rather than in coordinating and maintaining cooperative divergence. In other words, by incentivizing group earnings, between-dyad competition largely eliminated the potential allure of individual maximization—a conclusion underlined by the finding that a substantial subset of dyads implemented choice divergence inequitably at the expense of one player.[10]

In this respect, the competitive focus manipulation relates to melioration accounts of individual choice in tasks where immediate satisfaction and long-term goals are conflicted (see Herrnstein & Prelec, 1991). Such paradigms aim to model ubiquitous real-world decisions that encompass delayed and indirect consequences—for example, whether to have that second piece of cake now or to resist for future health benefits. In our paradigm, this temporal conflict is extended to interactive choice, in that an immediately lucrative action is at odds with long-term cooperative behavior. Between-group competition, however, alleviated this conflict by highlighting (and incentivizing) the long-term benefit of cooperation—higher dyad earnings permitted players to win the between-group contest. The CRT scores obtained after the choice task support this conclusion. Cognitive reflection, as measured by the CRT, is associated with higher sensitivity to delayed rewards (Frederick, 2005). In line with this interpretation, we found positive relationships between CRT scores and long-term oriented cooperative divergence under between-group competition in Experiment 2a and Experiment 2b. By contrast, CRT scores and choice divergence were not correlated under within-dyad competition or in Experiment 1. This pattern points to potentially different mechanisms generating individual differences in choice divergence: Differences in high levels of choice divergence under between-dyad competition may have resulted from differences in sensitivity to delayed cooperative rewards (as suggested by the CRT results). By contrast, differences in low levels of choice divergence under within-dyad competition and in Experiment 1 may have stemmed from individual differences in strategy use. For example, a pair of probability maximizing participants would have a lower choice divergence than a pair of probability matching participants. These differences in strategy preference (and their consequences for choice divergence) are, however, unrelated to abandoning immediately profitable action in favor of efficient dyad cooperation.

In this article, we examined how people cope with the combined challenge of navigating social contexts and facing environmental uncertainty by extending a standard solitary choice paradigm to include another person. Taking individual risky choice as the starting point and then asking how adding social dimensions affects people's behavior offers a fruitful approach to extending a large body of individual decision making research across a variety of domains. We are not the first to explore this route. Phillips, Hertwig, Kareev, and Avrahami (2014), for instance, examined how the introduction of competition affects experience-based choice in a sampling paradigm and found that, relative to solitary choice, the balance of the explore–exploit tradeoff shifted in favor

[10] This finding resonates with related work in experimental economics showing that behavior in competitive experimental markets can converge on equilibrated but extremely unequal payoff distributions (Roth, Prasnikar, Okuno-Fujiwara, & Zamir, 1991). It is important to note, however, that the "inequality acceptance" observed under between-dyad competition in our task may in part be due to incomplete learning of outcome probabilities and how best to implement an efficient solution. Comparing dyads' choice trajectories between Experiment 2a and Experiment 2b gives some indication of this (see Figure 4 and Figure 6). In Experiment 2a, where outcome probabilities needed to be learned from experience, strategy crossovers that equalize pay between players often occurred only in the very last block of the experiment. In Experiment 2b, where probability learning requirements were removed, more high-performing dyads diverged equitably early on. Moreover, the majority of dyads under between-dyad competition in both experiments distributed their choices equally among the two choice options and yielded very low overall payoff differences. Thus, given an opportunity to complete more learning rounds, more dyads would likely arrive at an equitable solution.

of rapid exploitation under competition. Although the respective research objectives differ, many of our findings are reminiscent of issues also addressed within the adjacent framework of behavioral game theory. Thus, it is important to discuss the contributions of our work in light of the existing literature.

Here, we have examined what essentially presents as a coordination problem under environmental uncertainty. To achieve efficiency, two people consistently need to select different options; to achieve equitability, both need to coordinate their behavior by taking turns in exploiting the richer option. The emergence, evolution, and benefit of cooperative turn-taking have previously been studied in iterated forms of related games, including the chicken (or hawk–dove) game and the battle of the sexes (or hero) game (see, e.g., Bornstein, Budescu, & Zamir, 1997; Browning & Colman, 2004; Colman & Browning, 2009; Crowley, 2001; Lau & Mui, 2008). In these games, too, alternation can be difficult to elicit but typically emerges as learning progresses. The payoff structure of these games is, however, characterized by dual pure strategy Nash equilibria and is thus distinctly different from our paradigm, in that both players are motivated *not* to converge on the same outcome. Moreover, the study of cooperative turn-taking in these games rarely considers environments in which the allocation of payoffs is stochastically determined.

Another related class of strategic games involving coordination of market entry by multiple players has been examined under simultaneous outcome and strategic uncertainty (see Rapoport, Seale, & Ordónez, 2002). In the stochastic form of these games (termed *endogenously determined lotteries* by Rapoport et al., 2002), players can opt to enter a lottery or not; the outcome probabilities of the lottery depend on the number of players electing to participate. Players who refrain from entering the lottery receive deterministically determined payoffs. As such, this game represents a social extension of a common individual choice paradigm involving repeated decisions between a safe and a risky option (see, e.g., Barron & Erev, 2003). Unlike the paradigm of interest here, however, this game allows for a line of action that evades the simultaneous challenge of social and environmental uncertainty. Examining the choices of large groups of people, Rapoport et al. (2002) found that, compared with market entry games that lack outcome uncertainty (see, e.g., Erev & Rapoport, 1998), choice behavior deviated systematically from the equilibrium solution. These results are in line with findings obtained with prisoner's dilemma and interdependent security games showing that outcome uncertainty affects how people learn and how they respond in experimental games otherwise studied under deterministic outcomes (e.g., Bereby-Meyer & Roth, 2006; Kunreuther, Silvasi, Bradlow, & Small, 2009; Shafran, 2012).

Which factors besides between-group competition might also facilitate optimal choices in our paradigm? In an effort to closely align our interactive task with the individual choice paradigm on which it is based, we did not allow explicit verbal or nonverbal communication between players. Instead, participants could only signal choice strategies through their decisions. The finding that optimal choice arose even under these constraints is remarkable and highlights the scope of the between-group competition manipulation. Nevertheless, an obvious alternative catalyst for effective choice in this paradigm is preplay communication. Costless preplay communication markedly increases cooperation in social dilemmas (Deutsch, 1958; Sally, 1995) and can enhance efficiency in coordination games (see, e.g., Cooper, DeJong, Forsythe, & Ross, 1992). A recent meta-analytic review suggests that the beneficial impact of communication on cooperation in social dilemmas is moderated by the type of communication and the size of the group, with large groups and face-to-face communication yielding the best results (Balliet, 2010). It is conceivable that pretask or task-concurrent discussions between players could promote cooperative divergence in the present paradigm in two ways. First, preplay communication may aid coordination and maintenance of cooperative divergence because players could verbalize commitment to cooperate prior to the task. Second, preplay communication may promote the availability of cooperative divergence as an efficient strategy: Two people discussing the task may be more likely to identify the optimal task solution without a prior experience of the game. This notion is in line with findings from related group choice paradigms showing that pretask small group discussions increase task performance relative to that in individual choice settings (Schulze & Newell, 2015).

Another potential method to facilitate efficient choice could involve directly manipulating the availability of cooperative divergence as an appropriate response strategy—for instance, by providing a strategy hint to one or both players. A similar approach has been taken to rebalance asymmetric strategy availability in related individual choice paradigms. For instance, Newell et al. (2013) gave participants a description of applicable strategies before they completed a binary choice task. The availability of probability maximizing (the superior strategy) was thus equated with that of probability matching (the inferior, but potentially intuitively more available, strategy). They found that the availability of probability maximizing increased the use of this strategy early in the task (see also Koehler & James, 2010). The availability of cooperative divergence might also be influenced more indirectly—for instance, by manipulating the choices of one player in a computerized version of this paradigm. Having a computerized second player signal cooperative divergence with varying degrees of patience could shed light on the determinants of strategy availability, coordination, and maintenance for single participants. Similar repeated-choice paradigms involving computerized opponents have found choice to be highly sensitive to the behavior of these competitors (see, e.g., Schulze & Newell, 2015). Thus, introducing virtual competitors within the tasks discussed here offers promising avenues for the study of human interactive choice under uncertainty.

To conclude, giving people an opportunity to exploit an uncertain environment by cooperating with another person made a simple game against nature more difficult. Alleviating this added complexity by aligning the task with competitive earning goals, however, quickly established efficient choice via cooperation—even when communication was severely limited. These results illustrate how individual choice under risk and uncertainty can benefit from a social perspective and link the concepts of individual risky choice and strategic decision making under both environmental and social uncertainty.

References

Arkes, H. R., Dawes, R. M., & Christensen, C. (1986). Factors influencing the use of a decision rule in a probabilistic task. *Organizational Behavior*

and Human Decision Processes, 37, 93–110. http://dx.doi.org/10.1016/0749-5978(86)90046-4

Avrahami, J., Güth, W., & Kareev, Y. (2005). Games of competition in a stochastic environment. *Theory and Decision, 59,* 255–294. http://dx.doi.org/10.1007/s11238-005-2883-y

Balliet, D. (2010). Communication and cooperation in social dilemmas: A meta-analytic review. *The Journal of Conflict Resolution, 54,* 39–57. http://dx.doi.org/10.1177/0022002709352443

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16,* 215–233. http://dx.doi.org/10.1002/bdm.443

Bereby-Meyer, Y., & Roth, A. E. (2006). The speed of learning in noisy games: Partial reinforcement and the sustainability of cooperation. *The American Economic Review, 96,* 1029–1042. http://dx.doi.org/10.1257/aer.96.4.1029

Berger, L. A., & Hershey, J. C. (1994). Moral hazard, risk seeking, and free riding. *Journal of Risk and Uncertainty, 9,* 173–186. http://dx.doi.org/10.1007/BF01064184

Bornstein, G. (2003). Intergroup conflict: Individual, group, and collective interests. *Personality and Social Psychology Review, 7,* 129–145. http://dx.doi.org/10.1207/S15327957PSPR0702_129-145

Bornstein, G., Budescu, D., & Zamir, S. (1997). Cooperation in intergroup, N-person, and two-person games of chicken. *Journal of Conflict Resolution, 41,* 384–406. http://dx.doi.org/10.1177/0022002797041003003

Bornstein, G., Gneezy, U., & Nagel, R. (2002). The effect of intergroup competition on group coordination: An experimental study. *Games and Economic Behavior, 41,* 1–25. http://dx.doi.org/10.1016/S0899-8256(02)00012-X

Browning, L., & Colman, A. M. (2004). Evolution of coordinated alternating reciprocity in repeated dyadic games. *Journal of Theoretical Biology, 229,* 549–557. http://dx.doi.org/10.1016/j.jtbi.2004.04.032

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.

Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012). Measuring risk literacy: The Berlin Numeracy Test. *Judgment and Decision Making, 7,* 25–47.

Colman, A. M., & Browning, L. (2009). Evolution of cooperative turn-taking. *Evolutionary Ecology Research, 11,* 949–963.

Cooper, R., DeJong, D. V., Forsythe, R., & Ross, T. W. (1992). Communication in coordination games. *The Quarterly Journal of Economics, 107,* 739–771. http://dx.doi.org/10.2307/2118488

Crowley, P. H. (2001). Dangerous games and the emergence of social structure: Evolving memory-based strategies for the generalized hawk–dove game. *Behavioral Ecology, 12,* 753–760. http://dx.doi.org/10.1093/beheco/12.6.753

Deutsch, M. (1958). Trust and suspicion. *The Journal of Conflict Resolution, 2,* 265–279. http://dx.doi.org/10.1177/002200275800200401

Erev, I., & Rapoport, A. (1998). Coordination, "magic," and reinforcement learning in a market entry game. *Games and Economic Behavior, 23,* 146–175. http://dx.doi.org/10.1006/game.1997.0619

Estes, W. K. (1964). Probability learning. In A. W. Melton (Ed.), *Categories of human learning* (pp. 89–128). New York, NY: Academic Press.

Fantino, E., & Esfandiari, A. (2002). Probability matching: Encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale, 56,* 58–63. http://dx.doi.org/10.1037/h0087385

Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives, 19,* 25–42. http://dx.doi.org/10.1257/089533005775196732

Fretwell, S. D. (1972). *Populations in a seasonal environment*. Princeton, NJ: Princeton University Press.

Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition, 109,* 416–422. http://dx.doi.org/10.1016/j.cognition.2008.09.007

Goldstone, R. L., & Ashpole, B. C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review, 11,* 508–514. http://dx.doi.org/10.3758/BF03196603

Goldstone, R. L., & Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science, 1,* 412–438. http://dx.doi.org/10.1111/j.1756-8765.2009.01038.x

Greenhouse, S. W., & Geisser, S. (1959). On methods in the analysis of profile data. *Psychometrika, 24,* 95–112. http://dx.doi.org/10.1007/BF02289823

Herrnstein, R. J., & Prelec, D. (1991). Melioration: A theory of distributed choice. *The Journal of Economic Perspectives, 5,* 137–156. http://dx.doi.org/10.1257/jep.5.3.137

Koehler, D. J., & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition, 38,* 667–676. http://dx.doi.org/10.3758/MC.38.6.667

Kraft, J. R., & Baum, W. M. (2001). Group choice: The ideal free distribution of human social behavior. *Journal of the Experimental Analysis of Behavior, 76,* 21–42. http://dx.doi.org/10.1901/jeab.2001.76-21

Kunreuther, H., Silvasi, G., Bradlow, E. T., & Small, D. (2009). Bayesian analysis of deterministic and stochastic prisoner's dilemma games. *Judgment and Decision Making, 4,* 363–384.

Lau, S. H. P., & Mui, V. L. (2008). Using turn taking to mitigate coordination and conflict problems in the repeated battle of the sexes game. *Theory and Decision, 65,* 153–183. http://dx.doi.org/10.1007/s11238-008-9100-8

Madden, G. J., Peden, B. F., & Yamaguchi, T. (2002). Human group choice: Discrete-trial and free-operant tests of the ideal free distribution. *Journal of the Experimental Analysis of Behavior, 78,* 1–15. http://dx.doi.org/10.1901/jeab.2002.78-1

Morey, R. D., & Rouder, J. N. (2014). *BayesFactor (Version 0.9.9)* [Computer software]. Retrieved from http://bayesfactorpcl.r-forge.r-project.org

Nash, J. (1951). Non-cooperative games. *The Annals of Mathematics, 54,* 286–295. http://dx.doi.org/10.2307/1969529

Newell, B. R., Koehler, D. J., James, G., Rakow, T., & van Ravenzwaaij, D. (2013). Probability matching in risky choice: The interplay of feedback and strategy availability. *Memory & Cognition, 41,* 329–338. http://dx.doi.org/10.3758/s13421-012-0268-3

Newell, B. R., & Rakow, T. (2007). The role of experience in decisions from description. *Psychonomic Bulletin & Review, 14,* 1133–1139. http://dx.doi.org/10.3758/BF03193102

Peterson, C. R., & Ulehla, Z. J. (1965). Sequential patterns and maximizing. *Journal of Experimental Psychology, 69,* 1–4. http://dx.doi.org/10.1037/h0021597

Phillips, N. D., Hertwig, R., Kareev, Y., & Avrahami, J. (2014). Rivals in the dark: How competition influences search in decisions under uncertainty. *Cognition, 133,* 104–119. http://dx.doi.org/10.1016/j.cognition.2014.06.006

Rapoport, A., Seale, D. A., & Ordónez, L. (2002). Tacit coordination in choice between certain outcomes in endogenously determined lotteries. *Journal of Risk and Uncertainty, 25,* 21–45. http://dx.doi.org/10.1023/A:1016315329973

Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., & Zamir, S. (1991). Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. *The American Economic Review, 81,* 1068–1095.

Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology, 56,* 356–374. http://dx.doi.org/10.1016/j.jmp.2012.08.001

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review, 16,* 225–237. http://dx.doi.org/10.3758/PBR.16.2.225

Sally, D. (1995). Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992. *Rationality and Society, 7,* 58–92. http://dx.doi.org/10.1177/1043463195007001004

Schulze, C., & Newell, B. R. (in press). More heads choose better than one: Group decision making can eliminate probability matching. *Psychonomic Bulletin & Review*.

Schulze, C., van Ravenzwaaij, D., & Newell, B. R. (2015). Of matchers and maximizers: How competition shapes choice under risk and uncertainty. *Cognitive Psychology, 78,* 78–98. http://dx.doi.org/10.1016/j.cogpsych.2015.03.002

Shafran, A. P. (2012). Learning in games with risky payoffs. *Games and Economic Behavior, 75,* 354–371. http://dx.doi.org/10.1016/j.geb.2011.09.001

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making, 15,* 233–250. http://dx.doi.org/10.1002/bdm.413

Smith, J. M., & Price, G. R. (1973). The logic of animal conflict. *Nature, 246,* 15–18. http://dx.doi.org/10.1038/246015a0

Sokolowski, M. B. C., Tonneau, F., & Freixa i Baqué, E. (1999). The ideal free distribution in humans: An experimental test. *Psychonomic Bulletin & Review, 6,* 157–161. http://dx.doi.org/10.3758/BF03210824

Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys, 14,* 101–118. http://dx.doi.org/10.1111/1467-6419.00106

Wetzels, R., & Wagenmakers, E.-J. (2012). A default Bayesian hypothesis test for correlations and partial correlations. *Psychonomic Bulletin & Review, 19,* 1057–1064. http://dx.doi.org/10.3758/s13423-012-0295-x

Received January 22, 2015
Revision received May 25, 2015
Accepted June 22, 2015 ■