Reference dependence, cooperation, and coordination in games

Mark Schneider[*]   Jonathan W. Leland[†]

Abstract

The problems of how self-interested players can cooperate despite incentives to defect, and how players can coordinate despite the presence of multiple equilibria, are among the oldest and most fundamental in game theory. In this report, we demonstrate that a plausible and even natural specification of the reference outcome in a game simultaneously predicts systematic cooperation and defection in the Prisoner's Dilemma, as well as equilibrium selection and out-of-equilibrium play in coordination games. The predictions hold even if players are purely self-interested, there are no salient labels, the game is played only once, and there is no communication of any kind. Furthermore, the predictions are unique, as opposed to the multiplicity of equilibria in the infinitely repeated Prisoner's Dilemma and in coordination games. We apply experimental results to test the predictions of the model.

Keywords: prisoner's dilemma, coordination games, reference-dependent preferences.

1 Introduction

Two of the most fundamental challenges in the social sciences concern how groups, from dyads to firms to nations, achieve cooperation and coordination. The Prisoner's Dilemma, a situation in which two players each choose between a socially desirable (cooperative) strategy and an alternative strategy more aligned with their material self-interests, epitomizes the difficulty of achieving cooperation. It has served as a paradigm for understanding a broad range of social, economic and political phenomena from the pricing decisions of firms to arms races during the cold war (Rapoport, 1974). Games like Rousseau's Stag-hunt, exemplify the difficulty of achieving coordination in settings with multiple Nash equilibria. In this report, we show that a simple heuristic in which players have reference-dependent preferences predicts cooperation and defection in the Prisoner's Dilemma and identifies when and how coordination problems will be solved. Previous explanations of cooperation in the Prisoner's dilemma rely on the game being repeated (Mas-Colell, Whinston & Green, 1995) or on players having preferences for altruism or reciprocity (Camerer & Fehr 2006). Common explanations for coordinated behavior rely on communication or salient labels (Schelling 1960). In contrast, our results obtain with purely self-interested players in single-shot games, without salient labels or communication.

Classical game theory assumes a great deal of sophistication on the part of players. Such players reason using backward induction and common knowledge, and they assume that their opponents are as sophisticated as themselves. This sophistication in strategic thinking seems at odds with the Bayesian view of choice under uncertainty that permeates classical decision theory, in which agents assign probabilistic beliefs over their opponent's strategies and maximize their expected payoff given their subjective information. This Bayesian approach to game theory has been advocated prominently by Aumann (1987) and others. Still the assumption inherent in the Bayesian approach that agents have uniquely defined subjective beliefs which satisfy the laws of probability theory has also been viewed as a strong assumption (e.g., Gilboa, 2009).

In this paper, we take an approach rooted in the judgment and decision-making literature (e.g., Payne, Bettman, & Johnson, 1993; Gigerenzer & Todd, 1999) by assuming agents use simple heuristics when formulating strategy choice, without requiring the aid of either strategic or probabilistic sophistication.

We note that one strong property of Nash equilibrium is that players focus only on unilateral deviations, assuming others are playing their best responses. But why should a reasonable player not be permitted to consider, or at least entertain the possibility, of bilateral, or even multilateral deviations? The types of games we will focus on in this paper are primarily coordination games in which players can either choose a safe strategy or a riskier strategy that offers the possibility of gain from mutual coordination as well as the possibility of loss if coordination is not achieved.

*University of Connecticut, School of Business, 2100 Hillside Road Unit 1041. Storrs, CT 06269-1041. E-mail: markschneider@aya.yale.edu.

†National Science Foundation, Arlington VA. Email: JLeland@nsf.gov. This work was done while serving as Senior Fellow in the Consumer Financial Protection Bureau's Office of Research. The views expressed are those of the authors and do not necessarily represent those of the Consumer Financial Protection Bureau, the National Science Foundation, or the United States Government.

In this context, and when defining a "safe" strategy without requiring the existence of probabilistic beliefs, we return to a strategy from classical game theory in which an agent seeks to maximize his minimum gain. We refer to this as the maximin strategy and the minimum payoff in this strategy is called the player's maximin payoff. The maximin strategy is a natural analog to a riskless option since it guarantees a player at least as much as can be obtained with certainty, independent of the other players' actions. We propose that, in games involving coordination, a player's payoff at the maximin strategy profile serves as a natural reference point from which to evaluate gains or losses arising from success or failure to coordinate.

Most game theoretic treatments of issues of cooperation and coordination assume players choose strategies in accordance with the expected utility hypothesis. However, researchers have long questioned the descriptive validity of expected utility theory (e.g., Kahneman & Tversky 1979; Tversky & Kahneman 1981). One criticism concerns the assumption that final wealth levels are the carriers of value. Instead, abundant experimental evidence suggests that the carriers of value are gains and losses relative to some reference point. This observation has spawned an entire class of "reference-dependent" utility models, most notably Kahneman and Tversky's Prospect theory. Given the success of such models at explaining risky decisions, it seems natural to examine the consequences of reference dependence in strategic settings. While prior research has proposed that reference dependence can be applied to game theory (Shalev, 2000), no general specification for the reference point has been provided. For this purpose, the classical maximin strategy provides a natural definition of the reference outcome and default strategy in a game. Recent work shows that many of the classical expected utility paradoxes can be resolved in a model that is linear in probabilities for any choice set when the maximin payoff serves as an agent's reference point (Schneider, Day & Garfinkel, 2014). Here we consider the implications of a maximin reference strategy when applied to games.

For simplicity and illustration purposes, our analysis focuses on 2x2 normal form games. Our argument is general in the sense that the maximin strategy can be identified in any game, although in general it need not be unique. We introduce a heuristic for predicting when a player will play or deviate from the maximin strategy. In particular, we consider agents who anchor on the maximin strategy profile (i.e., the strategy profile where all players play the maximin strategy), and then decide whether to deviate from that strategy according to the following criterion (stated for Player 1, but analogous for Player 2):

1. If the gain to Player 1 from bilateral deviation exceeds his loss from unilateral deviation, and Player 2 also benefits from bilateral deviation, then deviate from the maximin strategy.

2. If the gain to Player 1 from unilateral deviation exceeds his loss from bilateral deviation, and Player 2 is worse off from bilateral deviation, then deviate from the maximin strategy.

3. Otherwise, play the maximin strategy.

Note that Player 1 strives for bilateral deviation from the maximin strategy profile only if both players benefit from that deviation. In contrast, Player 1 strives for unilateral deviation from the maximin strategy profile only when such deviation benefits Player 1, but a simultaneous deviation is harmful, and therefore unlikely, for Player 2.

We can formalize the heuristic as follows: In a two player game, denote Player i's payoff at strategy profile (s_i, s_j) by $x_i(s_i, s_j)$. Let (m_1, m_2) denote the strategy profile when both players play their maximin strategies. For a strategy profile (s_1, s_2) let $x_1^*(s_1, s_2)$ denote Player 1's gain from bilateral deviation from the maximin strategies (i.e., if both players deviate from (m_1, m_2) to strategy profile (s_1, s_2)). That is, $x_1^*(s_1, s_2) = max[x_1(s_1, s_2) - x_1(m_1, m_2), 0]$. Denote Player 1's loss from bilateral deviation by $\underline{x}_1(s_1, s_2) = max[x_1(m_1, m_2) - x_1(s_1, s_2), 0]$ with analogous notation for Player 2. In addition, denote Player 1's gain or loss from unilateral deviation by $x_1^*(s_1, m_2)$ or $\underline{x}_1(s_1, m_2)$ respectively. The *reference-dependent maximin criterion* (stated for Player 1) can be formalized as follows:

Reference-Dependent Maximin (RDM) Criterion:[1]

1. If $x_1^*(s_1, s_2) > \underline{x}_1(s_1, m_2)$ and $x_2(s_1, s_2) > x_2(m_1, m_2)$, play s_1.
2. If $x_1^*(s_1, m_2) > \underline{x}_1(s_1, s_2)$ and $x_2(s_1, s_2) < x_2(m_1, m_2)$, play s_1.
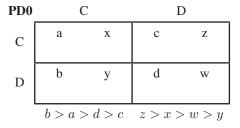3. Otherwise, play m_1.

We thus consider a decision rule in which a player deviates from the maximin strategy profile when it is advantageous to deviate, given the possible gains and losses from deviation and the incentives of the other player. The key assumptions of the RDM criterion are that players anchor on the maximin strategy profile, that gains and losses are measured relative to one's payoff at this strategy profile (which will always be the maximin payoff in any Prisoner's Dilemma or Stag-Hunt game, but may not be in other games), and that players consider both gains and losses due to unilateral deviation or bilateral deviation from the maximin strategy profile and whether their opponent is better or worse off in the event that both players deviate (and thus whether their opponent has an incentive to deviate to that outcome, relative to his payoff at the maximin strategy profile).

We apply the RDM criterion to experimental results from classic 2x2 games in the sections to follow.[2] Let x_i^* denote

---

[1]This strategy may be more broadly defined in terms of utility differences and may be further generalized to account for loss aversion.

[2]The data on the prisoner dilemma games are from Holt and Capra (2000). The data for the Stag-Hunt game are from Leland (2013). The data

Figure 1: The prisoner's dilemma. Here and in Figures 2–4, the rows are Player 1's choices, the columns, Player 2's. Each cell shows payoffs for Players 1 and 2, respectively.

$$b > a > d > c \quad z > x > w > y$$

the highest possible payoff to Player i from deviating from the maximin strategy, and let z_i denote the highest payoff to Player i from playing either strategy. We consider games where payoffs are scaled such that $x_i^* > cz_i$ for c in $(0, 1)$, where c is chosen so that x_i^* and z_i are not very different in magnitude.

2   Reference dependence in the prisoner's dilemma

We now consider the implications of the reference-dependent maximin (RDM) criterion in the context of the Prisoner's Dilemma. For the general game PD0 in Figure 1, and in all subsequent games, the left-most payoff in each cell corresponds to Player 1's payoff. In the experimental games in subsequent figures, Nash equilibria are highlighted in yellow in the bottom half of a cell, and the predictions of the RDM criterion are highlighted in blue in the top half of a cell. The modal outcome of the experiment is displayed in bold. In game PD0, each player chooses between cooperate (C) and defect (D). The game PD0 is a Prisoner's Dilemma if $b > a > d > c$ and $z > x > w > y$.

Payoffs d and w are the maximin payoffs for Players 1 and 2, respectively. If Player 1 follows the RDM criterion, the gain from switching from D to C, $a-d$, is compared with the potential loss from switching ($d-c$). Under the RDM criterion, Player 1 will cooperate in the Prisoner's Dilemma if and only if the possible gain from cooperating (relative to the maximin payoff) exceeds the possible loss from cooperating and $x > w$ (i.e., Player 2 benefits from bilateral deviation from the maximin strategy profile). More formally, we have the following:

Proposition 1: *If Player 1 and Player 2 each follow the RDM criterion then in a Prisoner's Dilemma (Game PD0) both players cooperate if and only if $d < 0.5(a + c)$, $w < 0.5(x + y)$.*

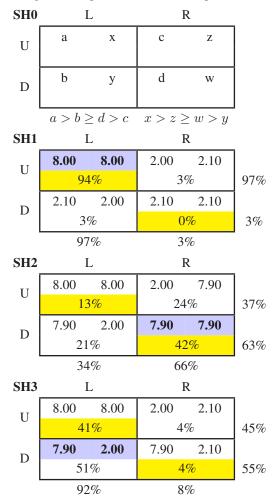Figure 2: Experimental test of Proposition 1.

Note that the RDM criterion also requires $x > w$ and $a > d$, but these conditions are already given as part of the definition of a Prisoner's Dilemma. As noted, the observation that players sometimes cooperate in the Prisoner's dilemma has been a theoretical challenge. The theory of repeated games can explain cooperation in the infinitely repeated prisoner's dilemma. However, in repeated games the same theory permits too many equilibria to reliably predict which strategy profile will be played and when cooperation will be observed. In addition, the theory of repeated games cannot explain the observation that players sometimes cooperate even when the game is played only once. Cooperation in the one-shot prisoner's dilemma can be explained by players with other-regarding preferences (Camerer & Fehr, 2006, Fehr & Schmidt, 1999). In contrast, Proposition 1 predicts that cooperation may arise even in a one-shot prisoner's dilemma with entirely self-interested players if the players use reference-dependent decision rules of the kind embodied in the RDM criterion. To test the necessary and sufficient conditions for cooperation, predicted in Proposition 1, first consider game PD1 reported in Holt and Capra (2000) and displayed in Figure 2.

Let the maximin payoff serve as a reference point for Player 1 and Player 2. Game PD1 tests the sufficient condition for equilibrium play in Proposition 1. In PD1, $a+c = 3$ and $d = 2$. Also, $x + y = 3$ and $w = 2$. By Proposition 1, the RDM criterion predicts both players to play D in PD1. In the experiment from Holt and Capra (2000) only 17% of players played C in PD1.

Game PD2 tests the necessary condition for equilibrium play in Proposition 1. In this game, $a + c = 8$ and $d = 2$. Also, $x + y = 8$ and $w = 2$. By Proposition 1, the RDM criterion predicts both players to play C in PD2. In the experiment, Holt and Capra (2000) found that 58% of players now played C in this game.

Figure 3: Experimental test of Proposition 2.

payoffs over larger riskier payoffs. It is widely recognized that play sometimes results in a payoff-dominant equilibrium, and sometimes in a security-minded equilibrium, but it is not clear how to systematically predict when each will be played. Formally, we have the following proposition:

Proposition 2: *If Players 1 and 2 follow the RDM criterion in a stag-hunt coordination game:*

1. *Both players will coordinate on the payoff-dominant Nash equilibrium if and only if $d < 0.5(a + c)$ and $w < 0.5(x + y)$.*
2. *Both players will coordinate on the security-minded Nash equilibrium if and only if $d > 0.5(a + c)$ and $w > 0.5(x + y)$.*
3. *Players' choices will produce one of the non-equilibrium outcomes if neither the conditions in* (1) *nor* (2) *hold.*

As before, the RDM criterion also requires $x > w$ and $a > d$ for both players to coordinate on the payoff dominant equilibrium, but these conditions are given as part of the structure of the Stag Hunt game. In Figure 3, games SH1, SH2, and SH3 test necessary and sufficient conditions for equilibrium play in Proposition 2. These games were played by experimental subjects (Leland, 2013). Of the two pure strategy Nash equilibria, UL is always payoff dominant, and DR is always security-minded. Equilibrium refinements predict that a Nash equilibrium will be played.
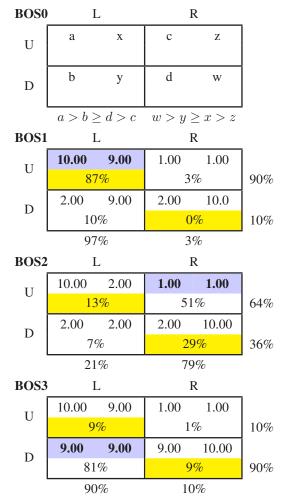
In SH1, $a + c = 10$ and $d = 2.10$. Also, $x + y = 10$ and $w = 2.10$. Thus, by Proposition 2, the payoff-dominant equilibrium, UL, should be played. As predicted by RDM, the experiment in Leland (2013) found that UL was the modal outcome in SH1, played 94% of the time.

In SH2, $a + c = 10$ and $d = 7.90$. Also, $x + y = 10$ and $w = 7.90$. Thus, by Proposition 2, the security-minded equilibrium, DR, should be played. As predicted by RDM, Leland (2013) found that DR was the modal outcome in SH2, played 42% of the time (No other outcome was played more than 25% of the time).

In SH3, $a + c = 10$, and $d = 7.90$. Also, $x + y = 10$ and $w = 2.10$. Thus, Proposition 2 predicts a non-equilibrium outcome to be played. More specifically, RDM predicts the particular disequilibrium outcome DL to be played. As predicted, DL was the modal outcome in SH3, played 51% of the time.

3 Reference-dependence in coordination games

We next consider implications of the RDM criterion for coordination games. Since Schelling (1960), coordination games have posed a fundamental challenge for game theory because it is not clear how to uniquely predict an outcome if there are multiple Nash equilibria.

3.1 The stag-hunt

Consider the stag-hunt coordination games in Figure 3. The general game SH0 is a stag-hunt coordination game if the payoffs for each player satisfy the inequalities specified in the figure. As before, d and w are the maximin payoffs for Players 1 and 2, respectively. This game has two pure strategy Nash equilibria, UL and DR. UL is the payoff dominant equilibrium. We refer to DR as the *security-minded equilibrium*, since it reflects a preference for smaller guaranteed

3.2 Battle of the sexes

A second classic coordination game is the battle of the sexes, illustrated in Figure 4. Game BOS0 is a generic battle of the sexes game if the payoffs for each player satisfy the inequalities specified in the figure. The game has two pure strategy equilibria, one which favors Player 1 (UL) and the other which favors Player 2 (DR). We refer to UL as the P1-preferred equilibrium and DR as the P2-preferred equilib-

Figure 4: Experimental test of Proposition 3.

BOS0

	L		R	
U	a	x	c	z
D	b	y	d	w

$$a > b \geq d > c \quad w > y \geq x > z$$

BOS1

	L		R	
U	**10.00**	**9.00**	1.00	1.00
	87%		3%		90%
D	2.00	9.00	2.00	10.0
	10%		0%		10%
	97%		3%	

BOS2

	L		R	
U	10.00	2.00	**1.00**	**1.00**
	13%		51%		64%
D	2.00	2.00	2.00	10.00
	7%		29%		36%
	21%		79%	

BOS3

	L		R	
U	10.00	9.00	1.00	1.00
	9%		1%		10%
D	**9.00**	**9.00**	9.00	10.00
	81%		9%		90%
	90%		10%	

rium. The RDM criterion provides a way to systematically predict when we will observe players coordinating on UL or DR, as well as when a non-equilibrium outcome will result.

Note first that the maximin strategy profile is DL and thus payoffs b and y serve as reference points for Players 1 and 2, respectively. Under the RDM criterion, we have the following result:

Proposition 3: *If Players 1 and 2 follow the RDM criterion in a battle-of-the-sexes game:*

1. *Both players will coordinate on the P1-preferred equilibrium if and only if $b < 0.5(a + c)$ and $y > 0.5(w + z)$.*

2. *Both players will coordinate on the P2-preferred equilibrium if and only if $b > 0.5(a + c)$ and $y < 0.5(w + z)$.*

3. *Players' choices will produce one of the non-equilibrium outcomes if neither the conditions in (1) nor (2) hold.*

Note that the RDM criterion also requires $b > c$ and $y > z$, but these are given as part of the structure of the battle-of-the-sexes game. Consider three instantiations of a battle of the sexes coordination game in Figure 4. Games BOS1, BOS2, and BOS3 test necessary and sufficient conditions for equilibrium play in Proposition 3. The games were played by experimental subjects (Leland and Schneider, 2014).

In BOS1, $a + c = 11$ and $b = 2$. Also, $w + z = 11$ and $y = 9$. By Proposition 3, the P1-preferred equilibrium, UL, should be played. As predicted by the RDM criterion, UL was the modal outcome in BOS1, played 87% of the time.

In BOS2, $a + c = 11$ and $b = 2$. Also, $w + z = 11$ and $y = 2$. By Proposition 3, we should observe one of the non-equilibrium outcomes to be played. In particular RDM predicts the disequilibrium UR to be played. As predicted by the RDM criterion, UR was the modal outcome in BOS2, played 51% of the time.

In BOS3, $a + c = 11$ and $b = 9$. Also, $w + z = 11$ and $y = 9$. Proposition 3 predicts a non-equilibrium outcome to be played. In particular, RDM predicts the disequilibrium outcome DL to be played. As predicted, DL was the modal outcome in BOS3, played 81% of the time.

4 Games with more than two strategies

In this section, we briefly illustrate how the RDM criterion may be extended to games with more than two strategies. Denote the set of strategies available for Players 1 and 2 by S^1 and S^2, respectively. We can generalize the RDM criterion to larger games as follows. Define:

$$y := max\{[x_1(s_1, s_2) - x_1(m_1, m_2),\ 0]\,|\,s_2 \text{ in } S^2\}$$
$$-max\{[x_1(m_1, m_2) - x_1(s_1, s_2'),\ 0]\,|\,s_2' \text{ in } S^2\}$$

Then the generalized RDM criterion (for Player 1) can be formalized as follows:

1. If strategy s_1 maximizes y over all strategies s_1 in S^1 (other than m_1), and if $x_2(s_1, s_2) > x_2(m_1, m_2)$, play s_1.

2. Otherwise play m_1.

The first part of Condition (1) states that Player 1 deviates from the maximin strategy to strategy s_1 only if the difference between the maximum possible gain and the maximum possible loss from playing s_1 (relative to her payoff at the maximin strategy profile) is greater than for any other strategy s_1 in S^1 (other than m_1) available to Player 1. The second part of Condition (1) states that Player 2 is better off at that new strategy profile than she would be at the maximin strategy profile (and thus views the change in payoffs as a gain). If no such strategy profile satisfies the properties in Condition (1), then Player 1 plays m_1.

4.1 A minimum effort coordination game

In this section, we illustrate how the generalized RDM criterion can be applied to a coordination game with more than two strategies. Goeree and Holt (2001) consider a coordination game as follows: Two players, P1 and P2, choose "effort" levels, e_1 and e_2 simultaneously. P1 receives $min(e_1, e_2) - ce_1$ where $c < 1$ is a coefficient indicating the cost of effort. Likewise, P2 receives $min(e_1, e_2) - ce_2$. Effort levels are integers in the interval [110, 170]. Any common effort level in this game is a Nash equilibrium and thus it is not clear how to select among the 61 different Nash equilbria, and whether players will be able to coordinate on an equilibrium at all.

In their experimental implementation, Goeree and Holt considered two variants of the game, one in which $c = 0.10$ and the other in which $c = 0.90$. Note that, for any $c > 0$, the maximin strategy is to choose an effort level of 110, which guarantees that player a payoff of $110(1 - c)$. For a given Player i, any other strategy admits the possibility of a lower payoff since any effort level $e_i > 110$ yields payoff $110 - ce_i$, whenever Player j chooses effort level 110.

For the case where $c = 0.10$, under the generalized RDM criterion, Player i deviates from the maximin strategy to the strategy that maximizes the expression in condition 1. The strategy that does so is to choose effort level 170. To see this, note that the maximum possible loss to Player 1 can occur only when $e_2 = 110$. Setting $e_2 = 110$ in computing Player 1's maximum loss from deviating from the maximin strategy, the generalized RDM criterion recommends the strategy for Player 1 that maximizes the expression, $min(e_1, e_2) - 2c(e_1 - 110)$ over all possible strategies for Player 2. For $c = 0.10$, Player 1 is predicted to choose the effort level that maximizes the expression, $min(e_i, e_2) - 0.2e_1$.

For the cases when $e_1 < e_2$ or $e_1 = e_2$, note that $0.8e_1$ is maximized at the highest possible value of e_1. Also note that for a fixed e_1, Player 1 is always worse off when $e_1 > e_2$ compared to when $e_1 < e_2$ or $e_1 = e_2$. Thus, the strategy profile which maximizes Player 1's payoff occurs when $e_1 = e_2 = 170$. Since Player 2 also benefits from deviating to this strategy, the generalized RDM criterion predicts players to choose the highest effort level of 170 when $c = 0.10$.

For the case where $c = 0.90$, the generalized RDM criterion recommends the strategy which maximizes the expression, $min(e_1, e_2) - 1.8e_1 + 88$.

For the cases when $e_1 < e_2$ or $e_1 = e_2$, note that $-0.8e_1$ is maximized at the lowest possible value of e_1. As before, note that for a fixed e_1, Player 1 is always worse off when $e_1 > e_2$ as compared to when $e_1 < e_2$ or $e_1 = e_2$. Thus, the strategy profile that maximizes Player 1's payoff occurs when $e_1 = e_2 = 110$. Hence, due to the high cost of effort, the downside from deviating from the maximin strategy outweighs the upside, and the generalized RDM criterion pre-

Table 1: Game SH3 with inequality-averse payoffs.

	L	R
U	8,8	$2 - 0.1\alpha_1,\ 2.1 - 0.1\beta_2$
D	$7.9 - 5.9\beta_1,\ 2 - 5.9\alpha_2$	$7.9 - 5.8\beta_1,\ 2.1 - 5.8\alpha_2$

dicts players to choose the lowest effort level (110) when $c = 0.90$.

As predicted, for $c = 0.10$, Goeree and Holt observed "*behavior is quite concentrated at the highest effort level of 170; subjects coordinate on the Pareto-dominant outcome. The high effort cost treatment ($c = 0.9$), however, produced a preponderance of efforts at the lowest possible level.*" (p. 1408).

5 Alternative models of behavior in games

We have seen that, at least for situations that admit the possibility of coordination or cooperation, the RDM criterion has some descriptive advantages over the Nash equilibrium. In the previous sections, the predictions made by the RDM criterion are unique, as compared to the multiplicity of Nash equilibria in coordination games. The RDM criterion also predicts experimentally observed out-of-equilibrium play in the stag-hunt and battle of the sexes, as well as when players will systematically cooperate and defect in the prisoner's dilemma.

More recently, a plethora of alternative models have emerged to predict behavior in games. Here we focus on two prominent examples—cognitive hierarchy models (e.g., Stahl & Wilson 1994, Camerer et al., 2002) and models of other-regarding preferences (e.g., Fehr and Schmidt, 1999). Models of boundedly rational behavior such as level-k thinking or cognitive hierarchy models postulate different levels of strategic thinking with higher level players best-responding assuming their opponents are less sophisticated than they are. One of the most successful implementations of this model for coordination games posits players who are level 1 boundedly rational and best respond assuming their co-players are level-0 players who play randomly. However, this model cannot explain cooperation in the prisoner's dilemma since D is always a dominant strategy, and thus C is never a best response under *any* probabilistic beliefs a player might have over his opponent's strategies. In addition, this model does not explain equilibrium selection in the minimum effort coordination game discussed in Section 4.1. If one treats each of his opponent's strategies as equally likely, he will choose an effort level of 164 when $c = 0.10$ and an effort level of 116 when $c = 0.90$, in contrast to the predominant equilibrium behavior at experimentally observed effort levels 170 when $c = 0.10$ and 110 when $c = 0.90$.

The classic model of other-regarding preferences (Fehr & Schmidt, 1999), postulates that in a two player game involving Players i and j, Player i will transform his payoffs according to the utility function

$$U_i(x) = x_i - \alpha_i max[x_j - x_i, 0] - \beta_i max[x_i - x_j, 0]$$

where Fehr and Schmidt (1999) assume that α_i and β_i are non-negative. Players then best respond according to Nash equilibrium strategies, given their transformed payoffs. For game SH3, the transformed payoffs are shown in Table 1.

Note that even under the transformed payoffs, players would never play DL in equilibrium, since Player 1 has an incentive to deviate to strategy U whenever α_1 and β_1 are non-negative. Thus, other-regarding preferences cannot explain the experimentally observed out-of-equilibrium play in game SH3 even when accounting for the two free parameters in the model. In contrast, the RDM criterion makes all of its predictions without any free parameters.

6   Conclusion

In this paper, we have introduced a player's maximin strategy as a plausible reference point in strategic settings as well as a criterion for predicting when a player will play or deviate from her maximin strategy in games involving coordination or cooperation. We have shown that this reference-dependent maximin criterion predicts experimentally observed systematic cooperation, coordination, and out-of-equilibrium play in classic games such as the Prisoner's Dilemma, the Stag Hunt, and the Battle of the Sexes. We also illustrated how the criterion may be generalized to games with more than two strategies and showed that this generalization predicts experimentally observed equilibrium selection in a minimum effort coordination game with 61 different Nash Equilibria. All of these results obtain even if players are purely self-interested, there are no salient labels, the game is played only once, and there is no communication of any kind. Furthermore, these predictions that follow from the RDM criterion are unique, in contrast to the multiplicity of equilibria which arise in coordination problems and infinitely repeated games.

In obtaining our results, the principle of reference dependence has been extended from individual decisions into the domain of strategic interactions via the maximin strategy of classical game theory. The Prisoner's Dilemma and the stag hunt are two of the most widely studied social dilemmas in game theory, as they epitomize the tension between social welfare and individual rationality. The results presented here provide a mechanism for predicting how this tension can be resolved.

References

Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica, 55*, 1–18.

Camerer, C. F., & Fehr, E. (2006). When does "economic man" dominate social behavior? *Science, 311*, 47–52.

Camerer, C., Ho, T., & Chong, J. K. (2002). A cognitive hierarchy theory of one-shot games: Some preliminary results. Manuscript. California Institute of Technology.

Fehr, E., & K. M. Schmidt. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics, 114*, 817–868.

Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Gilboa, I. (2009). *Theory of Decision under Uncertainty*. Econometric Society Monographs. New York: Cambridge University Press.

Goeree, J. K., & Holt, C. A. (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review, 91*, 1402–1422.

Holt, C. A., & Capra, M. (2000). Classroom games: A prisoner's dilemma. *Journal of Economic Education, 31*, 229–236.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–291.

Leland, J. W. (2013). Equilibrium selection, similarity judgments, and the "nothing to gain/nothing to lose" effect. *Journal of Behavioral Decision Making, 26*, 418–428.

Leland, J. W., & Schneider, M. (2014). Salience and strategy choice in 2x2 games. Manuscript.

Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic Theory*. New York: Oxford University Press.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The Adaptive Decision Maker.* New York: Cambridge University Press.

Rapoport, A. (Ed.) (1974). *Game theory as a theory of a conflict resolution* (No. 2). Springer Science and Business Media.

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

Schneider, M., Day, R., & Garfinkel, R. (2014). Endowment-adjusted utility functions and expected utility paradoxes. Manuscript.

Shalev, J. (2000). Loss aversion equilibrium. *International Journal of Game Theory, 29*, 269–287

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

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211*, 453–458.