University of Groningen

The rationale of rationality

Weesie, H.M.; Snijders, C.C.P.; Buskens, V.W.

Published in:
Rationality and Society

DOI:
[10.1177/1043463109103901](https://doi.org/10.1177/1043463109103901)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2009

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Weesie, H. M., Snijders, C. C. P., & Buskens, V. W. (2009). The rationale of rationality. *Rationality and Society*, 21(2), 249-277. <https://doi.org/10.1177/1043463109103901>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Rationality and Society

<http://rss.sagepub.com/>

The Rationale of Rationality

Jeroen Weesie, Chris Snijders and Vincent Buskens

Rationality and Society 2009 21: 249

DOI: 10.1177/1043463109103901

The online version of this article can be found at:

<http://rss.sagepub.com/content/21/2/249>

Published by:

<http://www.sagepublications.com>

Additional services and information for *Rationality and Society* can be found at:

Email Alerts: <http://rss.sagepub.com/cgi/alerts>

Subscriptions: <http://rss.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://rss.sagepub.com/content/21/2/249.refs.html>

THE RATIONALE OF RATIONALITY

Jeroen Weesie, Chris Snijders, and Vincent Buskens

ABSTRACT

Our research starts from the assumption that actors use a single decision theory to guide them on how to behave in all possible one-shot two-person encounters. To address which decision theories perform well, we let 17 theories compete in a large number of randomly selected symmetric 2×2 games. It turns out that the decision theory that optimizes its own payoff under the assumption that the other actor behaves randomly wins by a small margin. Second, we study the 'evolution of rationality.' In a quasi-biological setup where more successful strategies generate more offspring, the decision theory that always plays the behavior that belongs to the risk-dominant Nash equilibrium emerges as the long-term survivor from an initially mixed pool of decision theories. We also confront the decision theories with human experimental data. The decision theory that always aims for the highest possible payoff for itself performs best against humans.

KEYWORDS • decision theories • evolution • one-shot encounters
• rationality

1 Introduction

This article extends the discussion on whether rational behavior can be sustained or emerges as an evolutionary process from a situation in which initially most of the actors are not rational in the game-theoretic sense (e.g. see Güth and Kliemt 1998 on the evolution of rational behavior in trust situations). This question is also one of the main questions that is asked in the field of evolutionary economics. However, evolutionary economics uses evolutionary arguments more as equilibrium selection tools for situations with multiple equilibria. For example, the question is addressed on which equilibrium actors coordinate in a coordination game given that they start with a random decision and that behavior leading to a favorable outcome becomes more likely (e.g. Young 1993). Similarly,

Rationality and Society Copyright © 2009 Sage Publications. Los Angeles, London, New Delhi, Singapore and Washington DC, Vol 21(2): 249–277.
<http://rss.sagepub.com> DOI: 10.1177/1043463109103901

there are many studies on the evolution of cooperation in repeated Prisoner's Dilemmas or collective good problems (see Axelrod 1984; Macy and Flache 1995; Macy 1997).

An important limitation of the studies mentioned above is that they focus on just *one* type of interaction in which actors are involved repeatedly. Either a certain type of behavior in this specific game is reinforced because it provides higher benefits or certain preferences become more abundant in a society because these preferences are related to better outcomes. In one-shot encounters, behavior is just a specific choice in an interaction, while behavior in repeated interactions may include more complicated strategies that are related to experiences in the past (e.g. Tit-for-Tat in Axelrod 1984). This implies that these studies cannot explain the emergence of specific rationality principles that actors might use in a variety of interactions.

Heckathorn (1996) argued that collective good problems are related to a whole range of types of two-person games including Prisoner's Dilemmas, Chicken Games, and Assurance Games. We agree that two-person interactions are likely to take various forms in real-life. In addition, we argue that it is likely that decision principles will survive in an evolutionary context that do not just do well in one type of interaction, but that prevail when confronted with a range of different interactions. Also, Bednar and Page (2007) study emergence of strategies in situations where actors play different games. Their approach is complementary to ours. Although they use a more limited set of games (different versions of coordination games), they let their actors play repeatedly with actors in their local neighborhood. We use more different games but we restrict ourselves by letting people play with strangers all the time.

Formal decision theory has produced a long series of proposals for theories of rational behavior in situations of interdependent choice. Both among decision theory scholars and among applied scholars from a whole range of social sciences, there does not seem to be agreement with respect to the central question of what constitutes 'truly' rational behavior in situations of interdependent choice. In this article, we address this question by evaluating which decision theory survives in a world in which players can use different decision theories. It will not be possible to analyze this abstract and complicated problem in full generality. One important reason is that decision theories, especially those based on Nash equilibrium, are incomplete since many games have multiple equilibria. Harsanyi and Selten's (1988) theory of equilibrium selection has, as far as we know, not been confronted with a competing theory of a similar scope. Especially for many infinitely repeated games, the equilibrium selection problem is still

far from solved. Therefore, we decided for now to consider a restrictive set of games for which equilibrium selection is not problematic and for which multiple equilibrium selection theories have been proposed: *one-shot symmetric 2×2 games with complete information*.

An additional reason for taking this approach is that for these relatively simple games it is conceivable that people use a single decision theory across these games. Or, perhaps more adequately, it is conceivable that evolution will eventually take care of equipping people with a single decision theory for these situations. Therefore, the main question we analyze is: what kind of general decision theory is most successful in a world consisting entirely of one-shot symmetric 2×2 games. We will consider a range of decision theories that include one that chooses randomly one of the options in the 2×2 games, decision theories that play best replies assuming that the other player plays in a probabilistic way, as well as decision theories that play best replies assuming their opponents also try to play best replies. The last category of theories comprises the rational strategies in the game-theoretic sense.

Even for symmetric 2×2 games, a strictly formal analysis of the interactions between decision theories and the selection of an 'optimal' decision theory is a non-trivial matter analytically. For example, integration over the distribution of games might be straightforward for very simple distributions of payoffs such as the uniform distribution. However, assuming other distributions this becomes more complicated. Therefore, we apply Monte Carlo integration to compute expected values over the assumed distribution of games.

We proceed as follows. First, we fix a population of decision theories. Then we let these decision theories play a large number of one-shot symmetric 2×2 games against each other. We compare the average payoff that each decision theory receives in the encounters with all the other theories. This resembles the round robin tournament of Axelrod (1984). This analysis results in a ranking of strategies with respect to how well they perform in this population of strategies. We will see that relatively simple and optimistic strategies are performing quite well under these conditions, just like strategies that are based on Nash equilibrium. Subsequently, we add an evolutionary logic to this process: strategies that are doing well tend to become more frequent, while strategies that are doing badly tend to go extinct. The result of this is that the strategies that did not perform very well in the first analysis gradually disappear. Then, the strategies that performed best against the strategies that have disappeared are disappearing – in an evolutionary setting, lasting success should be in encounters with other successful strategies, not against

		Player 2	
		X	Y
Player 1	X	a, a	c, b
	Y	b, c	d, d

Figure 1. Notation for symmetric 2×2 games

unsuccessful ones. As it turns out, the Nash-based decision theories tend to survive this evolutionary process. Although we realize that this exercise is just a first step in a project on the evolution of decision rules in a dynamic environment, we still wanted to add an empirical twist to this article. Therefore, we end with a short investigation into how the artificial strategies perform against *human subjects* who were asked to make decisions in several 2×2 games. Here, it turns out that human subjects themselves obtain low average payoffs. Moreover, we find that strategies that did not perform very well in the tournament and in the evolutionary process are performing best against human beings. We outline the implications of these findings in the conclusion and discussion.

2 A Set of Decision Theories

Figure 1 shows a symmetric 2×2 game. Because we only consider decision theories that do not depend on the representation of the game, we can assume without loss of generality that $a > d$ in Figure 1.

We also assume that there are no ties in the payoffs. This implies there are 12 different possible payoff orderings, representing 12 different types of games (see Table 1). There are six payoff orderings in which both players have a dominant strategy (games 1, 3, 4, 8, 10, and 11). These games have a unique Nash equilibrium. Three of these six games (games 1, 3, and 4) have an equilibrium in which both players receive the maximal payoff in equilibrium. Two of the games (games 10 and 11) leave the players with the second highest payoff, and in the Prisoner's Dilemma (game 8) the equilibrium provides the players with the second lowest outcome. There are three types of games that have three symmetric equilibria (games 2, 5, and 6): one in which both receive a , one in which both get d , and one mixed equilibrium. The last three types (games 7, 9, and 12) also have three equilibria but two of these are asymmetric (one player gets b while the other gets c), whereas the third equilibrium is a symmetric

Table 1. Possible payoff orderings of symmetric 2×2 games

	<i>Order of payoffs</i>	<i>Equilibria (symmetric)</i>	<i>Name</i>
1	$a > b > c > d$	1 (1)	Unique Nash (a, a)
2	$a > b > d > c$	3 (3)	Assurance game
3	$a > c > b > d$	1 (1)	Unique Nash (a, a)
4	$a > c > d > b$	1 (1)	Unique Nash (a, a)
5	$a > d > b > c$	3 (3)	Assurance-like game
6	$a > d > c > b$	3 (3)	Assurance-like game
7	$b > a > c > d$	3 (1)	Chicken game
8	$b > a > d > c$	1 (1)	Prisoner's Dilemma
9	$b > c > a > d$	3 (1)	Chicken-like game
10	$c > a > b > d$	1 (1)	Unique Nash (a, a)
11	$c > a > d > b$	1 (1)	Unique Nash (a, a)
12	$c > b > a > d$	3 (1)	Chicken-like game (Battle of the sexes)

mixed equilibrium. Table 1 gives an overview in lexicographic ordering of the different games and their properties.

Table 2 briefly describes 17 decision theories that each generates unique predictions for behavior in symmetric 2×2 games. The decision theories are allowed to make randomized decisions. Most of these theories can be adequately defined by the assumption they make about the way in which the focal player (Ego) forms expectations about the behavior of Alter, supplemented with the assumption that players choose optimally given their expectations. In the first and simplest class of decision theories, we find theories where the expectation of the behavior of Alter is independent of the payoffs of Alter. For instance, the *Principle of Insufficient Reason* (PIR) assumes that both Alter's choices are equally likely, while *MaxiMin-pure* (MMPURE) supposes that Alter will try to minimize the payoffs of Ego ('people are out to get me', see Luce and Raiffa 1957: 284–286 for a historical account starting with Jakob Bernoulli). This implies that given what Ego plays, Ego assumes that Alter will play the strategy that corresponds with Ego's worst payoff. MMPURE plays the best move in *pure strategies* against this expectation. MAXIMAX expects that Alter will try to maximize the payoffs of Ego. We can conceive of these strategies as elements of a larger set of strategies $H(\alpha)$, where the 'H' stems from 'Hurwicz'. $H(\alpha)$ calculates the expected value of choosing X under the assumption that the opponent will choose the move that is best for Ego with probability α ($0 \leq \alpha \leq 1$) and the other move with probability $1 - \alpha$. Similarly, an $H(\alpha)$ player calculates the expected value of choosing Y . He then chooses the move (X or Y) with

Table 2. Four classes of decision theories for one-shot symmetric 2×2 games

1. Hurwicz: optimal behavior against simple expectations about Alter (5 strategies)

$H(\alpha)$	<i>Hurwicz</i> calculates the expected value of choosing X under the assumption that Alter will choose the move that is best for Hurwicz with probability α ($0 \leq \alpha \leq 1$) and the other move with probability $1 - \alpha$. Similarly, Hurwicz calculates the expected value of choosing Y . Hurwicz then chooses the move (X or Y) with the highest expected value. We include $H(0.25)$ and $H(0.75)$ in addition to the following three special cases.
$MMPURE$	<i>MaxiMin-pure</i> is equivalent to $H(0)$. It assumes that, no matter what, Alter will play the move that minimizes the Ego's payoff. $MMPURE$ is a best reply in pure strategies given the 'pessimistic' expectation that the worst will always happen.
PIR	<i>Principle of Insufficient Reason</i> is equivalent to $H(0.5)$. It chooses the move that is optimal given the assumption that Alter will play X and Y with equal probability. PIR is the best reply against $RANDOM$.
$MAXIMAX$	is equivalent to $H(1)$. It plays the best reply under the optimistic assumption that Alter will always choose the move that maximizes Ego's payoff.

2. Best replies against the Hurwicz strategies (5 strategies)

$bH(\alpha)$	assumes that Alter plays $H(\alpha)$, and plays the best reply against it. We include $bH(0.25)$ and $bH(0.75)$.
$bMMPURE$	assumes that Alter plays $MMPURE$, and plays the best reply against it.
$bPIR$	assumes that Alter plays PIR , and plays the best reply against it.
$bMAXIMAX$	assumes that Alter plays $MAXIMAX$, and plays the best reply against it.

3. Strategies using (refinements of) Nash equilibrium (2 strategies)

$PAYOFF$	selects the symmetric Nash equilibrium with the highest equilibrium payoff.
$RISK$	selects the risk-dominant symmetric Nash equilibrium.

4. Miscellaneous other strategies (5 strategies)

$MMMIXED$	<i>MaxiMin-Mixed</i> plays the (possibly mixed) equilibrium strategy for the game where the payoffs for Alter are replaced by minus the payoffs of Ego.
$SAVAGE$	plays the (possibly mixed) equilibrium strategy for the game where the payoffs R_{ij} for Ego are defined as $R_{ij} = M_{ij} - \max_h M_{hj}$, and the payoffs for Alter equal to minus the payoffs of Ego.
$PARETO$	assumes that both Ego and Alter play (X, Y) with probability $(p, 1 - p)$, and then chooses the value of p ($0 < p < 1$) that maximizes the expected payoff.
$NATURAL$	plays a dominant strategy if available. Otherwise it will play a Pareto-efficient equilibrium in pure strategies if it exists. If that does not exist, it applies the $MMPURE$ criterion.
$RANDOM$	plays X and Y with equal probability (0.5).

the highest expected value. The α parameter can be understood as an 'optimism index' (Luce and Raiffa 1957: 282–284): the larger α is, the more positive the expectations about the behavior of Alter. PIR equals $H(0.5)$, $MMPURE$ equals $H(0)$, and $MAXIMAX$ equals $H(1)$. We also included $H(0.25)$ and $H(0.75)$.

The second class of decision theories consists of the best replies against the five aforementioned $H(\alpha)$ strategies. The analysis in Appendix A shows that we need not worry about an infinite cycle of 'best replies against best replies against best replies,' because in fact the best replies against the best replies against the $H(\alpha)$ strategies, are the $H(\alpha)$ strategies themselves. Note that if we were to find out that someone plays according to PIR, this gives us no clue as to how 'deep' this person is actually thinking. Ego might be choosing PIR because he has no clue what Alter will do, or Ego might be playing a best reply against a best reply to PIR (which is PIR as well).

The third class of decision theories consists of refinements of the Nash equilibrium concept. We chose two relatively standard refinements that prescribe which equilibrium should be played if there are more. They differ only in which equilibrium is selected in case a game has more than one symmetric equilibrium. Payoff dominance (PAYOFF) as well as risk dominance (RISK) first calculates the Nash equilibria. When there is only one symmetric equilibrium (mixed or pure) then this equilibrium is played. When there are three (two is not possible), PAYOFF chooses the one that is *payoff dominant* (Harsanyi and Selten 1988: 80–82), while RISK chooses the one that is *risk dominant*, i.e., RISK plays the same as PIR in these cases, namely a best reply against the assumption that the other actor plays both options with equal probability (see Harsanyi and Selten 1988: 86–88; Güth 1992: 191–205).

The fourth and last class consists of decision theories not easily classified elsewhere. Maximin-mixed (MMMIXED) is the 'classic' maximin strategy where Ego plays the best reply against the expectation that Alter is out to get him. This corresponds to Ego playing the (possibly mixed) equilibrium in the zero-sum game that arises when the payoffs to Alter are replaced by minus the payoffs to Ego.

Savage's maximin-regret criterion is an interesting decision principle that combines a 'traditional' decision theory (MMPURE) with a *transformation* of the payoff structure. The transformed payoff structure – in the terminology of Kelley and Thibaut (1978), the 'effective matrix' – drives decisions, instead of just the 'objective' payoffs in the matrix. Savage's decision principle focuses on regret as a motivational force. It plays the (possibly mixed) equilibrium strategy for the game where the payoffs R_{ij} for Ego are defined as $R_{ij} = M_{ij} - \max_h M_{hj}$, and the payoffs for Alter equal minus the payoffs of Ego.¹ The payoffs reflect regret in the sense that they model that Ego feels worse when he could have had a higher payoff if he had chosen the other move. SAVAGE can be seen as utility maximizing behavior given that Ego assumes that the opponent is maximizing Ego's regret.

PARETO, the symmetric Pareto-efficient solution, is a cooperative decision theory. PARETO can be seen as the decision theoretic analogue of the

Kantian imperative: 'act as you want others to act' (cf. Hegselmann 1988). It assumes that the resulting behavior in any of the games is symmetric. That is, it assumes that Ego and Alter will both play (X, Y) with probability $(p, 1 - p)$. It then calculates for which p this gives the highest expected utility. Note that this implies that PARETO need *not* play dominant strategies – it chooses to cooperate in a Prisoner's Dilemma – and thus ignores incentives for individual improvement. In this category, we also include Rapoport's NATURAL solution (Rapoport and Guyer 1966), a boundedly rational version of Nash equilibrium in the sense that it plays pure strategy equilibria when they exist but does not allow for mixed strategies. Whenever there exists only one symmetric equilibrium in pure strategies, NATURAL plays the strategy related to this equilibrium. NATURAL plays the strategy related to the payoff dominant equilibrium if there are three symmetric equilibria. Otherwise, NATURAL mimics MMPURE. Finally, we added RANDOM as a decision theory that randomizes with probability 0.5 over both alternatives.

Whereas the first and the second class of decision theories are each other's best reply, and the third class of theories are their own best reply, we decided not to include the best replies to the fourth class strategies. As we show in Appendix A, it is actually possible to construct a 'closed' set of strategies for our set-up such that for each decision theory its best reply is also included. However, this would imply adding several theories (the best reply to the best reply to the best reply to SAVAGE, for instance) that are not very appealing intuitively. Still, we have experimented with adding these theories and found no results that vary substantially from the results we present. This provides some confidence that the results are to some extent robust for variations in the set of decision theories used. In addition, because in this set-up with best replies the most promising opponents to outperform another decision theory are included, it is unlikely that one can develop a new theory that harms the winners of the analyses in the paragraphs below.

Table 3 describes the behavior of each decision theory in each of the 12 types of games. The cell entries indicate the probability that a decision theory plays X in a given game (cf. Figure 1).

Almost without loss of generality we can assume that $a > d$. The definitions of the decision theories of Table 2 are incomplete in cases of certain types of indifference between the criteria that select the choice of alternatives. For instance, MMPURE behavior was left undefined in the case that the security level of the two behavioral alternatives are equal. Similarly, RISK deals only with the case of multiple strong equilibria. In our simulations we programmed the decision theories to play both alternatives with probability 1/2 in case of (near) indifference. In

Table 3. Probabilities to play X in the different types of games

<i>Game</i>	MMP_{PURE}	$H(\alpha)$	MAX/MAX	$BMMP_{PURE}$	$BH(\alpha)$	$BMAX/MAX$	P_{AYOFF}	$RISK$	$MMMIXED$	S_{WAGE}	P_{PARETO}	$NATURAL$
1	$a > b > c > d$	1	1	1	1	1	1	1	1	1	1	1
2	$a > b > d > c$	0	t_1	0	t_1	1	1	$t_1(\alpha = 0.5)$	0	$1 - r_1$	1	1
3	$a > c > b > d$	1	1	1	1	1	1	1	1	1	1	1
4	$a > c > d > b$	1	1	1	1	1	1	1	1	1	1	1
5	$a > d > b > c$	0	t_2	0	t_2	1	1	$t_2(\alpha = 0.5)$	r_2	$1 - r_1$	1	1
6	$a > d > c > b$	1	1	1	1	1	1	1	t_2	$1 - r_1$	1	1
7	$b > a > c > d$	1	t_1	0	t_4	1	r_1	r_1	1	r_1	r_3	1
8	$b > a > d > c$	0	0	0	0	0	0	0	0	0	1	0
9	$b > c > a > d$	1	t_3	0	t_5	1	r_1	r_1	r_2	r_1	r_3	1
10	$c > a > b > d$	1	1	1	1	1	1	1	1	1	r_3	1
11	$c > a > d > b$	1	1	1	1	1	1	1	1	1	1	1
12	$c > b > a > d$	1	1	0	0	0	r_1	r_1	r_2	r_1	r_3	1

$t_1 = 1$ if $\alpha a + (1 - \alpha) c > \alpha b + (1 - \alpha) d$, 0 otherwise.

$t_2 = 1$ if $\alpha a + (1 - \alpha) c > (1 - \alpha) b + \alpha d$, 0 otherwise.

$t_3 = 1$ if $(1 - \alpha) a + \alpha c > \alpha b + (1 - \alpha) d$, 0 otherwise.

$t_4 = 0$ if $\alpha a + (1 - \alpha) c > \alpha b + (1 - \alpha) d$, 1 otherwise.

$t_5 = 0$ if $(1 - \alpha) a + \alpha c > \alpha b + (1 - \alpha) d$, 1 otherwise.

$r_1 = (d - c)/(a - b - c + d)$.

$r_2 = (d - b)(a - b - c + d)$.

$r_3 = \min(1, (2d - b - c)/(2(a - b - c + d)))$.

fact, this specialized piece of code was only triggered for one game out of hundreds of millions of games that we simulated.

3 Decision Theories Play Round Robin

In a round robin tournament, our 17 contesting decision theories play a (large) number of symmetric 2×2 games with each other, and the decision theories are compared by the mean payoffs that they acquire across all their interactions. Clearly, the ranking of the decision theories may depend on the set of games that are being played. How should we proceed? The first possibility is the selection of a set of ‘interesting games’ by scanning the decision theoretic literature. Such an approach is still relatively ad hoc; what is interesting scientifically need not coincide with the kinds of games that are encountered regularly in real life. Selecting games at random is an obvious alternative procedure. Since no distribution of games seems of particular theoretical interest, we use a simple procedure: the payoffs a , b , c , and d of each game are drawn from a distribution F (and, if necessary, the payoffs are relabeled such that $a > d$). All drawings within and between games are independent.

The payoffs for two competing strategies can be obtained by integration over all possible games for a given distribution F . Instead of the tedious analytical way of calculating the scores for competing strategies, we simulated the round robin tournament.² We used four distributions: a uniform distribution $F = U[0, 1]$, a standard normal distribution $F = \text{Normal}(0, 1)$, an exponential distribution, and a log-normal distribution. To eliminate noise as much as possible, the number of randomly selected games was set at one hundred million.³ For each of these games, we determined the decisions made by all 17 theories; note that the decision theories may return mixed strategies. Next, we computed for each pair of decision theories (including the decision theories tied to them) the *expected payoffs* for a particular game.

Table 4 lists the average payoffs of the decision theories in the round robin tournament. For uniformly distributed payoffs, all payoffs are in the unit interval, so the average payoff of each decision theory is in the unit interval as well. For the normal distribution this clearly need not hold. The winner of the tournament, that is, the decision theory with the highest average payoff across the simulations, is the *Principle of Insufficient Reason*. However, note that PIR is the winner of the tournament by a very small margin on decision theories such as Risk. Also observe that the ranking of the decision theories is *identical* for the uniform and normal cases and only slightly different for the log-normal and exponential distributions. It seems that, for these latter distributions, it is profitable to be a little bit more optimistic.

Table 4. Average payoffs and the rank orderings of 17 decision theories in a round robin tournament of 10^8 randomly selected games with payoffs (a, b, c, d) generated independently from a uniform, a standard normal, log-normal, and exponential distribution

<i>Decision theory</i>	<i>Uniform</i>		<i>Normal</i>		<i>Log-normal</i>		<i>Exponential</i>	
PIR	0.64339	1	0.48741	1	2.55028	1	1.49390	1
H(0.75)	0.64056	2	0.47588	2	2.54683	2	1.49070	2
RISK	0.63844	3	0.47004	3	2.52043	4	1.47570	4
H(0.25)	0.63405	4	0.45086	4	2.51455	5	1.46583	5
MAXIMAX	0.62705	5	0.43392	5	2.53627	3	1.48085	3
PAYOFF	0.62570	6	0.42799	6	2.50873	6	1.46490	6
NATURAL	0.62092	7	0.41224	7	2.39625	8	1.41712	8
SAVAGE	0.61994	8	0.41003	8	2.43727	7	1.42171	7
BPIR	0.61879	9	0.40087	9	2.32797	10	1.38478	9
BH(0.75)	0.61622	10	0.39065	10	2.32415	11	1.38136	11
BH(0.25)	0.61587	11	0.38993	11	2.31746	13	1.37651	12
MMPURE	0.61405	12	0.38910	12	2.14307	15	1.30190	15
BMAXIMAX	0.60965	13	0.37027	13	2.31858	12	1.37623	13
BMMPURE	0.60890	14	0.36882	14	2.20543	14	1.32475	14
PARETO	0.60221	15	0.34280	15	2.35744	9	1.38470	10
MMMIXED	0.58786	16	0.29620	16	1.93807	16	1.19876	16
RANDOM	0.50891	17	0.03020	17	1.74638	17	1.04985	17

The average payoffs of most of the decision theories seem to be quite close. One of the reasons is that a symmetric game with randomly generated payoffs has a dominant strategy with probability 0.5. This dominant strategy is played by most of the decision theories. Nevertheless, if we restrict the analysis to games without dominant strategies, the conclusions are similar. More specifically, the top 7 of the ranking remain the same. The main change is that PARETO is moving to a higher position especially in the cases of the uniform and the normal distribution.

To understand where the differences in rankings between the different decision theories come from we examined (1) the average payoff of every decision theory against every other strategy and (2) the average payoff for every decision theory for each of the twelve types of games as indicated in Table 1. The average payoff matrix over the one hundred million games with payoffs drawn from a uniform distribution is listed in Table 5. For future reference we denote this matrix by U . So $U(T_i, T_j)$ is the expected average payoff of decision theory T_i in interactions with decision theory T_j .

It turns out that the success of PIR (and the other non-extreme Hurwicz strategies) is mainly due to its success against the strategies that are designed as best replies against the Hurwicz strategies. In addition, the

Table 5. Average payoffs of decision theories (times 1000) over 10^8 randomly selected symmetric 2×2 games against one other decision theory, with payoffs generated independently from a uniform distribution $U[0,1]$

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1	MMPURE	600	600	600	600	633	633	633	633	633	608	608	600	600	655	600	600
2	H(0.25)	613	613	613	613	621	671	671	671	671	627	627	613	613	683	638	613
3	PIR	617	617	617	617	600	650	700	700	700	642	642	617	617	704	667	617
4	H(0.75)	613	613	613	613	571	621	671	721	721	652	627	613	613	719	688	613
5	MAXIMAX	600	600	600	600	533	583	633	683	733	658	608	600	600	726	700	600
6	BMPPURE	667	642	617	592	567	533	558	583	608	608	608	600	617	668	667	583
7	BH(0.25)	642	667	642	617	592	533	558	583	608	627	627	613	600	678	667	583
8	BPIR	617	642	667	642	617	533	558	583	608	633	642	617	583	686	667	583
9	BH(0.75)	592	617	642	667	642	533	558	583	608	633	652	627	613	567	692	583
10	BMAXIMAX	567	592	617	642	667	533	558	583	608	633	658	608	600	550	697	583
11	PAYOFF	585	591	602	616	635	552	595	635	670	702	658	608	600	585	717	685
12	RISK	618	624	635	624	618	585	629	668	679	685	642	642	617	700	668	610
13	MM MIXED	603	603	603	603	603	570	607	636	657	670	627	610	603	685	653	603
14	SAVAGE	583	583	583	583	583	567	575	583	592	600	592	583	583	635	600	583
15	NATURAL	555	570	587	606	626	535	570	603	664	641	591	582	552	691	655	576
16	PARETO	550	575	600	625	650	583	608	633	658	683	658	608	600	567	705	600
17	RANDOM	500	500	500	500	500	450	475	500	525	525	500	500	500	576	550	500

Hurwicz strategies are doing pretty well among each other and also against the third and fourth class of strategies. For instance, PIR is a best reply against RISK and performs just as well as RISK against PAYOFF. MAXIMAX is a best reply against PAYOFF. The best replies against the various Hurwicz strategies lose the battle because their average payoff against each other is low. PAYOFF and RISK are doing well over the whole range of opponents, although PAYOFF loses against the more pessimistic strategies such as MMPURE and best replies against these pessimistic strategies as well. In general, we can conclude that being too pessimistic is not profitable in this population of strategies. Being too optimistic is also harmful, but not as bad as being too pessimistic. One could wonder what the optimal α for optimism would be in this population of decision theories, but this question is not that interesting given the subsequent analyses. Below we will learn that the Hurwicz strategies are not likely to survive RISK in the long run if other decision theories against which they perform well disappear and that there does not exist an optimal α to outperform RISK.

A deeper insight into the causes of these results is obtained by studying the results for the twelve types of games separately. First, as mentioned before, it should be noted that most decision theories play the dominant strategy if such a strategy is available. Only RANDOM and sometimes PARETO do not play the dominant strategy. Consequently, there is hardly any differentiation between decision theories in the six types of games for which there is a dominant strategy. PIR is better than or at least as good as all the other strategies for all other types of games except one in which PIR is outperformed by H(0.75) with a margin of less than 0.1%. The relatively low positions of PAYOFF, NATURAL, and PARETO are mainly due to their low scores in the games in which a is the highest and c is the lowest payoff, that is, the Assurance game and one of the Assurance-like games. In addition, the position of PARETO is especially low because of a large loss in the Prisoner's Dilemma in which the others are defecting. RISK and PAYOFF both lose against PIR in the games with three Nash equilibria from which only one is symmetric, for instance in the Chicken game.

4 The Evolution of Rationality

The round robin tournament gives a first impression of which kinds of strategies tend to perform best. However, one of the typical problems with such a tournament is the lack of robustness with respect to the inclusion of more or less well-performing decision theories such as random (erratic) choice, or other strategies that typically tend to perform poorly

(Binmore 1995: 198–201). Axelrod (1984) has rightly stressed that if successful types of behavior tend to grow in a population at the expense of unsuccessful types of behavior, behavior that is successful in the long run should depend for its success on other successful types of behavior. Types of behavior that thrive mainly because they are successful against types of behavior that are themselves unsuccessful, will likely become extinct in the medium time range: After the initially unsuccessful types of behavior have disappeared from the population, types of behavior that depended for their success on the now extinct types will likely disappear next. Thus, in the long run, success depends on success against other successful types rather than on exploitation of unsuccessful types.

Before turning to such an evolutionary analysis, we first consider a simple *two-stage game* with two players. In stage 1, the players independently decide to adopt one of the 17 decision theories. In stage 2, Nature selects a random game as described in the previous section, and the players use their selected decision theory to make a decision. How would a rational player select a decision theory in this two-stage game? Obviously, if player i uses decision theory T_i , the expected payoff of player 1 is $U(T_1, T_2)$ and of player 2 is $U(T_2, T_1)$.

Table 6 describes the symmetric Nash equilibria of the two-stage game for the four probability distributions, computed using Gambit (McKelvey et al. 2005). At first sight it may seem odd to allow players to choose between decision theories in stage 1, and afterwards compute the equilibria strictly based on one decision theory, in particular, Nash. The main reason is that the set-up we developed can be considered as a larger one-shot game in which the 17 decision theories are the available strategies and the matrix U is the payoff matrix. Therefore, the Nash equilibria of this larger game are precisely the candidates for survival in an evolutionary context. Given that only strategies that are part of a symmetric Nash equilibrium in symmetric games can be evolutionarily stable (Weibull 1995: chapter 2), we restrict ourselves to the symmetric equilibria in this two-stage game (for the different distributions F the Gambit analyses produced between 50 and 70 asymmetric equilibria). We find 7 symmetric Nash equilibria for the uniform, exponential and normal distribution, and 5 for the log-normal distribution. Obviously, we find two pure symmetric Nash equilibria. By definition, a best reply against RISK is RISK, while a best reply to PAYOFF is PAYOFF itself. In addition, we find a number of mixed equilibria. First, for all four probability distributions F , we find that $H(0.25)$ and $\text{b}H(0.25)$ form a mixed equilibrium, with somewhat different mixing probabilities for different F . The same holds for $H(0.75)$ and $\text{b}H(0.75)$, with the exception that this last equilibrium does not occur

Table 6. Symmetric Nash equilibria of the two-stage decision theory game with payoffs generated from various probability distributions. Decision theories that do not occur in any equilibrium are omitted. Columns do not always add up to 1 due to rounding. The attraction of an equilibrium is the fraction (times 1000) of randomly generated initial distributions for which the replicator dynamics converges to this equilibrium

Decision theory								Uniform distribution								Exponential distribution							
H(0.25)								0	0	0	0.67	0	0	0.46	0	0	0	0.72	0	0	0.55		
PIR								0	0	0	0	0	0.10	0.10	0	0	0	0	0	0.13			
H(0.75)								0	0	0.67	0	0.21	0.46	0	0	0.74	0	0.67	0.57				
BH(0.25)								0	0	0	0.33	0	0	0.25	0	0	0	0.28	0	0.24			
BPIR								0	0	0	0	0	0	0	0	0	0	0	0				
BH(0.75)								0	0	0.33	0	0.10	0.25	0	0	0.26	0	0.24	0.20				
PAYOFF								0	1	0	0	0.68	0	0	1	0	0	0	0.24				
RISK								1	0	0	0	0	0.19	0.19	1	0	0	0.09	0				
NATURAL								0	0	0	0	0.02	0	0	0	0	0	0	0				
Attraction								955	2	40	2	0	0	0	850	141	9	0	0	0			
Decision theory								Normal distribution								Log-normal distribution							
H(0.25)								0	0	0	0.68	0	0	0.46	0	0	0	0.77	0	0.67			
PIR								0	0	0	0	0	0.13	0.13	0	0	0	0	0.09				
H(0.75)								0	0	0.68	0	0.21	0.46	0	0	0	0	0.20	0				
BH(0.25)								0	0	0	0.32	0	0	0.25	0	0	0.23	0	0.22				
BPIR								0	0	0	0	0	0	0	0	0	0	0	0				
BH(0.75)								0	0	0.32	0	0.09	0.25	0	0	0	0	0.06	0				
PAYOFF								0	1	0	0	0.66	0	0	0	1	0	0.46	0.03				
RISK								1	0	0	0	0	0.17	0.17	1	0	0.29	0					
NATURAL								0	0	0	0	0.03	0	0	0	0	0	0	0				
Attraction								949	3	46	2	0	0	0	830	170	0	0	0				

under the log-normal distribution. There are some other equilibria that are mixtures of the equilibria mentioned above in which occasionally PIR and NATURAL are included.

We further specify in the rows labeled 'Attraction' which equilibria of the ones listed in Table 6 are more appealing from an evolutionary perspective. It is easy to show that the two pure equilibria are evolutionarily stable. Although there exist alternative best reply strategies against them, these strategies are doing worse against themselves. The equilibria that are mixtures with RISK or PAYOFF are not evolutionarily stable strategies. This follows from the theorem that if an evolutionarily stable strategy has support on a given set of pure strategies, then any strategy with support on a subset of this given set cannot be a Nash equilibrium strategy (Weibull 1995: 41, proposition 2.2). Therefore, the combinations of H(0.25) and H(0.75) with their best replies are the only other potential candidates of evolutionarily stable strategies. Because PIR only occurs in combination with PAYOFF and/or RISK, it will not survive in an evolutionary context. So now we investigate the relative importance of the various (potential) evolutionarily stable strategies. The central element of a dynamic analysis of the composition of a population of decision theories is a feedback mechanism between the payoff of decision theories and their share π in the population. We focus on a biological model that interprets average payoff of a decision theory as fitness. The fitness of a decision theory is proportional to the probability of getting 'offspring' of the decision theory. Assuming that the total size of the population is fixed, the share π_i of decision theory i in the population increases if and only if the payoff $(U \pi)_i$ of i is larger than the average payoff $\pi' U \pi$ in the population (see Hofbauer and Sigmund 1988: 108–137). In continuous time, this common evolutionary model, known as the replicator dynamics, can be described by the differential equation

$$\frac{d\pi_i(t)}{dt} = \rho \pi_i(t) ((U \pi(t))_i - \pi(t)' U \pi(t)) \text{ for } i = 1, \dots, 17.$$

Deriving analytical solutions of these differential equations is neither feasible nor very informative. Therefore, numerical solutions were computed in Matlab using the second and third order Runge-Kutta method. In Figure 2 we display these solutions for two initial distributions of the 17 decision theories for the payoff matrix U that is associated with uniformly distributed payoffs of 2×2 games. In the left panel, the initial population consists of equal shares of the 17 decision theories. As we see in the graph, there are three types of time-paths for decision theories:⁴

Figure 2. Population dynamics of the 17 decision theories. In the left graph, the initial distribution consists of all 17 decision theories in equal shares. In the right graph, RANDOM occupies 99% of the initial distribution with the remaining 1% equally shared by the other 16 decision theories.

- The proportion in the population of RANDOM, bH(0.25), MMPURE, BMMPURE, MMMIXED, bMAXIMAX, and PARETO monotonically decrease to zero.
- The proportion of MAXIMAX, PIR, bPIR, H(0.25), H(0.75), bH(0.75), SAVAGE, PAYOFF, and NATURAL increases in the initial phase of the population dynamics, and then decreases to zero.
- The proportion of RISK increases monotonically to fixation.

The above division of decision theories corresponds with the lower and higher performing theories in the round robin tournament. This is understandable because in the round robin tournament all decision theories also have the same number of representatives. Only after some time, it becomes more important that a decision theory does well against other well-performing strategies. For this, the results of the round robin tournament are not very informative anymore. After 100 time units, the population consists of 15% PIR, 6% bPIR, and 79% RISK. By tracing the time-paths for longer periods, we found that RISK does indeed occupy the whole population in the long run (fixation).

In the right panel of Figure 2 we consider the same dynamics from an initial population that consists of 99% RANDOM with the remainder 1% shared equally among the remaining 16 decision theories. Qualitatively, the time-paths of the decision theories are very similar, after a first phase in which RANDOM behavior is being wiped out. Most of the decision theories die out quickly. At time 120, only three decision theories have survived in a non-negligible proportion in approximately the same shares as in the first case.

We have indicated above that there are multiple Nash equilibria that could possibly be evolutionarily stable. Stability, however, is a local property, meaning that if the distribution is 'close enough,' the system converges to the stable state. To get a better understanding of the importance of the different equilibria, we therefore calculated the equilibrium to which the population converges for 5000 randomly chosen starting values, separately for each of the four probability distributions for payoffs that we use in this study. Table 6 lists the proportion of times (multiplied by 1000) that the process converged to each of the Nash equilibria, which can be interpreted as the sizes of the basins of attraction of these equilibria. As we mentioned earlier, if an equilibrium is stable, it cannot be part of a stable equilibrium with additional decision theories. This explains why we never observed convergence to the last three equilibria in case of the uniform, exponential, and normal distribution, and the last two in case of the log-normal distribution. We do observe convergence to the equilibria that are mixtures of $H(0.25)$ and $bH(0.25)$ and of $H(0.75)$ and $bH(0.75)$. These equilibria have rather small basins of attraction, at most 5% in all cases considered. The big attraction zones are for the Nash decision theories, especially for RISK. For the uniform and the normal distribution, in 95% of the cases starting from a random initial distribution, we end up in RISK. RISK is also doing very well with the other two probability distributions, but here PAYOFF has a basis of attraction of 14–17%. We do not have an explanation for this difference between probability distributions. Overall, however, RISK is the clear 'dominant' decision theory from the evolutionary perspective in the sense that starting from some mixture of the 17 decision strategies, it is by far the most likely to emerge in the end.

A last worry that we address related to the evolutionary stability of RISK and PAYOFF is whether one can design an 18th decision theory that might invade a population that consists of one type of actors, either RISK or PAYOFF. Such a decision theory should have two properties: first, it should be a best reply against RISK or PAYOFF. Second, it should perform better against itself than RISK or PAYOFF perform against this new decision theory (see Weibull 1995: 37). This implies that such a decision theory should copy RISK or PAYOFF in all games (see Table 3) in which RISK and PAYOFF play a pure strategy. Only in the games with one symmetric equilibrium that is a mixed equilibrium, such a new decision theory can choose whatever it wants still being a best reply against RISK and PAYOFF. However, in Theorem 2 in Appendix B, we show that it is impossible to choose behavior in these games in such a way that RISK or PAYOFF can be invaded.

Table 7. Payoffs in the Prisoner’s Dilemmas

	<i>red</i>	<i>green</i>	<i>red</i>	<i>green</i>	<i>red</i>	<i>green</i>	<i>red</i>	<i>green</i>
red	6	0	8	0	6	0	8	0
green	10	2	10	2	10	4	10	4

5 Dealing with Humans

In the previous sections, we studied symmetric interactions between 17 decision theories. We now confront these decision theories with the way in which human subjects choose between alternatives in symmetric 2×2 games in an experimental setting. The data were collected via booklets distributed as leave-behinds after a national household survey in the Netherlands (HIN) in 1995. Two booklets were left behind with each household. The booklets consisted of fourteen pages with various (abstract and scenario-like) experiments on interpersonal decision making. Three pages of this booklet described abstract 2×2 games, presented as ‘color games’ in which the respondent had to imagine that he or she was playing with an unknown other respondent of the survey. Both participants had to choose between the colors red and green, without knowing the other participant’s choice, with monetary outcomes depending on both participants’ choices. Thirty-two versions of the booklets with different experimental conditions (outcomes) were produced; subjects received a randomly selected version. Out of a total of $32 \times 3 = 96$ color games, we use the 48 color games that were symmetric 2×2 games (the other 48 color games were asymmetric 2×2 games). The outcomes used in the symmetric 2×2 are (0,2,6,10), (0,2,8,10), (0,4,6,10) and (0,4,8,10). These four outcomes were crossed with the 12 types of symmetric games.

For instance, we used four symmetric Prisoner’s Dilemmas (see Table 7), four variants of Chicken Games, 4 variants of Assurance Games, etc. Respondents were requested to fill out the booklets individually, stressing that there were no right or wrong answers. A recall was sent if the booklet was not returned within two weeks. After the completion of the data collection, six lucky participants were randomly chosen. They received a prize of 1 000 Dutch guilders (= Euro 450 = \$ 500) and the winners were announced in a national newspaper on a prescheduled date. The whole procedure was made clear in the booklet. Ultimately, 2 283 of 3 354 (68%) booklets were returned, with only small differences in the participation rates in terms of gender, age, education, religiosity, and political preference (see Bruins and Weesie 1996: section 1.6 for

Table 8. Total payoffs of decision theories against human behavior in 48 symmetric 2×2 games from the HIN-data. Payoffs are normalized so that the smallest payoff is 0, the largest is 1. The maximum score across all games is therefore 48

<i>Decision theory</i>	<i>All</i>	<i>Dominant strategy*</i>	<i>Prisoner's dilemma</i>	<i>Assurance-like</i>	<i>Chicken-like</i>
H(0.75), MAXiMAX	36.43	17.43	2.42	9.53	7.05
PAYOFF	36.09	17.43	2.42	9.53	6.72
bH(0.75), bMAXiMAX	36.01	17.43	2.42	9.53	6.63
bPIR	35.84	17.43	2.42	8.90	7.09
bH(0.25), bMMPURE	35.50	17.43	2.42	7.82	7.83
PIR	35.34	17.43	2.42	8.90	6.59
RISK	35.23	17.43	2.42	8.67	6.72
NATURAL	35.22	17.43	2.42	9.53	5.86
PARETO	34.58	17.39	1.22	9.53	6.44
SAVAGE	34.32	17.43	2.42	7.77	6.72
HUMAN	33.58	16.54	1.93	8.76	6.35
H(0.25), MMPURE	33.52	17.43	2.42	7.82	5.86
MMMIXED	32.65	17.43	2.42	6.41	6.40
RANDOM	26.81	11.22	1.82	6.93	6.84

*I.e., the dominant strategy games excluding the Prisoner's Dilemma.

details). Partial non-response on the color games was 1.6%. Since the subjects filled out the booklets without supervision, we were obviously concerned whether subjects had worked seriously and individually. Six booklets were marked as suspicious because we had some doubts as to whether subjects had worked seriously in at least part of the booklet. By comparing handwriting in the booklets received from households, we identified two pairs of booklets that we thought were filled out by the same person; these four booklets were marked as suspicious as well.

With these data it is possible to compare how our 17 strategies would fare if they were matched against the population of individuals who filled out the booklets. Table 8 shows the results of this comparison. We show results for all 48 games together, for the games with a dominant strategy excluding the Prisoner's Dilemma, for the Prisoner's Dilemma, for Assurance-like games, and for Chicken-like games. Besides the 17 strategies, we also include the 'decision theory' of the humans themselves (HUMAN in Table 8).

Let us first consider the columns in Table 8 that refer to games with a dominant strategy (excluding the Prisoner's Dilemma) and the Prisoner's Dilemma. In these columns we see that there is hardly any difference

between the strategies, as one could have anticipated, given that most of the 17 strategies choose to play the dominant strategy. The decision theory that suffers the most is PARETO. The main differences between strategies are found in the Assurance-like games and the Chicken-like games. We see that, perhaps somewhat surprisingly, relatively optimistic strategies tend to perform well. MAXIMAX, H(0.75), bH(0.75), and bMAXIMAX all perform well. That is, it pays to be optimistic about what humans will play in symmetric 2×2 games (MAXIMAX and H(0.75)), and it also pays to have an optimistic view about the strategies others are using and then play best reply against it (bH(0.75) and bMAXIMAX). Finally, the standard game-theoretic solution PAYOFF also ranks high.

When we restrict ourselves to behavior in Chicken-like games, we see that the best reply against humans seems to be to assume that others are playing MMPURE, and play a best reply against that. The situation is less clear for Assurance games; many decision theories seem to work (including NATURAL and PARETO).⁵

All in all, we can conclude that the differences in games without a dominant strategy drive the results and that PAYOFF and the relatively optimistic strategies tend to perform well against humans. Humans do not perform very well against themselves in our set-up. They tend to rank among the strategies with the lowest scores. Of course, it does not follow from this that humans are actually not very well able to deal with strangers. Our set-up is rather specific and the distribution of games in the real world might be quite different from the distribution analyzed here. In addition, the 'human' decision theory is not one decision theory. Some people apply different strategies than others and in this analysis we just assume the average behavior of humans per game to be the human decision theory.

We now compare the strategies to human behavior, to find out which of the strategies resembles average human behavior of respondents in Dutch households the most. For each game in the data we can derive the choice of each decision theory, and compare it to actual behavior in the data. We then calculate a difference-score as the absolute difference between the actual choice by the human and the prediction of the decision theory, and average this difference across individuals. The results are displayed in Table 9.⁶

We find that NATURAL is the decision theory that resembles human behavior in the data best. The differences between decision theories mostly arise in games without dominant strategies. PAYOFF, PARETO, MAXIMAX, and bMAXIMAX perform similarly well in the Assurance-like games, but NATURAL gains momentum in the Chicken-like games. Apparently, precisely in these kinds of games humans and NATURAL tend to follow a maximin

Table 9. Difference of decision theories with human behavior in 48 symmetric 2×2 games from the HIN data. Lower scores represent lower difference and thus higher resemblance

<i>Decision theory</i>	<i>All</i>	<i>Dominant strategy</i>	<i>Prisoner's dilemma</i>	<i>Assurance-like</i>	<i>Chicken-like</i>
NATURAL	0.184	0.074	0.430	0.217	0.252
PAYOFF	0.229	0.074	0.430	0.217	0.435
PARETO	0.231	0.079	0.570	0.217	0.386
PIR	0.239	0.074	0.430	0.346	0.348
bH(0.75), bMAXiMAX	0.246	0.074	0.430	0.217	0.504
H(0.25), MMPURE	0.249	0.074	0.430	0.488	0.252
H(0.75), MAXiMAX	0.249	0.074	0.430	0.217	0.496
RISK	0.260	0.074	0.430	0.346	0.435
SAVAGE	0.281	0.074	0.430	0.431	0.435
MMMiXED	0.311	0.074	0.430	0.595	0.398
bPIR	0.314	0.074	0.430	0.346	0.652
bH(0.25), bMMPURE	0.372	0.074	0.430	0.488	0.748
RANDOM	0.500	0.500	0.500	0.500	0.500

strategy. A bootstrap analysis suggests that differences between theories of about 0.026 are statistically significant (at 95%). Combining these findings with the findings from Table 8, it is logical that the best reply against NATURAL – a decision theory that we did not included in our original set – would have been the best reply against the humans in the data. And indeed, bNATURAL scored 37.21 against humans, outperforming every other decision theory in our theory set.

6 Conclusions and Discussion

We analyzed the relative performance of formal decision theories for one-shot symmetric 2×2 games, both in simulated interactions among the formal decision theories and between the theories and observed human behavior. It is hard to tell from looking at the winner of our round robin tournament of decision theories, PIR, but in general it seems that in such a tournament it pays if a decision theory is relatively optimistic, or at least not pessimistic. In an evolutionary analysis RISK mostly gets the upper hand eventually: virtually all simulations converge to a population of decision theories consisting of RISK only. In addition, we have shown that RISK as well as PAYOFF cannot be invaded by any other decision theory including decision theories that are not in our set of theories. In this sense, our

work can be conceived as providing additional arguments in favor of the idea that in the long run evolutionary forces tend to favor (game-theoretic) rationality in a world of symmetric 2×2 games. Or, more broadly speaking, our findings support that when dealing with strangers (i.e., in one-shot interaction) rational behavior might evolve in the long run.

When we compare the behavior of the decision strategies against humans in one-shot 2×2 games, we find that both *PAYOFF* and the relatively optimistic strategies tend to perform well against humans. Surprisingly, humans themselves tend to perform relatively poorly against themselves. This could be conceived as support for the claim that humans do not have a very adequate view of the 'average behavior' they are likely to encounter in the population (or, that they have an adequate view of the population, but do not know how, or do not want, to take advantage of it). When we look at this more closely, we find that human behavior resembles *NATURAL* the most. The best reply against *NATURAL* is in fact the decision theory that promises the highest revenue when pitted against a population of humans, or at least a population of respondents of Dutch households. When we relate this to the findings as mentioned in the previous paragraph, it seems that even though there is hope that rational behavior in one-shot 2×2 encounters will occur in the long run, our data do not support that humans use rational expectations. Of course, this claim should not be interpreted too strongly because we assume a specific distribution of games in our experiment, which might not correspond with the distribution of interactions that humans have in real life. In addition, we consider humans as a homogeneous group and we did not go into differences between different types of humans that might use different strategies.

An interesting extension of our analyses would be to analyze larger classes of games. A logical follow-up would be the asymmetric 2×2 games, because it would enable us to distinguish between strategies better. A main advantage is that most decision theories apply to asymmetric games as well. The only exceptions are *PARETO* (Pareto-optimal play is in general not unique in asymmetric games), and *PAYOFF* (the equilibria of 2×2 games are generally not ordered via payoff-dominance, so equilibrium selection does not work here). Since both these decision theories did not perform very well for symmetric 2×2 games, we would not expect that dropping them from the analysis changes our conclusions significantly. Or, it is possible to modify *PARETO* in asymmetric games in such a way that a Pareto-efficient strategy vector is selected by cooperative Nash bargaining with *MMMIXED* or another solution concept as the status quo. Empirically, such data are available both in the

HIN-data we used here, but also in Frenkel (1976; reported in Rapoport et al. 1976). Frenkel collected a large amount of data on subject's choices in a variety of asymmetric 2×2 games. It is therefore feasible to perform an analysis for asymmetric 2×2 games in a similar fashion: comprising both interaction between the formal theories and interaction between the formal theories and human subjects.

The evolutionary analysis of the emergence of rationality that we have presented here is incomplete in at least one respect: the set of competing types (decision theories) are explicitly defined as inputs in the evolutionary analysis. This type of analysis is also referred to as an 'ecological analysis' focusing only on selection, rather than an 'evolutionary analysis' that encompasses selection and creation of new types. That is, rationality is not being created in the evolutionary process, we have to include some rationality (although only a little bit) to make rationality flourish in the long run. However, a more elaborate analysis would be possible using *genetic algorithms* on a suitably defined set of functions from the payoffs of symmetric 2×2 games into the set of decisions. Since payoffs are cardinal entities, this only involves a (subset of the) set of functions $f: R^2 \rightarrow [0,1]$. Alternatively, a coding of such functions that seems to lend itself well to applying the genetic algorithms involves LISP-like symbolic expressions (Koza 1992). It would certainly add to our analyses when it can be shown that a genetic algorithm will generate (close approximations to) the long-term survivor theories such as equilibrium selection via risk dominance. It is probably not that difficult to generalize such an analysis to more general games. We hope to report on such an analysis in a future article.

Appendix A – A Note on Best Replies

We defined a best reply informally as the decision theory that always plays a best answer to a given other decision theory. This does not give a unique description of a strategy in every interaction, since if a decision theory plays a mixed strategy that leaves the partner indifferent, the best reply can be any pure or mixed strategy. More formally, we define best reply as follows.

Let T be a decision theory defined on one-shot symmetric 2×2 games with four different payoffs. Let bT be the best reply against T .

- If Alter's choice based on T does not leave Ego indifferent, Ego applying bT chooses the move that maximizes his outcome given that Alter uses T .

- If Alter's choice based on T leaves Ego indifferent, Ego applying bT chooses the move that is in Nash equilibrium with the move of Alter.

The second part of the definition ensures a unique definition because (1) given a move prescribed by T , there is at most one possible strategy by Ego that forms a Nash equilibrium with the move prescribed by T , and (2) if the move prescribed by T makes Ego indifferent between the two pure strategies, this implies T prescribes randomization and, therefore, bT prescribes the corresponding randomization that leaves Alter indifferent as well. In the last situation, the strategies are thus in equilibrium.

Theorem 1. For two theories T_1 and T_2 defined on a set of games Γ , we say that $T_1 = T_2$ if T_1 and T_2 prescribe the same behavior for all games in Γ . For any given decision theory T defined on one-shot symmetric 2×2 games, there are three possibilities:

1. $T = bT$,
2. $T = bbT$, or
3. $bT = bbbT$.

The set of decision theories for which the first equality holds is a subset of the set of decision theories for which the second equality holds. The third equality holds for all decision theories.

Proof. If T prescribes *symmetric* Nash equilibrium behavior, we are in case 1: the best reply to T is T itself.

If T always prescribes a strategy for which there is a best reply that constitutes a Nash equilibrium in combination with the strategy prescribed by T , the best reply will by definition prescribe the other part of the Nash equilibrium. Therefore, bbT again prescribes the same behavior as T .

Assume that we have a decision theory T that belongs not to case 1 or 2. Then, T prescribes behavior that is not part of an equilibrium combination of strategies for some 2×2 games. There is always a unique best reply in pure strategies in such a game, because if T prescribes behavior that makes the partner indifferent, there is always a reply that is in equilibrium with T . In symmetric 2×2 games with unequal payoffs, both players either have a dominant strategy or both pure strategies are part of an equilibrium combination. Thus, bT always plays a part of an equilibrium combination of strategies. This implies that the best reply of T fulfills case 2, and therefore, $bT = bbbT$. Clearly, $T = bT$ implies $bbT = bT = T$ and the conditions for case 2 both imply $bT = bbbT$.

Corollary 1. $BBH(A) = H(A)$ in symmetric 2×2 games.

The argument is that the Hurwicz decision theories always prescribe a pure strategy and they play the dominant strategy if it exists. Therefore, they belong to case 2 above.

For the other decision theories, we do not provide formal proofs, but it can be derived that PIR and bPIR are also best replies (in the informal sense, i.e. that not necessarily the condition on playing a Nash equilibrium in case of randomization is fulfilled) against SAVAGE and RISK. MAXIMAX is the best reply against PARETO and in the informal sense also against PAYOFF. bMAXIMAX is also in the informal sense a best reply against PAYOFF. PAYOFF and RISK are their own best replies. Using this logic of best replies we see that not all best replies to the decision theories we chose are in our set of theories; i.e. our set of decision theories is not 'closed' and in principle vulnerable to the addition of these best reply decision theories. The two strategies for which we did not include a best reply are MMMIXED and NATURAL. The best reply to the best reply of MMMIXED is also not in our experiment. From the theorem it follows that the 'triple' best reply of MMMIXED is the same as the best reply of MMMIXED. Because NATURAL always plays a pure strategy and it plays the dominant strategy if it exists, bBNATURAL is NATURAL itself. Finally, the best reply against SAVAGE in the strict sense is also not included.

Appendix B – The Impossibility to Invade Risk and Payoff

Although RISK and PAYOFF are their own best replies, this does not imply that there cannot be other decision theories that can invade them in an evolutionary process. Now we state the formal theorem on the impossibility to invade RISK and PAYOFF.

Theorem 2. For any decision theory T , it is impossible to invade a population that consists completely of actors using either RISK or PAYOFF.

Proof. If T would like to invade RISK or PAYOFF, it has to fulfill two conditions. First, it has to be a best reply against RISK or PAYOFF and, second, it has to perform better against itself, than RISK or PAYOFF performs against this new decision theory (see Weibull 1995: 37). This already implies that the new decision theory copies RISK or PAYOFF in all games in which those theories do not randomize and, therefore, the best reply is uniquely determined. Only in the games in which RISK and PAYOFF randomize, the new decision theory can choose a probability to

fulfill the second condition. We know that RISK and PAYOFF randomize in these games with probability $r_1 = \frac{d-c}{a-b-c+d}$. Assume the new decision theorem randomizes with probability r_0 . Then, the second condition implies that

$$r_0^2 a + (1 - r_0)^2 d + r_0(1 - r_0)(b + c) > r_0 r_1 a + (1 - r_0)(1 - r_1)d + r_0(1 - r_1)b + (1 - r_0)r_1 c,$$

which is equivalent to

$$r_0(r_0 - r_1)a - (1 - r_0)(r_0 - r_1)d - r_0(r_0 - r_1)b + (1 - r_0)(r_0 - r_1)c > 0,$$

Now, we need to distinguish two cases: $r_0 > r_1$ and $r_0 < r_1$. Assume $r_0 > r_1$, then the equation above is equivalent to $r_0 a - (1 - r_0)d - r_0 b + (1 - r_0)c > 0$ or, because $a - b - c + d < 0$ in all the relevant games, $r_0 < \frac{d-c}{a-b-c+d} = r_1$. This contradicts the assumption. Similarly, the assumption $r_0 < r_1$ leads to a contradiction. Therefore, there does not exist an r_0 that would allow an invasion of RISK or PAYOFF.

Acknowledgments

Helpful suggestions from Andreas Diekmann, Norman Braun, Werner Raub, and two anonymous reviewers are gratefully acknowledged. Weesie acknowledges funding by the Netherlands Organization for Scientific Research (NWO) for the research project 'The Management of Matches' (S 96–168 and PGS 50–370). Buskens' contribution is part of the UU High-Potential program 'Dynamics of Cooperation, Networks and Institutions.'

NOTES

1. A whole family of decision theories can actually be constructed by combining some kind of transformation from a given into an effective matrix with one of the decision principles. For instance, one could combine equilibrium selection via payoff dominance with a value transformation in which the utility of an actor is assumed to depend on the outcomes to self and to the other actor (altruism, envy, etc.).
2. This is equivalent to numerical integration across the possible games.
3. For the uniform distribution, we compared the simulated results with the analytical results. We are confident that the simulation results deviate less than 0.1% from the exact values.

4. From other starting values other time-paths are possible. For instance, decision theories may alternate a few times between becoming more and less frequent, before finally decreasing to 0.
5. Robustness of the results in Table 8 was analyzed by bootstrapping. This gives some mixed results. When we choose to bootstrap across the 48 games we get 95% confidence intervals that run from about 2.5 above and below the presented values. This implies that the confidence intervals from the best and the worst decision theory (excluding RANDOM) overlap. However, bootstrapping at the respondent level gives confidence intervals of about 0.5.
6. When dealing with a decision theory that prescribes mixed strategies, we do the following. Suppose that we have a decision theory that prescribes that the probability to choose a certain alternative equals 0.4 and we have an individual that indeed chooses this alternative. In this case the difference score is taken to be $10.4 - 11 = 0.6$.

REFERENCES

- Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Bednar, Jenna and Scott Page. 2007. 'Can Game(s) Theory Explain Culture? The Emergence of Cultural Behavior within Multiple Games.' *Rationality and Society* 19: 65–97.
- Binmore, Ken. 1995. *Game Theory and the Social Contract. Volume 1: Playing Fair*. Cambridge: MIT Press.
- Bruins, Joost and Jeroen Weesie. 1996. *Codebook of Hin95Exp. A Set of Experiments Conducted in Connection with Households in the Netherlands 1995*, Utrecht, ISCORE paper 77.
- Güth, Werner. 1992. *Spieltheorie und ökonomische (Bei)Spiele*. Berlin: Springer Verlag.
- Güth, Werner and Hartmut Kliemt. 1998. 'The Indirect Evolutionary Approach: Bridging the Gap Between Rationality and Adaptation.' *Rationality and Society* 10: 377–399.
- Harsanyi, John C. and Reinhard Selten. 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.
- Heckathorn, Douglas D. 1996. 'The Dynamics and Dilemmas of Collective Action.' *American Sociological Review* 61: 250–277.
- Hegselmann, Rainer. 1988. 'Wozu könnte Moral gut sein? oder: Kant, das Gefangenendilemma und die Klugheit.' *Grazer Philosophische Studien* 31: 1–28.
- Hofbauer, Josef and Karl Sigmund. 1988. *The Theory of Evolution and Dynamical Systems. Mathematical Aspects of Selection*. (English translation of German original 1984). Cambridge: Cambridge University Press.
- Kelley, Harold H. and John W. Thibaut. 1978. *Interpersonal Relations*. New York: Wiley.
- Koza, John R. 1992. *Genetic Programming. On the Programming of Computers by Means of Natural Selection*. Cambridge: MIT Press.
- Luce, R. Duncan and Howard Raiffa. 1957. *Games and Decisions: Introduction and Critical Review*. Wiley: New York.
- Macy, Michael W. 1997. 'Identity, Interest, and Emergent Rationality: An Evolutionary Synthesis.' *Rationality and Society* 9: 427–448.
- Macy, Michael W., and Andreas Flache. 1995. 'Beyond Rationality in Models of Choice.' *Annual Review of Sociology* 21: 73–91.
- McKelvey, Richard D., Andrew M. McLennan, and Theodore L. Turocy. 2005. *Gambit: Software Tools for Game Theory, Version 0.2005.12.12*. <http://econweb.tamu.edu/gambit>.

- Rapoport, Anatol and Melvin Guyer. 1966. 'A Taxonomy of 2×2 Games.' *General Systems* 11: 203–214.
- Rapoport, Anatol, Melvin J. Guyer, and David G. Gordon. 1976. *The 2×2 Game*. Ann Arbor, MI: Michigan Press.
- Weibull, Jorgen W. 1995. *Evolutionary Game Theory*, Cambridge: MIT Press.
- Young, H. Peyton. 1993. 'The Evolution of Conventions.' *Econometrica* 61: 57–84.

JEROEN WEESIE is an associate professor at the Department of Sociology/ICS, Utrecht University. His research interests focus on social networks, cooperative relations, social dilemmas, and statistics.

ADDRESS: Department of Sociology, Utrecht University, Heidelberglaan 2, 3584 CS Utrecht [email: j.weesie@uu.nl]

CHRIS SNIJDERS is Professor of Sociology of Technology and Innovation at the Eindhoven University of Technology. His research interests center on model building and simulation in the social sciences, trust and cooperation problems between individuals and organizations, and online forms of cooperation.

ADDRESS: Eindhoven University of Technology, PO Box 513, 5600 MB Eindhoven [email: c.c.p.snijders@tue.nl]

VINCENT BUSKENS is an associate professor at the Department of Sociology/ICS, Utrecht University. His research interests focus on dynamics of social networks, cooperative relations, and trust.

ADDRESS: Faculty of Technology Management, Department of Sociology, Utrecht University, Heidelberglaan 2, 3584 CS Utrecht [email: v.buskens@uu.nl]