

Evolution of State-Dependent Risk Preferences

PATRICK ROOS, J. RYAN CARR, and DANA S. NAU

University of Maryland

Researchers have invested much effort in constructing models of the state-dependent (sometimes risk-averse and sometimes risk-prone) nature of human decision-making. An important open question is how state-dependent risk behavior can arise and remain prominent in populations. We believe that one part of the answer is the interplay between risk-taking and sequentiality of choices in populations subject to evolutionary population dynamics. To support this hypothesis, we provide simulation and analytical results for evolutionary lottery games, including results on evolutionary stability. We consider a parameterized class of *imitation dynamics* in which the parameter $0 \leq \alpha \leq 1$ yields the *replicator dynamic* with $\alpha = 1$ and the *imitate-the-better* dynamic with $\alpha = 0$. Our results demonstrate that for every population dynamic in this class except for the replicator dynamic, the interplay between risk-taking and sequentiality of choices allows state-dependent risk behavior to have an evolutionary advantage over expected-value maximization.

Categories and Subject Descriptors: I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*; I.6.3 [Simulation and Modeling]: Applications; J.4 [Computer Applications]: Social and Behavioral Sciences

Additional Key Words and Phrases: Evolutionary games, population dynamics, risk, decision theory

1. INTRODUCTION

Empirical evidence shows that human decision-making, rather than conforming to the decision-theoretic notion of expected-value maximization, is *state-dependent*: the decisions are sometimes risk-averse and sometimes risk-seeking, depending on the decision maker's circumstances. Much effort has been invested into describing and modeling such behavior, but these efforts have largely lacked an explicit investigation of what evolutionary pressures might have influenced the behavior's spread. Thus an important open question is why state-dependent risk behavior is so prevalent.

In this paper, we use tools from evolutionary game theory to investigate how

Authors' address: Laboratory for Computational Cultural Dynamics, Department of Computer Science, Institute for Advanced Computer Studies, and Institute for Systems Research, University of Maryland, College Park, MD 20742; email: Patrick Roos (roos@cs.umd.edu), Ryan Carr (carr2@cs.umd.edu), Dana Nau (nau@cs.umd.edu).

This work was supported in part by AFOSR grant FA95500610405 and ARO grant W911NF0920072. The opinions expressed in this paper are those of the authors and do not necessarily reflect the opinions of the funders.

© ACM, (2010). This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in ACM Transactions on Intelligent Systems and Technology, {1, 1, (2010)}
<http://doi.acm.org/10.1145/1858948.1858954>

agents' risk behavior relates to different *population dynamics* (i.e., rules governing changes in the number of agents of each kind). In particular, we are interested in *imitation dynamics*, which model cultural evolution as a product of social learning by imitation. The best-known imitation dynamics are the replicator dynamic and the imitate-the-better dynamic, but there also are many others.

We consider a parameterized class of *imitation dynamics* in which the parameter $0 \leq \alpha \leq 1$ yields the *replicator dynamic* with $\alpha = 1$ and the *imitate-the-better* dynamic with $\alpha = 0$. Our study includes (1) a detailed analysis of how different imitation dynamics can affect risk behavior when agents make sequential choices, and (2) simulations, using several different imitation dynamics, of simple evolutionary lottery games in which agents make sequential choices among lotteries that have equal expected value but different risks. Our results demonstrate that for every population dynamic in this class except for the replicator dynamic, the interplay between risk-taking and sequentiality of choices allows state-dependent risk behavior to have an evolutionary advantage over expected-value maximization.

Our study provides a starting point for further investigation of how population dynamics influence risk behavior in evolutionary game environments. We anticipate that state-dependent risk behavior will outperform expected-value-maximizing strategies in a large variety of evolutionary game environments involving sequential choices of different risks.

The remainder of this paper is organized as follows: Section 2 provides a brief background on human decision making under risk and evolutionary game theory. Section 3 introduces our evolutionary game model, including the lottery game we use to study risk behavior and the imitation dynamics we are interested in. Section 4 provides theoretical analysis and predictions of how we expect different risk behavior to perform in different environments. Section 5 presents our experimental simulations and results. Finally, Section 6 provides concluding remarks.

2. BACKGROUND

2.1 Human Decision Making Under Risk

Human decision making under risk is the subject of much research effort in the social sciences. In most of the existing literature on models of human decision making under risk, the construction of such models is approached primarily through the analysis of a decision maker's choices among lotteries that have different payoff distributions, and thus potentially different risks. Under the most traditional model of decision making, expected utility theory, a rational agent's preferences can be modeled by assigning to each possible outcome a number called the outcome's *utility*; and a rational choice is one that maximizes the expected utility of the outcomes [von Neumann and Morgenstern 1944]. Empirical evidence of human decision making under risk shows that humans are sometimes risk-averse, sometimes risk-seeking, and even behave in ways that systematically violate the axioms of expected utility [Kahneman and Tversky 1979]. Expected utility theory can account for different attitudes towards risk through certain von Neumann-Morgenstern utility functions (e.g. [Friedman and Savage 1948]). Such risk propensities can differ greatly from simple expected-value considerations on prospective outcomes.

Researchers have invested much effort into constructing utility functions that

appropriately model human decision making under risk under the expected-utility model (e.g. [Friedman and Savage 1948; Arrow 1971; Rabin 2000]). Related efforts in economics have aimed to describe the preferences of humans over inter-temporal lotteries, recognizing the effects of temporally successive lotteries on risk preferences [Epstein and Zin 1989; 1991]. Other studies define utility functions that take into account inter-personal or population comparisons [Abel 1990]. Yet other researchers have constructed alternative descriptive theories of decision making that claim to correspond more closely to how humans make decisions involving risk. Among the most popular of these models are prospect theory [Kahneman and Tversky 1979; Tversky and Kahneman 1992], regret theory [Loomes and Sugden 1982], and SP/A (Security-Potential/Aspiration) theory [Lopes 1987; 1990; Lopes and Oden 1999]. One advantage of these models is that they more explicitly or perhaps more naturally model some of the mechanics involved in human decision making processes. For example, state-dependent attitudes toward risk are modeled in prospect theory by using a reference point with respect to which prospective outcomes can be interpreted as potential gains or losses, and are modeled in SP/A theory by including an aspiration level as an additional decision criterion in decisions involving risk. A common theme of both Prospect theory and SP/A theory is that agents are risk-averse when they have done well relative to some reference point, and risk-seeking when they have not done well relative to the reference point.

2.2 Evolutionary Game Theory and Risk Behavior

Evolutionary game theory has been used to study and explain a great variety of social and cultural phenomena. Examples of such phenomena studied through evolutionary games include altruism [Axelrod and Hamilton 1981; Riolo et al. 2001; Bowles and Gintis 2004], trust [Fang et al. 2002; Kimbrough 2005], fairness [Binmore 1998], empathy [Page and Nowak 2002], and social learning [Schotter and Sopher 2003]. The social-science literature is filled with examples of empirical studies showing that humans violate the strong rationality assumptions of classical game theory. Evolutionary game theory is thus an attractive framework under which to model human behavior in such domains, because the assumption of perfectly rational agents is not required. Furthermore, most game-theoretic work in economics, the focus of our evolutionary-game-theoretic approach is not to define internal preferences and whether or not agents satisfy these. Rather, the focus is to evaluate the preferences according to their evolutionary fitness (which is external).

Evolutionary game theory studies a population of agents making choices according to an assigned strategy in a game that models a situation of interest. After agents have played the game, they reproduce into the next generation according to a reproduction function or population dynamic that, generally speaking, increases the frequency of the types of agents that were successful in the current generation. To study risk behavior in this framework, we can model the situation of interest as a game in which agents are faced with choices among lotteries of different risks. As described in [Alexander 2009], these lotteries dispense resources that are considered to be an objective quantity of which 1) agents always want more than less and 2) interpersonal comparisons are meaningful. The reproduction function defining the dynamics of strategies in the population then acts directly on these resources.

2.3 Imitation Dynamics

Imitation dynamics are a class of population dynamics commonly used to model the evolution of behaviors in societies [Hofbauer and Sigmund 1998; Nowak and May 1992; Huberman and Glance 1993; Nowak et al. 1994; Eguluz et al. 2005]. The general framework for imitation dynamics is stated by Hofbauer and Sigmund [1998] as follows:

We shall suppose that occasionally a player is picked out of the population and afforded the opportunity to change his strategy. He samples another player at random, and adopts his strategy with a certain probability.

In what follows, we will call these players the “observer” and the “observed agent,” respectively.

Important theoretical studies have been done of two specific imitation dynamics. One of these is the *replicator dynamic* [Taylor and Jonker 1978; Schlag 1998; 1999; Hofbauer and Sigmund 1998; Gintis 2000], in which the probability that the observer adopts the strategy of the observed agent is proportional to how much more successful the observer was than the observed. The other is the *imitate-the-better* dynamic [Blume 1993; Vega-Redondo 1997; Szabó and Tóke 1998; Riolo et al. 2001; Hauert and Szabo 2005; Dawid 2007; Traulsen and Hauert 2009; Traulsen et al. 2007],¹ in which the observer always adopts the observed agent’s strategy if it was more successful than the observer’s strategy.²

Several experimental studies investigating social learning through imitation between humans have found experimental support for Vega-Redondo’s model [Huck et al. 1997; 2000; Offerman et al. 2002]. Experiments on human imitation reported by Apesteguia et al. [2007] indicate that the difference in observed payoff to an agent’s own payoff does affect imitation (the higher the difference, the more likely imitation occurs). This is more in line with Schlag’s model. Due to this evidence for both imitation models, in this paper we explore a parameterized range of imitation dynamics based on a definition in [Hofbauer and Sigmund 1998], that includes the replicator dynamic, the imitate-the-better dynamic, and a spectrum of other dynamics in between those two.

3. EVOLUTIONARY GAME MODEL

In previous work we have proposed the sequential lottery game, a class of games that we use to investigate risk behavior under evolutionary pressures [Roos and Nau 2009; 2010b]. In this section we describe this *sequential lottery game* and the particular range of imitation dynamics under which we explore the evolution of risk behavior.

3.1 Sequential Lottery Game

We consider a game in which agents acquire payoffs dispensed by lotteries. In each generation, each agent must make a sequence of n choices, where each choice is

¹The imitate-the-better dynamic is sometimes called tournament selection [Riolo et al. 2001].

²Vega-Redondo [1997] generalizes the imitate-the-better dynamic by allowing the observer to observe a collection of agents and adopt the strategy of the most successful agent in C .

between two lotteries: a *safe* lottery, whose payoff is always 4, and a *risky* lottery, which one can *win* (a payoff of 8) with probability p , or *lose* (a payoff of 0) with probability $1 - p$.³ Note that if $p = 0.5$, both lotteries have expected value 4.

Our population consists of agents that follow strategies chosen from the set $\mathcal{S} = \{s_1, \dots, s_k\}$ of all possible pure strategies for the sequence of lottery choices. In any generation, a vector $\mathbf{x} = (x_1, \dots, x_k)$ gives the state of the population, where each x_i is the proportion of agents in the population using strategy s_i . We let $\pi(i)$ denote the payoff accumulated in a generation from the n lottery choices by agents of type i (i.e. agents following strategy s_i).

3.2 Imitation Dynamics in Our Evolutionary Games

As we discussed in Section 2.3, we want to explore a range of population dynamics that includes the replicator dynamic, imitate-the-better dynamic, and other dynamics intermediate between these two extremes. Hofbauer and Sigmund [1998] give the following parameterized formula for these population dynamics:

$$\dot{x}_i = x_i \sum_j x_j |\pi(i) - \pi(j)|^\alpha \text{sign}(\pi(i) - \pi(j)), \quad (1)$$

where x_i is the current proportion in the population of agents of type i , \dot{x}_i is the change in x_i over time,⁴ and $\alpha \geq 0$ is a parameter that determines the particular imitation dynamic. Our formulation is based on theirs, but incorporates the following changes:

- We are interested in population dynamics based on payoff comparisons among individuals, as in [Blume 1993; Szabó and Tóke 1998; Riolo et al. 2001; Hauert and Szabo 2005; Traulsen and Hauert 2009; Traulsen et al. 2007; Dawid 2007]. To model payoff comparisons among individuals, we must take into account the stochastic variability in the payoffs to individual agents. We do this by treating $\pi(i)$ and $\pi(j)$ as discrete random variables representing the distributions of payoffs that an agent of type i and an agent of type j receive from their lottery choices.
- The imitate-the-better dynamic and the replicator dynamic correspond to $\alpha = 0$ and $\alpha = 1$, respectively. Since the imitation dynamics that interest us are these two and the ones that are intermediate between them, we only consider $0 \leq \alpha \leq 1$.

If we let r and s be any possible payoff values acquired by agents of type i and j , and let $p(r, s)$ the probability of obtaining this pair of values, then our modified version of Eq. (1) is

$$\dot{x}_i = x_i \sum_j x_j \nabla_\alpha(i, j), \quad (2)$$

³The choice of payoffs here is arbitrary. The results in this paper (see Sections 4 and 5) would be qualitatively similar if we had used any fixed value v for the safe lottery's payoff and $v + \delta$ and $v - \delta$ (for fixed δ) for the risky lottery's payoffs.

⁴As is common in the evolutionary game theory literature, Hofbauer and Sigmund approximate the current population as a real-valued function $x_i(t)$ where t is the current time, so that $\dot{x}_i = dx_i/dt$.

Table I. All of the possible pure strategies in our lottery game when $n = 2$.

Strategy	1st lottery	2nd lottery
<i>SS</i>	choose safe	choose safe
<i>RR</i>	choose risky	choose risky
<i>SR</i>	choose safe	choose risky
<i>RS</i>	choose risky	choose safe
<i>RwS</i>	choose risky	choose safe if 1st lottery was won, risky otherwise
<i>RwR</i>	choose risky	choose risky if 1st lottery was won, safe otherwise

Table II. Payoff distributions of the six pure strategies in the sequential lottery game with $n = 2$ and $p = 0.5$. All six strategies have the same expected payoff, 8, but differing distributions. In this table and others in this section, we use boldface numbers to denote payoff values, and non-boldface to denote probabilities.

	RwS			RwR			SR		RS		SS	RR		
Payoff V	12	8	0	16	8	4	12	4	12	4	8	16	8	0
Prob. $P(V)$.5	.25	.25	.25	.25	.5	.5	.5	.5	.5	1	.25	.5	.25

where $\nabla_\alpha(i, j) = \sum_r \sum_s |(r - s)|^\alpha \text{sign}(r - s) \cdot p(r, s)$ is the *switching rate* between two agent types i and j . This *switching rate* determines the effect that a pairing for imitation between agents of type i and j has, on average, on type i 's growth rate in the population. For example, if $\nabla_\alpha(i, j)$ is positive, a pairing for imitation between agents of type i and j (which happens with non-zero probability if agents of type i and j exist in the population) on average has a positive effect on i agents' growth rate. If this is the case, we say i has an evolutionary advantage over j .

4. THEORETICAL PREDICTIONS

In the two-choice sequential lottery game, there are six possible pure strategies. These are listed in Table I.⁵

Table II gives each each strategy, its possible numeric payoffs, and the probabilities of these payoffs for the case where p , the probability of winning the risky lottery, is 0.5. Note that in this case, all six strategies have the same expected payoff value of 8, but they have differing probabilities of being above or below 8. For example, $P[\pi(\text{RWS}) > 8] = 0.5$ and $P[\pi(\text{RWS}) < 8] = 0.25$.

In this section, we examine how these strategies will perform against one another for different population dynamics (Sections 4.1, 4.2, and 4.3) and for different expected values of the lotteries (Section 4.4), culminating in an evolutionary stability result for the RWS strategy (Section 4.5). Finally, Section 4.6 briefly discusses cases where $n \neq 2$.

⁵For simplicity, we have restricted our study to pure strategies. For lottery games like the ones we are considering, the reproductive fitness of a mixed strategy should be intermediate among the pure strategies in the mixed strategy's support, hence the inclusion of mixed strategies should not have a substantial affect on our results.

If the agents were interacting in a nonzero-sum game rather than making simple choices among lotteries, then the inclusion of mixed strategies could make a significant difference in the results; but such games are beyond the scope of this paper.

4.1 The Replicator Dynamic

We now show that under the replicator dynamic ($\alpha = 1$), the switching rate between any two agent types is equal to the difference in their expected payoffs.

PROPOSITION 4.1. *Under the population dynamics given by Eq. (2) when $\alpha = 1$, $\nabla_\alpha(i, j)$ (the switching rate between any two agent types i and j) is the difference between the expected payoffs of agents of type i and j (given by the discrete random variables $\pi(i)$ and $\pi(j)$).*

PROOF. Let r and s be possible payoff values acquired by agents of type i and j , and let $p(r, s)$ be the probability of obtaining this pair of values. With $\alpha = 1$,

$$\begin{aligned}\nabla_\alpha(i, j) &= \sum_r \sum_s |(r - s)|^\alpha \text{sign}(r - s) \cdot p(r, s) \\ &= \sum_r \sum_s (r - s) \cdot p(r, s) = EV(\pi(i) - \pi(j)).\end{aligned}$$

Assuming independence between payoffs, $\nabla_\alpha(i, j) = EV(\pi(i) - \pi(j)) = EV(\pi(i)) - EV(\pi(j))$, and the proposition follows. \square

Since all strategies have the same expected payoff in this environment, Proposition 4.1 tells us that the switching rate between any two strategies will be 0. Consequently, all six of the above strategies will perform equally well evolutionarily.

4.2 The Imitate-the-Better Dynamic

If we use the imitate-the-better dynamic ($\alpha = 0$), then in Eq. (2), only the sign of the payoff difference between two paired agents plays a role in determining the switching rate; the magnitude of the difference is irrelevant. We can compute the switching rate between two strategies by using the probabilities in Table II to calculate the probability of each pair of payoffs occurring. Table III shows these values for the RWS vs. RR pairing. We can then use these probabilities as the values of $p(r, s)$ in Eq. (2) to calculate the switching rate, as follows:

$$\begin{aligned}\nabla_\alpha(\text{RWS}, \text{RR}) &= \text{sign}(0 - 0) \cdot 0.0625 + \text{sign}(8 - 0) \cdot 0.0625 + \text{sign}(12 - 0) \cdot 0.125 \\ &\quad + \text{sign}(0 - 8) \cdot 0.125 + \text{sign}(8 - 8) \cdot 0.125 \\ &\quad + \text{sign}(12 - 8) \cdot 0.25 + \text{sign}(0 - 16) \cdot 0.0625 \\ &\quad + \text{sign}(8 - 16) \cdot 0.0625 + \text{sign}(12 - 16) \cdot 0.125 \\ &= 0.0625 + 0.125 - 0.125 + 0.25 - 0.0625 - 0.0625 - 0.125 \\ &= 0.0625.\end{aligned}$$

Using similar calculations, we see that $\nabla_0(\text{RWS}, \text{SS}) = 0.25$, $\nabla_0(\text{RWS}, \text{RWR}) = 0.0625$, and $\nabla_0(\text{RWS}, \text{RS}) = \nabla_0(\text{RWS}, \text{SR}) = 0$. This suggests that RWS will be able to consistently win an evolutionary competition against RR, RWR, or SS and remain stable with SR and RS in this environment. The experimental results in Section 5 verify this prediction.

Table III. Payoff distribution for RWS vs. RR in the sequential lottery game with $n = 2$ and $p = 0.5$. Entries in the table show give the probabilities that each pair of payoffs occurs.

		RWS		
		12	8	0
RR	16	0.125	0.0625	0.0625
	8	0.25	0.125	0.125
	0	0.125	0.0625	0.0625




Fig. 1. Switching rate between RWS and each of the other five pure strategies for $0 \leq \alpha \leq 1$. A positive switching rate indicates an evolutionary advantage of RWS against the other strategy (in a population made up of solely the two strategies). We see that RWS has such an advantage over all other strategies when $0 < \alpha < 1$, and all but SR and RS when $\alpha = 0$.

4.3 Arbitrary α

For $0 \leq \alpha \leq 1$, we can calculate the switching rate in a method similar to the previous section, combining the probabilities from Table II to get the values of $p(r, s)$ for Eq.(2). However, since both the sign and magnitude of the payoff differences are now important, the calculation is slightly more complex. For instance, the switching rate for the RWS vs. RR pairing is now as follows:

$$\begin{aligned}
 \nabla_{\alpha}(\text{RWS}, \text{RR}) &= |0 - 0|^{\alpha} \text{sign}(0 - 0) \cdot 0.0625 + |8 - 0|^{\alpha} \text{sign}(8 - 0) \cdot 0.0625 \\
 &\quad + |12 - 0|^{\alpha} \text{sign}(12 - 0) \cdot 0.125 + |0 - 8|^{\alpha} \text{sign}(0 - 8) \cdot 0.125 \\
 &\quad + |8 - 8|^{\alpha} \text{sign}(8 - 8) \cdot 0.125 + |12 - 8|^{\alpha} \text{sign}(12 - 8) \cdot 0.25 \\
 &\quad + |0 - 16|^{\alpha} \text{sign}(0 - 16) \cdot 0.0625 + |8 - 16|^{\alpha} \text{sign}(8 - 16) \cdot 0.0625 \\
 &\quad + |12 - 16|^{\alpha} \text{sign}(12 - 16) \cdot 0.125 \\
 &= 4^{\alpha}(0.125) + 8^{\alpha}(-0.125) + 12^{\alpha}(0.125) + 16^{\alpha}(-0.0625).
 \end{aligned}$$

Figure 1 shows how the switching rate between RWS and the other strategies varies with α . We see that RWS has an advantage over *all* other strategies for $0 < \alpha < 1$, suggesting that RWS should be able to win any evolutionary competition in these environments. Again, this prediction is supported by the simulation results in Section 5.

Table IV. Payoff distributions of the six pure strategies in the sequential lottery game with $n = 2$ and $0 \leq p \leq 1$.

		RwS			RwR		
Payoff V		12	8	0	16	8	4
Prob. $P(V)$		p	$(1-p)p$	$(1-p)^2$	p^2	$(1-p)p$	$(1-p)$

		RS, SR		SS	RR	
Payoff V		12	4	8	16	8
Prob. $P(V)$		p	$(1-p)$	1	p^2	$2(1-p)p$
						$(1-p)^2$

 Table V. Payoff distribution for RWS vs. RR in the sequential lottery game with $n = 2$ and arbitrary p . Entries in the table show give the probabilities that each pair of payoffs occurs.

		RwS		
		12	8	0
16		p^3	$(1-p)p^3$	$(1-p)^2p^2$
8		$2(1-p)p^2$	$2(1-p)^2p^2$	$2(1-p)^3p$
0		$(1-p)^2p$	$(1-p)^3p$	$(1-p)^4$

4.4 Arbitrary p

We now consider the case where p , the probability of winning the risky lottery, is any number between 0 and 1 (whence the risky lottery's expected value is between 0 and 8). Table IV gives the probability distributions for each pure strategy.

We can also construct a new probability matrix for each pairing, such as the one for RWS vs. RR shown in Table V. We can then compute the switching rate for our pairing as before. For example, the switching rate for RWS vs. RR is now

$$\begin{aligned}
 \nabla_{\alpha}(\text{RwR}, \text{RR}) &= |0 - 0|^{\alpha} \text{sign}(0 - 0) \cdot (1 - p)^4 \\
 &\quad + |8 - 0|^{\alpha} \text{sign}(8 - 0) \cdot (1 - p)^3 p \\
 &\quad + |12 - 0|^{\alpha} \text{sign}(12 - 0) \cdot (1 - p)^2 p \\
 &\quad + |0 - 8|^{\alpha} \text{sign}(0 - 8) \cdot 2(1 - p)^3 p \\
 &\quad + |8 - 8|^{\alpha} \text{sign}(8 - 8) \cdot 2(1 - p)^2 p^2 \\
 &\quad + |12 - 8|^{\alpha} \text{sign}(12 - 8) \cdot 2(1 - p)p^2 \\
 &\quad + |0 - 16|^{\alpha} \text{sign}(0 - 16) \cdot (1 - p)^2 p^2 \\
 &\quad + |8 - 16|^{\alpha} \text{sign}(8 - 16) \cdot (1 - p)p^3 \\
 &\quad + |12 - 16|^{\alpha} \text{sign}(12 - 16) \cdot p^3 \\
 &= 4^{\alpha}(2(1 - p)p^2 - p^3) + 8^{\alpha}((1 - p)^3 p - 2(1 - p)^3 p - (1 - p)p^3) \\
 &\quad + 12^{\alpha}(1 - p)^2 p + 16^{\alpha}(1 - p)^2 p^2
 \end{aligned}$$

Notice that when $p > 0.5$, the risky lottery has a higher expected value than the safe lottery, and the opposite is true when $p < 0.5$. Thus, RR has the highest expected value when $p > 0.5$, and SS has the highest expected value when $p < 0.5$.

Surprisingly, even though RWS has a suboptimal expected value when $p \neq 0.5$, by examining the switching rates we can see that it still has an evolutionary advantage over both SS and RR for many values of p and α . Figure 2 shows the values of p and α for which $\nabla_{\alpha}(\text{RWS}, \text{RR}) > 0$ and $\nabla_{\alpha}(\text{RWS}, \text{SS}) > 0$, meaning that for




Fig. 2. The shaded area indicates values of p and α for which the switching rate between RWS and the expected-value maximizing strategy (i.e. RR if $p \geq 0.5$, SS otherwise) is greater than 0. RWS is at a disadvantage in terms of expected utility when $p \neq 0.5$, but it still manages to retain its evolutionary advantage for a wide range of values.

these values of p and α , RWS has an evolutionary advantage over expected-value maximizing strategies. Under the imitate-the-better dynamic ($\alpha = 0$), the range is surprisingly large. For example, RWS outperforms SS even when $p = 0.4$, and SS has a significantly higher expected value than RWS. As α increases, the range shrinks at a roughly linear rate, disappearing at $\alpha = 1$ (i.e., the replicator dynamic).

4.5 Evolutionary Stability of State-Dependent Risk Behavior

In this section we discuss whether RWS is an *evolutionarily stable strategy* (ESS) in the 2-lottery game when $p = 0.5$ and $0 < \alpha < 1$. We first give the classical definition of an ESS, explain why it cannot be directly applied to n -lottery games when $\alpha \neq 1$, and propose an intuitive modification to make it applicable. We then show that RWS fits our modified definition of an ESS.

4.5.1 Evolutionary Stability under Imitation Dynamics. In an evolutionary game, a population of agents using an *evolutionarily stable strategy* (ESS) is resilient against a small invading population using any other strategy [Smith 1982]. According to Maynard Smith, strategy S is an ESS if for every strategy $T \neq S$, one of the following conditions holds:

- $E(S, S) > E(T, S)$;
- $E(S, S) = E(T, S)$ and $E(S, T) > E(T, T)$;

where $E(X, Y)$ is the expected payoff an agent receives by playing strategy X against strategy Y [Smith 1982].

This definition does not apply directly to the n -lottery game, because it assumes we are using the replicator dynamic, for which the expected payoff of a strategy pairing is all that is necessary to determine whether one strategy will perform better than the other. As shown in Section 4.4, with population dynamics other than the replicator dynamic, the expected value of a pairing is *not* sufficient to determine which strategy will perform better in the n -lottery game. Thus, we will use the switching rate $\nabla_\alpha(X, Y)$ rather than $E(X, Y)$ when defining evolutionary stability for the n -lottery game under imitation dynamics. We believe this is appropriate

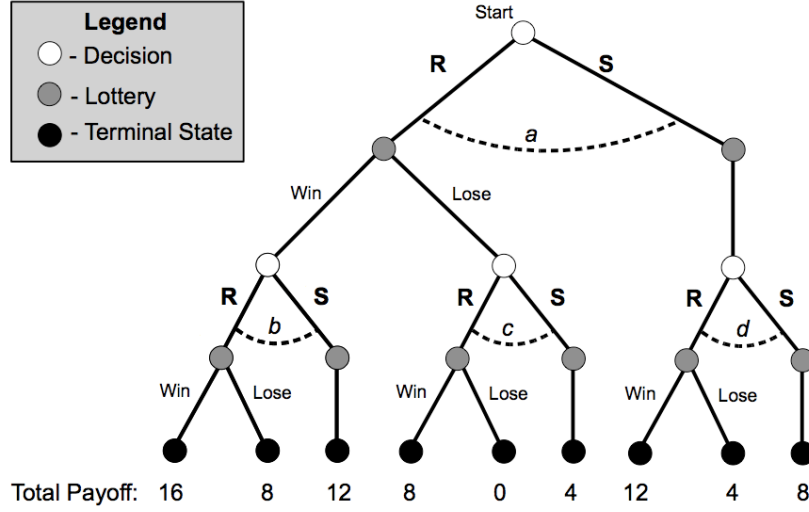


Fig. 3. Representation of the 2-lottery game as a game tree consisting of decision nodes, in which the agent chooses between a risky (**R**) and safe (**S**) lottery, lottery nodes, and terminal nodes, which assign the agent its total payoff. Any strategy in this game can be represented as $S_{a,b,c,d}$, where $a, b, c,$ and d give the probabilities of choosing the risky lottery at each of the four corresponding decision nodes.

because if there are many agents using X and few using Y , then

- $-\nabla_{\alpha}(X, X) > \nabla_{\alpha}(Y, X)$ implies that X will grow faster playing against itself than Y will grow playing against X , so Y will not be able to gain population and will eventually die off;
- $-\nabla_{\alpha}(X, X) = \nabla_{\alpha}(Y, X)$ and $\nabla_{\alpha}(X, Y) > \nabla_{\alpha}(Y, Y)$ implies that X and Y grow at the same rate when playing against X , but X grows faster than Y when playing against Y , so Y will still not be able to gain population and will eventually die off.

and these scenarios correspond to the two conditions for X to be an ESS in the classical definition.

Replacing expected value with switching rates in the definition above gives us the following two conditions for S to be an ESS:

- $-\nabla_{\alpha}(S, S) > \nabla_{\alpha}(T, S)$ or
- $-\nabla_{\alpha}(S, S) = \nabla_{\alpha}(T, S)$ and $\nabla_{\alpha}(S, T) > \nabla_{\alpha}(T, T)$

for all $T \neq S$. Since $\nabla_{\alpha}(S, S) = 0$ and $\nabla_{\alpha}(S, T) = -\nabla_{\alpha}(T, S)$ for all S and T , we can combine these conditions and define ESS in the n -lottery game as follows:

DEFINITION 1. A strategy S is an evolutionarily stable strategy (ESS) in the n -lottery game if, for any strategy $T \neq S$, $\nabla_{\alpha}(S, T) > 0$.

4.5.2 RWS is Evolutionarily Stable. In this section we show that RWS is an ESS by Definition 1 for the 2-lottery game with $p = 0.5$ and $0 < \alpha < 1$. To do this, we must show that it has a positive switching rate with an arbitrary strategy.

Therefore, our first step must be to devise a method for representing an arbitrary strategy for the 2-lottery game. We have found that, if the decisions an agent makes and the possible lottery outcomes are arranged into a game tree as shown in Figure 3, then any strategy can be expressed as $S_{a,b,c,d}$, where a, b, c , and d give the agent's probability of choosing the risky lottery at each of the four decision nodes indicated in the figure. For instance, the pure strategies we have been dealing with thus far can be represented as follows (here, a “-” in place of one of the four probabilities indicates that any value is acceptable, since the decision corresponding to that probability is never reached):

- SS is $S_{0,-,-,0}$
- RR is $S_{1,1,1,-}$
- RS is $S_{1,0,0,-}$
- SR is $S_{0,-,-,1}$
- RWS is $S_{1,0,1,-}$
- RWR is $S_{1,1,0,-}$

We can now calculate the switching rate between RWS and an arbitrary strategy $S_{a,b,c,d}$ in terms of a, b, c, d , and α in the same way we calculated $\nabla_\alpha(\text{RWR}, \text{RR})$ and $\nabla_\alpha(\text{RWR}, \text{SS})$ in Section 4.3. This comes out to:⁶

$$\begin{aligned} \nabla_\alpha(\text{RWR}, S_{a,b,c,d}) &= \left[\frac{1}{8}(2 * 8^\alpha - 12^\alpha - 4^\alpha) \right] (a(1-c) + (1-a)d) \quad (3) \\ &\quad + \left[\frac{1}{4}(2 * 4^\alpha - 8^\alpha) \right] (1-a)(1-d) \\ &\quad + \left[\frac{1}{16}(2 * 12^\alpha + 2 * 4^\alpha - 2 * 8^\alpha - 16^\alpha) \right] ab. \end{aligned}$$

The three bracketed terms are all strictly greater than 0 for $0 < \alpha < 1$; in fact, they are identical to the curves in Figure 1 for SR/RS, SS, and RR, respectively. Intuitively, this says that the switching rate for the arbitrary strategy is just the probability that the strategy follows the pure strategies SS, SR, RS, and RR, times the switching rate between RWS and each of those strategies. Given that the bracketed terms are strictly positive, it is easy to see that $\nabla_\alpha(\text{RWR}, S_{a,b,c,d}) > 0$ unless $a = 1$, $b = 0$, and $c = 1$. Since $S_{1,0,1,-}$ is equivalent to RWS, this means that $\nabla_\alpha(\text{RWR}, S_{a,b,c,d}) > 0$ for all $S_{a,b,c,d} \neq \text{RWS}$. Thus, RWS is an evolutionarily stable strategy by Definition 1.

4.6 Other Values of n

With the exception of Section 4.5.1, our theoretical development has been largely restricted to the case $n = 2$. We now discuss briefly what happens for other values of n .

The case $n = 1$ is relatively trivial: there are only two pure strategies, both are unconditional, and both perform equally well (for more details, see [Roos and Nau 2010a]).

⁶We give a full derivation in the Appendix

The case $n > 2$ is very hard to analyze, because the number of pure strategies is super-exponential in n . However, intuition suggests that the behavior pattern exhibited by RWS for $n = 2$, namely to play safe when having done well and risky otherwise, should also have an advantage when $n > 2$. We discuss some pilot experiments that support this in Section 5.3.

5. SIMULATIONS

To further investigate the population dynamics in a population consisting of all pure strategies, we ran computer simulations of agent-based models playing our two-choice evolutionary lottery game. These simulations and results are described in the next section.

5.1 Setup and Implementation

Our simulations for the two-choice lottery game environment explore population evolution under a variety of parameter combinations of α (the imitation parameter) and p (the probability of winning the risky lottery). The types of agents included were the six pure strategies for the two-choice game described earlier. All simulations started with an initial population of 1000 agents for each agent type.

To model the imitation dynamics given by Eq. (2) in our finite population agent-based model, we used a pairwise comparison process [Traulsen and Hauert 2009] to model the transmission of strategies among agents. Each generation, after all agents have received payoffs from chosen lotteries, each agent i compares its (individual) payoff $\rho(i)$ to that of a randomly drawn agent from the population j (with payoff $\rho(j)$) and adopts the strategy of this agent with a probability q if $\rho(j) > \rho(i)$. In order to achieve the parameterized dynamics given by Eq. (1) in our agent-based model, we use

$$q = [|\rho(j) - \rho(i)|/\Delta\rho]^\alpha, \quad 0 \leq \alpha \leq 1,$$

where $\Delta\rho$ is the highest possible difference in payoff. Figure 4 provides pseudocode on how this pairwise comparison imitation process was implemented.

5.2 Results

Figure 5(a,b,c) shows the results for simulations with $p = 0.5$ for each $\alpha = 1$ (the replicator dynamic), $\alpha = 0$ (the imitate-the-better dynamic), and $\alpha = 0.5$. Each plot is an average over 20 simulation runs (the amount of variation from one run to another was quite small). These experiments confirm our analysis from Section 4.3, which shows that RWS has an evolutionarily advantageous risk behavior under any $0 \leq \alpha < 1$.

As predicted by our analysis, RWS outperformed the other strategies evolutionarily except when $\alpha = 1$:

- For $\alpha = 1$, all of the strategies performed equally well and remained at their initial population counts.
- For $\alpha = 0$, the state-dependent strategy RWS outperformed the other strategies. RWS rose in population proportion relatively quickly to comprise the majority ($> 2/3$) of the population and remained there throughout subsequent generations. Furthermore, the two unconditional strategies SR and RS remained, comprising the proportion of the population not taken over by RWS.

```

LOOP for each agent in population
  Select current agent (imitator) from population
  Select random agent (observed) from pop
  Let  $\rho_{im}$  = payoff achieved by imitator
  Let  $\rho_{obs}$  = payoff achieved by observed
  Let  $\Delta\rho$  = highest possible payoff difference
  Let  $r$  = random number in  $[0, 1)$ 
  Let  $q = [|\rho_{im} - \rho_{obs}|/\Delta\rho]^\alpha$ 
  // Imitate observed agent with probability  $q$  based on payoff comparison
  IF  $\rho_{im} < \rho_{obs}$  AND  $r \leq q$  THEN
    Place an agent following the strategy of observed into next generation
  ELSE
    Place an agent following the strategy of imitator into next generation
  END IF
END LOOP

```

Fig. 4. Pseudocode for the pairwise comparison imitation process used to reproduce agents into the next generation in our evolutionary simulations. Takes parameter $0 \leq \alpha \leq 1$.

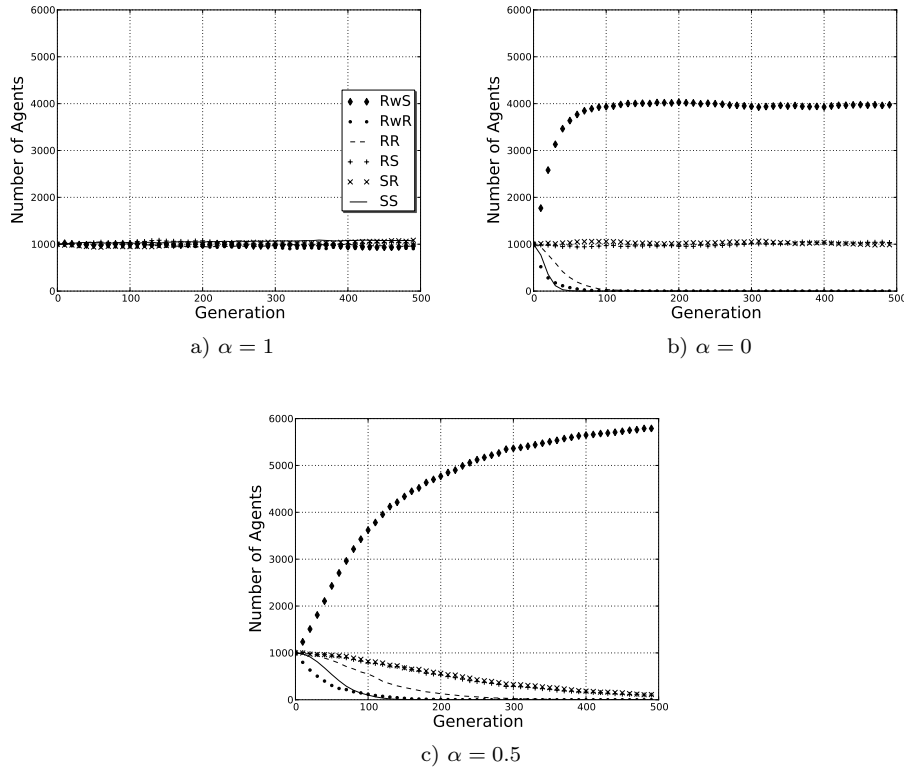


Fig. 5. Results of simulations for different α values for the population dynamic and $p = 0.5$. Each plot shows the number of agents of each type over the course of evolution for 500 generations. Simulation runs for $\alpha = 0.2, 0.4, 0.6,$ and 0.8 produced qualitatively identical results to the case of $\alpha = 0.5$.

- For $\alpha = 0.5$, the RWS agent population grew similarly as for $\alpha = 0$, but here RWS also had an advantage against SR and RS (as indicated by Figure 1) and thus continued to grow to comprise 100% of the population.

We also ran simulations with $\alpha = 0.2, 0.4, 0.6$, and 0.8 . The results for these α values are all essentially equal to the case of $\alpha = 0.5$. The only difference is that the rate at which RWS grows to take over the population is inversely related to α (i.e. for larger α values, it takes longer for RWS to take over the population).

In order to explore lottery games in which the risky lottery has a different expected value than the safe lottery, we also ran experiments with $p = 0.3, 0.4, 0.55$, and 0.7 . These values were chosen because for $\alpha = 0$, two of them lie within ($p = 0.4, 0.55$) the shaded area of Figure 2 and the other two ($p = 0.3, 0.7$) lie outside that area. Recall that the shaded area is the region for which our analysis predicts that RWS has an evolutionary advantage over the expected-value-maximizing strategy, and the unshaded areas are regions for which our analysis predicts the reverse.

As shown in Figure 6(a,b,c,d), the simulation results confirm the theoretical predictions. More specifically:

- For $p = 0.2$ (Figure 6(a)), SS is the expected-value maximizing strategy and it takes over the population.
- For $p = 0.7$ (Figure 6(d)), RR is the expected-value maximizing strategy and it takes over the population.
- Even though SS is the expected-value-maximizing strategy for $p = 0.4$ (Figure 6(b)) and RR is for $p = 0.55$ (Figure 6(c)), in both cases RWS has an evolutionary advantage and takes over the population.

In Figures 6(b,c,d), some fluctuations occur before stabilization. These occur because of the differing amounts of evolutionary advantage that different strategies have over others. For example, a strategy a may grow in number temporarily because it has an advantage over another strategy b . But once b becomes extinct (or sufficiently small in number), a will diminish because some other strategy c has an advantage over a .

5.3 Simulations for $n > 2$

In Section 4.6, we hypothesized that RWS's behavior pattern, namely to be risk-averse when it has done well and risk-seeking when it has done badly, may be advantageous in lottery games with $n > 2$. To test this hypothesis, we ran a pilot experiment for the case $n = 4$, using the six strategies shown in Table VI.

For our simulations we used an initial population of 1000 agents of each type, and the parameters $p = 0.5$ for the risky lottery and $\alpha = 0$ for the population dynamic. The results were qualitatively the same as the ones in Figure 5(c): the RRWSS strategy dominated the other strategies and grew to comprise 100% of the population. This would seem to confirm our hypothesis—but since there are hundreds of pure strategies when $n = 4$ and we only looked at six of the simpler ones, the result should be regarded as purely preliminary.

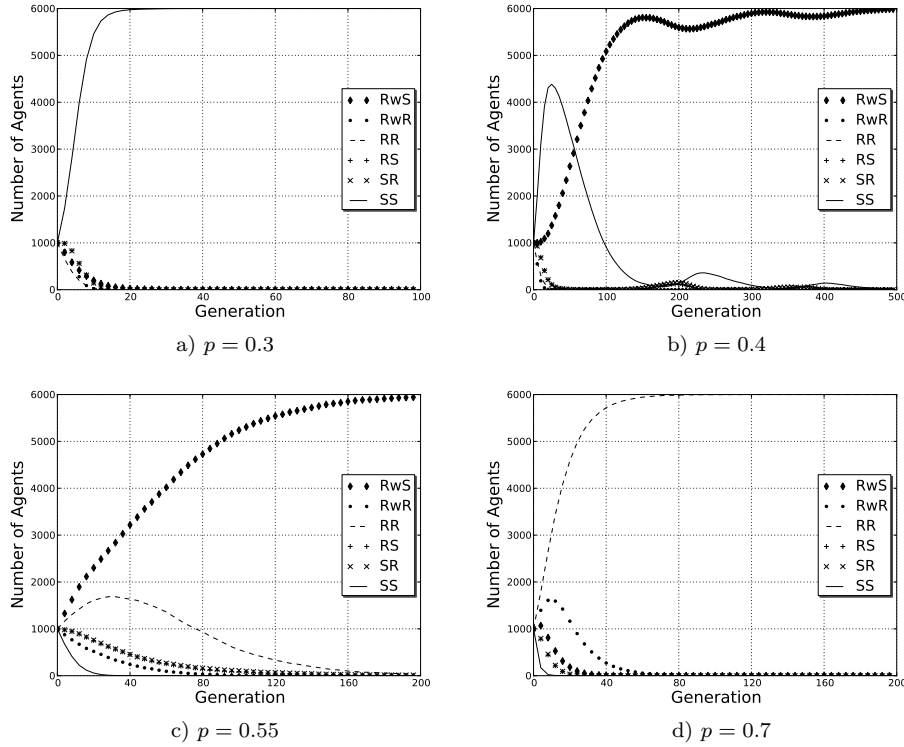


Fig. 6. Results of simulations with four different p values for the risky lottery and $\alpha = 0$. Each plot shows the number of agents of each type over the course of evolution for 500 generations.

Table VI. Pure strategies included in our lottery game simulations for $n = 4$.

Strategy	Behavior
SSSS	always choose safe
RRRR	always choose risky
SSRR	choose safe in lotteries 1 & 2, then choose risky in 3 & 4
RRSS	choose risky in lotteries 1 & 2, then choose safe in 3 & 4
RRwSS	choose risky in lotteries 1 & 2, then choose safe in 3 & 4 only if 1 & 2 were won, else choose risky
RRwRR	choose risky in lotteries 1 & 2, then choose risky in 3 & 4 only if 1 & 2 were lost, else choose safe

6. CONCLUSIONS

Our results show that for imitation dynamics other than the pure replicator dynamic, there are evolutionary game environments in which the RWS strategy has an evolutionary advantage over expected-value maximization. Since RWS's risk-taking behavior is similar to the risk preferences captured in several prominent models of human decision making, this suggests that population dynamics other than the replicator dynamic may model an important mechanism for the emergence of those risk preferences.

There are several ways in which our work could be extended:

- Our study was restricted to simple lottery games, in order to avoid confounding factors that may be present in more complex social games (e.g. trust, reputation, kinship, etc.). It would be interesting to explore the relationships among reproduction dynamics, risk preferences, and social interactions such as cooperation. We have already begun some work along this line: [Roos and Nau 2010a] shows that risk behavior similar to RWS can promote the evolution of cooperation in a situation where the cooperation requires a risky decision (namely, choosing to cooperate).
- It will be important to examine other population dynamics in which a strategy's reproductive success is not always proportional to its expected payoff. For example, if a death rate (e.g. [Nowak and Sigmund 1993]) is implemented as a payoff-dependent threshold function, we might expect risk propensities to differ depending on whether an agent is above or below that threshold, in a manner similar to behavior above or below an aspiration level in SP/A theory.
- Our study focused on the case $n = 2$, i.e., in each generation the agents made two decisions. We believe that state-dependent risk preferences like those of RWS should also have an advantage when $n > 2$, and Section 5.3 discussed some pilot experiments that support this intuition. Conducting more extensive studies will be a topic for future work.
- We assumed a well-mixed population in which every agent was able and equally likely to imitate any other. It would be interesting to explore the possible effects of social or physical structures (that may guide or constrain imitation) on the evolution of risk behavior.
- Finally, it may be useful to conduct empirical studies of which type of imitation dynamics best models human imitation propensities. The insights of this paper combined with such knowledge have potential application in domains where human decision making under risk is of interest, including economics, medicine, military and public policy.

Acknowledgements

We wish to thank the anonymous reviewers for their helpful suggestions. This work was supported in part by AFOSR grant FA95500610405 and NAVAIR contract N6133906C0149. The opinions in this paper are those of the authors and do not necessarily reflect the opinions of the funders.

REFERENCES

- ABEL, A. 1990. Asset prices under habit formation and catching up with the Joneses. *The American Economic Review* 80, 2, 38–42.
- ALEXANDER, J. M. 2009. Evolutionary game theory. In *The Stanford Encyclopedia of Philosophy*, Fall 2009 ed., E. N. Zalta, Ed.
- APESTEGUIA, J., HUCK, S., AND OECHSSLER, J. 2007. Imitation—theory and experimental evidence. *Journal of Economic Theory* 136, 1 (Sept.), 217–235.
- ARROW, K. J. 1971. *Essays in the theory of risk-bearing*. Markham, Chicago.
- AXELROD, R. AND HAMILTON, W. D. 1981. The evolution of cooperation. *Science* 211, 4489 (March), 1390–1396.
- BINMORE, K. 1998. The evolution of fairness norms. *Rationality and Society* 10, 275–302.

- BLUME, L. 1993. The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5, 3, 387–424.
- BOWLES, S. AND GINTIS, H. 2004. The evolution of strong reciprocity: cooperation in heterogeneous populations. *Theoretical Population Biology* 65, 1 (Feb.), 17–28.
- DAWID, H. 2007. Evolutionary game dynamics and the analysis of agent-based imitation models: The long run, the medium run and the importance of global analysis. *Journal of Economic Dynamics and Control* 31, 6 (June), 2108–2133.
- EGULUZ, V. M., ZIMMERMANN, M. G., CELA-CONDE, C. J., AND MIGUEL, M. S. 2005. Cooperation and emergence of role differentiation in the dynamics of social networks. *American Journal of Sociology* 110, 977 (Feb.).
- EPSTEIN, L. AND ZIN, S. 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica: Journal of the Econometric Society*, 937–969.
- EPSTEIN, L. AND ZIN, S. 1991. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of Political Economy*, 263–286.
- FANG, C., KIMBROUGH, S. O., PACE, S., VALLURI, A., AND ZHENG, Z. 2002. On adaptive emergence of trust behavior in the game of stag hunt. *Group Decision and Negotiation* 11, 6 (Nov.), 449–467.
- FRIEDMAN, M. AND SAVAGE, L. J. 1948. The utility analysis of choices involving risk. *The Journal of Political Economy* 56, 4 (Aug.), 279–304.
- GINTIS, H. 2000. *Game Theory Evolving*. Princeton University Press, Princeton, New Jersey.
- HAUERT, C. AND SZABO, G. 2005. Game theory and physics. *American Journal of Physics* 73, 5, 405–414.
- HOFBAUER, J. AND SIGMUND, K. 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, United Kingdom.
- HUBERMAN, B. A. AND GLANCE, N. S. 1993. Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences of the United States of America* 90, 16 (Aug.), 7716–7718.
- HUCK, S., NORMANN, H., AND OECHSSLER, J. 2000. Does information about competitors' actions increase or decrease competition in experimental oligopoly markets? *International Journal of Industrial Organization* 18, 1, 39–57.
- HUCK, S., NORMANN, H. T., AND OECHSSLER, J. 1997. Learning in cournot oligopoly: An experiment. *SSRN eLibrary*.
- KAHNEMAN, D. AND TVERSKY, A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–291.
- KIMBROUGH, S. O. 2005. Foraging for trust: Exploring rationality and the stag hunt game. In *Trust Management*. 1–16.
- LOOMES, G. AND SUGDEN, R. 1982. Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal* 92, 805–824.
- LOPES, L. L. 1987. Between hope and fear: The psychology of risk. *Advances in Experimental Social Psychology* 20, 255–295.
- LOPES, L. L. 1990. Re-modeling risk aversion. In *Acting under uncertainty: Multidisciplinary conceptions*, G. M. von Furstenberg, Ed. Boston: Kluwer, 267–299.
- LOPES, L. L. AND ODEN, G. C. 1999. The role of aspiration level in risky choice: A comparison of cumulative prospect theory and sp/a theory. *Journal of Mathematical Psychology* 43, 286–313.
- NOWAK, M. A., BONHOEFFER, S., AND MAY, R. M. 1994. Spatial games and the maintenance of cooperation. *Proceedings of the National Academy of Sciences of the United States of America* 91, 11 (May), 4877–4881.
- NOWAK, M. A. AND MAY, R. M. 1992. Evolutionary games and spatial chaos. *Nature* 359, 6398 (Oct.), 826–829.
- NOWAK, M. A. AND SIGMUND, K. 1993. A strategy of win-stay, lose-shift that outperforms tit for tat in the prisoner's dilemma game. *Nature* 364, 56–58.

- OFFERMAN, T., POTTERS, J., AND SONNEMANS, J. 2002. Imitation and belief learning in an oligopoly experiment. *The Review of Economic Studies* 69, 4 (Oct.), 973–997. ArticleType: primary_article / Full publication date: Oct., 2002 / Copyright 2002 The Review of Economic Studies Ltd.
- PAGE, K. AND NOWAK, M. 2002. Empathy leads to fairness. *Bulletin of Mathematical Biology* 64, 6, 1101–1116.
- RABIN, M. 2000. Risk aversion and Expected-Utility theory: A calibration theorem. *Econometrica* 68, 5 (Sept.), 1281–1292.
- RIOLO, R. L., COHEN, M. D., AND AXELROD, R. 2001. Evolution of cooperation without reciprocity. *Nature* 411, 441–443.
- ROOS, P. AND NAU, D. S. 2009. Conditionally risky behavior vs expected value maximization in evolutionary games. In *Sixth Conference of the European Social Simulation Association (ESSA 2009)*.
- ROOS, P. AND NAU, D. S. 2010a. Risk preference and sequential choice in evolutionary games. *Advances in Complex Systems*. To appear.
- ROOS, P. AND NAU, D. S. 2010b. State-dependent risk preferences in evolutionary games. In *2010 International Conference on Social Computing, Behavioral Modeling, and Prediction*.
- SCHLAG, K. 1999. Which one should I imitate? *Journal of Mathematical Economics* 31, 4, 493–522.
- SCHLAG, K. H. 1998. Why imitate, and if so, how? a bounded rational approach to Multi-Armed bandits. *Journal of Economic Theory* 78, 1, 130–156.
- SCHOTTER, A. AND SOPHER, B. 2003. Social learning and coordination conventions in intergenerational games: An experimental study. *Journal of Political Economy* 111, 3.
- SMITH, J. M. 1982. *Evolution and the theory of games*. Cambridge University Press, Cambridge, United Kingdom.
- SZABÓ, G. AND TÓKE, C. 1998. Evolutionary prisoner’s dilemma game on a square lattice. *Phys. Rev. E* 58, 1 (Jul), 69–73.
- TAYLOR, P. D. AND JONKER, L. B. 1978. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences* 40, 1-2, 145 – 156.
- TRAUlsen, A. AND HAUERT, C. 2009. Stochastic evolutionary game dynamics. In *Reviews of Nonlinear Dynamics and Complexity*, H.-G. Shuster, Ed. Wiley-VHC, New York.
- TRAUlsen, A., PACHECO, J. M., AND NOWAK, M. A. 2007. Pairwise comparison and selection temperature in evolutionary game dynamics. *Journal of Theoretical Biology* 246, 3, 522–529.
- TVERSKY, A. AND KAHNEMAN, D. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5, 297–323.
- VEGA-REDONDO, F. 1997. The evolution of walrasian behavior. *Econometrica* 65, 2 (Mar.), 375–384. ArticleType: primary_article / Full publication date: Mar., 1997 / Copyright 1997 The Econometric Society.
- VON NEUMANN, J. AND MORGENSTERN, O. 1944. *Theories of games and economic behavior*. Princeton University Press Princeton, NJ.

7. APPENDIX A

In this section we present the derivation of $\nabla_{\alpha}(\text{RWS}, S_{a,b,c,d})$ used in Section 4.5. Using Table IV and Figure 3, we can determine the probability of each pair of

payoffs occurring and use them for the values of $p(r, s)$ as follows:

$$\begin{aligned}
\nabla_{\alpha}(\text{RwS}, S_{a,b,c,d}) = & |12 - 8|^{\alpha}(p)(apb(1-p) + apc(1-p) + (1-a)(1-d)) \\
& + |12 - 4|^{\alpha}(p)(a(1-p)(1-c) + (1-a)d(1-p)) \\
& + |12 - 0|^{\alpha}(p)(a(1-p)c(1-p)) \\
& + |8 - 4|^{\alpha}(p(1-p))(a(1-p)(1-c) + (1-a)d(1-p)) \\
& + |8 - 0|^{\alpha}(p(1-p))(a(1-p)c(1-p)) \\
& - |16 - 12|^{\alpha}p(apbp) \\
& - |16 - 8|^{\alpha}(1-p)p(apbp) \\
& - |16 - 0|^{\alpha}(1-p)^2(apbp) \\
& - |12 - 8|^{\alpha}(1-p)p(ap(1-b) + (1-a)dp) \\
& - |12 - 0|^{\alpha}(1-p)^2(ap(1-b) + (1-a)dp) \\
& - |8 - 0|^{\alpha}(1-p)^2(apb(1-p) + a(1-p)cp + (1-a)(1-d)) \\
& - |4 - 0|^{\alpha}(1-p)^2(a(1-p)(1-c) + (1-a)d(1-p))
\end{aligned}$$

Since we are considering the case where $p = 0.5$, we can collect terms as though $p = (1-p)$ (for ease of exposition we will wait to substitute 0.5 for p):

$$\begin{aligned}
\nabla_{\alpha}(\text{RwS}, S_{a,b,c,d}) = & |16|^{\alpha}(0 - abp^4) + |12|^{\alpha}(acp^3 - a(1-b)p^3 - (1-a)dp^3) \\
& + (|4|^{\alpha}p - |8|^{\alpha}p^2)(abp^2 + acp^2 + (1-a)(1-d)) \\
& + (|8|^{\alpha}p + |4|^{\alpha}p^2 - |4|^{\alpha}p^2)(a(1-c)p + (1-a)dp) \\
& + |8|^{\alpha}(acp^4 + abp^4) \\
& + |4|^{\alpha}(0 - abp^3 - a(1-b)p^3 - (1-a)dp^3),
\end{aligned}$$

which yields

$$\begin{aligned}
\nabla_{\alpha}(\text{RwS}, S_{a,b,c,d}) = & - |16|^{\alpha}(abp^4) + |12|^{\alpha}(acp^3 - a(1-b)p^3 - (1-a)dp^3) \\
& + |8|^{\alpha}(acp^4 + a(1-c)p^2 + (1-a)dp^2 \\
& - abp^4 - abp^4 - acp^4 - (1-a)(1-d)p^2) \\
& + |4|^{\alpha}(abp^3 + acp^3 + (1-a)(1-d)p - abp^3 - a(1-b)p^3 - (1-a)dp^3),
\end{aligned}$$

which yields

$$\begin{aligned}
\nabla_{\alpha}(\text{RwS}, S_{a,b,c,d}) = & - |16|^{\alpha}(abp^4) + |12|^{\alpha}(acp^3 - a(1-b)p^3 - (1-a)dp^3) \\
& + |8|^{\alpha}(a(1-c)p^2 + (1-a)dp^2 - abp^4 - abp^4 - (1-a)(1-d)p^2) \\
& + |4|^{\alpha}(acp^3 + (1-a)(1-d)p - a(1-b)p^3 - (1-a)dp^3).
\end{aligned}$$

Recollecting terms gives us

$$\begin{aligned}
\nabla_{\alpha}(\text{RwS}, S_{a,b,c,d}) = & (12^{\alpha} + 4^{\alpha})acp^3 + 8^{\alpha}a(1-c)p^2 \\
& + 4^{\alpha}(1-a)(1-d)p + 8^{\alpha}(1-a)dp^2 \\
& - (4^{\alpha} + 12^{\alpha})(1-a)dp^3 - 8^{\alpha}(1-a)(1-d)p^2 \\
& - (16^{\alpha} + 2 * 8^{\alpha})abp^4 - (12^{\alpha} + 4^{\alpha})a(1-b)p^3.
\end{aligned}$$

Substituting $p = \frac{1}{2}$ and expanding the final term, we get

$$\begin{aligned}\nabla_{\alpha}(\text{RwS}, S_{a,b,c,d}) &= \frac{1}{8}(12^{\alpha} + 4^{\alpha})ac + \frac{1}{4}8^{\alpha}a(1-c) \\ &\quad + \frac{1}{4}(2 * 4^{\alpha} - 8^{\alpha})(1-a)(1-d) + \frac{1}{8}(2 * 8^{\alpha} - 4^{\alpha} - 12^{\alpha})(1-a)d \\ &\quad - \frac{1}{16}(16^{\alpha} + 2 * 8^{\alpha})ab - \frac{1}{8}(12^{\alpha} + 4^{\alpha})a + \frac{1}{8}(12^{\alpha} + 4^{\alpha})ab,\end{aligned}$$

which yields

$$\begin{aligned}\nabla_{\alpha}(\text{RwS}, S_{a,b,c,d}) &= \frac{1}{8}(12^{\alpha} + 4^{\alpha})(ac - a) + \frac{1}{4}8^{\alpha}a(1-c) \\ &\quad + \frac{1}{4}(2 * 4^{\alpha} - 8^{\alpha})(1-a)(1-d) + \frac{1}{8}(2 * 8^{\alpha} - 4^{\alpha} - 12^{\alpha})(1-a)d \\ &\quad + \frac{1}{16}(2 * 12^{\alpha} + 2 * 4^{\alpha} - 2 * 8^{\alpha} - 16^{\alpha})ab.\end{aligned}$$

Since $(ac - a) = -a(1 - c)$, this yields

$$\begin{aligned}\nabla_{\alpha}(\text{RwS}, S_{a,b,c,d}) &= \frac{1}{8}(2 * 8^{\alpha} - 4^{\alpha} - 12^{\alpha})a(1-c) \\ &\quad + \frac{1}{4}(2 * 4^{\alpha} - 8^{\alpha})(1-a)(1-d) + \frac{1}{8}(2 * 8^{\alpha} - 4^{\alpha} - 12^{\alpha})(1-a)d \\ &\quad + \frac{1}{16}(2 * 12^{\alpha} + 2 * 4^{\alpha} - 2 * 8^{\alpha} - 16^{\alpha})ab.\end{aligned}$$

Finally, recollecting terms gives us

$$\begin{aligned}\nabla_{\alpha}(\text{RwR}, S_{a,b,c,d}) &= \frac{1}{8}(2 * 8^{\alpha} - 12^{\alpha} - 4^{\alpha})(a(1-c) + (1-a)d) \\ &\quad + \frac{1}{4}(2 * 4^{\alpha} - 8^{\alpha})(1-a)(1-d) \\ &\quad + \frac{1}{16}(2 * 12^{\alpha} + 2 * 4^{\alpha} - 2 * 8^{\alpha} - 16^{\alpha})ab,\end{aligned}$$

which matches Equation 3.