

Smart Representations: Rationality and Evolution in a Richer Environment

Paolo Galeazzi and Michael Franke*†

Standard applications of evolutionary game theory look at a single game and focus on the evolution of behavior for that game alone. Instead, this article uses tools from evolutionary game theory to study the competition between choice mechanisms in a rich and variable multigame environment. A choice mechanism is a way of subjectively representing a decision situation, paired with a method for choosing an act based on this subjective representation. We demonstrate the usefulness of this approach by a case study that shows how subjective representations in terms of regret that differ from the actual fitness can be evolutionarily advantageous.

1. Introduction. If agents deal with a rich and variable environment, they have to face many different choice situations. Standard evolutionary game models frequently simplify reality in at least two ways. First, the environment is represented as a fixed stage game; second, the focus of evolutionary selection is behavior for that stage game alone. In contrast, some argue for

Received February 2016; revised June 2016.

*To contact the authors, please write to: Paolo Galeazzi, ILLC, University of Amsterdam, The Netherlands; e-mail: pagale87@gmail.com. Michael Franke, Department of Linguistics, University of Tübingen, Germany; e-mail: mchfranke@gmail.com.

†The research of Paolo Galeazzi leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no. 283963. Michael Franke's work was supported by the Institutional Strategy of the University of Tübingen (DFG, ZUK 63). We are also grateful to Gerhard Jäger for funding this work. We want to thank all the participants in Decisions, Games and Logic 2015 workshop at London School of Economics, in LogiCIC workshop 2014 at the University of Amsterdam, and in Bridges 2 workshop at Rutgers University for the valuable feedback. Special thanks go to Pierpaolo Battigalli, Massimo Marinacci, Valerio Nispi Landi, Tommaso Orlando, Eric Pacuit, Erik Quaeghebeur, Rasmus Rendsvig, Teddy Seidenfeld, Peter Wakker, and Kevin Zollman for their attention and comments. We are also grateful for the helpful comments we received by two anonymous reviewers.

Philosophy of Science, 84 (July 2017) pp. 544–573. 0031-8248/2017/8403-0008\$10.00
Copyright 2017 by the Philosophy of Science Association. All rights reserved.

studying the evolutionary competition of general choice mechanisms in a rich and variable environment (e.g., Hammerstein and Stevens 2012; Fawcett, Hamblin, and Giraldeau 2013; McNamara 2013). In response to this and adding to recent like-minded approaches, this article introduces a general meta-game model that conservatively extends the scope of evolutionary game theory to deal with evolutionary selection of choice mechanisms in variable environments (see also Harley 1981; Bednar and Page 2007; Rayo and Becker 2007; Zollman 2008; Skyrms and Zollman 2010; Zollman and Smead 2010; Smead and Zollman 2013; O'Connor 2015).¹

A choice mechanism associates decision situations with action choices. A crucial part of a choice mechanism is the subjective representation of the decision situation, in particular the manner of forming preferences and beliefs about a possibly uncertain world. To show the usefulness of the meta-game approach, this article asks: which preference and belief representations are ecologically valuable and lead to high fitness? The evolution of preferences has been the subject of recent interest in theoretical economics (e.g., Dekel, Ely, and Ylankaya 2007; Robson and Samuelson 2011; Alger and Weibull 2013). Here, we argue that questions of preference evolution should take variability in uncertainty representation into account as well. We demonstrate that if agents have imprecise probabilistic beliefs (e.g., Levi 1974; Gärdenfors and Sahlin 1982; Walley 1996), faithful and objective representations in terms of true evolutionary fitness can be outperformed by subjective (e.g., regret-based) preference representations that deviate from the true fitness that natural selection operates on.

The article is organized as follows. Section 2 sets the scene by reviewing different perspectives on rational choice. Section 3 introduces the meta-game approach. In doing so, it covers key notions such as choice mechanisms, decision rules and subjective representations, all with an eye toward the evolutionary application of section 4. Section 5 contains the main results for that application, and section 6 discusses some interesting extensions. Finally, section 7 concludes.

1. Some of these contributions are closely related to ours. Bednar and Page (2007) use a multigame framework, composed of a fixed selection of six possible games, to study the emergence of different cultural behaviors, and model agents as finite-state automata playing games from the fixed selection. Zollman (2008) explains seemingly “irrational” fair behavior in social dilemmas (like the Ultimatum game) by means of a model where agents have to play the Ultimatum game together with the Nash bargaining game, but they are constrained to choose the same strategy for both games. Finally, Rayo and Becker (2007) consider, in a more decision-theoretic setting, what subjective utility function a cognitively limited agent should be endowed with in order to maximize her evolutionary fitness. Our framework can then be viewed as a generalization of those models, mainly in that here players do not necessarily have any specific cognitive limitations, and we allow for larger and possibly variable classes of games.

2. Rationality and Subjective Representations. The standard textbook definition of rationality in economics and decision theory traces back to the seminal work by de Finetti (1937), von Neumann and Morgenstern (1944), and Savage (1954). It says that a choice is rational only if it maximizes (subjective) expected utility. Expected utility is subjective in the sense that it is a function of subjective beliefs and subjective preferences of the decision maker (DM). To wit, a choice can be rational, that is, the best choice from the DM's point of view, even if based on peculiar beliefs and/or aberrant preferences.

If beliefs and preferences are subjective, there is room for rationalization or redescriptionism of observable behavior. For example, in the case of social decision making, including considerations of fairness allows us to describe as rational empirically observed behavior, such as in experimental Prisoner's Dilemmas or public goods games, that might otherwise appear irrational (e.g., Fehr and Schmidt 1999; Charness and Rabin 2002).

The main objection to redescriptionism is that, without additional constraints, the notion of rationality is likely to collapse, as it seems possible to deem rational almost everything that is observed, given the freedom to adjust beliefs and preferences at will. Normativism therefore emphasizes that there are many ways in which ascriptions of beliefs and preferences should be constrained by normative considerations of rationality as well: for example, subjective beliefs should reflect objective chance where possible; subjective preferences should be oriented toward tracking objective fitness. For instance, profit maximization seems a necessary requirement for evolution in a competitive market because only firms behaving according to profit maximization will survive in the long run (e.g., Alchian 1950; Friedman 1953).

An alternative view on rationality of choice is adaptationism (e.g., Anderson 1991; Chater and Oaksford 2000; Hagen et al. 2012). Adaptationism aims to explain rational behavior by appealing to evolutionary considerations: DMs have acquired choice mechanisms that have proved to be adaptive with respect to the variable environment where they have evolved. A choice mechanism can be a set of distinct heuristics (the DM's adaptive toolbox) that have little in common (e.g., Gigerenzer and Goldstein 1996; Tversky and Kahnemann 1981; Scheibehenne, Rieskamp, and Wagenmakers 2013). But to closely relate to the literature on evolution of preferences and to the philosophical debate about the nature of rational choice, we here suggest thinking of a choice mechanism as a map from choice situations to action choices that includes an explicit level of subjective representation of the situation. Specifically, a subjective representation is a general way of forming preferences and beliefs about the choice situation. We are most interested in the question of which subjective representations, and which choice mechanisms in general, are better than others from an evolutionary point of view.

3. Choice Mechanisms and Meta-Games. We view a choice mechanism as the combination of three different things: a subjective utility (or preference), a subjective belief, and a decision rule. In general, the agent’s action choice will depend both on the agent’s utility at different possible outcomes of the choice situation and on the agent’s beliefs about the realization of these outcomes. The decision rule then combines the agent’s subjective utility and belief, and dictates how the agent should act: a decision rule is a function that associates an action choice with the agent’s utility and beliefs:

$$\text{Decision Rule: Utility} \times \text{Beliefs} \rightarrow \text{Actions.}$$

The subjective utility of an agent can be formally expressed by a function $u: W \times A \rightarrow \mathbb{R}$, where A stands for a (finite) set of actions available to the agent and W is a (finite) set of possible states of the world. There are many different ways to describe beliefs, but for concreteness of later applications we here assume that the agent’s beliefs are represented in terms of a (possibly singleton) convex compact set of probability functions $\Gamma \subseteq \Delta(W)$ over the possible states of the world. Given a utility u and a belief Γ , examples of well-known decision rules from the literature that we will encounter later are:

1. **Maxmin:**

$$a^*(u, \Gamma) = \operatorname{argmax}_{a \in A} \min_{\mu \in \Gamma} \sum_{w \in W} u(w, a) \mu(w),$$

2. **Maximax:**

$$a^*(u, \Gamma) = \operatorname{argmax}_{a \in A} \max_{\mu \in \Gamma} \sum_{w \in W} u(w, a) \mu(w),$$

3. **Laplace rule:**

$$a^*(u, \Gamma) = \operatorname{argmax}_{a \in A} \sum_{w \in W} \frac{1}{|W|} u(w, a),$$

4. **Expected utility maximization** (for Γ singleton):

$$a^*(u, \Gamma) = \operatorname{argmax}_{a \in A} \sum_{w \in W} u(w, a) \mu(w).$$

It is worth noticing that both maxmin and maximax boil down to expected utility maximization when the set Γ is a singleton, and in turn expected utility maximization reduces to the Laplace rule when the belief μ is a uniform probability over the states.

As mentioned previously, for a choice mechanism to prescribe an action, the decision rule needs to be given a specific utility u and belief Γ as input. We call the pair (u, Γ) a subjective representation of the decision situation.

In the following, we investigate the evolutionary fitness of general and systematic ways of forming such subjective representations across many different decision situations.

A fitness game is an interactive decision situation. For a given fitness game $G = \langle N, (A_i, \pi_i^G)_{i \in N} \rangle$, let us denote the evolutionary payoff, or fitness, of player i by the function $\pi_i^G : \prod_{i \in N} A_i \rightarrow \mathbb{R}$, where A_i is player i 's (finite) set of actions. For simplicity of exposition we assume that all games that are played are symmetric two-player games where $N = \{1, 2\}$, $A_1 = A_2$, and $\pi_1^G(a_1, a_2) = \pi_2^G(a_2, a_1) = \pi^G(a, a')$.² The fitness of a choice mechanism c with decision rule a_c^* and subjective representation (u_c, Γ_c) is measured in terms of the expected evolutionary payoff of c . Formally, the fitness of choice mechanism c against choice mechanism c' in a symmetric two-player game $G = \langle \{1, 2\}, A, \pi^G \rangle$ is given by:³

$$F_G(c, c') = \pi^G(a_c^*(u_c^G, \Gamma_c), a_{c'}^*(u_{c'}^G, \Gamma_{c'})).$$

Given the game-theoretic setting, the subjective utility u_c^G is now a function $u_c^G : A \times A \rightarrow \mathbb{R}$, and the subjective belief Γ_c is a set of probability functions over the coplayer's actions, $\Gamma_c \subseteq \Delta(A)$.

Going beyond a single fixed fitness game, we consider a class of possible games. For concreteness, let \mathcal{G} be a class of two-player symmetric games, together with a probability measure $P_G(G')$ for the occurrence probability of game $G' \in \mathcal{G}$. Intuitively, the probability P_G encodes the statistical properties of the environment. A meta-game is then a tuple $MG = \langle CM, \mathcal{G}, P_G, F \rangle$, where CM is a set of choice mechanisms, \mathcal{G} is a class of possible games, $P_G(G')$ is the probability of game G' to occur, and $F : CM \times CM \rightarrow \mathbb{R}$ is the (meta-)fitness function, defined as:

$$F(c, c') = \int F_G(c, c') dP_G(G). \tag{1}$$

Hence, $F(c, c')$ determines the evolutionary payoff of choice mechanism c against c' in the meta-game. The set CM can be thought of as the set of choice mechanisms that are present within a given population playing the games from the class \mathcal{G} . Consequently, it is possible to compute the average fitness of c against the population, that is given by:

2. Since payoff functions are symmetric, we simply write $\pi^G(a, a')$ for $\pi_1^G(a_1, a_2)$ and $A = A_1 = A_2$, as usual. However, notice that all definitions and results can be extended to more general cases.

3. Whenever a choice mechanism would not select a unique action, we assume that the player chooses one of the equally optimal actions at random. That is, $F_G(c, c') = \sum_{a \in a_c^*(u_c^G, \Gamma_c)} \sum_{a' \in a_{c'}^*(u_{c'}^G, \Gamma_{c'})} (1/|a_c^*(u_c^G, \Gamma_c)|)(1/|a_{c'}^*(u_{c'}^G, \Gamma_{c'})|)\pi^G(a, a')$.

$$F(c) = \int F(c, c') dP_c(c') = \iint F_G(c, c') dP_c(c') dP_G(G), \tag{2}$$

where $P_c(c')$ is the probability of encountering a coplayer with choice mechanism c' .

Meta-games are then abstract models for the evolutionary competition between choice mechanisms in interactive decision-making contexts. Standard notions of evolutionary game theory apply to meta-games as well. For example, a choice mechanism c is a strict Nash equilibrium if $F(c, c) > F(c', c)$ for all c' ; it is evolutionarily stable if for all c' either (i) $F(c, c) > F(c', c)$ or (ii) $F(c, c) = F(c', c)$ and $F(c, c') > F(c', c')$; it is neutrally stable if for all c' either (i) $F(c, c) > F(c', c)$ or (ii) $F(c, c) = F(c', c)$ and $F(c, c') \geq F(c', c')$ (Maynard Smith 1982). Similarly, evolutionary dynamics can be applied to meta-games. Later we will also turn toward a dynamical analysis in terms of replicator dynamics (Taylor and Jonker 1978) and replicator mutator dynamics (e.g., Nowak 2006).

4. Evolution of Preferences. To demonstrate the usefulness of a meta-game approach, we compare a selection of general ways of forming belief and preference representations against each other. As for subjective preferences, consider initially:

1. the objective utility, defined by: for all $G \in \mathcal{G}$,

$$\text{obj}^G(a, a') = \pi^G(a, a');$$

2. the regret, defined by: for all $G \in \mathcal{G}$,

$$\text{reg}^G(a, a') = \pi^G(a, a') - \max_{a'' \in A} \pi^G(a'', a').$$

As motivation for this comparison, it is to be stressed that regret minimization is one of the main alternatives to utility (or value) maximization in the literature on decision criteria (see also Bleichrodt and Wakker 2015). For a start, the subjective beliefs that we take into consideration are also two:

1. *prc*: a precise uniform belief $\bar{\mu}$ such that $\bar{\mu}(a) = 1/|A|$ for all $a \in A$;
2. *imp*: a maximally imprecise belief $\bar{\Gamma} = \Delta(A)$.

Although a thorough discussion of this issue goes beyond the scope of this work, let us say that these two kinds of belief underlie two different and alternative views on uncertainty. Faced with uncertain events, a strict Bayes-

ian will always form a precise belief, specified by a single probability μ . In the absence of any information about future uncertain events, the Bayesian would mostly invoke the principle of insufficient reason and accordingly choose a uniform probability over the possible outcomes. In contrast, others have argued against the obligation of representing a belief by means of a single probability measure, opposite to the Bayesian paradigm (e.g., Gilboa and Marinacci 2013). They argue instead in favor of a more encompassing account, according to which uncertainty can be unmeasurable and represented by a (convex and compact) set of probabilities (e.g., Gilboa and Schmeidler 1989). This line of thought has its origin in decision theory, motivated by Ellsberg's famous paradoxes (Ellsberg 1961), and appears extremely relevant in game-theoretic contexts too. Indeed, in a recent paper Battigalli et al. (2015, 646) write: "Such [unmeasurable] uncertainty is inherent in situations of strategic interaction. This is quite obvious when such situations have been faced only a few times."

In evolutionary game theory, for instance, players obviously face uncertainty about the composition of the population that they are part of, and consequently about the (type of) coplayer that they are randomly paired with at each round and about the coplayer's action. In case of complete lack of information about the composition of the population, a non-Bayesian player would thus entertain maximal unmeasurable uncertainty, that is, a maximally imprecise belief.⁴ As already anticipated, we will see that the way agents form beliefs, and the possibility of holding imprecise beliefs in particular, can have a fundamental impact on their evolutionary success.

As for the decision rule, we assume that players use the maxmin rule. This is in line with many representation results of decision making under unmeasurable uncertainty (e.g., Gilboa and Schmeidler 1989; Ghirardato and Marinacci 2002) and seems corroborated by empirical findings too. Ellsberg's paradoxes are prominent examples (Ellsberg 1961), and evidence from experimental literature suggests that agents are generally averse to unmeasurable uncertainty (e.g., Trautmann and Kuilen 2016).

Finally, note that when the maxmin rule acts on subjective representations of type (obj, imp), that is, objective preferences and imprecise beliefs, the generated behavior corresponds to the classic maxmin strategy (von Neumann and Morgenstern 1944). When the maxmin rule acts on subjective representation (reg, imp), the agent's behavior is known as regret minimiza-

4. Such a radical uncertainty could ensue, for example, if agents have no conception of their coplayer or her preferences. Unsophisticated agents, as considered in evolutionary game theory, might be entirely unaware of the fact that they are engaged in social decision making (see Heifetz, Meier, and Schipper [2013] for game-theoretic models of unawareness). It is therefore not ludicrous to consider radical uncertainty first and tend to more sophisticated ways of forming beliefs later (more on this below).

tion.⁵ Two facts follow from these observations. The first is related to our focus on different types of uncertainty that players may entertain.

Fact 1 For any precise (Bayesian) belief μ , maximization of expected (objective) utility based on μ and minimization of expected regret based on μ are behaviorally equivalent.

The second fact highlights another behavioral equivalence, which we will make use of shortly in the following section.

Fact 2 In the class of 2×2 symmetric games, the acts selected by the Laplace rule are exactly the acts selected by regret minimization.

Here is a simple example that shows these choice mechanisms in action. Consider the coordination fitness game G depicted in figure 1*a*. Since the game is symmetric, it suffices to specify the evolutionary payoffs for the row player. Figure 1*a* also represents the objective utility obj^G , since $\text{obj}^G = \pi^G$ by definition, whereas figure 1*b* pictures the representation of G in terms of regret-based utilities. While classic maxmin is indifferent between *I* and *II* (fig. 1*a*), regret minimization uniquely selects *II* (fig. 1*b*).

5. Results

5.1. Simulation Results. Since for now we keep the decision rule fixed to maxmin, a player's choice mechanism will only depend on the player's subjective representation (u, Γ) . For brevity, from now on we will refer to the pair (u, Γ) , like (reg, imp) or (obj, prc), as the type of the player. Sometimes we will also distinguish types by referring to the subjective utility only, for instance (reg, imp) and (reg, prc) are regret types.

As observed earlier, meta-games factor in statistical properties of the environment. For particular empirical purposes, one could consult a specific class of games \mathcal{G} with appropriate, maybe empirically informed probability P_G in order to match the natural environment of a given population. For our present purposes, let \mathcal{G} be a set of symmetric two-player fitness games with two acts for a start. Each game $G \in \mathcal{G}$ is then individuated solely by its payoff function, that is, by a quadruple of numbers $G = (a, b, c, d)$. As for the

5. The notion of regret in decision theory dates back at least to the work by Savage (1951) and has later been developed by Bell (1982), Fishburn (1982), and Loomes and Sugden (1982) independently. Recently, Halpern and Pass (2012) showed how the use of regret minimization can give solutions to game-theoretic puzzles (like the Traveller's dilemma and the Centipede game) in a way that is closer to everyday intuition and empirical data. In this article the notion of regret defined earlier is the same as in Halpern and Pass (2012).

	I	II
I	1	0
II	0	2

(a) Coordination game G .

	I	II
I	0	-2
II	-1	0

(b) Regret-based representation of G .Figure 1. A coordination game (*left*) and the associated regret representation (*right*).

occurrence probability $P_G(G)$ of game G , we imagine that the values a , b , c , d are independently and identically distributed (i.i.d.) random variables sampled from the set $\{0, \dots, 10\}$ according to uniform probability P_V . Using Monte Carlo simulations, we can then approximate the values of equation (1) to construct meta-game payoffs. Results based on 100,000 randomly sampled games are given in table 1.⁶

Simulation results obviously reflect fact 2 in that all encounters in which types (reg, imp), (reg, prc), or (obj, prc) are substituted for one another yield identical results. More interestingly, table 1 shows that (obj, imp), the maxmin strategy, is strictly dominated by the three other types: in each column (i.e., for each type of coplayer), the maxmin strategy is strictly worse than any of the three competitors. This has a number of interesting consequences.

If we restrict attention to subjective representations with imprecise beliefs only, then a monomorphic state in which every agent has regret-based preferences is the only evolutionarily stable state. More strongly, since (obj, imp) is strictly dominated by (reg, imp), we expect selection that is driven by (expected) fitness to invariably weed out maxmin players (obj, imp) in favor of (reg, imp), regret minimization. In terms of choice rules, this means that regret minimization is evolutionarily better than maxmin over the class of games considered. In terms of subjective preferences, it shows that players using the objective representation that directly looks at fitness (possibly money, or profit) are outperformed by nonveridical (regret) representations, when players' beliefs are imprecise.

Next, if we look at the competition between all four types represented in table 1, (reg, imp) is no longer evolutionarily stable. Given behavioral equivalence (fact 2), types (reg, imp), (reg, prc), and (obj, prc) are all neutrally stable (Maynard Smith 1982). But since (obj, imp) is strictly dominated and so disfavored by fitness-based selection, we are still drawn to conclude that

6. Concretely, 100,000 games were sampled repeatedly by choosing independently four integers between 0 and 10 uniformly at random. For each game, the action choices of all four choice mechanisms were determined and payoffs from all pairwise encounters recorded. The number in each cell of table 1 is the average payoff for the choice mechanism listed in the row when matched with the choice mechanism in the column.

TABLE 1. AVERAGE EVOLUTIONARY FITNESS FROM MONTE CARLO SIMULATIONS OF 100,000 SYMMETRIC 2 × 2 GAMES

	reg, imp	obj, imp	reg, prc	obj, prc
reg, imp	6.663	6.662	6.663	6.663
obj, imp	6.486	6.484	6.486	6.486
reg, prc	6.663	6.662	6.663	6.663
obj, prc	6.663	6.662	6.663	6.663

maxmin behavior is weeded out in favor of a population with a random distribution of the remaining three types.

Simulation results of the (discrete time) replicator dynamics (Taylor and Jonker 1978) indeed show that random initial population configurations are attracted to states with only three player types: (reg, imp), (reg, prc), and (obj, prc). The relative proportions of these depend on the initial shares in the population. This variability fully disappears if we add a small mutation rate to the dynamics. Take a fixed, small mutation rate ϵ for the probability that a player’s subjective utility or her subjective belief changes to another utility or belief. The probability that a player’s subjective representation randomly mutates into a completely different representation with altogether different utility and belief would then be ϵ^2 . With these assumptions about “component-wise mutations,” numerical simulations of the (discrete time) replicator mutator dynamics (Nowak 2006) show that already for very small mutation rates almost all initial population states converge to a single fixed point in which the majority of players have regret-based utility. For instance, with $\epsilon = 0.001$, almost all initial populations are attracted to a final distribution with proportions:

(reg, imp)	(obj, imp)	(reg, prc)	(obj, prc)
.289	.021	.398	.289

What this suggests is that, if biological evolution selects behavior-generating mechanisms, not behavior as such, it need not be the case that behaviorally equivalent mechanisms are treated equally all the while. If mutation probabilities are a function of individual components, it can be the case that certain components of such behavior-generating mechanisms are more strongly favored by a process of random mutation and selection. This is exactly the case with regret-based preferences. Since regret-based preferences are much better in connection with imprecise beliefs than veridical preferences are, the proportion of expected regret minimizers, (reg, prc), in the attracting state is substantially higher than that of expected utility maximizers, (obj, prc), even though these types are behaviorally equivalent.

5.2. Analytical Results. Results based on the single meta-game in table 1 are not fully general and possibly spoiled by random fluctuations in the sampling procedure. Fortunately, for the case of 2×2 symmetric games, the main result that maxmin types (obj, imp) are strictly dominated by regret minimizers can also be shown analytically for considerably general conditions.

Proposition 1 Let \mathcal{G} be the class of 2×2 symmetric games $G = (a, b, c, d)$ generated by i.i.d. sampling a, b, c, d from a set of values with at least three elements in the support. Then, (reg, imp) strictly dominates (obj, imp) in the resulting meta-game.

Proof All proofs are in the appendix.

Corollary 1 Let \mathcal{G} be as in proposition 1. If we only consider imprecise belief types, (obj, imp) and (reg, imp), then the unique evolutionarily stable state is a monomorphic population of (reg, imp) players.

The result shows that there is support for the main conceptual point that we wanted to make: objective preference representations are not necessarily favored by natural selection; objective preferences are outperformed by non-veridical regret preferences if agents have imprecise beliefs. This tells us that the main conclusions drawn in the previous section based on the approximated meta-game of table 1 hold more generally for arbitrary 2×2 symmetric games with i.i.d. sampled payoffs.

This result presupposes at least occasional imprecise beliefs. The assumed imprecise beliefs do not need to be maximally uncertain, however. Let the uncertainty held by a player be defined by a convex compact set of probabilities $[s, t] \subseteq \Delta(A)$ over the coplayer's actions, where s is the lower probability and t is the upper probability of action I . We can then prove the following proposition, which is the analogue of proposition 1 for any possible (not necessarily maximal) degree of uncertainty $[s, t]$, with $s \neq t$. There is only one difference: we are now going to require i.i.d. drawing of a continuous random variable with uniform distribution. This is due to the fact that, for arbitrarily small intervals $[s, t]$, objective players (obj, $[s, t]$) and regret players (reg, $[s, t]$) can behave as holding a unique probability measure (precise belief) if the underlying payoff space is not dense. The reason for this technical requirement will become clearer from the proof.

Proposition 2 Let \mathcal{G} be the class of symmetric 2×2 games generated by i.i.d. drawing of a continuous random variable with uniform distribution over any set of values, and fix any imprecise belief $[s, t]$. Then the only evolutionarily stable state of a population with regret players (reg, $[s, t]$) and objective players (obj, $[s, t]$) is a monomorphic state of (reg, $[s, t]$) players.

This tells us that regret-based preferences can outperform objective preference representations when agents are also capable of learning or otherwise restricting their assumptions about the coplayer's behavior as long as there is, at least on occasion, some imprecision in their beliefs. We will enlarge on the issue of belief formation after having covered some more relevant extensions in the next section.

6. Extensions. How do the basic results from the previous section carry over to richer models? Section 6.1 first introduces further conceptually interesting subjective representations that have been considered in the literature. Section 6.2 then addresses the case of symmetric two-player $n \times n$ games for $n \geq 2$. Finally, section 6.3 ends with a brief comparison to the case of solitary decision making.

6.1. More Preference Types. The space of possible preference types is enormous, and we have only compared regret and objective types so far. Let us now look at two other types of subjective preferences that have been investigated, especially in behavioral economics and in evolutionary game theory. A famous example is the altruistic preference (e.g., Becker 1976; Bester and Güth 1998), summoned to explain the possibility of altruistic behavior. At the other end of the spectrum, the competitive preference is located. The two subjective utilities are defined as follows:

1. Altruistic utility: for all $G \in \mathcal{G}$, $\text{alt}^G(a, a') = \pi^G(a, a') - \pi^G(a', a)$;
2. Competitive utility: for all $G \in \mathcal{G}$, $\text{com}^G(a, a') = \pi^G(a, a') - \pi^G(a', a)$.⁷

Table 2 shows results of Monte Carlo simulations that approximate the expected fitness in the relevant meta-game with all the subjective representations considered so far. These results confirm basic intuitions about altruistic and competitive types: everybody would like to have an altruistic coplayer and nobody would like to play against a competitive player. Perhaps more surprisingly, (alt, imp) comes up strictly dominated by (com, imp), but competitive types themselves are worse off against all types except against maxmin players (obj, imp) than any of the behaviorally equivalent types (reg, imp), (obj, prc), and (reg, prc). It is thus easy to see that the previous re-

7. A more general formulation would be to define α -altruistic utility, for $\alpha \in [0, 1]$, $u_\alpha^G(a, a') = \pi^G(a, a') + \alpha\pi^G(a', a)$. Since we are not interested in the evolution of degrees of altruism, here we simply fix $\alpha = 1$. Analogously for α -competitive utilities too. Other possible generalizations could also take into account combinations of all these preferences with different decision rules. Maximax, minimax, and minimin, for example, would be possible rules for choice. Here we opted for maxmin because we were specifically interested in comparing maxmin and regret minimization, as these are two major alternatives for decision making.

TABLE 2. AVERAGE EVOLUTIONARY FITNESS FROM MONTE CARLO SIMULATIONS OF 100,000 SYMMETRIC 2×2 GAMES

	(reg, imp)	(obj, imp)	(com, imp)	(alt, imp)	(reg, prc)	(obj, prc)	(com, prc)	(alt, prc)
(reg, imp)	6.663	6.662	5.829	7.105	6.663	6.663	5.829	7.489
(obj, imp)	6.486	6.484	6.088	6.703	6.486	6.486	6.088	6.875
(com, imp)	6.323	6.758	5.496	6.977	6.323	6.323	5.496	7.149
(alt, imp)	5.949	5.722	5.326	6.396	5.949	5.949	5.326	6.568
(reg, prc)	6.663	6.662	5.829	7.105	6.663	6.663	5.829	7.489
(obj, prc)	6.663	6.662	5.829	7.105	6.663	6.663	5.829	7.489
(com, prc)	6.323	6.758	5.496	6.977	6.323	6.323	5.496	7.149
(alt, prc)	6.331	5.893	5.497	6.566	6.331	6.331	5.497	7.152

sults still obtain for the larger meta-game in table 2: (reg, imp), (obj, prc), and (reg, prc) are still neutrally stable; simulation runs of the (discrete-time) replicator dynamics on the 8×8 meta-game from table 2 end up in population states consisting of only these three types in variable proportion. In sum, the presence of other subjective representations, such as those based on altruistic or competitive utilities, does not undermine, but rather strengthens, our previous results.

6.2. More Actions. Results from section 5 relied heavily on fact 2, which is no longer true when we look at arbitrary $n \times n$ games. Table 3 gives approximations of expected fitness in the class of $n \times n$ symmetric games. Concretely, the numbers in table 3 are averages of evolutionary payoffs obtained in 100,000 randomly sampled symmetric games, where each fitness game G was sampled by first picking a number of acts $n^G \in \{2, \dots, 10\}$ uniformly at random, and then filling the necessary $n^G \times n^G$ payoff matrix with i.i.d. sampled numbers, as before.

TABLE 3. AVERAGE EVOLUTIONARY FITNESS FOR 100,000 RANDOMLY GENERATED $N \times N$ SYMMETRIC GAMES WITH N RANDOMLY DRAWN FROM $\{2, \dots, 10\}$

	(reg, imp)	(obj, imp)	(com, imp)	(alt, imp)	(reg, prc)	(obj, prc)	(com, prc)	(alt, prc)
(reg, imp)	6.567	6.570	5.650	6.992	6.564	6.564	5.593	7.409
(obj, imp)	6.476	6.483	5.896	6.818	6.484	6.484	5.850	7.124
(com, imp)	6.468	6.647	5.512	7.169	6.578	6.578	5.577	7.354
(alt, imp)	5.968	5.923	5.363	6.685	5.975	5.975	5.086	6.973
(reg, prc)	6.908	6.918	5.988	7.456	6.929	6.929	5.934	7.783
(obj, prc)	6.908	6.918	5.988	7.456	6.929	6.929	5.934	7.783
(com, prc)	6.529	6.680	5.445	7.276	6.542	6.542	5.521	7.440
(alt, prc)	6.450	6.337	5.772	6.978	6.457	6.457	5.479	7.500

The most important result is that the regret-minimizing type (reg, imp) is strictly dominated by (reg, prc) and by (obj, prc) in the meta-game from table 3. This means that while simple regret minimization can thrive in some evolutionary contexts, there are also contexts where it is demonstrably worse off. While this may be bad news for regret-minimizing types (reg, imp), it is not the case that regret types as such are weeded out by selection. Since, by fact 1, (reg, prc) and (obj, prc) are behaviorally equivalent in general, it remains that selection based on meta-games constructed from $n \times n$ games will still not eradicate regret preferences.

On the other hand, there are plenty of ways in which the basic insights from propositions 1 and 2 can make for situations in which evolution would favor regret types, even in $n \times n$ games. If, for example, the belief of a player is a trait that biological evolution has no bite on, but rather something that the particular choice situation would exogenously give us (possibly because of the different amount of information available in different choice situations), then regret-based preferences can again drive out veridical preferences altogether. For example, suppose that only preference representations compete and that agents' beliefs are exogenously given, in such a way that both players hold precise (Bayesian) uniform beliefs with probability p and they both have maximally imprecise beliefs otherwise. This transforms the meta-game from table 3 into a simpler 4×4 meta-game in which the payoff obtained by a subjective preference is the weighted average over the payoffs of the subjective representations including that preference in table 3. Setting $p = .98$ for illustration, we get the meta-game in table 4. The only evolutionarily stable state of this meta-game is again a monomorphic population of regret types. Accordingly, all our simulation runs of the (discrete-time) replicator dynamics converge to monomorphic regret-type populations. The reason why regret-based utilities prosper is because they have a substantial fitness advantage when paired with imprecise beliefs (propositions 1 and 2). If unmeasurable uncertainty is exogenously given as something that happens to agents because of the information available in some choice situations, and even if that happens only very infrequently (i.e., for rather low p), regret preferences will outperform objective preferences, as well as competitive and altruistic preferences.

TABLE 4. META-GAME FOR THE EVOLUTIONARY COMPETITION BETWEEN SUBJECTIVE UTILITIES WHEN BELIEFS ARE EXOGENOUSLY GIVEN (SEE MAIN TEXT)

	reg	obj	com	alt
reg	6.926	6.926	5.942	7.757
obj	6.924	6.924	5.948	7.751
com	6.566	6.570	5.481	7.434
alt	6.463	6.461	5.478	7.469

6.3. Solitary Decisions. To see how different choice mechanisms behave in evolutionary competition based on solitary decision making, we approximated, much in the spirit of meta-games, average accumulated fitness obtained in randomly generated solitary decision problems. For our purposes, a decision problem $D = \langle W, A, \pi^D \rangle$ consists of a set of states of the world W , a set of acts A , and a payoff function $\pi^D : W \times A \rightarrow \mathbb{R}$. We generate arbitrary decision problems by selecting, uniformly at random, numbers of states and acts $n_w^D, n_a^D \in \{2, \dots, 10\}$ and then filling the payoff table, so to speak, by i.i.d. samples for each $\pi^D(w, a) \in \{0, 10\}$. Unlike with two-player games, we need to also sample the actual state of the world, which we selected uniformly at random from the available states in the current decision problem. Accordingly, the fitness of choice mechanism c in decision problem D is given by:

$$F_D(c) = \pi^D(w, a_c^*(u_c^*, \Gamma_c)) \bar{\mu}(w),$$

with $\bar{\mu}(w) = 1/n_w^D$ for all w . As subjective representations, we considered the original cast of four from table 1, since altruistic and competitive types are meaningless in solitary decision situations. As before, the relevant fitness measure, defined in equation (3), was approximated by Monte Carlo simulations, the results of which are given in table 5.

$$F(c) = \int F_D(c) dP_D(D). \quad (3)$$

Facts 1 and 2 still apply: (reg, prc) and (obj, prc) are behaviorally equivalent in general, and (reg, imp) is behaviorally equivalent to the former two in decision problems with two states and two acts. This shows in the results from table 5 in that the averages for (reg, prc) and (obj, prc) are identical. But since we included decision problems with more acts and more states as well, the average for regret minimizers (reg, imp) is not identical to the one of (reg, prc) and (obj, prc). It is, in fact, lower, but again not as low as that of (obj, imp).

This means that every relevant result we have seen about game situations is also borne out for solitary decisions. Evolutionary selection based on objective fitness will not select against regret preferences, as these are indistinguishable from veridical preferences when paired with precise beliefs. But when paired with imprecise beliefs, regret-based utilities outperform objective utilities. Consequently, if there is a chance, however small, that

TABLE 5. EXPECTED FITNESS OF CHOICE MECHANISMS APPROXIMATED FROM 100,000 SIMULATED SOLITARY DECISION PROBLEMS (SEE MAIN TEXT)

(reg, imp)	(obj, imp)	(reg, prc)	(obj, prc)
6.318	6.237	6.661	6.661

agents fall back on imprecise beliefs, evolution will actually positively select for nonveridical regret-based preferences.

6.4. Sophisticated Beliefs. Since one of our main purposes was to illustrate the usefulness of a meta-game approach by the case study of objective and regret preferences, we have partially neglected an important and interesting issue, namely, the evolution of ways of forming beliefs about coplayers' behavior or the actual state of the world. For reasons of space we must, unfortunately, leave a deeper exploration of belief type evolution to another occasion. Two remarks are in order nonetheless. First, belief type evolution can be studied without conceptual hurdles in the meta-game framework, so that there is no principled argument against the main methodological contribution of this article. Second, our results regarding the comparison between regret and objective types remain to be informative, even if we allow agents to learn or reason strategically.⁸ This is because we know from fact 1 that regret and objective preferences come up behaviorally equivalent when paired with precise probabilistic beliefs (given identical decision rule). This holds no matter what the content of that belief is. So, if learning, reasoning, or statistical knowledge about a recurrent situation can be brought to bear, this will not make evolution select against regret-based preferences. If, on the other hand, agents resort to imprecise beliefs at least occasionally (e.g., when they are unaware of the coplayer or her utilities or when strategic reasoning cannot reduce all uncertainty about the coplayer's choice), then regret-based preferences can be favored by natural selection over objective preferences.

7. Conclusion. The assumption that players and decision makers maximize their (subjective) utility is central through all economic literature, and the maximization of actual (objective) payoffs is often justified by appealing to evolutionary arguments and natural selection. In contrast to the standard view, we showed the existence of player types with subjective utilities different from the actual evolutionary payoffs that can outperform types whose subjective utilities coincide with the evolutionary payoffs. Here the claim is not that regret preferences are the best on the market, but rather that utilities that perfectly mirror evolutionary fitness can be outclassed by subjective utilities that differ from the objective fitness. While the literature on evolution of preferences has focused on fixed games, we have adopted a more general approach here. We suggested that attention to "meta-games" is crucial, because what may be a good subjective representation in one type of game (e.g., cooperative preferences in the Prisoner's Dilemma) need not be generally beneficial.

8. Some research has recently been done along these lines. See, in particular, Mengel (2012), Mohlin (2012), and Robalino and Robson (2016).

Appendix

Proofs

The proof of proposition 1 relies on a partition of \mathcal{G} , and on some lemmas. For brevity, let us denote the regret minimizer (reg, imp) by R and the maximinimizer (obj, imp) by M . Following equation (1), let $F_{\mathcal{G}}(X, Y)$ denote the expected payoff of choice mechanism X against choice mechanism Y on the possibly restricted class of fitness games \mathcal{G} .

Proof of Proposition 1. By definition of strict dominance, we have to show that in the class \mathcal{G} of symmetric 2×2 games with payoffs sampled from a set of i.i.d. values with at least three elements in the support, it holds that:

- (i) $F_{\mathcal{G}}(R, R) > F_{\mathcal{G}}(M, R)$;
- (ii) $F_{\mathcal{G}}(M, M) < F_{\mathcal{G}}(R, M)$.

To show this we use the following partition of \mathcal{G} , based on payoffs parametrized as follows:

	I	II
I	a	b
II	c	d

1. Coordination games \mathcal{C} : $a > c$ and $d > b$;
2. Anticoordination games \mathcal{A} : $a < c$ and $d < b$;
3. Strong dominance games \mathcal{S} : aut ($a > c$ and $b > d$) aut ($a < c$ and $b < d$);
4. Weak dominance games \mathcal{W} : aut $a = c$ aut $b = d$;
5. Boring games \mathcal{B} : $a = c$ and $b = d$.

Before proving the lemmas, it is convenient to fix some notation. Let us call x, y, z the three elements in the support, and without loss of generality suppose that $x > y > z$. We denote by C a coordination game in \mathcal{C} with payoffs $a_C, b_C, c_C,$ and d_C ; similarly for games $A \in \mathcal{A}, S \in \mathcal{S}, W \in \mathcal{W},$ and $B \in \mathcal{B}$. Let us denote by I_{RC} the event that a R -player plays action I in the game C ; and similarly for action II, for player M , and for games $A, S, W,$ and B . We first consider the case of i.i.d. sampling with finite support.

Lemma 1

R and M perform equally well in \mathcal{S} and in \mathcal{B} .

Proof

By definition of regret minimization and maxmin it is easy to check that whenever in a game there is a strongly dominant action a^S , then a^S is both the max-

min action and the regret-minimizing action. Then, for all the games in \mathcal{S} , R chooses action a if and only if M chooses action a . Consequently, R and M always perform equally (well) in \mathcal{S} . In the case of \mathcal{B} it is trivial to see that all the players perform equally. QED

Lemma 2

In \mathcal{W} , R strictly dominates M .

Proof

Assume without loss of generality that $b = d$, and that $a > c$. There are two cases that we have to check: (i) $c < b = d$ and (ii) $c \geq b = d$. In the first case it is easy to see that R and M perform equally: act I is the choice of both R and M . In the case of (ii) instead we have that I is the regret-minimizing action, whereas both actions have the same minimum and M plays $((1/2)I; (1/2)II)$, since both I and II maximize the minimal payoff. Consider now a population of R and M playing games from the class \mathcal{W} . Whenever (i) is the case R and M perform equally well. But suppose $W \in \mathcal{W}$ and (ii) is the case. Then, $\pi_W(R, R) = a > (1/2)a + (1/2)c = \pi_W(M, R)$, whereas

$$\pi_W(M, M) = \frac{1}{4}a + \frac{1}{4}b + \frac{1}{4}c + \frac{1}{4}d < \frac{1}{2}a + \frac{1}{2}b = \pi_W(R, M).$$

Hence, we have that in general $F_W(R, R) > F_W(M, R)$, and $F_W(M, M) < F_W(R, M)$. QED

Since it is not difficult to see that both (R, R) and (M, M) are strict Nash equilibria in \mathcal{C} , and that (R, R) and (M, M) are not Nash equilibria in \mathcal{A} , the main part of the proof will be to show that R strictly dominates C in the class $\mathcal{C} \cup \mathcal{A}$, that is:

- (i') $F_{\mathcal{C} \cup \mathcal{A}}(R, R) > F_{\mathcal{C} \cup \mathcal{A}}(M, R)$,
- (ii') $F_{\mathcal{C} \cup \mathcal{A}}(M, M) < F_{\mathcal{C} \cup \mathcal{A}}(R, M)$.

This part needs some more lemmas to be proven, but first we introduce the following bijective function ϕ between coordination and antcoordination games.

Definition 3 (ϕ)

The permutation $\phi(a, b, c, d) = (c, d, a, b)$ defines a bijective function $\phi : \mathcal{C} \rightarrow \mathcal{A}$ that for each coordination game $C \in \mathcal{C}$ with payoffs (a_C, b_C, c_C, d_C) gives the antcoordination game $A \in \mathcal{A}$ with payoffs $(a_A, b_A, c_A, d_A) = (c_C, d_C, a_C, b_C)$. Essentially, ϕ swaps rows in the payoff matrix.

Lemma 4

Occurrence probability of C equals that of $\phi(C)$: $P(\phi(C)) = P(C)$.

Proof

By definition, each game $C \equiv (a_C, b_C, c_C, d_C)$ is such that $a_C > c_C$ and $d_C > b_C$, and each game $A \equiv (a_A, b_A, c_A, d_A)$ is such that $a_A < c_A$ and $d_A < b_A$. Given that a, b, c, d are i.i.d. random variables and that a sequence of i.i.d. random variables is exchangeable, it is clear that the probability of (a_C, b_C, c_C, d_C) equals the probability of (c_C, d_C, a_C, b_C) . Hence, $P(\phi(C)) = P(C)$. QED

Lemma 5

Let $P(E)$ be the probability of event E , for example, $P(I_{RC})$ is the probability that a random R -player plays act I in coordination game C , which is either 0, .5, or 1. It then holds that:

- $P(I_{RC}) = P(II_{R\phi(C)})$, and $P(II_{RC}) = P(I_{R\phi(C)})$;
- $P(I_{MC}) = P(II_{M\phi(C)})$, and $P(II_{MC}) = P(I_{M\phi(C)})$.

Proof

It is easy to check that if $b_C - d_C > c_C - a_C$, an R -player plays action I in C ; that if $b_C - d_C < c_C - a_C$, R plays II; and that if $b_C - d_C = c_C - a_C$, an R -player is indifferent between I and II in C , and so randomizes with $((1/2)I; (1/2)II)$. Similarly, if $a_A - c_A > d_A - b_A$, an R -player plays action I in A ; if $a_A - c_A = d_A - b_A$, R plays II; and if $a_A - c_A < d_A - b_A$, an R -player is indifferent between I and II in A , and randomizes with $((1/2)I; (1/2)II)$. Consequently, if $b_C - d_C > c_C - a_C$, then $P(I_{RC}) = 1$, and by definition of ϕ we have $P(II_{R\phi(C)}) = 1$. Likewise, if $b_C - d_C < c_C - a_C$, then $P(II_{RC}) = 1 = P(I_{R\phi(C)})$; and if $b_C - d_C = c_C - a_C$, then $P(I_{RC}) = P(II_{RC}) = 1/2 = P(II_{R\phi(C)}) = P(I_{R\phi(C)})$.

In the same way, in coordination games we have that if $b_C > c_C$, an M -player plays I; if $c_C > b_C$, an M -player plays II; and if $b_C = c_C$, M is indifferent between I and II and plays $((1/2)I; (1/2)II)$. In antcoordination games instead, if $a_A > d_A$, M plays I; if $a_A = d_A$, M plays II; if $a_A < d_A$, M plays $((1/2)I; (1/2)II)$. By definition of ϕ : $P(I_{MC}) = 1 = P(II_{M\phi(C)})$ if $b_C > c_C$; $P(II_{MC}) = 1 = P(I_{M\phi(C)})$ if $c_C > b_C$; and $P(I_{MC}) = P(II_{MC}) = 1/2 = P(II_{M\phi(C)}) = P(I_{M\phi(C)})$ if $b_C = c_C$. QED

Lemma 6

It holds that:

- $a_C > d_C \rightarrow (I_{MC} \subseteq I_{RC})$;
- $a_C = d_C \rightarrow I_{MC} = I_{RC}$.
- $a_C < d_C \rightarrow (II_{MC} \subseteq II_{RC})$;

Proof

The event that R plays action I, I_{RC} , with positive probability is the event that $b_C - d_C \geq c_C - a_C$: if $b_C - d_C > c_C - a_C$, R plays I, and if $b_C - d_C = c_C - a_C$, R plays $((1/2)I; (1/2)II)$. Similarly, the event that I_{MC} has positive occurrence is the event that $b_C \geq c_C$: if $b_C > c_C$, M plays I, and if $b_C = c_C$, M plays $((1/2)I; (1/2)II)$. Then, I_{RC} implies that $b_C - d_C \geq c_C - a_C$, and I_{MC} implies that $b_C \geq c_C$. Moreover, on the assumption that $a_C > d_C$, it is easy to check that $b_C \geq c_C$ implies $b_C - d_C > c_C - a_C$. Hence, in any C with $a_C > d_C$ it holds that I_{MC} implies I_{RC} , that is, $a_C > d_C \rightarrow (I_{MC} \subseteq I_{RC})$. Instead, it is possible that $a_C > d_C$, $b_C - d_C > c_C - a_C$, and $b_C < c_C$ hold simultaneously, so that $I_{MC} \not\subseteq I_{RC}$. By a symmetric argument it can be shown that $a_C > d_C \rightarrow (\Pi_{MC} \subseteq \Pi_{RC})$ too. Finally, when $a_C = d_C$ it holds that: $b_C - d_C > c_C - a_C$ iff $b_C > c_C$; $b_C - d_C < c_C - a_C$ iff $b_C < c_C$; and $b_C - d_C = c_C - a_C$ iff $b_C = c_C$. Hence, $a_C = d_C \rightarrow I_{MC} = I_{RC}$. QED

We are now ready to prove that $F_{C \cup A}(R, R) > F_{C \cup A}(M, R)$. With notation like $P(I_{RC} \cap I_{RC})$ denoting the probability that a random R -player plays I and another R -player plays I as well in game C , rewrite the inequality as:

$$\begin{aligned} & \sum_{C \in \mathcal{C}} P(C) [P(I_{RC} \cap I_{RC}) \cdot a_C + P(\Pi_{RC} \cap \Pi_{RC}) \cdot d_C + P(I_{RC} \cap \Pi_{RC}) \cdot b_C \\ & \quad + P(\Pi_{RC} \cap I_{RC}) \cdot c_C] + \sum_{A \in \mathcal{A}} P(A) [P(I_{RA} \cap I_{RA}) \cdot a_A + P(\Pi_{RA} \cap \Pi_{RA}) \\ & \quad \cdot d_A + P(I_{RA} \cap \Pi_{RA}) \cdot b_A + P(\Pi_{RA} \cap I_{RA}) \cdot c_A] \\ & > \sum_{C \in \mathcal{C}} P(C) [P(I_{RC} \cap I_{MC}) \cdot a_C + P(\Pi_{RC} \cap \Pi_{MC}) \cdot d_C + P(I_{RC} \cap \Pi_{MC}) \\ & \quad \cdot c_C + P(\Pi_{RC} \cap I_{MC}) \cdot b_C] + \sum_{A \in \mathcal{A}} P(A) [P(I_{RA} \cap I_{MA}) \cdot a_A \\ & \quad + P(\Pi_{RA} \cap \Pi_{MA}) \cdot d_A + P(I_{RA} \cap \Pi_{MA}) \cdot c_A + P(\Pi_{RA} \cap I_{MA}) \cdot b_A]. \end{aligned}$$

By lemma 4 and lemma 5, we can express everything in terms of C only:

$$\begin{aligned} & \sum_C P(C) [P(I_{RC} \cap I_{RC}) \cdot a_C + P(\Pi_{RC} \cap \Pi_{RC}) \cdot d_C + P(I_{RC} \cap \Pi_{RC}) \cdot b_C \\ & \quad + P(\Pi_{RC} \cap I_{RC}) \cdot c_C + P(\Pi_{RC} \cap \Pi_{RC}) \cdot c_C + P(I_{RC} \cap I_{RC}) \cdot b_C \\ & \quad + P(\Pi_{RC} \cap I_{RC}) \cdot d_C + P(I_{RC} \cap \Pi_{RC}) \cdot a_C] \\ & > \sum_C P(C) 2[P(I_{RC} \cap I_{MC}) \cdot a_C + P(\Pi_{RC} \cap \Pi_{MC}) \cdot d_C + P(I_{RC} \cap \Pi_{MC}) \cdot c_C \\ & \quad + P(\Pi_{RC} \cap I_{MC}) \cdot b_C + P(\Pi_{RC} \cap \Pi_{MC}) \cdot c_C + P(I_{RC} \cap I_{MC}) \cdot b_C \\ & \quad + P(\Pi_{RC} \cap I_{MC}) \cdot a_C + P(I_{RC} \cap \Pi_{MC}) \cdot d_C] \end{aligned}$$

This simplifies to:

$$\begin{aligned} & \sum_C P(C) [a_C \cdot (P(I_{RC} \cap I_{RC}) + P(I_{RC} \cap \Pi_{RC})) + b_C \cdot (P(I_{RC} \cap \Pi_{RC}) \\ & \quad + P(I_{RC} \cap I_{RC})) + c_C \cdot (P(\Pi_{RC} \cap I_{RC}) + P(\Pi_{RC} \cap \Pi_{RC})) + d_C \\ & \quad \cdot (P(\Pi_{RC} \cap \Pi_{RC}) + P(\Pi_{RC} \cap I_{RC}))] \\ & > \sum_C P(C) [a_C \cdot (P(I_{RC} \cap I_{MC}) + P(\Pi_{RC} \cap I_{MC})) + b_C \\ & \quad \cdot (P(\Pi_{RC} \cap I_{MC}) + P(I_{RC} \cap I_{MC})) + c_C \\ & \quad \cdot (P(I_{RC} \cap \Pi_{MC}) + P(\Pi_{RC} \cap \Pi_{MC})) + d_C \\ & \quad \cdot (P(\Pi_{RC} \cap \Pi_{MC}) + P(I_{RC} \cap \Pi_{MC}))]. \end{aligned}$$

Now let us split into $a > d$ and $a < d$, and consider $a > d$ first. Notice that, by lemma 6, the case $a = d$ is irrelevant in order to discriminate between R and M . If $a > d$, by lemma 6 we can eliminate the cases where R plays Π and M plays I :

$$\begin{aligned} & \sum_{C_{a>d}} P(C) [a_C \cdot (P(I_{RC} \cap I_{RC}) + P(I_{RC} \cap \Pi_{RC}) - P(I_{RC} \cap I_{MC})) \\ & \quad + b_C \cdot (P(I_{RC} \cap \Pi_{RC}) + P(I_{RC} \cap I_{RC}) - P(I_{RC} \cap I_{MC})) \\ & \quad + c_C \cdot (P(\Pi_{RC} \cap I_{RC}) + P(\Pi_{RC} \cap \Pi_{RC}) - P(I_{RC} \cap \Pi_{MC}) - P(\Pi_{RC} \cap \Pi_{MC})) \\ & \quad + d_C \cdot (P(\Pi_{RC} \cap \Pi_{RC}) + P(\Pi_{RC} \cap I_{RC}) - P(\Pi_{RC} \cap \Pi_{MC}) - P(I_{RC} \cap \Pi_{MC}))] \\ & > 0. \end{aligned}$$

We now distinguish between two cases: (1) $a - c = d - b$ and (2) $a - c \neq d - b$. Notice that $P(I_{RC} \cap \Pi_{RC}) \neq 0$ if and only if case (1) obtains, and that $a > d$ and (1) imply Π_{MC} . Then, from (1) we have:⁹

$$\begin{aligned} & \sum_{C_{a>d}} P(C) \left[a_C \cdot \left(\frac{1}{4} + \frac{1}{4} \right) + b_C \cdot \left(\frac{1}{4} + \frac{1}{4} \right) + c_C \cdot \left(\frac{1}{4} + \frac{1}{4} - \frac{1}{2} - \frac{1}{2} \right) \right. \\ & \quad \left. + d_C \cdot \left(\frac{1}{4} + \frac{1}{4} - \frac{1}{2} - \frac{1}{2} \right) \right] > 0. \end{aligned}$$

9. Note that when we have only 3 elements in the support it is not guaranteed that case (1), together with $a > d$, may arise in a coordination game, whereas it is guaranteed that case (2), together with $a > d$, occurs with some positive probability. If we take for instance $x = 5, y = 2, z = 1$, then case (1) cannot obtain, whereas if we take $x = 3, y = 2, z = 1$, both (1) and (2) may obtain ($a = 3, b = 1, c = 2, d = 2$ for case (1), and $a = 3, b = 1, c = 2, d = 2$ for case (2)). Moreover, under the assumption that $a > d$, having three elements in the support is a necessary and sufficient condition for case (2) to have positive occurrence in a coordination game. As it will be clear in the following, a positive occurrence of case (2) only is enough for the theorem to hold.

Since we have assumed $a - c = d - b$, the last inequality is not satisfied. We have instead:

$$\sum_{C_{a>d}} P(C) \left[\frac{1}{2} a_C + \frac{1}{2} b_C - \frac{1}{2} c_C - \frac{1}{2} d_C \right] = 0.$$

This means that where $a_C > d_C$ and where (1) is the case, R and M are equally fit. This changes when we turn to (2). In that case, since $a_C > d_C \rightarrow (I_{MC} \subset I_{RC})$ by lemma 6, we have that $P(I_{RC} \cap I_{RC}) - P(I_{RC} \cap I_{MC}) = P(I_{RC} \cap \Pi_{MC})$. Moreover, when $a_C > d_C, b_C \geq c_C$ implies $b_C - d_C > c_C - a_C$ (see lemma 6). Consequently, when M plays either I or $((1/2)I; (1/2)II)$, R always plays I. Hence, whenever $a_C > d_C$ and (2) obtain, it also holds that $P(\Pi_{RC} \cap \Pi_{MC}) = P(\Pi_{RC} \cap \Pi_{RC})$. In this case we can simplify to:

$$\sum_{C_{a>d}} P(C) [P(I_{RC} \cap \Pi_{MC}) \cdot (a_C + b_C - c_C - d_C)] > 0.$$

We know that I_{RC} implies that $a_C - c_C \geq d_C - b_C$. Since we have assumed that $a_C - c_C \neq d_C - b_C$, we have that $a_C - c_C > d_C - b_C$. Hence, the inequality

$$\sum_{C_{a>d}} P(C) [P(I_{RC} \cap \Pi_{MC}) \cdot (a_C + b_C - c_C - d_C)] > 0$$

is satisfied. So, when $a_C > d_C$, R strictly dominates M . Symmetrically, from $a < d$ and by distinguishing between the two cases (1) and (2) as before, in the end we get:

- (1) $\sum_{C_{a>d}} P(C) \left[-\frac{1}{2} a_C - \frac{1}{2} b_C + \frac{1}{2} c_C + \frac{1}{2} d_C \right] = 0$; and
- (2) $\sum_{C_{a>d}} P(C) [P(\Pi_{RC} \cap I_{MC}) \cdot (-a_C - b_C + c_C + d_C)] > 0$.

Hence, we can conclude that R strictly dominates M in the class $\mathcal{C} \cup \mathcal{A}$.

It remains to be shown that $F_{\mathcal{C} \cup \mathcal{A}}(M, M) < F_{\mathcal{C} \cup \mathcal{A}}(R, M)$. As before, spell this out as:

$$\begin{aligned} & \sum_C P(C) [P(I_{MC} \cap I_{MC}) \cdot a_C + P(\Pi_{MC} \cap \Pi_{MC}) \cdot d_C \\ & \quad + P(I_{MC} \cap \Pi_{MC}) \cdot b_C + P(\Pi_{MC} \cap I_{MC}) \cdot c_C] \\ & + \sum_A P(A) [P(I_{MA} \cap I_{MA}) \cdot a_A + P(\Pi_{MA} \cap \Pi_{MA}) \cdot d_A \\ & \quad + P(I_{MA} \cap \Pi_{MA}) \cdot b_A + P(\Pi_{MA} \cap I_{MA}) \cdot c_A] \\ & < \sum_C P(C) [P(I_{RC} \cap I_{MC}) \cdot a_C + P(\Pi_{RC} \cap \Pi_{MC}) \cdot d_C \\ & \quad + P(I_{RC} \cap \Pi_{MC}) \cdot b_C + P(\Pi_{RC} \cap I_{MC}) \cdot c_C] \\ & + \sum_A P(A) [P(I_{RA} \cap I_{MA}) \cdot a_A + P(\Pi_{RA} \cap \Pi_{MA}) \cdot d_A \\ & \quad + P(I_{RA} \cap \Pi_{MA}) \cdot b_A + P(\Pi_{RA} \cap I_{MA}) \cdot c_A]. \end{aligned}$$

Similarly to the above derivation, we consider $a > d$ first, and we now distinguish between (1) $b = c$, (2) $b > c$, and (3) $b < c$. Notice that either (1) or (2), together with $a > d$, implies I_{RC} . Then we obtain:¹⁰

$$\begin{aligned}
 (1) \quad & \sum_{C_{a>d}} P(C) \left[-\frac{1}{2}a_c - \frac{1}{2}b_c + \frac{1}{2}c_c + \frac{1}{2}d_c \right] < 0; \\
 (2) \quad & \sum_{C_{a>d}} P(C) [a_c \cdot (P(I_{MC} \cap I_{MC}) - P(I_{RC} \cap I_{MC})) + b_c \\
 & \quad \cdot (P(I_{MC} \cap I_{MC}) - P(I_{RC} \cap I_{MC}))] = 0; \\
 (3) \quad & \sum_{C_{a>d}} P(C) [a_c \cdot (-P(I_{RC} \cap \Pi_{MC})) + b_c \cdot (-P(I_{RC} \cap \Pi_{MC})) \\
 & \quad + c_c \cdot (P(\Pi_{MC} \cap \Pi_{MC}) - P(\Pi_{RC} \cap \Pi_{MC})) + d_c \\
 & \quad \cdot (P(\Pi_{MC} \cap \Pi_{MC}) - P(\Pi_{RC} \cap \Pi_{MC}))] \leq 0.
 \end{aligned}$$

When $a < d$, the derivation proceeds symmetrically and we get:

$$\begin{aligned}
 (1) \quad & \sum_{C_{a<d}} P(C) \left[\frac{1}{2}a_c + \frac{1}{2}b_c - \frac{1}{2}c_c - \frac{1}{2}d_c \right] < 0; \\
 (2) \quad & \sum_{C_{a<d}} P(C) [a_c \cdot (P(I_{MC} \cap I_{MC}) - P(I_{RC} \cap I_{MC})) + b_c \\
 & \quad \cdot (P(I_{MC} \cap I_{MC}) - P(I_{RC} \cap I_{MC})) + c_c \cdot (-P(\Pi_{RC} \cap I_{MC})) \\
 & \quad + d_c \cdot (-P(\Pi_{RC} \cap I_{MC}))] \leq 0; \\
 (3) \quad & \sum_{C_{a<d}} P(C) [c_c \cdot (P(\Pi_{MC} \cap \Pi_{MC}) - P(\Pi_{RC} \cap \Pi_{MC})) + d_c \\
 & \quad \cdot (P(\Pi_{MC} \cap \Pi_{MC}) - P(\Pi_{RC} \cap \Pi_{MC}))] = 0.
 \end{aligned}$$

Finally, we can conclude that $F_{C \cup A}(M, M) < F_{C \cup A}(R, M)$.

When we have i.i.d. sampling with continuous support, games in \mathcal{W} and \mathcal{B} never occur, and the proof for the other cases reduces to the proof of proposition 2 for $s = 0$ and $t = 1$.

Proof of Proposition 2. As shown in figure A1a, given a game (a, b, c, d) , action I corresponds to the line $a + (b - a)x$, while action II corresponds to the line $c + (d - c)x$. The slope of action I is then $(b - a)$, and the slope of action II is $(d - c)$. Action I is steeper than action II if $|b - a| > |d - c|$, and action II is steeper than action I if the reverse of the last inequality

10. Note that here, when we only have three elements in the support, case (2) is impossible, but cases (1) and (3) can occur with positive probability, and this enough for our purpose.




Figure A1. Examples of coordination game C and corresponding anticoordination game $\psi(C)$.

holds. Given a belief $\Gamma = [s, t]$, let us define

$$\begin{aligned}
 a' &:= (1 - s)a + sb = a + s(b - a) \\
 b' &:= (1 - t)a + tb = a + t(b - a) \\
 c' &:= (1 - s)c + sd = c + s(d - c) \\
 d' &:= (1 - t)c + td = c + t(d - c).
 \end{aligned}
 \tag{4}$$

Next, type R is indifferent between the two acts if

$$\frac{c - a}{c - a + b - d} - s = t - \frac{c - a}{c - a + b - d}$$

and prefers I over II if

$$\frac{c - a}{c - a + b - d} - s < t - \frac{c - a}{c - a + b - d}.$$

For succinctness, let us abbreviate $Z := c - a + b - d$. Whenever $|d - c| > |b - a|$ and $[(c - a)/Z] - s \leq t - [(c - a)/Z]$, it is the case that $a' > d'$, so that M prefers I over II. Indeed, when $s = t = (c - a)/Z$, we have that $a' = d' = (cb - ad)/Z$. When we enlarge the interval Γ by moving s to the left of $(c - a)/Z$ and t to the right of $(c - a)/Z$ by the same extent, such that $[(c - a)/Z] - s = t - [(c - a)/Z]$, we get that $a' > d'$, since a' moved by

$$(a - b) \left(\frac{c - a}{Z} - s \right)$$

while d' moved by

$$(d - c) \left(t - \frac{c - a}{Z} \right).$$

Consequently, in a game where action II is steeper than action I, the only possible way in which the two types can differ is M playing I and R playing II, since we will never observe M playing II and R playing I in such games. By a similar argument, when action I is steeper than action II, the only possible way in which the two types can differ is M playing II and R playing I. Hence, whenever one of the actions is steeper than the other, the two types can differ in only one of the two possible ways. Finally, when action I and action II are equally steep, $|d - c| = |b - a|$, then one type strictly prefers one action if and only if the other type does too. We say that a game is relevant if the two types play different actions.

Since the two cases are symmetric, consider the case $|d - c| = |b - a|$. Suppose that c and d have been drawn such that $c < d$. For $|d - c| = |b - a|$ to hold, the game has to be a coordination game, otherwise we would have $b = d$ and $c = a$, so $|d - c| < |b - a|$.

If the game is a coordination game, the two types choose differently if and only if

$$s < \frac{c - a}{Z} < \frac{s + t}{2}$$

and

$$b' > c'.$$

Consider all the points $P \in (s, (s + t)/2)$. Each of these points can be expressed as a linear combination:

$$\frac{ks + (n - k)t}{n},$$

for $k > (n/2)$. By simple algebra, for each point in $P \in (s, [(s + t)/2])$ and $c < d$, the point

$$c + \frac{ks + (n - k)t}{n} (d - c)$$

expresses the expected value of action II for

$$P = \frac{ks + (n - k)t}{n},$$

that is, it is the y -value of the line corresponding to action II when

$$x = \frac{ks + (n - k)t}{n}.$$

Then, given a two-dimensional point (x_0, y_0) , the sheaf of lines passing through that point is defined by all the equations

$$y - y_0 = m(x - x_0)$$

for $m \in \mathbb{R}$. (The vertical line $m = \infty$ is excluded from the sheaf, but it is not relevant for the proof.) Consequently, given the two-dimensional point

$$\left(\frac{ks + (n - k)t}{n}, c + \frac{ks + (n - k)t}{n}(d - c) \right),$$

the sheaf of lines passing through that point is defined by the set of equations, for $m \in \mathbb{R}$:

$$y = m \left(x - \frac{ks + (n - k)t}{n} \right) + c + \frac{ks + (n - k)t}{n}(d - c).$$

If, for each equation in the set, we define

$$\begin{aligned} a^\circ &:= m \left(-\frac{ks + (n - k)t}{n} \right) + c + \frac{ks + (n - k)t}{n}(d - c) \\ b^\circ &:= m \left(1 - \frac{ks + (n - k)t}{n} \right) + c + \frac{ks + (n - k)t}{n}(d - c) \end{aligned}, \tag{5}$$

then each equation corresponds to a possible game

	I	II
I	a°	b°
II	c	d

such that

$$\frac{c - a^\circ}{c - a^\circ + b^\circ - d} = \frac{ks + (n - k)t}{n}.$$

By algebraic computations, the condition $d - c > |a^\circ - b^\circ|$ is equivalent to

$$|m| < d - c.$$

Moreover, among the coordination games such that $d - c > |a^\circ - b^\circ|$, the relevant ones are those that also satisfy $b' > c'$, otherwise type M would not (strictly) prefer I over II. If we rewrite a', b', c', d' as in equation (4), then the inequality $b' > c'$ reduces to

$$m \left(\frac{k}{k - n} \right) < d - c.$$

By symmetric arguments, whenever $c > d$, the only possible relevant games for the same interval $[s, t]$ are anticonoordination games for which

$$\frac{s + t}{2} < \frac{c - a}{Z} < t,$$

and such that $c - d > |a - b|$, and that $d' > d'$. Similarly, for $k < n/2$, these correspond to all games

$$y = m \left(x - \frac{ks + (n - k)t}{n} \right) + c + \frac{ks + (n - k)t}{n} (d - c)$$

such that

$$|m| < c - d$$

and

$$m \left(\frac{k - n}{k} \right) > d - c.$$

Consider now the following bijective function $\psi : C \rightarrow \mathcal{A}$ between coordination and anticonoordination games, that, for $d > c$, associates the coordination game

$$y = m \left(x - \frac{ks + (n - k)t}{n} \right) + c + \frac{ks + (n - k)t}{n} (d - c)$$

with the anticonoordination game

$$y = -m \left(x - \frac{(n - k)s + kt}{n} \right) + d + \frac{(n - k)s + kt}{n} (c - d).$$

Essentially, ψ changes c to d , m to $-m$, and k to $n - k$. In particular, note that ψ is a bijection that, for a fixed interval $[s, t]$, sends relevant coordination games to relevant anticonoordination games. Figure A1 gives a graphical example of the bijection.

We can then pair these two games and consider the average fitness in $\{C, \psi(C)\}$ of (reg, $[s, t]$) against (reg, $[s, t]$), and then compare it to the fitness of (obj, $[s, t]$) against (reg, $[s, t]$), denoted respectively by M and R henceforth. In the pair of relevant games C and $\psi(C)$, R strictly dominates M if

$$F_{\{C, \psi(C)\}}(R, R) > F_{\{C, \psi(C)\}}(M, R)$$

and

$$F_{\{C, \psi(C)\}}(R, M) > F_{\{C, \psi(C)\}}(M, M).$$

Consider the first inequality. Since both C and $\psi(C)$ are relevant with respect to the interval $[s, t]$, it implies that $F_{\{C, \psi(C)\}}(R, R) = d + c$, and $F_{\{C, \psi(C)\}}(M, R) = b^\circ + \psi(b^\circ)$. Therefore, the first inequality is equivalent to

$$\frac{d + c}{2} > \frac{b^\circ + \psi(b^\circ)}{2},$$

which can be spelled out as

$$\begin{aligned} d + c &> m \left(1 - \frac{ks + (n - k)t}{n} \right) + c + \frac{ks + (n - k)t}{n} (d - c) \\ &\quad - m \left(1 - \frac{(n - k)s + kt}{n} \right) + d + \frac{(n - k)s + kt}{n} (c - d). \end{aligned}$$

After some computations, the previous inequality boils down to

$$d - c > m,$$

which we know is the case, since we have seen that the condition $d - c > |a^\circ - b^\circ|$ is equivalent to $d - c > |m|$.

Finally, from the previous argument it follows that, for any given interval $[s, t]$, if we consider the set of all relevant coordination games, is denote it by C^r , and the set of all relevant anticoordination games \mathcal{A}^r , then it holds that

$$F_{\{C \cup \mathcal{A}\}}(R, R) > F_{\{C \cup \mathcal{A}\}}(M, R).$$

Let us now check that the second inequality for R to strictly dominate M also holds. In $\{C, \psi(C)\}$, that is equivalent to

$$d + c > a^\circ + \psi(a^\circ),$$

which amounts to $m < d - c$. As before, it then follows that

$$F_{\{C \cup \mathcal{A}\}}(R, M) > F_{\{C \cup \mathcal{A}\}}(M, M).$$

Therefore, R strictly dominates M .

REFERENCES

Alchian, Armen. 1950. "Uncertainty, Evolution and Economic Theory." *Journal of Political Economy* 58:211–21.

Alger, Ingela, and Jörgen W. Weibull. 2013. "Homo Moralis: Preference Evolution under Incomplete Information and Assortative Matching." *Econometrica* 81 (6): 2269–2302.

Anderson, John R. 1991. "Is Human Cognition Adaptive?" *Behavioral and Brain Sciences* 14 (3RA): 471–517.

Battigalli, Pierpaolo et al. 2015. "Self-Confirming Equilibrium and Model Uncertainty." *American Economic Review* 105 (2): 646–77.

- Becker, Gary S. 1976. "Altruism, Egoism, and Genetic Fitness: Economics and Sociobiology." *Journal of Economic Literature* 14:817–26.
- Bednar, Jenna, and Scott Page. 2007. "Can Game(s) Theory Explain Culture? The Emergence of Cultural Behavior within Multiple Games." *Rationality and Society* 19 (1): 65–97. <http://rss.sagepub.com/content/19/1/65.full.pdf+html>.
- Bell, David E. 1982. "Regret in Decision Making under Uncertainty." *Operations Research* 30 (5): 961–81.
- Bester, Helmut, and Werner Güth. 1998. "Is Altruism Evolutionarily Stable?" *Journal of Economic Behavior and Organization* 34:193–209.
- Bleichrodt, Han, and Peter P. Wakker. 2015. "Regret Theory: A Bold Alternative to the Alternatives." *Economic Journal* 125 (583): 493–532.
- Charness, Gary, and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117 (3): 817–69.
- Chater, Nick, and Mike Oaksford. 2000. "The Rational Analysis of Mind and Behavior." *Synthese* 122:93–131.
- de Finetti, Bruno. 1937. "La prevision: Ses lois logiques, ses sources subjectives." *Annales de l'Institute Henri Poincare* 7:1–68.
- Dekel, Eddie, Jeffrey C. Ely, and Okan Ylankaya. 2007. "Evolution of Preferences." *Review of Economic Studies* 74 (3): 685–704.
- Ellsberg, Daniel. 1961. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics* 75 (4): 643–69.
- Fawcett, Tim W., Steven Hamblin, and Luc-Alain Giraldeau. 2013. "Exposing the Behavioral Gambit: The Evolution of Learning and Decision Rules." *Behavioral Ecology* 24 (1): 2–11.
- Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics* 114:817–68.
- Fishburn, Peter C. 1982. "Nontransitive Measurable Utility." *Journal of Mathematical Psychology* 26 (1): 31–67.
- Friedman, Milton. 1953. *Essays in Positive Economics*. Chicago: University of Chicago Press.
- Gardenfors, Peter, and Nils-Eric Sahlin. 1982. "Unreliable Probabilities, Risk Taking, and Decision Making." *Synthese* 53:361–86.
- Ghirardato, Paolo, and Massimo Marinacci. 2002. "Ambiguity Made Precise: A Comparative Foundation." *Journal of Economic Theory* 102:251–89.
- Gigerenzer, Gerd, and Daniel G. Goldstein. 1996. "Reasoning the Fast and Frugal Way: Models of Bounded Rationality." *Psychological Review* 103 (4): 650–69.
- Gilboa, Itzhak, and Massimo Marinacci. 2013. "Ambiguity and the Bayesian Paradigm." In *Advances in Economics and Econometrics*, vol. 1, ed. Daron Acemoglu, Manuel Arellano, and Eddie Dekel, 179–242. Cambridge: Cambridge University Press.
- Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin Expected Utility with Non-unique Prior." *Journal of Mathematical Economics* 18:141–53.
- Hagen, Edward H., et al. 2012. "Decision Making: What Can Evolution Do for Us?" In *Evolution and the Mechanisms of Decision Making*, ed. Peter Hammerstein and Jeffrey R. Stevens, 97–126. Cambridge, MA: MIT Press.
- Halpern, Joseph Y., and Rafael Pass. 2012. "Iterated Regret Minimization: A New Solution Concept." *Games and Economic Behavior* 74:184–207.
- Hammerstein, Peter, and Jeffrey R. Stevens. 2012. "Six Reasons for Invoking Evolution in Decision Theory." In *Evolution and the Mechanisms of Decision Making*, ed. Peter Hammerstein and Jeffrey R. Stevens. Cambridge, MA: MIT Press.
- Harley, Calvin B. 1981. "Learning the Evolutionarily Stable Strategy." *Journal of Theoretical Biology* 89 (4): 611–33.
- Heifetz, Aviad, Martin Meier, and Burkhard C. Schipper. 2013. "Dynamic Unawareness and Rationalizable Behavior." *Games and Economic Behavior* 81:50–68.
- Levi, Isaac. 1974. "On Indeterminate Probabilities." *Journal of Philosophy* 71 (13): 391–418.
- Loomes, Graham, and Robert Sugden. 1982. "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty." *Economic Journal* 92 (368): 805–24.
- Maynard Smith, John. 1982. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.

- McNamara, John M. 2013. "Towards a Richer Evolutionary Game Theory." *Journal of the Royal Society Interface* 10 (88): 1–9.
- Mengel, Friederike. 2012. "Learning across Games." *Games and Economic Behavior* 74 (2): 601–19.
- Mohlin, Erik. 2012. "Evolution of Theories of Mind." *Games and Economic Behavior* 75 (1): 299–318.
- Nowak, Martin A. 2006. *Evolutionary Dynamics: Exploring the Equations of Life*. Cambridge, MA: Harvard University Press.
- O'Connor, Cailin. 2015. "Evolving to Generalize: Trading Precision for Speed." *British Journal for the Philosophy of Science*. doi:10.1093/bjps/axv038.
- Rayo, Luis, and Gary S. Becker. 2007. "Evolutionary Efficiency and Happiness." *Journal of Political Economy* 115:302–37.
- Robalino, Nikolaus, and Arthur Robson. 2016. "The Evolution of Strategic Sophistication." *American Economic Review* 106 (4): 1046–72.
- Robson, Arthur, and Larry Samuelson. 2011. "The Evolutionary Foundations of Preferences." In *Handbook of Social Economics*, vol. 1A, ed. J. Benhabib, M. Jackson, and A. Bisin. Amsterdam: North-Holland.
- Savage, Leonard Jimmie. 1951. "The Theory of Statistical Decision." *Journal of the American Statistical Association* 46:55–67.
- . 1954. *The Foundations of Statistics*. New York: Dover.
- Scheibehenne, Benjamin, Jörg Rieskamp, and Eric-Jan Wagenmakers. 2013. "Testing the Adaptive Toolbox Models: A Bayesian Hierarchical Approach." *Philosophical Review* 120 (1): 39–64.
- Skyrms, Brian, and Kevin J. S. Zollman. 2010. "Evolutionary Considerations in the Framing of Social Norms." *Politics, Philosophy and Economics* 9 (3): 265–73. <http://ppe.sagepub.com/content/9/3/265.full.pdf+html>.
- Smead, Rory, and Kevin J. S. Zollman. 2013. "The Stability of Strategic Plasticity." Unpublished manuscript, Carnegie Mellon University.
- Taylor, Peter D., and Leo B. Jonker. 1978. "Evolutionary Stable Strategies and Game Dynamics." *Mathematical Bioscience* 40 (1–2): 145–56.
- Trautmann, Stefan T., and Gijs van de Kuilen. 2016. "Ambiguity Attitudes." In *The Wiley Blackwell Handbook of Judgment and Decision Making*, ed. G. Keren and G. Wu. Chichester: Wiley-Blackwell.
- Tversky, Amos, and Daniel Kahnemann. 1981. "The Framing of Decisions and the Psychology of Choice." *Science* 211 (4481): 453–58.
- von Neumann, John, and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Walley, Peter. 1996. "Inferences from Multinomial Data: Learning about a Bag of Marbles." *Journal of the Royal Statistical Society* 58 (1): 3–57.
- Zollman, Kevin J. S. 2008. "Explaining Fairness in Complex Environments." *Politics, Philosophy and Economics* 7 (1): 81–97.
- Zollman, Kevin J. S., and Rory Smead. 2010. "Plasticity and Language: An Example of the Baldwin Effect?" *Philosophical Studies* 147 (1): 7–21.