



The ecological rationality of decision criteria

Paolo Galeazzi^{1,2}  · Alessandro Galeazzi³

Received: 6 November 2019 / Accepted: 7 July 2020 / Published online: 1 August 2020
© The Author(s) 2020

Abstract

Standard evolutionary game theory investigates the evolutionary fitness of alternative behaviors in a fixed and single decision problem. This paper instead focuses on decision criteria, rather than on simple behaviors, as the general behavioral rules under selection in the population: the evolutionary fitness of classic decision criteria for rational choice is analyzed through Monte Carlo simulations over various classes of decision problems. Overall, quantifying the uncertainty in a probabilistic way and maximizing expected utility turns out to be evolutionarily beneficial in general. Minimizing regret also finds some evolutionary justifications in our results, while maximin seems to be always disadvantaged by differential selection.

Keywords Ecological rationality · Evolutionary selection · Decision criteria · Population games

1 Introduction

One of the most fundamental questions that sparked the development of probability and decision theory is how to evaluate and choose from different uncertain prospects, such as lottery tickets or business ventures. Among the candidates that have been advanced in the literature, the decision criterion dictating to maximize a subjectively expected utility enjoys the largest favor at present, especially among economists.

In the wake of behavioral experiments, though, subjective expected utility maximization has been objected to not only on descriptive but also on normative grounds. Criticisms have particularly targeted its probabilistic sophistication: for expected util-

The Center for Information and Bubble Studies is funded by the Carlsberg Foundation.

✉ Paolo Galeazzi
pagale87@gmail.com

¹ University of Bayreuth, Bayreuth, Germany

² University of Copenhagen, CIBS, Copenhagen, Denmark

³ University of Brescia, Brescia, Italy

ity maximization to be applicable in a given decision problem, the decision maker is required to hold probabilistic beliefs about the state of the world and/or the actions of the opponents. In technical terms, all uncertainty must be measurable, namely, all uncertainty must be reduced to a probability measure. Decision-theoretic experiments have disproved the descriptive accuracy of such a view on uncertainty (e.g., Allais (1953), Ellsberg (1961)), but objections regarding the rightfulness of denying all non-probabilistic uncertainty have been raised from a normative standpoint too (Levi (1974), Gärdenfors (1979), Gärdenfors and Sahlin (1982), Levi (1982), Levi (1985), Gilboa et al. (2009), Gilboa et al. (2012)).

Two of the main alternative decision criteria that have been proposed for guiding rational choice in situations of uncertainty are expected utility maximization and regret minimization. Unlike expected utility maximization, these two criteria allow both probabilistic and non-probabilistic uncertainty. Maximin can be traced back at least to the work in Wald (1939), von Neumann and Morgenstern (1944), and Wald (1945). It has been later proposed as prominent alternative to expected utility maximization in numerous other works (e.g., Gärdenfors (1979), Gärdenfors and Sahlin (1982), Levi (1982)) and famously axiomatized by Gilboa and Schmeidler (1989). It prescribes to pick the option that maximizes the utility in the worst-case scenario. Regret minimization, firstly advanced by Savage (1951), has been recently axiomatized by Hayashi (2008) and Stoye (2011), and it dictates to choose the option that minimizes the loss in the worst-case scenario.

In the presence of multiple candidates, the meta-problem of selecting among different criteria for rational decision making arises. There are at least two different approaches that can shed some light on this meta-problem and guide our choices. The first and traditional one is axiomatic: some decision-theoretic axioms are considered healthy rationality principles (e.g., transitivity, Mas-Colell et al. (1995)), and decision criteria are compared and evaluated on the basis of the axioms they satisfy or do not satisfy.

The second approach is game-theoretic: it thinks of different criteria as general behavioral rules, or higher-order strategies, and it analyzes and compares their evolutionary fitness. This second approach is thus ecological: fitness and performance of each criterion have to be expected to vary depending on the characteristics of environment (e.g., the existing competitors, the decision problems at hand, and the like).

Philosophical as well as biological research has largely profited from the game-theoretic methodology (Skyrms (1996), Bergstrom and Godfrey-Smith (1998), Skyrms (2004), Okasha (2007), Zollman (2008), Skyrms (2010), Huttegger and Zollman (2010), Huttegger and Zollman (2013)). Different from classic evolutionary game theory, however, the ecological approach proposed here views decision criteria, rather than simple behaviors, as the types under selection.¹ In order to study the differential selection of general behavioral rules, the standard single-game model of evolutionary game theory is thus expanded to include a variety of different games, and therefore named multigame model. To exemplify, consider an environment consisting of two different games, the Prisoner's Dilemma and a coordination game:

¹ To the best of our knowledge, evolutionary competitions where the types under selection are general decision criteria have not been yet investigated, with only a few exceptions (Galeazzi and Franke (2017), Klein et al. (2018)).

| | | | | | |
|-----------|----------|-----------|-----------|----------|-----------|
| PD | <i>I</i> | <i>II</i> | CG | <i>I</i> | <i>II</i> |
| <i>I</i> | 1;1 | 3;0 | <i>I</i> | 1;1 | 1;0 |
| <i>II</i> | 0;3 | 2;2 | <i>II</i> | 0;1 | 3;3 |

Suppose that these two games are played in a population with two agent types differing in the decision criterion: the expected utility maximizer with uniform prior, that assigns equal probability to all the actions of the other and picks the option that maximizes his or her own expected utility, and the maximinimizer, that simply chooses the option that maximizes the utility in the worst-case scenario. In the Prisoner's Dilemma, both these criteria prescribe action *I*, since it leads to the highest expected utility ($1/2 + 3/2 = 2 > 0/2 + 2/2 = 1$) as well as to the highest minimal utility ($\min\{1, 3\} = 1 > 0 = \min\{0, 2\}$). In the coordination game, on the other hand, action *II* is associated with the highest expected utility ($1/2 + 1/2 = 1 < 0/2 + 3/2 = 3/2$), while action *I* is associated with the highest minimum ($\min\{1, 1\} > \min\{0, 3\}$).

The two decision criteria therefore dictate the same action in the Prisoner's Dilemma, but different actions in the coordination game. Had the Prisoner's Dilemma been the only game in the environment, the evolutionary pressure would have been the same on both types in the population. By introducing the coordination game, the two subpopulations may instead be subject to differential selection. Adding the coordination game to the environment thus allowed us to make a distinction we would otherwise not have been able to make.

Richer environments would permit finer discriminations: for instance, in an environment consisting only of the Prisoner's Dilemma and the coordination game above, the maximin decision criterion is not behaviorally distinguishable from the decision criterion that always prescribes action *I*. These two criteria would instead dictate different actions in many other games which can be part of richer environments. Intuitively, moreover, the richer in different games, the less biased against or in favor of certain criteria the environment will be.²

The core of this paper compares different decision criteria in terms of their evolutionary fitness. The investigation includes the main candidates for rational choice under uncertainty proposed in the literature. In particular, we contrast the maximization of expected utility with maximin expected utility and regret minimization.

Special attention is given to regret minimization, as recent axiomatizations have focused only on one of the (at least) two possible versions of regret minimization for decision making under unmeasurable uncertainty (Hayashi 2008; Stoye 2011). Here, instead, we consider both these ways of minimizing regret and compare their evolutionary performances.

The structure of the paper is as follows. Section 2 introduces the main characters by defining the decision criteria that we will encounter throughout the paper. Section 3 sets the stage by presenting the model we use in our investigations, whose results for interactive and individual decision making are in Sects. 5 and 6, respectively. Finally, Sect. 7 concludes by relating to the existing literature and the closing remarks.

² The model introduced below, however, allows the converse too: one could purposely select a specific set of games when interested in studying the evolution in a specific environment.

2 Decision criteria

In this section, we formally define the decision criteria which are studied in the following sections. As game theory can be seen as interactive decision theory, the relevant decision criteria are introduced in a decision-theoretic setting. In doing so, however, we pay attention to use formal notation that is consistent with the multigame setting developed in the next sections.

The main alternatives investigated in this paper are four: maximization of expected utility, maximin expected utility, and two forms of regret minimization. In the following analysis, the term rationality is uniquely understood and used in its ecological acceptance, in the sense of being advantageous and beneficial for the individual to survive and thrive.

By and large, decision criteria are functions associating a decision problem with an action choice. A decision problem can be formalized as a tuple (S, A, Z, c) where S is a set of possible states of the world, A is the agent's set of actions,³ Z is a set of outcomes, and c is the outcome function

$$c : S \times A \rightarrow Z.$$

On the top of the objective structure (S, A, Z, c) of the decision problem, there are two crucial subjective components: the agent's subjective utility u and the agent's subjective belief B . The subjective utility is a real-valued function associating each possible outcome with a numerical subjective utility

$$u : Z \rightarrow \mathbb{R}.$$

The subjective belief of the agent is encoded by a set B , whose elements can be objects of different nature, depending on the belief representation in the model. For instance, in a qualitative approach, the belief set B may be just a subset of possible states of the world $B \subseteq S$ which the agent considers possible or most plausible. In a more quantitative approach, the belief set may be a set of probability measures $B \subseteq \Delta(S)$, where $\Delta(S)$ denotes the set of all probability measures p over the set S .

All decision criteria considered here can be thus thought of as functions associating a subjective utility and a subjective belief with an action choice:

$$\text{Decision Criterion: Utilities} \times \text{Beliefs} \rightarrow \text{Actions}.$$

In order to output an action, the decision criterion must be able to resolve the kind of uncertainty encoded by the belief set B . For instance, if the belief set B is not a singleton set containing a unique probability measure, but rather a set of states $B \subseteq S$, actions cannot be associated with expected utilities, and the maximization of expected utility is hence unfit to solve the choice problem. In those cases, the agent has to resort to other criteria, such as maximin or regret minimization.

³ Throughout this paper, S and A are finite sets. This assumption can be dropped by accommodating the usual measure-theoretic issues.

In the following, we assume that B is a nonempty set of probability measures. The advantage of this choice is that it allows to encompass the main approaches to the formal representation of beliefs mentioned above. The models where the belief of the agent is given in terms of a set of states $B = S' \subseteq S$ can be alternatively expressed by the set of probability measures $B = \Delta(S')$. The models assuming Bayesian agents with probabilistic beliefs are the special cases where the set $B \subseteq \Delta(S)$ is a singleton.

2.1 Expected utility maximization

Of all decision criteria advanced in the literature, the maximization of a subjective expected utility has gained most favor for defining rational choice. Given a decision problem where the belief set B is represented by a single probability measure $p \in \Delta(S)$, the subjective expected utility of each action $a \in A$ is

$$E_p[u|a] = \sum_{s \in S} u(c(s, a)) \cdot p(s).$$

According to expected utility maximization, the rational choice to make is to pick an action a^* such that

$$a^* \in \operatorname{argmax}_{a \in A} E_p[u|a].$$

Whenever the uncertainty of the agent is not measurable, i.e., whenever the set B is not representable by a single probability measure, each action a is thus associated with an expected utility $E_q[u|a]$ for each probability measure $q \in B$. Under unmeasurable uncertainty, also known as ambiguity, the maximization of subjective expected utility is thus unable to assign each option a unique expected-utility value and to prescribe a course of action accordingly. When faced with unmeasurable uncertainty, the agent necessitates more general decision criteria to evaluate her options.

2.2 Maximin expected utility

The most famous criterion for decision making under ambiguity is probably the *maximinimization* of subjective expected utility, or just maximin expected utility. This criterion dictates to rank the actions according to the minimal expected utility and then choose among the top-ranked:

$$a^* \in \operatorname{argmax}_{a \in A} \min_{p \in B} E_p[u|a].$$

Note that maximin expected utility is adequate for choice under both unmeasurable and measurable uncertainty. When the uncertainty is measurable, that is when B is reduced to a singleton, maximin expected utility reduces to the maximization of expected utility.

2.3 Linear regret minimization

Another important criterion for decision under both measurable and unmeasurable uncertainty is regret minimization, which comes in at least two different forms. The more common version of regret minimization is the one that we call linear, and has recent axiomatizations in Hayashi (2008) and Stoye (2011).

To formalize linear regret minimization, let us first define the linear regret of action a given probability measure p as the following quantity:

$$r_L(a, p) := E_p \left[\max_{a' \in A} u(c(s, a')) - u(c(s, a)) \right] = \sum_{s \in S} p(s) \cdot \left(\max_{a' \in A} u(c(s, a')) - u(c(s, a)) \right).$$

Linear regret minimization prescribes to pick an action a^* such that

$$a^* \in \operatorname{argmin}_{a \in A} \max_{p \in B} r_L(a, p).$$

Linear regret minimization also prescribes the same action choice as expected utility maximization in case of measurable uncertainty (see Appendix A).

2.4 Nonlinear regret minimization

An equally reasonable version of regret minimization is the one we call nonlinear regret minimization. Let us define the nonlinear regret of action a given probability measure p as the quantity

$$r_N(a, p) := \max_{a' \in A} E_p[u|a'] - E_p[u|a] = \max_{a' \in A} \sum_{s \in S} p(s) \cdot u(c(s, a')) - \sum_{s \in S} p(s) \cdot u(c(s, a)).$$

Nonlinear regret minimization accordingly prescribes to opt for an action a^* such that

$$a^* \in \operatorname{argmin}_{a \in A} \max_{p \in B} r_N(a, p).$$

Similarly to linear regret minimization, nonlinear regret minimization too agrees with expected utility maximization on the action that shall be chosen in case of measurable uncertainty (see Appendix A).

Example 1 To better appreciate the behavioral difference between the two ways of understanding and minimizing regret, consider the following example. A bag contains ten marbles, which are either blue or red. Seven marbles are red, one is blue, and nothing else is known about the remaining two (the corresponding belief set can be simply expressed as $\{0.1, 0.2, 0.3\}$, where it suffices to state the probability of a blue marble only, since drawing a red marble is the complementary event). A marble will be drawn from the bag and an agent has the chance to bet on its color according to the

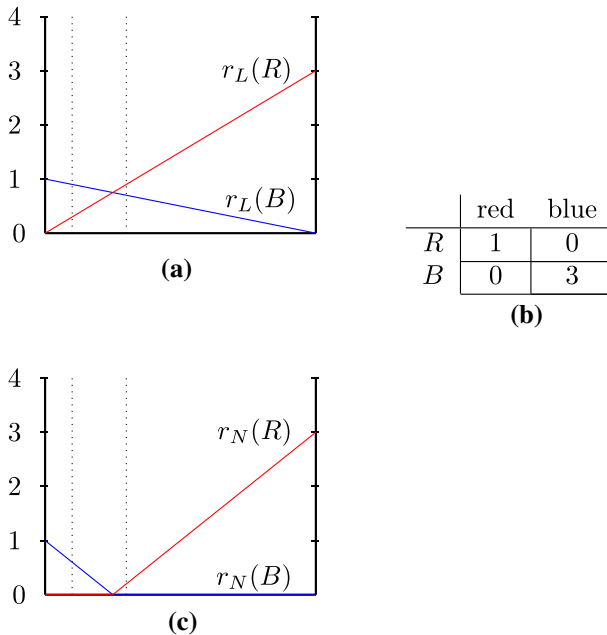


Fig. 1 An example where the choices prescribed by linear regret and nonlinear regret differ. Figure 1a, c picture, respectively, the linear and the nonlinear regret of the two bets. The uncertainty about the bag is represented by the interval between the two vertical dotted lines, with the left line corresponding to a probability of a blue marble of 0.1 (i.e., the remaining two marbles are red) and the right line to a probability of 0.3 (i.e., the remaining marbles are all blue). (Color figure online)

payoffs given in Fig. 1b: a winning bet on blue yields 3, a winning bet on red yields 1, and losing bets cost nothing.

A linear regret minimizer will look at Fig. 1a and pick the action corresponding to the line whose highest point within the dotted interval is lower. A linear regret minimizer is hence indifferent between *R* and *B*, as the two lines reach equally high maxima within the dotted interval (where the maximum of $r_L(R)$ is reached when the probability of a blue marble is 0.3, and the maximum of $r_L(B)$ when that probability is 0.1). On the contrary, a nonlinear regret minimizer will prefer to bet on red: as depicted in Fig. 1c, the maximum reached by $r_N(R)$ within the dotted interval is lower than the maximum of $r_N(B)$.

It is worth stressing that in order to resolve the same uncertainty about the bag of marbles through expected utility maximization, a single probability distribution is needed. Given the knowledge on the composition of the bag (7 red marbles, 1 blue marble, 2 either red or blue marbles), it is reasonable for an expected utility maximizer to discard all probability distributions that are out of the interval delimited by the two vertical dotted lines and opt for a distribution that assigns probability of a blue marble between 0.1 and 0.3. In what follows, we assume that the expected utility maximizer reaches a probabilistic belief by averaging the points of the belief set, e.g., $(0.1 + 0.2 + 0.3)/3 = 0.2$ in the case of the bag of marbles. The expected utility maximizer may thus be viewed as resolving unmeasurable uncertainty by means

of the *principle of insufficient reason*, or principle of indifference, which prescribes to assign equal probabilistic weight to all possible alternatives. In the following, by expected utility maximizer we hence mean a decision maker employing the principle of insufficient reason first and the maximization of expected utility next.

All candidate criteria for rational choice that we have seen agree in all those cases where the belief of the agent is represented by a unique probability measure. The disagreement about the course of action to be taken arises when the agent is unmeasurably uncertain. Maximin expected utility and both versions of regret minimization may all prescribe different behaviors, whereas the maximization of expected utility is of no help unless the uncertainty is reduced to a probabilistic belief.

3 The model

In the multigame model, the environment consists of a set of symmetric games with n possible actions. The multigame model is defined by a repeated loop with four steps. First, n^2 fitness values are randomly drawn from a set of i.i.d. possible values $\{0, \dots, \bar{v}\}$, and a new fitness game is thereby generated.⁴

| | | | |
|-----|-----------|-----|-----------|
| | I | ... | n |
| I | $v_{1,1}$ | ... | $v_{1,n}$ |
| ... | ... | ... | ... |
| n | $v_{n,1}$ | ... | $v_{n,n}$ |

Second, a belief set B is selected at random by the following procedure: a set $\{p_1, \dots, p_m\}$ of m i.i.d. points is randomly drawn from a uniform distribution over the simplex $\Delta(\{I, \dots, n\})$. Each point p_i is thus a probability measure over the actions of the opponent, and we set $B = \{p_1, \dots, p_m\}$.⁵ Third, the action prescribed by each decision criterion is computed and all decision criteria in the population are paired to play the game. Fourth and finally, the fitness achieved by each criterion against the others is recorded, and the process repeats.

Three remarks on the model are in order. First, the players hold the same beliefs independent of their criterion. The reason is that we aim to study the ecological value of different decision criteria *ceteris paribus*. It would otherwise be possible that one criterion outperforms another in terms of evolutionary fitness just because of the difference in beliefs, e.g., because one player happens to have better information than the other, rather than because one criterion is actually fitter than the other. We hence vary only the decision criteria between players, all else being equal.

Second, no constraints on beliefs of rationality are imposed. By construction of the model, one agent may in fact believe that the other will play a dominated action with positive probability.

⁴ Since the games are symmetric, it suffices to write the values for the row player only.

⁵ For maximin and the two regret minimizations, this is equivalent to set $B = \mathcal{C}(\{p_1, \dots, p_m\})$, where $\mathcal{C}(\{p_1, \dots, p_m\})$ denotes the convex hull of the set of points $\{p_1, \dots, p_m\}$. (See Appendix A.) This is relevant because those three decision criteria are usually represented by means of a convex compact set of probability measures (Gilboa and Schmeidler 1989; Hayashi 2008; Stoye 2011).

It is beyond the scope of this paper to compare alternative decision criteria when further assumptions on belief of rationality are put in place. To reiterate, here we limit ourselves to a merely ecological understanding of rationality, in the sense of being evolutionarily successful, and therefore do not impose any epistemic constraints, except for the fact that the players hold the same beliefs independent of their criterion. We leave the study of the interplay between evolutionary success and epistemically sophisticated agents to future work.

Third, in order for expected utility maximization to produce a decision, the possibly unmeasurable uncertainty has to be reduced to measurable uncertainty. Given the uncertainty encoded in $B = \{p_1, \dots, p_m\}$, we assume that a Bayesian agent assigns equal probability $1/m$ to all the points in B . This gives rise to the probabilistic belief

$$\bar{p}(a) = \frac{1}{m} \sum_{i=1}^m p_i(a)$$

for all $a \in A$.⁶

4 Method

We analyze the evolutionary competition between different criteria through Monte Carlo simulations.⁷ The parameters of the model are the following:

- n : the number of actions in the symmetric games, $A = \{a_1, \dots, a_n\}$
- m : the number of points in the belief set, $B = \{p_1, \dots, p_m\}$
- \bar{v} : the maximum of the possible fitness values $\{0, \dots, \bar{v}\}$

We test the model for the parameters $n = \{2, 3, 5, 7, 9, 12, 15\}$, $m = \{2, 3, 7\}$, and $\bar{v} = \{50, 100\}$. Ten simulation runs are performed for each possible combination (n, m, \bar{v}) of parameter values. Each simulation run represents a multigame consisting of 10000 randomly generated games.

5 Results

Each simulation run outputs a table with the average fitnesses accumulated over the 10000 randomly generated games by each decision criterion:

| | <i>LRm</i> | <i>NRm</i> | <i>Mm</i> | <i>EU</i> |
|------------|------------|------------|-----------|-----------|
| <i>LRm</i> | 32,228 | 32,26 | 32,248 | 32,27 |
| <i>NRm</i> | 32,229 | 32,175 | 32,198 | 32,15 |
| <i>Mm</i> | 31,788 | 31,771 | 31,847 | 31,769 |
| <i>EU</i> | 32,356 | 32,303 | 32,379 | 32,286 |

The table above, for instance, shows the average accumulated fitnesses of the four decision criteria (linear regret minimization *LRm*, nonlinear regret minimiza-

⁶ As already stressed in Example 1, this may be seen as the Bayesian applying the *principle of insufficient reason* to the belief set in order to quantify the uncertainty in a probabilistic way.

⁷ The source code is available from the authors upon request.

Table 1 Number of times *EU* strictly dominates all other criteria for all combinations of *m* (in the rows), *n* (in the columns), and \bar{v} ($\bar{v} = 50$ in Table 1a, and $\bar{v} = 100$ in Table 1b)

| (a) | | | | | | | |
|-----|---|----|----|----|----|----|----|
| | 2 | 3 | 5 | 7 | 9 | 12 | 15 |
| 2 | 0 | 0 | 4 | 8 | 7 | 8 | 8 |
| 3 | 0 | 1 | 10 | 10 | 10 | 10 | 10 |
| 7 | 0 | 7 | 10 | 10 | 10 | 10 | 10 |
| (b) | | | | | | | |
| | 2 | 3 | 5 | 7 | 9 | 12 | 15 |
| 2 | 0 | 0 | 5 | 5 | 7 | 6 | 5 |
| 3 | 0 | 0 | 9 | 10 | 10 | 10 | 10 |
| 7 | 0 | 10 | 10 | 10 | 10 | 10 | 10 |

tion *NRM*, maximin expected utility *Mm*, and expected utility maximization *EU*) for $n = 3$, $m = 7$ and $\bar{v} = 50$. Specifically, the number in each cell is the average accumulated fitness of a player making choices according to the criterion of the corresponding row when paired with a player who complies with the prescriptions by the criterion of the corresponding column.

5.1 Qualitative analysis

Imagining a population of agents whose decision criteria are the competing traits in the population, the evolutionary selection of the decision criteria would then be driven by the fitness accumulated in the *multigame*. Looking at the results for the tested parameters, a few things can be noted. Table 1a, b show that expected utility maximization *EU* turns out to be strictly better than the other decision criteria in many but not all cases.

This comes from the fact that, out of the ten multigames for each combination of parameters, the average accumulated fitnesses are most of the times highest in the row corresponding to *EU*. When this happens, expected utility maximization is the only evolutionarily stable criterion, and a monomorphic state of *EU*-types is the only evolutionarily stable state of the population (under replicator-mutator dynamics, Nowak (2006)). That is, the sole monomorphic population that cannot be invaded by mutants using a different criterion is a population consisting entirely of expected utility maximizers, and any initial population configuration will lead to a monomorphic state of *EU*-types, provided that there is a nonzero chance of expected utility maximizers to spring by mutation from the other types.

There are, however, exceptions to the evolutionary supremacy of expected utility maximization. Under maximal unmeasurable uncertainty, expected utility maximization, linear regret minimization and nonlinear regret minimization are provably equivalent when $n = 2$, and therefore are all evolutionarily stable, whereas maximin is provably inferior (see Galeazzi and Franke (2017)). Furthermore, the nonlinear version of regret remains equivalent to utility maximization under nonmaximal uncertainty too, for $n = 2$ and $m = 2$.

Table 2 Number of times LRm strictly dominates all other criteria for all combinations of m (in the rows), n (in the columns), and \bar{v} ($\bar{v} = 50$ in Table 2a, and $\bar{v} = 100$ in Table 2b)

| (a) | | | | | | | |
|-----|----|---|---|---|---|----|----|
| | 2 | 3 | 5 | 7 | 9 | 12 | 15 |
| 2 | 10 | 5 | 1 | 0 | 0 | 0 | 0 |
| 3 | 10 | 0 | 0 | 0 | 0 | 0 | 0 |
| 7 | 8 | 0 | 0 | 0 | 0 | 0 | 0 |
| (b) | | | | | | | |
| | 2 | 3 | 5 | 7 | 9 | 12 | 15 |
| 2 | 10 | 5 | 0 | 0 | 0 | 0 | 0 |
| 3 | 10 | 1 | 0 | 0 | 0 | 0 | 0 |
| 7 | 9 | 0 | 0 | 0 | 0 | 0 | 0 |

These considerations partly explain why expected utility maximization turns out never strictly dominant in Tables 1a, b for $m = 2$. Moreover, in these and the remaining cases where EU is not strictly dominant, linear regret minimization comes out undominated by expected utility maximization and sometimes even as the strictly dominant criterion, as shown in Table 2a, b.

5.2 Statistical analysis

Pooling together all ten Monte Carlo simulations for each of the parameter combinations, we also perform statistical analyses of the differences in fitness between different criteria. Figure 2 for instance shows the performances of the decision criteria against linear regret minimization and nonlinear regret minimization.

In order to spare space, Fig. 2 pictures only the fitnesses against two of the four criteria for $m = 2$ and $\bar{v} = 50$, but all four cases look alike. The horizontal axis represents the parameter n , i.e., the number of actions in the games, and average accumulated fitnesses are given by the values in the vertical axis. For each fitness value, the related 95% confidence interval is also reported by means of a vertical segment limited above and below by small horizontal bars.

As already observed in Sect. 5.1 (Table 2a, b), the red and purple lines, corresponding respectively to nonlinear regret minimizers and expected utility maximizers, are initially indistinguishable from each other in both graphs of Fig. 2 at $n = 2$, while the blue line corresponding to linear regret minimization lies above them. It is possible to appreciate from the picture, however, that the difference in fitness between linear regret minimization and expected utility maximization is never statistically significant.

As for the case of $m = 2$ in general, Fig. 2 shows also that maximinimizers are always the worst off for all n , even though the difference in fitness with respect to the other criteria is statistically significant solely in games with two actions. Moreover, expected utility maximization becomes superior to the other criteria in games with five or more actions for $m = 2$, although the difference is never statistically significant.

Figure 3a depicts a representative picture for the case of $m = 3$, i.e., $B = \{p_1, p_2, p_3\}$.

Fig. 2 First, the average fitnesses, as functions of the number n of actions in the games, accumulated by the four criteria when paired with linear regret minimization for $m = 2$ and $\bar{v} = 50$; second, the same average accumulated fitnesses when the four criteria are paired with nonlinear regret minimization







Fig. 3 **a** The average fitnesses, as functions of the number n of actions in the games, of the four criteria when paired with linear regret minimization for $m = 3$ and $\bar{v} = 50$. **b** The average fitnesses, as functions of the number n of actions in the games, of the four criteria when paired with linear regret minimization for $m = 7$ and $\bar{v} = 50$



The situation is qualitatively similar to the case $m = 2$, but important statistical differences are also found. First of all, nonlinear regret minimization and utility maximization are no longer equivalent for $n = 2$, and both ways of minimizing regret turn out to have initially higher accumulated fitness than utility maximization. For $n = 3$ instead, linear regret minimization and expected utility maximization perform almost identically, while the accumulated fitness achieved by nonlinear regret minimizers is slightly lower. When the number n of actions in the games increases to five and more, expected utility maximizers are the best off. Importantly, the difference in fitness between expected utility maximization and the other criteria is always statistically significant in games with seven or more actions (Fig. 3a). It is also noticeable that the performance by utility maximizers is always statistically distinguishable from that by maximinimizers for all parameters n , while the difference between maximin and the two types of regret minimization becomes statistically insignificant as the number of actions increases.

These tendencies are even more pronounced as the number m of points in the belief set goes up. Figure 3b shows the situation for $m = 7$. There, the small initial advantage of linear regret minimization is already overtaken by expected utility maximizers at $n = 3$ (although the confidence intervals of EU , LRm , and NRm are still overlapping for games with only two or three actions, $n = 2, 3$). For games with five or more actions, the gap between utility maximization and the others becomes even bigger than in previous cases ($m = 2, 3$), and always statistically significant.

In general, as m increases, the accumulated fitnesses of every criterion increase as well, suggesting an evolutionary benefit in having a larger belief set.⁸ Further, the expected utility maximizers seem to profit more than the others from larger parameters m , in that their fitness peaks at $n = 5$ (rather than at $n = 3$) for $m = 7$. Nonlinear regret minimization, instead, suffers from the combination of high m and high n more than the competitors and is even outperformed by maximin, sometimes in a statistically significant way (see Fig. 3b at $n = 15$ for instance).

Overall, it turns out that expected utility maximization is largely favored in interactive decision problems where the possible alternatives are many. The difference in fitness between utility maximizers and the others becomes always statistically significant when a large number n of actions is combined with a large number m of points in B . With few possible actions, instead, linear regret minimization seems to perform slightly better than the other criteria, but rarely in a statistically significant way. When one considers the rationality of different criteria in ecological terms, therefore, rationality may become a context-dependent concept, which cannot be studied in vacuum and independent of the environment in which decisions take place. This is what emerges from our multigame model for interactive decision making. Next, we analyze the case of individual decision making.

⁸ While it is reasonable to expect that higher values of m correspond in the long run to belief sets with larger area, volume or hypervolume, it is also possible that belief sets with lower m be larger than belief sets with higher m in some cases. Take, for instance, the case of games with 3 actions (the simplex is hence a two-dimensional object): a belief set generated by 4 points ($m = 4$) has *on average* larger area than one generated by 3 points, but the 3-points convex hull may well be larger than the 4-points convex hull in some cases.

6 Individual decisions

A little modification of the previous model allows us to investigate the evolutionary fitness of the criteria introduced above in the context of individual decision making too. In individual decision problems, all the uncertainty about the outcome of an action is caused by the uncertainty about the actual state of the world, as no other player is simultaneously involved in the same decision problem. Let us denote by k the number of possible states of the world, $|S| = k$. We then assume that the realization of a state $s \in S$ obeys a certain probability distribution $p \in \Delta(S)$, extracted at random from a uniform distribution on $\Delta(S)$, which hence determines the actual expected fitness of each possible action in the decision problem at hand. The rest of the model remains unchanged: at each time step, $n \cdot k$ fitness values are randomly drawn, together with a belief set B , to generate the single-agent decision problem.⁹ Therefore, the set $B = \{p_1, \dots, p_m\} \subseteq \Delta(S)$ now encodes the agent's uncertainty about the actual probability p , as in the example with the bag of marbles in Sect. 2. Differently than the game-theoretic framework of Sect. 3, the modeler has thus the possibility to distinguish between two instructive cases and analyze them separately: the case where the actual probability p is in the convex hull of the belief set, $p \in \mathcal{C}(\{p_1, \dots, p_m\})$, termed the case of good beliefs; and the case where $p \in \Delta(S) \setminus \mathcal{C}(\{p_1, \dots, p_m\})$, termed bad beliefs.

6.1 Good beliefs

Figures 4 and 5a picture the average accumulated fitnesses of the criteria under consideration. Similarly to the game-theoretic scenario, expected utility maximization turns out to be the most profitable criterion. In fact, the advantage is even more marked in individual decisions than in games: apart from the case of $n = 2$ and $m = 2$, where EU and NRm are still equivalent, utility maximizers have strictly higher accumulated fitness than the competitors for all other tested parameters. Moreover, this difference in fitness is statistically significant in most of the cases.

A notable difference with respect to the results of Sect. 5 is that the accumulated fitness is here mostly increasing in n , whereas it was mainly decreasing in interactive situations (it was increasing from $n = 2$ to $n = 3$ and decreasing thenceforth). A related difference is that, while the average accumulated fitness in individual decisions is close to that in interactive decisions for $n = 2$, the two substantially diverge as the number of actions increases. Overall, hence, the accumulated fitnesses in individual decision making with good beliefs are higher for all criteria. Notice, however, that only EU keeps improving its fitness as n goes up for all m , while the other criteria experience a small decrease in fitness when a large number n of actions is combined with a large number m of points in B . Moreover, in Figs. 4 and 5a, the average accumulated fitness decreases as m increases, contrary to the interactive setting: in individual decisions with good beliefs, larger belief sets seem to be detrimental.

⁹ The values that we tested also remained the same as in Sect. 3 for all parameters \bar{v} , n and m , while $k \in \{2, 3, 7\}$.

Fig. 4 First, the average fitnesses, as functions of the number n of actions in the individual decisions, accumulated by the four criteria in single-agent decision problems where the agent holds good beliefs for $m = 2$ and $\bar{v} = 50$; second, the same average accumulated fitnesses for $m = 3$ and $\bar{v} = 50$







Fig. 5 a The average fitnesses, as functions of the number n of actions in the individual decisions, accumulated by the four criteria in single-agent decision problems where the agent holds good beliefs for $m = 7$ and $\bar{v} = 50$. **b** The average fitnesses, as functions of the number n of actions in the individual decisions, accumulated by the four criteria in single-agent decision problems where the agent holds bad beliefs for $m = 2$ and $\bar{v} = 50$



(a)



(b)

Nonlinear regret minimization comes out here as the second-best criterion. This is different from the results obtained for interactive decision making in Sect. 5, where linear regret minimization was mostly superior to nonlinear regret minimization. In the next section, the investigation on the case of individual decisions under bad beliefs seems to confirm that both this last observation and the general improvement in accumulated fitness experienced by all criteria depend on the fact that beliefs are often bad in the game-theoretic setting.

6.2 Bad beliefs

Figures 5b and 6 show the situation where $p \in \Delta(S) \setminus \mathcal{C}(\{p_1, \dots, p_m\})$. These graphs and those of Sect. 5 look extremely alike, thus supporting the conclusion that the major differences in fitness between Sects. 5 and 6.1 are due to the agents holding bad rather than good beliefs.

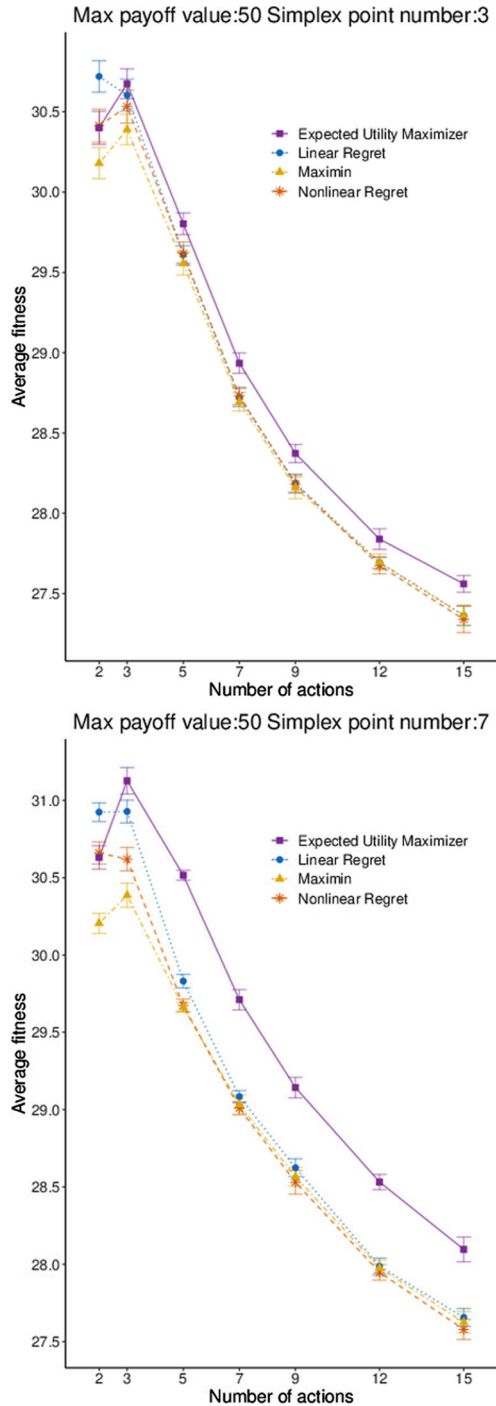
Some observations made in Sect. 5 with regard to the optimality of linear regret minimization should then be reconsidered. In light of the last findings, the cases in favor of the ecological rationality of linear regret minimization stem from situations in which the agents' beliefs are bad. On the one hand, it is hence possible to find some settings where linear regret minimization would be ecologically sustainable. On the other hand, if one enlarges the scope of rationality also to more intentional and introspective acceptations, it would then be difficult to justify the rationale of adopting linear regret minimization when making decisions: one should go for the option that minimizes linear regret only when believing his or her own beliefs bad. Yet, this broader account of rationality may also imply that one updates to new beliefs in order to include the actual probability p whenever recognizing his or her own beliefs as bad, possibly prior to any decision. We leave this last observation for future development, though, in that it implies a broader notion of rationality than the merely ecological one that we investigate in this paper.

Finally, a general remark is in order. Overall, reducing the uncertainty to a Bayesian belief and maximizing expected utility turns out to be evolutionarily beneficial. In this respect, another possible interpretation of the results is that the decision criterion does not really matter when one sticks to Bayesianism, in that all criteria are equivalent under probabilistic beliefs, as we have seen in Sect. 2. Our findings would thus become especially relevant for those situations where the agent is unable to reduce the uncertainty to a probability measure, and has to resort to some criterion fit for unmeasurable uncertainty, such as maximin or regret minimization. The analysis of the accumulated fitnesses in both games and individual decisions then seems to favor the nonlinear version of regret minimization in such situations, while maximin does not receive any convincing evolutionary support from these models.

7 Concluding remarks

Besides the axiomatic representations, investigations about the evolutionary fitness of different decision criteria may prove to be instructive and shed new light on ratio-

Fig. 6 The average fitnesses, as functions of the number n of actions in the individual decisions, accumulated by the four criteria in single-agent decision problems where agents hold bad beliefs for $m = 3$ and $\bar{v} = 50$ (first), and for $m = 7$ and $\bar{v} = 50$ (second)



nal decision making. Rather than considering various decision criteria as alternative philosophical theories, we think of them as high-level strategies that may coexist in the same environment and play off against each other. Allowing different criteria to coexist and interact in the same model may then help to better understand the biological and ecological advantages and disadvantages of each of them. In the multi-game model, representing uncertainty by a single probability measure and maximizing expected utility turns out to be largely advantageous in terms of evolutionary fitness. Furthermore, whenever this reduction is not possible to the agent, regret minimization almost always outperforms the other prominent criterion for choice under unmeasurable uncertainty, i.e., maximin expected utility. In this respect, this paper extends the results in Galeazzi and Franke (2017) by focusing on $n \times n$ games in general and by distinguishing between two types of regret minimization, and it shows that the ecological advantage of nonlinear regret minimization in particular is statistically significant especially when the agent's beliefs are not completely off.

The idea that the concept of rationality hinges on the interplay between the structure of the environment and the computational limitations of the agent can be traced back at least to the work by Herbert Simon (e.g., Simon (1955, 1990)) and is still widely exploited in the research on heuristics by Gerd Gigerenzer and colleagues (e.g., Gigerenzer and Goldstein (1996), Goldstein and Gigerenzer (2002)). Heuristics, however, are normally understood as simple rules of thumb to tackle specific tasks, differently from decision criteria that are supposed to be general principles fit for all sorts of decision problems.

The need of elevating the analysis from expressed behavior to underlying psychological mechanisms has also been advocated in theoretical biology and behavioral ecology (e.g., Fawcett et al. (2013)). In economics, similar thoughts have led evolutionary game theorists to the study of the evolution of preferences, where subjective utility functions rather than simple behaviors become the types under selective pressure (e.g., Robson and Samuelson (2011), Heifetz et al. (2007b), Heifetz et al. (2007a), Dekel et al. (2007), Alger and Weibull (2013)). An orthogonal but related research direction investigates the evolution of learning: keeping utility function and decision criterion fixed, the goal is to understand which learning procedures are more beneficial than others (e.g., Blume and Easley (2006)).

We view the results presented here as complementary to the evolution of preferences and the evolution of learning. A specific utility function, a specific learning rule (which updates the agent's beliefs over time), and a specific decision criterion can be seen as three different individual traits, jointly determining the choices of the agent bearing them. A general model, which we shall leave to future investigation, should hence study the evolutionary selection of all these traits in combination.

Acknowledgements Open Access funding provided by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted

by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A. Appendix

A.1 Equivalence under Measurable Uncertainty

To see the equivalence between linear regret minimization and expected utility maximization under measurable uncertainty, note that from

$$a^* \in \operatorname{argmin}_{a \in A} \max_{p \in B} r_L(a, p)$$

it follows that, for all $a \in A$,

$$\max_{p \in B} E_p \left[\max_{a' \in A} u(c(s, a')) - u(c(s, a^*)) \right] \leq \max_{p \in B} E_p \left[\max_{a' \in A} u(c(s, a')) - u(c(s, a)) \right].$$

Since E_p is a linear operator, it follows from B containing a single measure that

$$E_p \left[\max_{a' \in A} u(c(s, a')) \right] - E_p [u|a^*] \leq E_p \left[\max_{a' \in A} u(c(s, a')) \right] - E_p [u|a]$$

and therefore $a^* \in \operatorname{argmax}_{a \in A} E_p [u|a]$.

Similarly, the equivalence between nonlinear regret minimization and expected utility maximization under measurable uncertainty directly follows from the definition of $r_N(a, p)$. If $B = \{p\}$ and a^* is the action minimizing nonlinear regret, then, for all $a \in A$,

$$\begin{aligned} & \max_{a'} \sum_s p(s) \cdot u(c(s, a')) - \sum_s p(s) \cdot u(c(s, a^*)) \\ & \leq \\ & \max_{a'} \sum_s p(s) \cdot u(c(s, a')) - \sum_s p(s) \cdot u(c(s, a)) \end{aligned}$$

and hence

$$- \sum_{s \in S} p(s) \cdot u(c(s, a^*)) \leq - \sum_{s \in S} p(s) \cdot u(c(s, a)).$$

Equivalence between B and $\mathcal{C}(B)$

We show that the optimal choices for a belief set $B = \{p_1, \dots, p_m\}$ always coincide with the optimal choices for the convex hull of that belief set $\mathcal{C}(B)$ for maximin and the two versions of regret minimization. To see this, consider for instance the game and the projection of the points from the belief set B onto the corresponding expected utilities of I pictured in Fig. 7.




Fig. 7 A symmetric game on the left; on the right, a possible belief set B consisting of five points from the simplex $\Delta(\{I, II, III\})$ and the relative convex hull $\mathcal{C}(B)$. The projections onto the plane representing action I correspond to the expected utilities associated with those five points

Since both the utility function u and the linear regret function r_L are linear in probability, and linear functions on a convex compact set reach minimum and maximum at some extreme point, it follows that interior points of $\mathcal{C}(B)$ play no role in determining the optimal choice for both the maximinimizer and the linear regret minimizer.

The same holds true for nonlinear regret minimization, even though the nonlinear regret function r_N is not linear in probability. Notice that we can rewrite the nonlinear regret $r_N(a_j, p)$ of action $a_j \in A$ given p as

$$\max \{E_p[u|a_1] - E_p[u|a_j], \dots, E_p[u|a_n] - E_p[u|a_j]\}.$$

From the functions $E_p[u|a_i] - E_p[u|a_j]$ being linear, and hence convex, in p for all $i \in \{1, \dots, n\}$, it follows that $r_N(a_j, p)$ is also convex in p , since it is the pointwise maximum function of a set of convex functions. A convex function on a convex compact set $\mathcal{C}(B)$ reaches its maximum at some extreme point. The maximum of $r_N(a_j, p)$ is therefore attained at some point in B .

References

Alger, I., & Weibull, J. W. (2013). Homo moralis: Preference evolution under incomplete information and assortative matching. *Econometrica*, 81(6), 2269–2302.

Allais, M. (1953). Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’école américaine. *Econometrica*, 21(4), 503–546.

Bergstrom, C. T., & Godfrey-Smith, P. (1998). On the evolution of behavioral heterogeneity in individuals and populations. *Biology and Philosophy*, 13(2), 205–231.

Blume, L., & Easley, D. (2006). If you’re so smart, why aren’t you rich? Belief selection in complete and incomplete markets. *Econometrica*, 74(4), 929–966.

Dekel, E., Ely, J. C., & Ylankaya, O. (2007). Evolution of preferences. *The Review of Economic Studies*, 74(3), 685–704.

- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75(4), 643–669.
- Fawcett, T. W., Hamblin, S., & Giraldeaub, L.-A. (2013). Exposing the behavioral gambit: The evolution of learning and decision rules. *Behavioral Ecology*, 24(1), 2–11.
- Galeazzi, P., & Franke, M. (2017). Smart representations: Rationality and evolution in a richer environment. *Philosophy of Science*, 84(3), 544–573.
- Gärdenfors, P. (1979). Forecasts, decisions and uncertain probabilities. *Erkenntnis*, 14(2), 159–181.
- Gärdenfors, P., & Sahlin, N.-E. (1982). Unreliable probabilities, risk taking, and decision making. *Synthese*, 53, 361–386.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650–669.
- Gilboa, I., Postlewaite, A., & Schmeidler, D. (2009). Is it always rational to satisfy savage's axioms? *Economics and Philosophy*, 25(3), 285–296.
- Gilboa, I., Postlewaite, A., & Schmeidler, D. (2012). Rationality of belief or: Why Savage's axioms are neither necessary nor sufficient for rationality. *Synthese*, 187(1), 11–31.
- Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18, 141–153.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1), 75–90.
- Hayashi, T. (2008). Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139(1), 242–268.
- Heifetz, A., Shannon, C., & Spiegel, Y. (2007a). The dynamic evolution of preferences. *Economic Theory*, 32(2), 251–286.
- Heifetz, A., Shannon, C., & Spiegel, Y. (2007b). What to maximize if you must. *Journal of Economic Theory*, 133, 31–57.
- Huttenberger, S. M., & Zollman, K. J. S. (2010). Dynamic stability and basins of attraction in the Sir Philip Sidney game. *Proceedings of the Royal Society of London B: Biological Sciences*, 277(1689), 1915–1922.
- Huttenberger, S. M., & Zollman, K. J. S. (2013). Methodology in biological game theory. *British Journal for the Philosophy of Science*, 64(3), 637–658.
- Klein, D., Marx, J., & Scheller, S. (2018). Rationality in context. *Synthese*, 197, 209–232.
- Levi, I. (1974). On indeterminate probabilities. *The Journal of Philosophy*, 71(13), 391–418.
- Levi, I. (1982). Ignorance, probability and rational choice. *Synthese*, 53(3), 387–417.
- Levi, I. (1985). Imprecision and indeterminacy in probability judgment. *Philosophy of Science*, 52(3), 390–409.
- Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic theory*. Oxford: Oxford University Press.
- Nowak, M. A. (2006). *Evolutionary dynamics: Exploring the equations of life*. Cambridge: Harvard University Press.
- Okasha, S. (2007). Rational choice, risk aversion, and evolution. *Journal of Philosophy*, 104(5), 217–235.
- Robson, A., & Samuelson, L. (2011). The evolutionary foundations of preferences. In J. Benhabib, M. Jackson, & A. Bisin (Eds.), *Handbook of social economics*, 1A. North-Holland: Elsevier.
- Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, 46, 55–67.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41(1), 1–20.
- Skyrms, B. (1996). *Evolution of the social contract*. Cambridge: Cambridge University Press.
- Skyrms, B. (2004). *The stag hunt and the evolution of social structure*. Cambridge: Cambridge University Press.
- Skyrms, B. (2010). *Signals: Evolution, learning, and information*. Oxford: Oxford University Press.
- Stoye, J. (2011). Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146, 2226–2251.
- von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, New Jersey: Princeton University Press.
- Wald, A. (1939). Contributions to the theory of statistical estimation and testing hypotheses. *The Annals of Mathematical Statistics*, 10(4), 299–326.

- Wald, A. (1945). Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, 46, 265–280.
- Zollman, K. J. S. (2008). Explaining fairness in complex environments. *Politics, Philosophy and Economics*, 7(1), 81–97.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.