Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Emergence of Human Decision-Making Behavior Using Evolutionary Selection of an Ecology of Learning Neural Networks

Master's thesis

Anna Schmidt-Rohr
D-INFK, ETH Zürich

Revised version: July 20th, 2021
(Initial version submitted on June 11th, 2021)

Supervisor: Prof. Dr. Didier Sornette
Co-Supervisor: Giuseppe M. Ferro

Chair of Entrepreneurial Risks, D-MTEC, ETH Zürich

Abstract

In this work we aim to use an evolutionary argument to explain experimentally observed decision-making choices that are considered irrational or inconsistent through the lens of classical expected utility theory. Our hypothesis is that human decision-making reflects the complex and uncertain environments our ancestors adapted to, which result in sub-optimal decisions in abstract and oversimplified laboratory setups. To investigate the effect of different evolutionary environments on decision-making behavior, we implement an agent-based model of the evolutionary process and analyse the decision-making behavior of the trained agents. The simulation results confirm the aforementioned hypothesis and indicate that the more ambiguous the evolutionary environment, the worse is the decision-making performance in simpler non-ambiguous environments. We also detect expected utility violations (using the Allais paradox experiment) across a wide range of evolutionary training environments. We find considerable heterogeneity in ambiguity preferences within evolutionary environments.

Acknowledgements

I would like to express my gratitude to my supervisor Professor Didier Sornette for giving me the opportunity to write my master's thesis at the Chair of Entrepreneurial Risks. I admire your broad interest and approach to science which is reflected in the diversity of topics researched in your group. This appreciation for interdisciplinary research enriched my experience working on this project, as I was encouraged to draw inspiration from many fields, such as economics, evolutionary theory and computer science.

I would also like to thank my co-supervisor, PhD candidate Giuseppe M. Ferro, for giving me the freedom to explore my own ideas, but at the same time always giving me constructive and insightful feedback and advice when needed. I very much appreciate the time you invested to discuss my thesis, ranging from high-level philosophical ideas to specific design choices.

Contents

Chapter 1

Introduction

Decision-making is perhaps the most fundamental action affecting every aspect of human life. On an individual level everyone must make all types decisions, small to large and irrelevant to important, every second of one's conscious life. Each and any event one experiences is preceded by some form of decision-making.

Due to the evident importance of decision-making, scientists from various fields (including economics, data science, psychology and biology) have used laboratory experiments to gather empirical data on decisions under uncertainty. When looking at the experimentally observed decision-making choices through the lens of classical expected utility theory (and therefrom derived theories), they are considered irrational or inconsistent. Examples of such consistently observed paradoxes are the Allais paradox and the Ellsberg paradox [1, 2]. Previous work on decision-making can only explain a subset of these paradoxes and fails to take into account how the characteristics of decision-making in a laboratory setting differ from real-world decision-making. The focus of this project is to address these shortcomings by approaching the decision-making process through an evolutionary perspective, by simulating how evolution has shaped human decision-making.

In this work we use an evolutionary argument to justify experimentally observed decision-making paradoxes. Our hypothesis is that humans decision-making behavior reflects the complex and uncertain environments our ancestors adapted to, presumably by developing some form of heuristics, resulting in optimal behavior in such complex real-world scenarios, but sub-optimal decisions in abstract and oversimplified laboratory setups. To simulate the evolutionary process we implement an agent based model, where each agent is modeled as a neural network and a population of agents is trained using a genetic algorithm. We incorporate the relevant aspects of (the evolution of) human decision-making into the model i.e. the complexity and uncertainty of decision environments as well as competition and cooperation mechanisms.

Our model builds on previous exploratory work done at the chair of Entrepreneurial Risk at ETH Zürich [3]. Key differences of our work, are that our model allows for stochastic decision-making and incorporates a cooperation mechanism. Also, our data analysis is independent of other decision-making theories. We do not report if the trained agents behave according to predictions of dominant decision-theories, but directly measure performance, expected utility theory violations and ambiguity preferences. Specifically we investigate the effect of the modeled evolutionary environment in regards to (1) the types of choices (i.e. high/low stake, high/low ambiguity, high/low gains) and (2) cooperation dynamics on the decision-making behavior in simpler environments.

The remainder of this report is organized as follows: First, an overview of preliminary notions of established decision theories and their shortcomings is given in Section 2. Section 3 examines the motivations and approaches to modeling decision-making behavior from an evolutionary perspective. Section 4 describes the implemented agent-based model. Finally, in Section 5 and 6, we present and discuss the results of our simulations. Three appendices complement the main text.

Chapter 2

Established Theories of Decision Theory

Decision theory is the study of (mainly) human choices. This field focuses on *choice under uncertainty*, where an agent is faced with a number of actions, each of which could give rise to more than one possible outcome with different probabilities, not necessarily known to the agent. Many decisions humans make are indeed choices under uncertainty, as the complexity and interactivity of real-world decisions rarely allow for complete certainty.

Note, however, that not all decisions can be modeled as static choices under uncertainty. For example, in inter-temporal choices, different actions lead to outcomes that are realised at different stages over time and complex decisions describe choices that are difficult simply due to their complexity.

2.1 Decisions under Uncertainty

Most models of choice under uncertainty characterize uncertain (i.e risky) prospects in terms of lotteries: a discrete probability distribution on a set of possible outcomes. A decision-maker then needs to choose between multiple lotteries. Often, these outcomes represent monetary payoffs, but this need not be the case. An outcome could also be a sports team winning, a law passing or a student finishing their work in time for a deadline. If $\mathcal{X}$ is the set of outcomes, a lottery $L : \mathcal{X} \to [0, 1]$ with n possible outcomes is written as:

$$L = (x_1, p_1; \ldots; x_n, p_n), \tag{2.1}$$

where the lottery yields outcome x_i with probability p_i. Clearly, for a lottery to be valid $p_1 + p_2 + \cdots + p_n = 1$ must hold.

For instance, the prospect that France and Germany are equally likely to win the FIFA World Cup and one of them will definitely win, can be represented as a lottery $L_0 = ('France\ wins', 1/2; 'Germany\ wins', 1/2)$. If the winning probabilities depend on who is chosen as the new coach for the German national team, there would be multiple lotteries, such as L_0 and $L_1 = ('France\ wins', 2/3; 'Germany\ wins', 1/3)$, depending on who is coaching the team.

It is important to differentiate between two concepts of uncertainty: *risk* and *ambiguity*. Risk refers to decision-making situations, where all potential outcomes and their probabilities of occurring are known to the decision-maker, while ambiguity refers to situations, where either the outcomes and/or their probabilities are not fully known to the decision-maker. Note that the terms - uncertainty, risk and ambiguity - are not used consistently in the decision theory literature: some authors use our definitions given above, while others use the term uncertainty to refer to the concept we defined as ambiguity. The formal definition of a lottery presented above characterizes a risky prospect because both the set of outcomes and their probabilities are given. One can extend the presented lottery definition to characterize decisions under ambiguity. One of the simplest ways to model ambiguity is by using *lower envelope lotteries* [4], which are defined as:

$$L = (x_1, \underline{p}_1; \ldots; x_n, \underline{p}_n), \tag{2.2}$$

where $\underline{p}_i$ specifies lower bounds on probability p_i. Equivalently to Definition 2.1, the lower envelope lottery yields outcome x_i with probability p_i, however, the lottery includes ambiguity because the actual value of p_i is not revealed. Here, a straightforward definition of ambiguity presents itself, namely, the unassigned probability mass $y = 1 - \sum_{i=1}^{n} \underline{p}_i$. As an example, consider an urn with 100 balls where 20 are known to be black, 30 are known to be red and the remaining 50 are either black or red in unknown quantities. With this knowledge, the prospect of drawing a ball from the urn can be modeled by the following lower envelope lottery $L_2 = ('Black\ ball\ drawn', 0.2, 'Red\ ball\ drawn', 0.3)$, with $y = 0.5$. Thus, while the actual probabilities p_i of the lottery L_2 depend on the actual unknown distribution of the red and black balls, with the given information we can bound the actual probabilities as such: $0.2 \leq p_1 \leq 0.7$ and $0.3 \leq p_2 \leq 0.8$, while $p_1 + p_2 = 1$.

2.2 Expected Utility Theory

The first theory of decision-making we will discuss is *expected utility theory* (EUT), which has been the major paradigm in decision-making since the mid 20$^{\text{th}}$ century. The expected utility hypothesis is standard in economic modeling largely because of its simplicity and convenience. However, a growing body of empirical evidence challenging the validity of EUT is beginning to undermine its influence.

2.2.1 Expected Value

We begin by discussing the expected value of lotteries, as this is where the motivation of EUT stems from. Using the lottery notation introduced in 2.1, the expected value $EV(L)$ of a lottery L is given by:

$$EV(L) = \sum_{i=1}^{n} p_i x_i \,. \tag{2.3}$$

The simplest decision-making rule mathematically, would be to select the lottery L that has the highest expected value $EV(L)$. While this rule seems reasonable for decisions that are repeatedly made, this is not necessarily the case for choices with high stake lotteries that are played only once. For example, take a choice between two lotteries:

L_1: a certain outcome of \$1 million

L_2: an uncertain option with an 50% chance of 5 million and 50% chance of receiving nothing

The expected value calculation prescribes that one should choose the second lottery because $EV(L_2) = \$2.5\text{ million} > \$1\text{ million} = EV(L_1)$. In this example, however, one would expect most people to choose the safe first option of receiving \$1 million. Empirical studies confirm that people do not make decisions according to the expected value rule, especially when the outcome values of the lotteries are large (e.g. \$1 million) [5].

There are two intuitively convincing concepts that explain why people do not base their decisions solely on the expected value of prospects:

1. *Risk aversion*: People tend to choose outcomes with low uncertainty over outcomes with high uncertainty, even if the average outcome of the latter is equal to or higher in monetary value than the more certain outcome.
2. *Diminishing marginal utility*: Additional utility, for the moment meant as "pleasure" or "happiness", gained from an increase in consumption, decreases with each subsequent increase in the level of consumption.

In our example, risk aversion can explain why the first lottery L_1 appeals to people. Even though the expected value of L_1 is lower than for L_2, L_1 is definitely more certain. Moreover, diminishing marginal utility explains why receiving \$1 million (compared to receiving nothing) feels like a larger utility gain, than receiving \$5 million (compared to receiving \$1 million), even though in the latter comparison the difference in monetary value is four times larger. Furthermore, diminishing marginal utility can explain the fact that receiving \$1 million does not have the same value to a billionaire and a broke man. These observations led to the conclusion that people view monetary outcomes subjectively. To address this, the concept of utility was introduced into the expected value model to create EUT.

2.2.2 Expected Utility

Individual preferences over lotteries following Expected Utility Theory [6] can be succinctly stated as follows: *expected utility* EU(L):

$$\begin{aligned} L_1 \succeq L_2 \iff EU(L_1) \geq EU(L_2) \\ EU(L) = \sum_{i=1}^{n} p_i U(W_0 + x_i) \end{aligned} \tag{2.4}$$

where $L_1 \succeq L_2$ denotes weak preference for L_1 over L_2. $U(x)$ is the utility function of the decision-maker and W_0 is the initial wealth. Note that the utility levels are computed using the absolute levels of wealth of an outcome $W_0 + x_i$.

John von Neumann and Oskar Morgenstern provided an axiomatic foundation for EUT [6], which defines a *rational* decision-maker, i.e. they showed that if a decision-maker's preferences over lotteries comply to these axioms, there exists a utility function such that Eq. (2.4) holds. For later discussion, we state these axioms below:

1. Preferences for lotteries are *complete*, which means that for any choice between lotteries L_1 and L_2, a decision-maker must either prefer L_1 to L_2 (denoted $L_1 \succ L_2$), $L_2 \succ L_1$ or both are equally attractive (denoted $L_1 \sim L_2$).
2. Preferences for lotteries are *transitive*, which implies that if $L_1 \succsim L_2$ and $L_2 \succsim L_3$ then $L_1 \succsim L_3$.
3. Preferences for lotteries are *continuous*, which implies that if a decision-maker ranks three lotteries L_1, L_2, and L_3, they will be indifferent between the middle-ranked lottery and some probability mixture of the best- and worst-ranked lotteries. Formally, if $L_1, \succsim L_2 \succsim L_3$, then there exists a probability $p \in [0, 1]$, such that $p \cdot L_1 + (1 - p)L_3 \sim L_2$, where the notation on the left side refers to a compound lottery, in which L_1 is received with probability p and L_3 is received with probability $1 - p$.
4. Preferences for lotteries are *independent*, which means that if two lotteries have an identical probability and payoff branch, the levels of this payoff and probability should not affect a decision-makers choice between lotteries. Formally, if $L_1, \succsim L_2$ then for any lottery L_3 and probability $p \in [0, 1]$ it must hold that $p \cdot L_1 + (1 - p)L_3 \succsim p \cdot L_2 + (1 - p)L_3$.

If a decision-maker's preferences can be represented by a utility function $U(x)$, this function is unique up to positive linear transformations, i.e. if the function $U(x)$ represents a decision-maker's risk preferences, then so will $U'(x) = a \cdot U(x) + b$ for any a and b where $a > 0$.

Utility functions can take many forms. For monetary outcomes, one can assume receiving a larger amount yields greater utility and, thus, increasing utility functions are adopted. The curvature of a utility function describes the decision-maker's attitudes towards risk: a concave utility function implies that the decision-maker is risk averse, i.e. a sure amount is always preferred over a risky bet with the same expected value (See Figure 2.1). Reversely, a convex utility function represents a risk-prone decision-maker. The property of diminishing marginal utility corresponds to the mathematical property of concavity, and thus in fact implies risk aversion. An example of a simple utility function of a risk-averse decision-maker is the logarithmic function: $U(X) = \log(X)$.

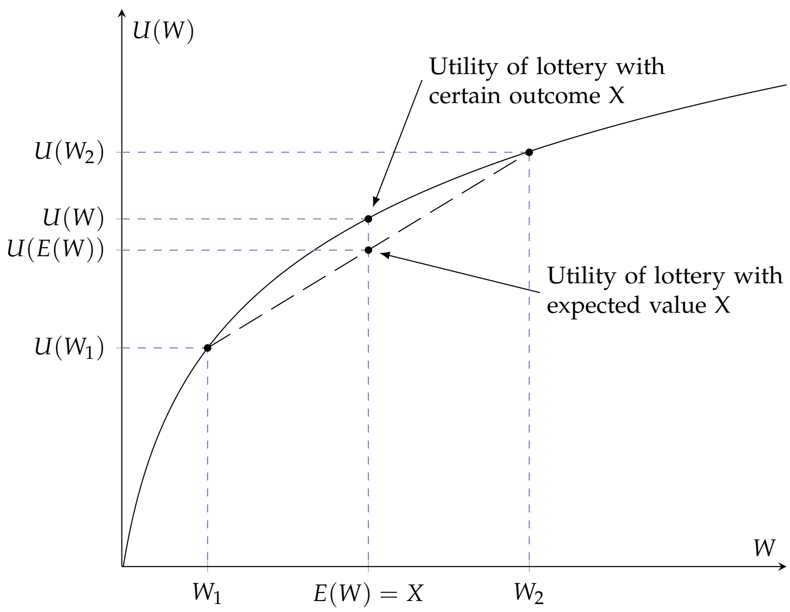

Figure 2.1: A utility function of a risk-averse decision-maker. Specifically, a lottery with payoffs W_1, W_2 and expected value $E(W)$, is displayed. The graph illustrates that the concavity of the utility function implies risk-aversion, because a certain payoff is always be preferred over a risky lottery with the same expected value, i.e. $U(X) > U(E(W) = X)$.

2.2.3 Subjective Expected Utility

Later, EUT was extended to *subjective expected utility* (SEU) to include choices under ambiguity, i.e. choices where there are no objectively known probabilities. Outcomes are assigned subjective probabilities π_i and SEU suggests that the decision-maker should select the lottery L with the highest expected subjective utility $SEU(L)$ [7]. Using the notation of Equation 2.4 we have:

$$SEU(L) = \sum_{i=1}^{n} \pi_i U(W_0 + x_i). \tag{2.5}$$

SEU is also applied to choices under risk, where one assumes that decision-makers weigh outcomes with subjective probabilities, possibly different from objective (exogenously given) ones.

2.3 Paradoxes and Empirical Violations of Expected Utility Theory

To determine if EUT is a theory of practical use in modeling human decision-making behavior, one needs to assess if the axioms EUT is based on (listed in Section 2.2.2) hold sufficiently. In this section, we will present several empirically reported violations of these axioms.

2.3.1 Allais Paradox

The Allais paradox is a classic example that illustrates how people violate the independence axiom of expected utility [1]. The paradox arises when comparing people's choices in two different experiments, each consisting of a choice between two lotteries. The outcomes for the lotteries in each experiment are as follows:

Experiment A:

L_{1a}: 100% chance to win \$1 million.

L_{2a}: 10% chance to win 5\$ million, 89% chance to win \$1 million, 1% chance to win nothing.

Experiment B:

L_{1b}: 11% chance to win \$1 million, 89% chance to win nothing.

L_{2b}: 10% chance to win 5\$ million, 90% chance to win nothing.

Studies have found that when presented with a choice between L_{1a} and L_{2a}, most people choose L_{1a} and when presented with a choice between L_{1b} and L_{2b}, most people choose L_{2b} [8]. However, to be in line with expected utility theory, one person would need to choose either L_{1a} and L_{1b} or L_{2a} and L_{2b}. Recall that, according to the independence axiom of EUT (axiom 4, Section 2.2.2), equal outcomes should cancel out. Concretely, this means one can disregard the 89% common consequence of winning \$1 million and winning nothing in Experiment A and B, respectively. Then the remaining part of the choices in both experiments are the same:

L_1: 11% chance of winning $1 million

L_2: 10% chance to win 5$ million

Thus, not choosing the same lottery (first or second) in both experiments is a direct violation of the EUT independence axiom. This paradox has led to the development of many alternative decision theories, some of them we discuss in Section 2.4.1.

2.3.2 Ellsberg Paradox

The Ellsberg paradox [2] is another paradox showing the descriptive inadequacy of Subjective Expected Utility. Similarly to the Allais paradox experiment, we compare people's decisions in two experiments, each consisting of a choice between two lotteries. Here the lotteries characterize the rewards for possible draws from a urn containing 90 balls, where 30 balls are red and the remaining 60 balls are either black or yellow in unknown proportions. The outcomes for the lotteries in each experiment are as follows:

Experiment A:

L_{1a}: win $100 if you draw a red ball

L_{2a}: win $100 if you draw a black ball

Experiment B:

L_{1b}: win $100 if you draw a red or yellow ball

L_{2b}: win $100 if you draw a black or yellow ball

Studies have found that in Experiment A, people tend to choose the first lottery L_{1a} and in Experiment B, most people choose the second lottery L_{2b} [9]. Subjective Expected Utility theory prescribes that when choosing between these lotteries in mind people behave as if they had a subjective probability of the non-red balls being yellow or black and then compute the expected utility of the two lotteries accordingly (see Section 2.2.3). Having formed a *unique* probabilistic belief over the composition of the urn, one person should choose the same lottery in both experiments (either L_{1a} and L_{2a} or L_{1b} and L_{2b}), depending on if they believe that drawing a red ball is more likely than drawing a black ball. This follows from a the independence axiom (axiom 4, Section 2.2.2), which allows us to disregard the 'common consequence' of winning $100 due to drawing a yellow ball in Experiment B. However, this choice-pattern is not the empirically observed and therefore an assumption of EUT is violated.

The choice-pattern can be explained by *ambiguity aversion*, which is the tendency to avoid options whose outcome probabilities are unknown. Concretely, in our experiment people prefer to choose the lotteries to which they can attach probabilities to their outcomes, i.e. in lottery L_{1a}, the probability of winning \$100 is known to be $\frac{1}{3}$, for the lottery L_{2b} the probability of winning \$100 is known to be $\frac{2}{3}$, while for the less preferred lotteries L_{2a} and L_{1b}, the probability is not known exactly. EUT fails to model this preference for "probabilized uncertainty".

2.4 Beyond Expected Utility Theory

So far we have presented EUT, the dominating normative theory that models the optimal decision-making behavior of rational agents. In this section, we present alternative decision theories, mainly developed to account for the above presented paradoxes (Section 2.3).

Differently from normative theories (which prescribe which choices should be made), descriptive theories aim to explain and predict people's actual decision-making behavior. They do so by condensing empirically found phenomena into simple mechanisms. Their focus is usually not on giving psychological explanations for the mechanisms they find. We present two descriptive theories, *prospect theory* and *rank-dependent utility theory* in Section 2.4.1 and Section 2.4.2, respectively. In contrast, computational theories examine the underlying cognitive processes of decision-making. They construct dynamic decision-making systems by connecting simple components taken from elementary principles of cognition. Although the properties of the components are simple, the emergent behavior of the ensemble system can become complex. We briefly outline the major computational theories in Section 2.4.3. An alternative approach is the so-called *quantum decision theory* (QDT), which is based on the mathematics of Hilbert spaces and describes decisions as an intrinsically stochastic event, in the same spirit as a quantum measurement [10].

2.4.1 Prospect Theory

The most popular descriptive decision-making theory is prospect theory, which was developed by Daniel Kahneman and Amos Tversky in 1979 [11]. Prospect theory splits the decision-making process into two distinct phases: an *editing phase* and an *evaluation phase*. In the editing phase, the decision-maker organizes and reformulates the available options to simplify the choice. This phase attempts to explain all framing effects. Once the choices have been framed for decision, the decision-maker enters the evaluation phase, where they maximize a utility measure based on the potential outcomes and their respective probabilities. Formally, the utility of a lottery in prospect theory is evaluated as:

$$PT(L) = \sum_{i=1}^{n} \pi_i v(x_i), \tag{2.6}$$

where $v(x)$ is the *value function* and $\pi_i = w(p_i)$ are the *decision weights*. Note that this equation is very similar to the SEU equation (Equation 2.5). The value function can be seen as the utility function in SEU: this is merely a semantic change. The decision weights are computed using a *weighting function* $w(p)$ that takes as its argument an objective probability.

Prospect theory focuses on describing characteristics of the probability weighting function $w(p)$ and the value function $v(x)$ that represent observed behavior. We present the properties of the value function $v(x)$ below and in Figure 2.2:

1. *Reference dependence*: People derive value from gains and losses, measured relative to some reference point, and not from absolute levels of wealth. Thus, the argument of $v(x)$ is x_i and not $W_0 + x_i$.

2. *Reflection effect*: People are risk averse in the gain domain and risk-seeking in the loss domain. Thus, $v(x)$ is concave in the gain region and convex in the loss region.

3. *Loss aversion*: People are more sensitive to losses than to gains of the same magnitude. Thus, $v(x)$ is steeper in the loss region than in the gain region.

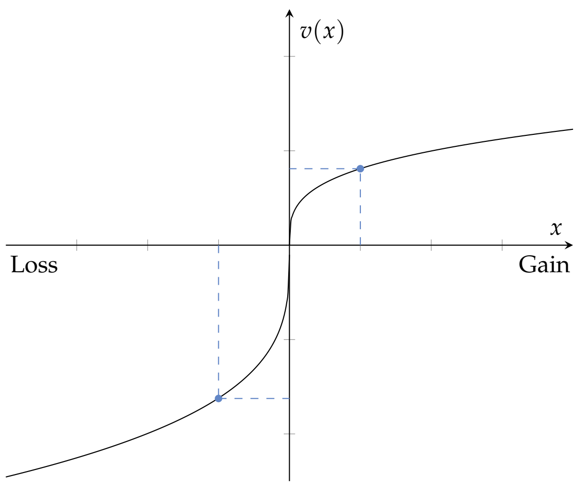

Figure 2.2: A hypothetical value function as prescribed by prospect theory. The function is (1) s-shaped to account for reflection effect and (2) asymmetrical, with a steeper slope for losses than gains, to account for loss-aversion.

Specifically, Kahneman and Tversky, propose a value function of the following form:

$$v(x) = \begin{cases} x^{\alpha}, & \text{if } x \geq 0 \\ -\lambda(-x^{\beta}), & \text{otherwise} \end{cases} \tag{2.7}$$

where their empirical data suggests $\alpha = \beta = 0.88$ and $\lambda = 2.25$ [12].

For the probability weighting function $w(p)$, prospect theory prescribes only one property, namely that people overweight low probabilities and underweight high probabilities (see Figure 2.3). Specifically, Kahneman and Tversky propose a probability weighting function of the following form:

$$w(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}} \ , \tag{2.8}$$

where their data suggests that γ equals 0.61 and 0.69 for gains and losses, respectively [12].

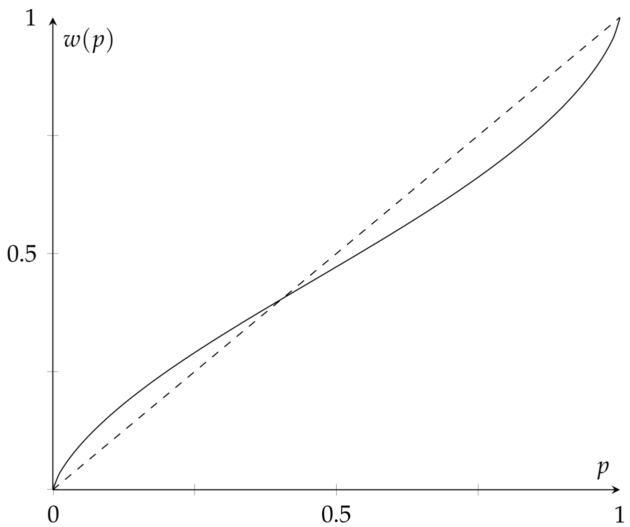

Figure 2.3: A hypothetical probability weighting function, which overweights low probabilities and underweights high probabilities, as prescribed by prospect theory. For reference we plot linear probability weighting using a dotted line.

The combination of the characteristics of the value function and the probability weighting function give rise to the *fourfold pattern of risk attitudes*, which predicts people to be risk-seeking over low-probability gains and high-probability losses and risk-averse over high probability gains low probability losses, as shown in Table 2.1.

	Gains	Losses
Low probability	risk seeking *(hope of gain)*	risk averse *(fear of loss)*
High probability	risk averse *(fear of missing gain)*	risk seeking *(hope to avoid loss)*

Table 2.1: The fourfold pattern of risk preferences prescribed by prospect theory.

Since the introduction of prospect theory, numerous studies have investigated and confirmed the proposed properties of the weighting function $w(p)$ and value function $v(x)$ [13, 14, 15]. Findings are especially robust for probability weighting. However, the existing empirical evidence for the fourfold pattern of risk attitudes is mixed, for a review see [16].

A closer look at the literature reveals a number of gaps and shortcomings in prospect theory. One shortcoming is that the theory is based on the assumption that the decision weights only depend on the objective probability and thus are independent of the value of the outcomes. Kahneman and Tversky themselves suspected that the decision weights may be sensitive to other properties, including the value of the outcomes and framing effects. Numerous studies have confirmed that this is the case [17]. This entanglement of probabilities and outcome values is assessed in rank-dependent theories, discussed in Section 2.4.2, and stochastic representations of decision theory [18]. Another gap in prospect theory concerns how to define the reference point used to measure gains and losses [19]. To define the reference point, one needs to determine what the decision-maker considers to be the 'neutral' outcome of an uncertain prospect. This is not always as simple as looking at a positive versus negative monetary payoff. Take for example the risky prospect of buying a stock. Here, potential definitions of a gain could be that the return on the stock was positive, or the return exceeded the risk-free rate or the return exceeded the value that the investor expected to earn. Another general critique of prospect theory is that it is difficult to consider prospect theory as a unified solution because the parameters (describing the weighting and value functions) are adjustable and adapted ad-hoc to fit the specific data set that is being considered. Being a descriptive theory, prospect theory also does not in any way attempt to explain the reasons behind the behavior it describes.

2.4.2 Rank-Dependent Theories

Rank-dependent theories are based on the assumption that the rank of the outcomes affects decision-making. In prospect theory this means that the probability weighting is affected by the outcome values. We present one rank-dependent theory here, *cumulative prospect theory* (CPT), which is a rank-dependent extension of prospect theory [12].

The general equation used to evaluate lotteries is the same as in original prospect theory (Equation 2.6). However, CPT predicts that not all unlikely events are overweighted equally, but unlikely extreme outcomes are especially overweighted. In general, two outcomes with the same objective probability need not to have the same subjective weight. This is implemented by applying probability weighting to the *cumulative* probability distribution, rather than to individual probabilities. Thus, we have for the decision weights:

$$\pi_i = w\left(\sum_{j=i}^{n} p_i\right) - w\left(\sum_{j=i+1}^{n} p_i\right), \tag{2.9}$$

where $x_1 \leq x_2 \cdots \leq x_n$ and the probability weighting function w is the same as described in Section 2.4.1 for prospect theory.

2.4.3 Computational Theories

The decision theories we have presented so far, are based on the maximization of utility functions. These theories do not claim to describe the actual decision-making process in nature, but rather are used as modeling tools. In fact, for many decisions following the presented theories' axioms would be computationally intractable [20]. For example, the assumption of the completeness of preferences (axiom 1, Section 2.2.2) can require huge cognitive resources (that may exceed peoples cognitive capabilities) if the set of options is very large. Computational theories, in contrast, focus on the underlying processes in nature and try to use biological motivations to justify their models.

Many computational theories are inspired by heuristics, which are simplifying processes that people rely on when making choices. For example, a popular heuristic is the *elimination by aspects* model, where the decision-maker chooses among a given set of prospects during a sequential process. At each stage, the decision-maker eliminates the prospects, which are inadequate in a particular attribute, until only one prospect remains [21]. This strategy is simple because only one aspect is considered at a time. However this also means that an option can be eliminated on the basis of a single attribute even if it might be the best option as a whole. Other computational theories, such as *the adaptive decision maker* or *the adaptive toolbox*, are based on adaptive heuristics. This means that a decision-maker has multiple strategies (not only one like the elimination by aspects theory), which they use adaptively depending on the task, context, and individual difference factors. [22, 23].

Another type of computational models are sequential sampling models, where a decision-maker gradually accumulates noisy information until a threshold of evidence is reached. One such theory, *decision field theory* (DFT), is implemented as follows: (1) The decision-maker makes a momentary evaluation of each prospect at each moment in time and integrates these across time to produce the overall (cumulative) preference state, (2) the cumulative preference states of all options are then compared to determine the valence for each option and (3) the first prospect who's valence reaches the decision threshold is chosen by the decision-maker [24]. Most importantly DFT incorporates the cognitive mechanism of attention. Specifically, the attention an individual devotes to each attribute of a prospect is assumed to fluctuate over time and is modeled as a random walk. Attributes that are in the focus of attention contribute more to the evaluation of a prospect. In the case of a choice between lotteries, attention fluctuates between the lotteries' possible outcomes and the amount of attention depends on the probabilities with which these outcomes occur i.e. outcomes with higher probability receive more attention. DFT is one of the few models that (unlike SEU and prospect theory) can explain violations of stochastic dominance, independence, and stochastic transitivity. Also, DFT can account for speed and accuracy trade-offs by assuming that the decision threshold varies across individuals and decision situations.

A recent model, inspired by DFT, is *stochastic representation decision theory* (SRDT) [18]. Similarly to DFT, in SRDT a stochastic process is assumed to mimic the deliberation process, resulting in a decision once a certain threshold is reached. The novel aspect of SRDT is that it takes into account the empirical observation that the probabilities and outcomes of a prospect are often not separable in the mind of the decision-maker. Instead of disentangling the probabilities and outcomes of a prospect to calculate a subjective expectation (as is done in almost all theories), in SRDT the probabilities and outcomes interact non-trivially as they are combined in a non-symmetric and non-separable way.

Chapter 3

Decision-Making Behavior as an Evolutionary Adaptation

In the previous chapter we outlined the major normative, descriptive and computational decision-making theories. We saw that normative and most descriptive theories dictate that people behave as if they were maximizing some sort of utility, while computational approaches study the emergence of decision-making behavior given the underlying cognitive and motivational processes [25].

Evolutionary theories focus on the underlying evolutionary processes, from which - they argue - the decision-making behavior emerges dynamically. We focus specifically on *evolutionary game theory* [26], a method to study evolutionary theories. In evolutionary games, individuals of a population repeatedly interact and through a modeled evolutionary process the population evolves. In this project, our agent-based model can be seen as an evolutionary game.

In this chapter we begin by briefly describing the process of evolution from an biological perspective in Section 3.1. Then, in Section 3.2, we review arguments about decision-making behavior as an evolutionary adaptation. Specifically, we discuss three aspects of evolution that may be the underlying cause of observed decision-making biases: (1) The maximization of the fitness function, (2) the evolutionary environment and (3) biological limitations. Finally, in Section 3.3, we introduce evolutionary algorithms, which can be used to model the evolutionary process.

3.1 Human Evolution

In nature, populations evolve and adapt to their environment through a process called *natural selection* or *survival of the fittest* [27]. Individuals in a population differ in regards to some heritable traits, where some traits are better suited to the environment than others. These traits are controlled by units called genes. The individuals with the favorable traits are considered to have higher fitness because they have an advantage in survival and reproduction. The fittest individuals survive, reproduce and pass on their fittest genes to their offspring during the sexual recombination process. Given enough generations, a constant environment and enough genetic variation, this process of natural selection then results in population's fitness increasing.

Crucial to the human evolutionary process were two concepts: *competition* and *cooperation* [28]. Competition for food and mates, two necessities to maximize one's chances to reproduce, was the underlying driving force of natural selection. At the same time, the vast majority of our hunter-gatherer ancestors lived in cooperative, small groups comprising several nuclear families [29]. The cooperation took the form of sharing food and child rearing responsibilities, these being the most important and energetically expensive tasks, where cooperation was advantageous for all group members. Food sharing, for example, was advantageous, because meat was often procured in large units but only sporadically obtained by a given hunter. This cooperative behavior gives rise to two separate forms of competition: inter-individual and inter-group competition.

3.2 Optimality and Rationality in the Context of Evolution

Various biases (see Section 2.3) are consistently found in humans and other primates [30]. These biases are considered irrational because they violate principles of economic rationality, i.e. the choice patterns do not respect the theoretic axioms of EUT. However, we argue that these decision-making biases evolved because they follow principles of ecological rationality, i.e the behavior is adapted to the environment in which humans act. We explore three aspects of ecological rationality that potentially explain the biases we find in studies. Firstly, in Section 3.2.1, we consider individuals maximizing their fitness from the perspective of natural selection. Secondly, in Section 3.2.2, we discuss the role of the evolutionary environment and thirdly, in Section 3.2.3, we take into account the computational limitations organisms face.

3.2.1 Fitness Optimization vs Rationality Optimization

To give a flavor of ecological rationality and how it differs from what is generally considered rational, we consider the phenomenon of overconfidence, i.e showing a bias towards overestimating one's capabilities, control over events and one's invulnerability to risk. Intuitively, one would assume that being overconfident and thus having an inaccurate view of the world would lead to faulty judgements and would not be able to compete with decisions based on accurate, unbiased beliefs. However, this is not generally the case. In fact, a robust finding in the psychology of judgment is that people are overconfident [31, 32, 33]. We now present two evolutionary models that have been used to explain the origin of such consistent presence of overconfidence.

Error Management Theory (EMT) [34] convincingly argues that biases, such as overconfidence, are effective decision-making strategies because of the asymmetric costs of false-positive and false-negative errors under uncertainty. For example, one type of overconfidence that has been researched is men's overperception of women's sexual intent [35, 36, 37]. In this scenario, the false-positive error of ancestral men falsely inferring a prospective mate's sexual intent, resulted in the fairly low costs of wasted time and energy on a failed sexual pursuit. In contrast, the false-negative error of men falsely inferring that a woman lacked sexual intent, resulted in the costs of losing a sexual opportunity and hence a reproductive opportunity. Because one primary factor limiting men's reproductive success over evolutionary history was their ability to gain sexual access to fertile women, missing out on such an opportunity would have a very high cost in the currency of natural selection. Thus, following the argumentation of EMT, men's overconfidence in women's sexual interest is not a mistake but adaptive behavior [38].

Another evolutionary argument explaining the presence of overconfidence is altruism: overconfidence encourages exploration, and thus overconfident individuals can provide valuable additional information benefiting the population as a whole [39]. As a result, when groups within a population compete, groups with some overconfident individuals have an evolutionary advantage over groups without such individuals. This argument assumes that there are two intertwined optimizations occurring in the evolutionary process: the fitness of the group and the fitness of the individuals. In this case, there is a trade-off between the fitness of the individuals that are overconfident and putting themselves at risk and the fitness of the benefiting group.

These evolutionary arguments provide mechanisms that can explain a type of rationality - ecological rationality - that differs from what we generally consider rational. The root of the difference is that ecological rationality assumes that:

1. organisms maximize for reproductive fitness (not utility) and
2. natural selection acts as the optimizing selection process.

The most important consequence of maximizing for fitness is that there is no need for the choices to be internally consistent, because the fitness value is influenced only by the outcomes and not the choices themselves. For example, one can argue that efficiency is more crucial to survival than choice consistency [20]. Then inconsistent biases can be optimal in the sense that they produce the best-case behavior from a biological perspective [40]. Another difference between fitness and utility is that the fitness value is a relative measure, i.e depends on the other individuals' performance in the population. In fact, an individual could choose differently between the same options in different points of time (which would traditionally be considered highly inconsistent) and still maximize their fitness value.

The fact that natural selection acts as the optimizing selection process means that, in contrast to economic rationality, ecological rationality is based on the internal decision-making process: Only the decision process is heritable and not the behavior directly, i.e. only the strategy will subject to any consistency conditions [41]. For example, a strategy could be to always choose the first option when making a decision. Here the decision-process is consistent, while the choices would be considered inconsistent from the standard rationality perspective.

3.2.2 Adaptation to the Evolutionary Environment

Another aspect that plays an important role in the evolutionary process is the environment. Because ecological rationality is optimized through natural selection, humans are only trained to act optimally in settings which they encounter naturally. However, a decision-making process that performs well in the environment in which humans evolved, may not perform well in novel environments. In fact, some economically irrational behavior has been justified by humans overgeneralizing rules that are reasonable in many evolutionary contexts [42].

The settings in laboratory experiments that study irrational behavior of humans, are not only novel but mostly very simple i.e. one-shot decisions. In contrast, in real world environments, decisions are complex and may depend on future expectations and/or interact with many other decisions. This simplicity of laboratory experiments could make these environments extremely different to natural environments and thus encourage the findings of economically irrational behavior. Real world environments also include contextual information, such as the energetic state of an organism, that a decision-maker needs to take into account [43]. These state dependencies can give rise to choice patterns that violate transitivity, even though they are maximizing the fitness of the individuals [44].

3.2.3 Biological Computational Limitations

A different biological explanation of human decision-making biases is that they are a result of the computational limitations of the human brain. Even if the human brain had the information-processing capacity to actually calculate various expected fitness consequences and apply calculus to find the optimal solution, it would be impractical apply this process for every decision due to a lack of time. Thus, biases may not be optimal for a decision-maker with perfect knowledge and infinite time, but optimal given the real biological constraints humans face. In fact it may even be ecologically rational for individuals to use heuristics, which occasionally result in sub-optimal choices, because they are less computationally intensive [23].

Biological limitations also imply that there will be some non-zero chance of decision-making errors. Consider a sub-optimal option being added to a decision. Because future errors could result in the sub-optimal option being wrongly chosen, this can affect the expected value of the choice even though the decision-maker does not plan to choose the sub-optimal option. Models that incorporate such error mechanisms have shown that this may result in violations of the transitivity and independence axioms [45].

Not only the human brain has computational limitations, but also the process of natural selection is limited. Natural selection does not create completely new traits out of nothing, but rather makes incremental changes building on previous traits. There may be trade-offs with other behaviors, which biologically cannot be combined. Thus, the starting point and evolutionary history of an individual constrains potential evolutionary trajectories [46, 47]. This is a further constraint on the decision-making optimization process.

3.3 Evolutionary Algorithms

Early analogies between the mechanism of natural selection and learning- and optimization processes led to the development of *evolutionary algorithms* (EAs) [48, 49], algorithms that are inspired by the biological evolutionary process. EAs were developed with two somewhat separate intentions: firstly, to solve specific optimization problems, where the mechanisms of natural adaptation are used solely for performance reasons and evolutionary plausibility is irrelevant. Secondly, EAs were developed to directly study the phenomenon of adaptation as it occurs in nature and by doing so uncovering adaptive behavior. In this report, we focus on the latter purpose of EAs, specifically we will focus on the most popular type of EAs, namely genetic algorithms (GAs) [49].

A GA maintains a population of N individuals, $P(t) = \{A_1, \ldots, A_N\}$ for iteration t. Each individual A_i represents a potential solution to the problem, which is manipulated by the algorithm. Traditionally a binary encoding has been adopted to store the individuals' genetic information [50], although the best representations reflect something about the problem being solved. In our model, for example, we chose as an appropriate encoding an array of numbers, which represent the weights of a neural network.

A GA consists of the following steps: First, an initial population is randomly generated. Then the fitness function $f : A \to \mathbb{R}$, is applied to each of the individuals, in order to measure the quality of the solution encoded by them: the higher the fitness value, the better the solution. This function can have various forms and is derived from the objective of the problem to be solved. Using the individuals' fitness values, a selection scheme is used to choose the individuals that will be the parents of the updated population. We discuss several prominent (probabilistic) selection schemes in Section 3.3.1. Finally, recombination and/or mutation operators are applied to the selected intermediate population to create the updated population:

$$P(t+1) = \mathsf{mutation}(\mathsf{recombination}(P_{\mathsf{selected}}(t))).$$

The basic operation of a GA is illustrated in Figure 3.1.

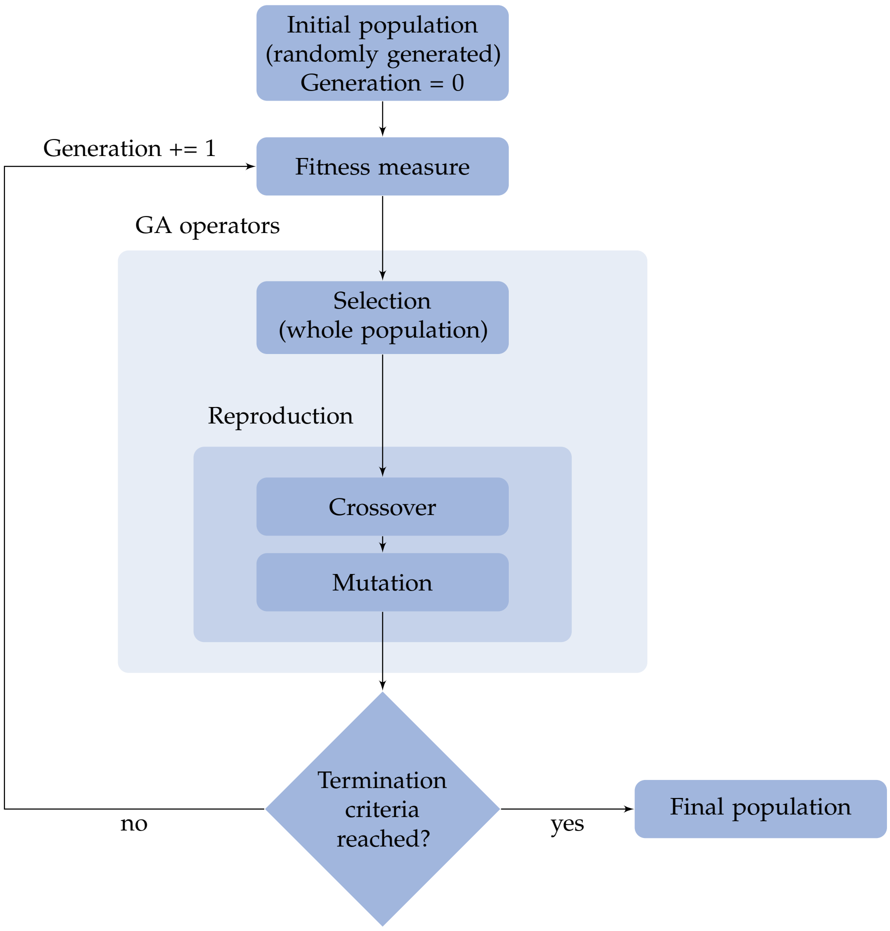

Figure 3.1: Flowchart of the basic operation of a genetic algorithm (GA)

Each iteration of this process is called a generation. The entire set of generations, usually in the hundreds, is called a run. At the end of a run, one expects the population to contain one or more individuals with high fitness. To ensure this is the case, one needs to define an appropriate termination criterion, such as: (1) a individual's genetic material satisfies some minimum criteria, (2) the highest ranking individual's fitness has reached a plateau or (3) a fixed number of generations has been reached. Due to the randomness involved in generating the initial population and in the selection and recombination schemes, two runs with different random-number seeds will generally produce different behavior.

In the following subsections, we will outline common implementations of the core components of GAs.

3.3.1 Selection Schemes

The selection operator chooses some individuals from a current population to create a new population - the mating pool - containing the parents to the individuals of the next population: $P_{mating}(t) = selection\_scheme(P(t))$. Some individuals may appear in population $P_{mating}(t)$ multiple times, while others may not be present at all. To ensure the population size stays constant, we require $|P_{mating}(t)| = |P(t)|$. The purpose of the selection operator is to increase the average quality (i.e. fitness value) of the population, by giving individuals of higher quality a higher probability to be copied into the next generation. The *selection pressure* is the degree to which the individuals of higher quality are favored. This component of the GA can be seen as the algorithmic implementation of the evolutionary mechanism of survival of the fittest.

The selection scheme of a GA plays a crucial rule in determining the convergence rate of a GA: increasing the selection pressure increases the convergence rate of the population. Too high of a selection pressure leads to a loss of population diversity, resulting in an increased chance of the GA prematurely converging to a sub-optimal local maximum. On the other hand, if the selection pressure is too low, the average fitness of the population does not increase sufficiently.

An ideal selection scheme should be easy and efficiently implementable and enable the adjustment of the selection pressure that is needed for the domain. Beyond selection pressure, selection schemes differ in their selected populations' expected (1) average fitness, (2) fitness variance, (3) loss of diversity and (4) selection variance. In the following subsections we present common selection schemes that we implemented in our model (See [51] for more details regarding the characteristics of these schemes).

Proportional Selection

In proportional selection - the original selection method proposed for genetic algorithms - the probability of an individual i to be selected is proportionate to its fitness value [49]. Thus, we have:

$$p_i = \frac{f_i}{\sum_{j=0}^{N} f_j}. \tag{3.1}$$

We provide pseudocode of the proportional selection algorithm in Algorithm 1.

Algorithm 1 Proportional selection

Input: The population $P(t) = \{A_1, \dots, A_N\}$
Output: The population after selection $P(t)' = \{A'_1, \dots, A'_N\}$

proportional($A_1, A_2, \dots, A_N$):
  $cum\_sum_i = 0$
  for $i \leftarrow 0$ **to** N **do**
    $cum\_sum_i \leftarrow cum\_sum_{i-1} + \frac{f_i}{\sum_{j=0}^{N} f_j}$
  end for
  for $i \leftarrow 0$ **to** N **do**
    $r \leftarrow \text{random}([0, cum\_sum_N])$
    $A'_i \leftarrow A_k$, where $cum\_sum_{k-1} \leq r < cum\_sum_k$
  end for
  return $\{A'_1, \dots, A'_N\}$

An advantage of proportional selection is that also individuals with low fitness may survive the selection process. This is beneficial because there is a chance that even individuals with low fitness have some characteristics, which could be successful in the recombination process. A great disadvantage of proportional selection is the fact that this method is not translation invariant [52]. As a result the selection probabilities strongly depend on the scaling of the fitness function. Take a population of ten individuals, with the following fitness values:

Population A:

$$f_{best} = 10 \text{ (one individual)}$$
$$f_{worst} = 1 \text{ (one individual)}$$
$$f_{other} = 5 \text{ (eight individuals)}$$

Population B:

$$f_{best} = 110 \text{ (one individual)}$$
$$f_{worst} = 101 \text{ (one individual)}$$
$$f_{other} = 105 \text{ (eight individuals)}$$

Note that the fitness values of population A added by 100 result in the fitness values of population B, which can be achieved by translating the fitness function by 100. This simple change in the fitness function can have very large effects on the selection probabilities, in fact in population B the selection probabilities of the best and worst individual are almost the same, while for population A the difference is almost 20%:

Population A:

$$p_{best} \approx 19.61\%$$
$$p_{worst} \approx 1.96\%$$
$$p_{other} \approx 9.80\%$$

Population B:

$$p_{best} \approx 10.47\%$$
$$p_{worst} \approx 9.61\%$$
$$p_{other} \approx 9.99\%$$

Proportional selection also inherently favors risk-averse behavior for finite populations. This results from the fact that the frequency of an agent with a specific strategy being selected is a concave function of their payoff. Jensen's inequality then dictates that strategies with lesser variance, which correspond to risk-averse strategies, will have an advantage [53].

Tournament selection

In tournament selection individuals are randomly chosen from the population to compete in a tournament, where τ is the size of the tournament [54]. The individual with the highest fitness of the competitors is copied into the mating pool. This process is repeated N times.

Algorithm 2 Tournament selection

Input: The population $P(t) = \{A_1, \dots, A_N\}$, tournament size $\tau \in \{1, \dots, N\}$
Output: The population after selection $P(t)' = \{A'_1, \dots, A'_N\}$

tournament($\tau, A_1, A_2, \dots, A_N$):
  for $i \leftarrow 0$ **to** N **do**
    $T_i \leftarrow \tau$ individuals sampled uniformly at random from $P(t)$ (with replacement)
  end for
  return $\{A'_1, \dots, A'_N\}$

The tournament selection scheme is translation and scaling invariant i.e. unaffected by the scale of the fitness values. Clearly, increasing the size of the tournaments τ increases the selection pressure.

Truncation Selection

In truncation selection with threshold τ, only the fraction τ best individuals can be selected into the mating pool and they all have the same selection probability [55].

Algorithm 3 Truncation selection

Input: The population $P(t) = \{A_1, \ldots, A_N\}$, truncation threshold $\tau \in [0,1]$
Output: The population after selection $P(t)' = \{A'_1, \ldots, A'_N\}$

truncation($\tau, A_1, A_2, \ldots, A_N$):
  $P^{sort} \leftarrow$ population $\{A_1, A_2, \ldots, A_N\}$ sorted decreasingly by fitness value
  for $i \leftarrow 0$ **to** N **do**
    $r \leftarrow$ random($\{1, 2, \ldots, \tau N\}$)
    $A'_i \leftarrow P^{sort}_r$
  end for
  return $\{A'_1, \ldots, A'_N\}$

Linear Ranking Selection

In linear ranking selection [56], the selection probability is linearly assigned to the individuals according to their rank. For individual i with rank r_i out of N individuals (the higher the value, the higher the rank), we present the following equation to assign the selection probability:

$$p_i = \frac{1}{N}\left(\tau^- + (\tau^+ - \tau^-)\frac{r_i - 1}{N-1}\right), \tag{3.2}$$

where $\frac{\tau^-}{N}$ and $\frac{\tau^+}{N}$ are the probabilities of selecting the individual with the lowest and highest fitness, respectively. To keep the population size constant we set $\tau^+ = 2 - \tau^-$.

Exponential Ranking Selection

Exponential ranking selection follows the same mechanism as linear ranking selection except that the probabilities of the ranked individuals are weighted exponentially (see Algorithm 4 and replace 3.2 with Equation 3.3).

We present the following equation to assign selection probabilities to individuals with rank $r_i \in \{1, \ldots, N\}$:

$$p_i = \frac{\tau^{N-r_i}}{\sum_{j=1}^{N} \tau^{N-j}} , \tag{3.3}$$

where the base of the exponent is set by the parameter $0 < \tau < 1$.

Algorithm 4 Linear Ranking selection

Input: The population $P(t) = \{A_1, \ldots, A_N\}$, reproduction rate of the worst individual $\tau^- \in [0, 1]$
Output: The population after selection $P(t)' = \{A'_1, \ldots, A'_N\}$

linear_ranking($\tau^-, A_1, A_2, \ldots, A_N$):
  $P^{sort} \leftarrow$ population $\{A_1, A_2, \ldots, A_N\}$ sorted increasingly by fitness value
  for $i \leftarrow 0$ **to** N **do**
    $r_i \leftarrow index(P^{sort} == A_i)$
  end for
  $cum\_sum_i = 0$
  for $i \leftarrow 0$ **to** N **do**
    $cum\_sum_i \leftarrow cum\_sum_{i-1} + p_i$ (Equation 3.2)
  end for
  for $i \leftarrow 0$ **to** N **do**
    $r \leftarrow$ random($[0, cum\_sum_N]$)
    $A'_i \leftarrow P^{sort}_k$, where $cum\_sum_{k-1} \leq r < cum\_sum_k$
  end for
  return $\{A'_1, \ldots, A'_N\}$

3.3.2 Reproduction Operators

The next step in the GA is to generate an updated population of individuals (i.e. solutions) from the selected mating-pool. To do this two main genetic operators are used: *crossover* and *mutation*.

Crossover refers to the combining of genetic information of two (or more) parent individuals to generate new offspring. The offspring that is created typically shares many of the characteristics of its parent individuals. Note that some GAs also use asexual reproduction, in which individuals can be carried over to the next generation without applying the crossover operator. The simplest type of crossover (mostly used for a bit array encoding of the genetic information) is *single point crossover*: A point on both parents' arrays is picked randomly as the crossover point, where bits to the right of that point are swapped between the two parents arrays. This results in two offspring, each carrying some genetic information from both parents. In our model, the genetic information of the individuals represent the weights of neural networks and thus different ways of implementing crossover are more applicable.

Mutation is conceptually analogous to biological mutation, altering values in an individuals genetic information according to some specified stochastic process. This process introduces new genetic material into the population and allows the GA to explore the complete solution space.

Chapter 4

The Agent-Based Model (ABM)

To simulate the evolutionary process, we implement an *agent-based model* (ABM), a flexible method commonly used to model competitive and cooperative behaviors [57]. In our ABM, each agent is modeled as a neural network and interacts with an environment by making choices under uncertainty and receiving payoffs according to their choices. Besides interacting with the environment, the agents interact with each other by cooperating and competing for payoff advantages (under a specific set of rules).

First, we describe the decision-making environment and the agent setup in Section 4.1 and Section 4.2, respectively. Then, in Section 4.3, we outline how the agents in the ABM evolve over time through the application of a GA.

4.1 Environment

The agents are confronted with choices, which are represented as a pair of binary lotteries L_1 and L_2, each with two possible outcomes x_1 and x_2. The probabilities of the lotteries' outcomes are not exactly known by the agents, who only receive partial information encoded by three parameters: the minimum probabilities p_{min}, the maximum probability p_{max} and the shape of the sampling distribution t. Thus, a choice under uncertainty between lotteries L_1 and L_2 can be represented as follows:

$$C(L_1, L_2) = (\underbrace{x_1, x_2, p_{min}, p_{max}, t}_{L_1}, \underbrace{x'_1, x'_2, p'_{min}, p'_{max}, t'}_{L_2}). \tag{4.1}$$

Note that the probabilities refer to the first outcome x_1. The probability interval of the second outcome can be inferred, as we only consider binary lotteries.

To explore the effect of varying evolutionary environments on decision-making behavior, we generate different types of lotteries in a controlled manner. We use the term *environment* to refer to the types of the lotteries used in the ABM. We define three characteristics of the lotteries to specify different types of environments, which we describe below:

1. *Outcome variance* $\mathcal{V}$: The outcome variance refers to the difference between the two possible outcomes of a binary lottery $|x_1 - x_2|$. The larger the outcome variance, the more important the probabilities of a lottery are in determining the payoff and thus "value" of a lottery. If $|x_1 - x_2| = 0$, then we have a (degenerate) lottery which yields outcome $x_1 = x_2$ for sure. We run simulations in environments with three outcome variance settings: (1) 'high': $0.4 \leq \mathcal{V} \leq 0.8$, (2) 'low': $0 \leq \mathcal{V} \leq 0.4$, and (3) 'high, low': where high and low $\mathcal{V}$ is equally likely.
2. *Probability range* $\mathcal{A}$: The information given regarding the probability p of an outcome is an interval in which the probability lies, i.e. the probability p_i of outcome x_i is represented by p^i_{max} and p^i_{min}. Lottery types are characterized by the range $p^i_{min} - p^i_{max}$ of the probability intervals they provide. The larger the probability range, the more uncertain the lotteries of that type are. We run simulations in environments with three probability range settings: (1) 'high': $0.4 \leq \mathcal{A} \leq 0.6$, (2) 'low': $0 \leq \mathcal{A} \leq 0.2$, and (3) 'high, low': where high and low $\mathcal{A}$ is equally likely.
3. *Probability distribution* $\mathcal{D}$: To generate the actual probability p used to instantiate the lottery, we sample from a probability distribution, which is non-zero only within the defined probability interval $[p_{min}, p_{max}]$. Each distribution is identifiable by an ordinal value of uncertainty t (uncertainty increases with increasing t). We use three settings, displayed in Figure 4.1, in which we sample p from a *bell-shaped*, *uniform* and *U-shaped* distribution. In the bell-shaped setting, the probability of a lottery can be determined with high likelihood i.e. p will mostly lie close to the median of the interval. In the uniform setting, the neutral setting, any probability in the interval is equally likely. In the U-shaped setting the probability is likely to be near to one of the extreme values of the interval p_{max} and p_{min}. Thus, in this setting the payoff of the lottery is the most difficult to predict.

   We use the the beta distribution $\mathcal{B}$ with different parameters α and β, defined on the interval [0, 1], to sample from each of the three probability distributions. The sampling process is implemented as follows:

   $$X \sim \mathcal{B}(\alpha, \beta)$$
   $$p = p_{min} + X \cdot (p_{max} - p_{min}).$$

   The parameters of the distributions are set as follows:

   (a) *Bell-shaped distribution (t = 0)* : $\alpha = 5, \beta = 5$

(b) *Uniform distribution (t = 0.5)*: $\alpha = 1, \beta = 1$

(c) *U-shaped distribution (t = 1)*: $\alpha = 0.5, \beta = 0.5$

When running simulations, 20% of the lotteries are always in the bell-shaped distribution, to ensure that the agents do not overspecialize and are given data to interpret the input t.

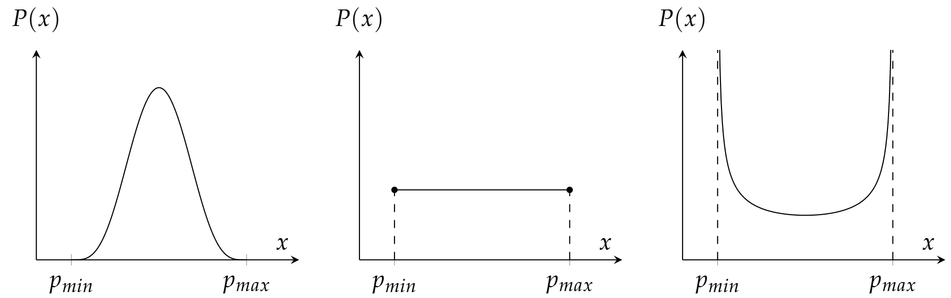

(a) Bell-shaped (b) Uniform (c) U-shaped

Figure 4.1: The probability density functions of the three probability distributions used to generate the instantiated probability p of a lottery given its probability interval $[p_{min}, p_{max}]$. The distributions are ordered from most certain to least certain (left to right). See Section 4.1 for details on the sampling procedure.

4.2 Agents' Neural Network Topology

Each agent of the ABM was modeled using a feed-forward artificial neural network (ANN), displayed in Figure 4.2. Formally, each agent a_i is represented as an array $W_i = (w_1^i, \ldots, w_X^i)$ that defines the weights of the agent's ANN. The used ANNs are simple multilayer perceptrons (MLPs) consisting of three layers with input layer (11 nodes), one hidden layer (10 nodes) and output layer (1 node). Additionally, there are bias inputs to the hidden and output layers. Thus, each ANN consists of 131 weights.The neurons in the hidden and output layer use the nonlinear activation functions sigmoid function and the hyperbolic tangent function, respectively.

The main consideration when setting the ANN structure was to keep the neural network simple, to effectively model the time and computing constraints of the human decision-making process. At the same time, the ANN needed to have a large enough capacity to learn the inputs meaning. To ensure this was the case, we trained the agents using a typical back propagation algorithm [58] to fit a dataset of choices and EV-conform targets. The ANNs with the chosen topology had the capacity to easily fit the expected value function,

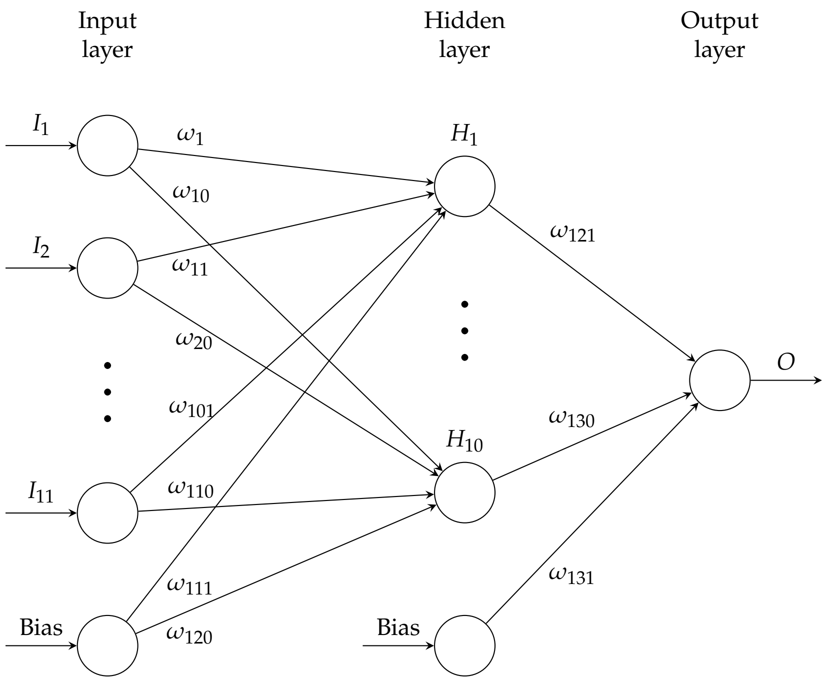

Figure 4.2: The neural network structure of the agents in the ABM. The composition of the input I is described in Equation 4.8.

suggesting that they are able to fit a wide variety of mapping functions.

4.3 The Genetic Algorithm

In this section, we present our GA, which maintains a population of agents, where each agent's ANN exhibits a specific decision-making behavior. As introduced in Section 3.3, GAs generate an initial population and then repeatedly (1) measure the fitness of the individuals in their population and (2) create an updated population based on the measured fitness values using the genetic operators selection, crossover and mutation. This process goes on until some termination criteria are met, as described in Section 3.3. We describe our implementations of these two steps in Section 4.3.1 and Section 4.3.2, respectively. Finally, in Section 4.3.3, we present the additional cooperation and competition mechanism, implemented as an extension of the classical GA. An overview of the implemented GA is provided in Figure 4.3.

4.3.1 Fitness Assessment

In a given generation, each agent is confronted with M choices $C(L_i, L_j)$ between two lotteries. The types of lotteries depend on the environment type the ABM is set in, which we described in Section 4.1. Note that, within one generation, each agent is confronted with the same choices. However, at each generation, the choice set is re-sampled.

The choices of an agent are determined by forward propagating the input data (representing a decision task) through the agent's ANN. The output of the forward propagation is between 0 and 1 and is interpreted as a probability that the first lottery is chosen. Thus, the agents' choices are made probabilistically. The chosen lottery is played and the (stochastic) outcome of the lottery is added to the score of the agent. The agents final score after receiving the payoffs of all their choices, is their fitness value.

4.3.2 Updating the Population

Initialization

The initial population was generated by initializing the agents' ANN weights randomly (u.a.r) to values from a closed interval $[-\alpha_0, +\alpha_0]$, where α_0 was set to 10.

Selection

We use the exponential ranking selection scheme described by Equation 3.3. We set the parameter τ, which determines the exponentiality of the selection method, to 0.99.

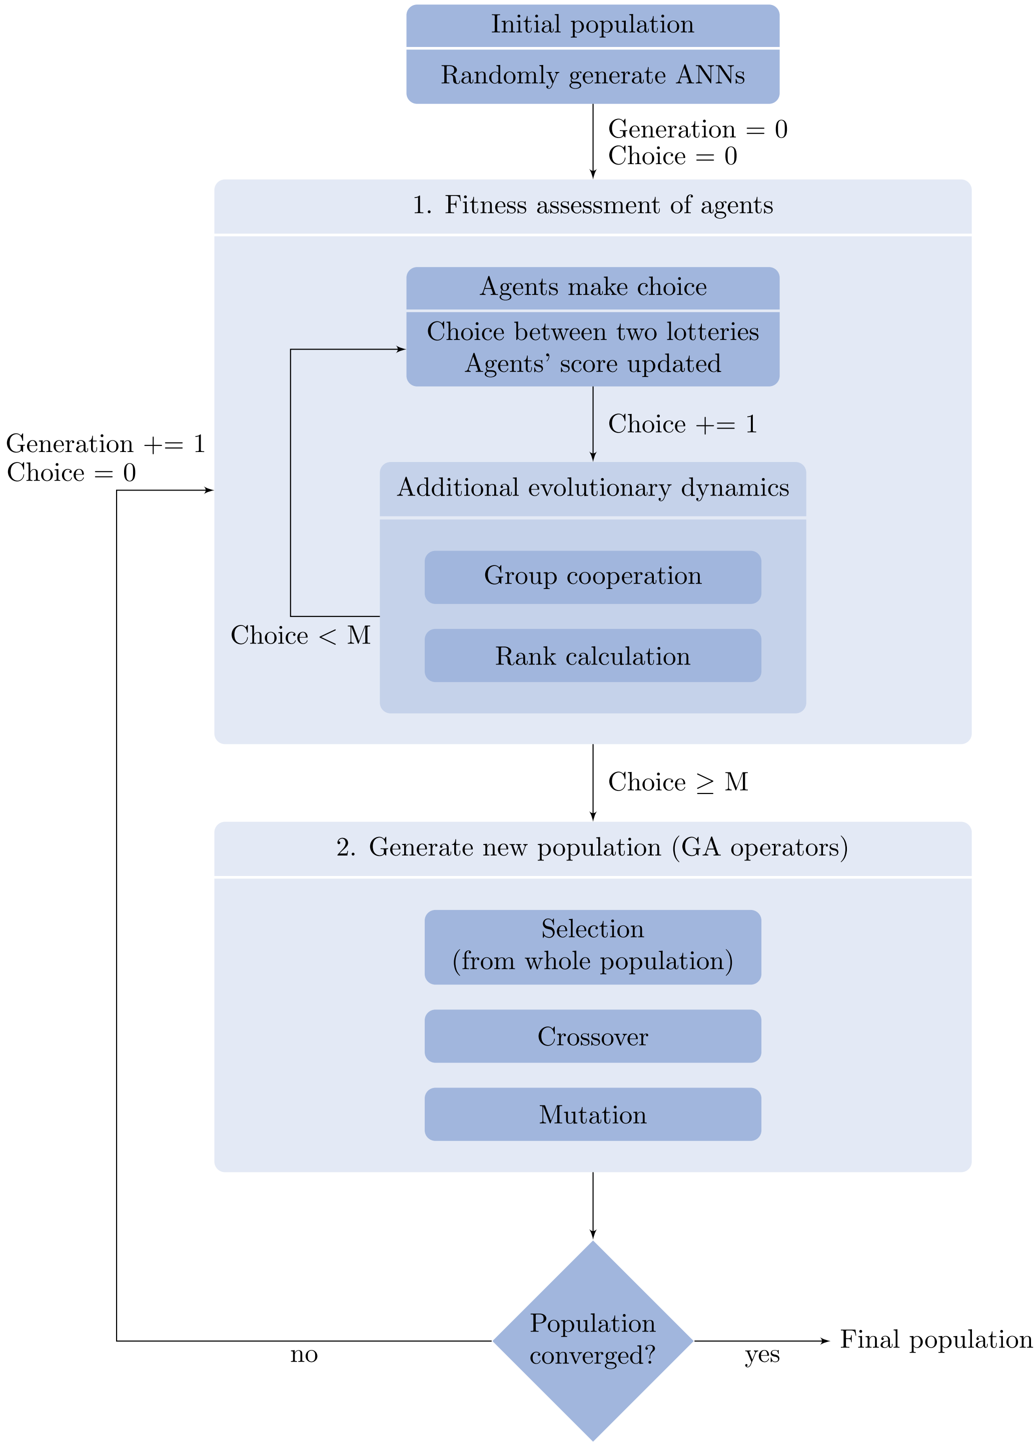

Figure 4.3: Flowchart of the implemented ABM (M represents the number of choices per generation).

Crossover

We use sexual recombination, modeled by an *arithmetical crossover operator*, in our GA. The arithmetical crossover operator combines the ANN's weights of two parent agents by computing two linear combinations of each weight of the two parent agents to create two new agents. Formally, given the two parent agents, a_1 with $W_1 = (w_1^1, \dots, w_X^1)$ and a_2 with $W_2 = (w_1^2, \dots, w_X^2)$, the weight arrays of the offspring agents, o_1 with $W_{o_1} = (w_1^{o_1}, \dots, w_X^{o_1})$ and o_2 with $W_{o_2} = (w_1^{o_2}, \dots, w_X^{o_2})$, are computed according to the following equations:

$$W_{o_1}(i) = \mu w_i^1 + (1-\mu) w_i^2, \tag{4.2}$$

$$W_{o_2}(i) = \mu w_i^2 + (1-\mu) w_i^1, \tag{4.3}$$

where $\mu \in [0.5, 1]$ is a user-specified constant set to 0.8 in our experiments. Because $\mu > 0.5$ the agents a_1 and a_2 are considered the dominant parents of offspring o_1 and o_2, respectively.

Mutation

The implemented mutation operator, a form of dynamic mutation, was inspired by previous work on mutation-based training of neural network weights [59, 60]. Formally, the updating rule for the i-th weight of agent a is:

$$w_i' = w_i + \mathcal{N}(0, \alpha(a)), \tag{4.4}$$

where $\mathcal{N}(0, \alpha(a))$ is the Gaussian perturbation with mean 0 and standard deviation $\alpha(a)$, which is adapted during the process to control the severity of the mutation.

Intuitively, agents with low fitness values should be mutated severely (i.e. have a high value of $\alpha(a)$), while those with high fitness values should only be mutated slightly. To ensure that the mutation rate responds to the rate of progress of the agent, we introduced an adaptive element in the mutation operator by using the ratio of the maximum possible fitness value and the current fitness value of an agent. The function $\alpha(a)$ is defined as follows:

$$\alpha(a) = \alpha_0 \cdot \left(1 - \frac{fitness(a)}{fitness_{max}}\right)^{\beta}, \tag{4.5}$$

where α_0 (the initial value of α) and β are set by user and $fitness_{max}$ is the maximum score that the agent could have achieved in its lifetime (one generation). Thus, if the agents are learning to adapt to their environment, $\alpha(a)$ should be high in the beginning of the GA and should reduce (for all agents) towards convergence of the agents' ANNs. The constant β must be chosen such that the the mutation severity is large enough to allow the GA to sufficiently explore the solution space, while small enough to ensure an acceptable rate of progress. We set β to 2.

Termination

A record is maintained of the number of consecutive generations showing no improvement, i.e the fitness score of the agents stagnates. If this reaches a predetermined number of generations then the process is terminated.

Note that we used this termination criterion for pre-testing and found that in all settings the GA terminated within 500 generations. For consistency purposes we ran all our reported simulations for 500 generations and report the behavior of the population of agents in the final generation.

4.3.3 Cooperative and Competitive Dynamics

Group-based Cooperation

We implement a form of dynamic group-based cooperation, where agents are organized into groups and redistribute a percentage of their score equally between all members of their group. The score of agent a_i in group G_k after the redistribution step, denoted by $s'(a_i)$, is given by the following equations:

$$R_{G_k} = \sum_{a_i \in G_k} c_k s(a_i) \tag{4.6}$$

$$s'(a_i) = s(a_i) + \frac{R_{G_k}}{|G_k|}, \qquad \text{for } \forall a_i \in G_k \tag{4.7}$$

where c_k and $|G_k|$ represent the redistribution rate and number of agents of group G_k, respectively. R_{G_k} is the amount of redistributed wealth within group G_k.

The number of groups and their respective redistribution percentages are set when initializing the algorithm. The agents in the initial population are randomly assigned to a group. When a new population is generated the agents' group memberships are assigned to the group of their dominant parent, i.e the parent they inherited the most genes (Definition in Section 4.3.2). The redistribution percentages of each group stay constant throughout the complete run of the GA. Thus, no new groups with different redistribution rates can emerge, but groups that perform poorly may die out.

This redistribution mechanism was inspired by the famous public goods game [61], where (1) subjects (secretly) choose how many of their private tokens to contribute to a public pot, (2) the tokens in this pot are multiplied by a factor (greater than one and less than the number of players, N) and finally (3) this 'public good' payoff is evenly divided among the players. Each subject also keeps the tokens they do not contribute. In the classic version of the public goods game, players can (and in fact have an incentive to) free-ride off of other players who are contributing to the common pool. This is not the case in our setup, because once an agent is part of a group the amount they contribute to the public pot is predetermined. Our implementation can be seen as a special version of the public goods game, where (1) the participants' contributions to the public pot are distributed asymmetrically among a group, i.e. a percentage of their wealth (score) is contributed, so successful agents contribute more than less successful agents and (2) the multiplication factor of the public good is 1.

It is also important to differentiate the implemented group mechanism to island-based or multi-population GAs that are primarily used to avoid premature convergence in classical GAs [62, 63]. In multi-population GAs sub-populations evolve in (semi-) isolation for generations because the migration rate of agents between the different sub-populations is restricted. In our implementation, the migration rate is not restricted, i.e. each agent in the selection pool regardless of group membership is equally likely to be paired up to reproduce. Also, the selection mechanism in our model is global, i.e. agents compete to be selected with all agents within and outside of their group. We do not aim to increase diversity, but use this group mechanism to best approximate the evolutionary process where individuals lived in small cooperative groups. Island-based approaches have been used with the same intent in previous work [64]. One already mentioned difference is that we make a different design choice regarding migration rates, which seems to be scientifically sound as recent studies of current hunter-gatherers show that group memberships were relatively fluid, i.e. new members joined groups continuously [29]. Most importantly, however, our model also incorporates a cooperation mechanism, which is crucial as cooperation was a key component of human evolution (as discussed in Section 3.1). Especially novel about our model is that it allows dominant cooperation rates to emerge depending on the performance of the agents in the groups. Thus, we can explore which cooperation rates are evolutionary superior and to what decision-making behavior they lead.

Rank Awareness

We implement a direct feedback loop to the agents about their current performance. Specifically, the agents' ranks are recomputed after each choice (after redistribution) and the rank is communicated to the agent by adding it to the next input. Thus, the input to agent a_i's ANN is given by the choice between two lotteries $C(L_1, L_2)$ and the agent's current rank $\mathbf{r_i}$:

$$I(L_1, L_2) = (C(L_1, L_2), \mathbf{r_i}) = (\underbrace{x_1, x_2, p_{min}, p_{max}, t}_{L_1}, \underbrace{x'_1, x'_2, p'_{min}, p'_{max}, t'}_{L_2}, \mathbf{r_i}), \tag{4.8}$$

where r_i is scaled between 0 and 1.

Feeding this contextual information to the agents adds real-world complexity to the GA and can be interpreted as the energetic state of an organism or the stakes of the choice. As discussed in Section 3.2.2, such contextual information affects decision-making and has been suggested as a cause of paradoxical decision-making behavior in simpler environments. This added awareness mechanism allows the ANNs to differentiate their decision-making behavior depending on the contextual information, possibly leading to the emergence of paradoxical decision-making behavior.

4.4 Setting the ABM Parameters

In general, our goal of using a ABM is to allow for the emergence of decision-making behavior and not to train the agents to behave in line with a predetermined strategy. Thus, there is no obvious performance indicator and it is difficult to define a performance metric to perform parameter optimization.

When setting the parameters of the ABM we focused on two criteria, which show that the agents are following a reasonable strategy, but do not impose decision-making biases. Specifically, we aimed for the agents to reach a sufficient level of *environment adaptation* and *order consistency*. We measured the adaptation to the environment by computing the average payoff received on a large dataset of choices of their environment. The order consistency was measured by computing the consistency between two inputs where the lottery order is reversed, i.e $I(L_1, L_2)$ and $I(L_2, L_1)$. Clearly a high level of order consistency is desirable, as it indicates that the agents have learnt to recognize the meaning of the input data.

For completeness, we list some observations we made while determining the parameter settings:

1. *Lottery similarity*: We found that for environments where most lotteries had very similar expected payoff, the agents very quickly converge to always preferring the first (or always preferring the second) lottery presented in the choice, indicating that they were unable to learn a strategy that performed better than random choices.

2. *Selection sensitivity*: We found that the overall population development was very sensitive to the selection pressure. Too high of a selection pressure resulted in premature convergence and almost no improvement in performance over random strategies. Too low of a selection sensitivity often resulted in the best agents dying out early on.

3. *Number of choices*: When setting the number of choices per iteration to very high values, the agents tended to develop towards choosing the lotteries with the highest expected value. However, from an evolutionary standpoint this set-up is problematic, as humans do not just make a large amount of low stake choices in their lifetime.

Chapter 5

Data Analysis Methods

We use four separate data analysis methods to assess the trained agents' decision-making behavior.

Our hypothesis is that humans display seemingly irrational behaviors in simple decision-making tasks tested in laboratories, due to the fact they have adapted their behavior to the very different complex and uncertain environments faced during evolution. We address this hypothesis by investigating the decision-making behavior of the agents in environments that are more simple (less ambiguous) than the ones in which they were trained in throughout our data analysis.

In the first and second approach, described in Section 5.1 and 5.2, we measure the degree of stochastic behavior and the general decision-making performance of the agents, respectively. The other two methods investigate more specific patterns of decision-making behavior to enable us to gain a better understanding of the strategies the agents learn. We directly measure the agents' choice-patterns in the Allais paradox and ambiguity preferences (explaining the Ellsberg paradox) using methods described in Section 5.3 and 5.4, respectively.

5.1 Measuring Stochastic Behavior

The decision-making process of the agents is modeled through an ANN and allows for stochastic behavior. The output $O_{L_1,L_2} \in [0,1]$ of the ANN for a choice between $L1$ and L_2, is interpreted as the probability that the first lottery is chosen. We investigate the effect of differing training environments on the extent of stochastic behavior the agents display.

We define the stochasticity of a choice with output O_{L_1,L_2} as follows:

$$\mathsf{stoch}(O_{L_1,L_2}) = \begin{cases} 1 - O_{L_1,L_2} & \text{for } O_{L_1,L_2} \geq 0.5 \\ O_{L_1,L_2} & \text{otherwise} \end{cases} \tag{5.1}$$

Thus, if the choice is deterministic i.e. the choice probability is 0 or 1, the stochasticity is 0. If the choice is maximally stochastic i.e. L_1 is equally likely to be chosen as L_2, the stochasticity is 0.5.

We calculate the average stochasticity of an agent on a set of M choices $C = \{(L_1^1, L_2^1), ..., (L_1^M, L_2^M)\}$ as follows:

$$\overline{\mathsf{stoch}}(C) = \frac{1}{M} \sum_{i=1}^{M} \mathsf{stoch}(O_{L_1^i, L_2^i}) \tag{5.2}$$

The average stochasticity of the population is simply the average stochasticity of the agents in the population.

5.1.1 Fuzzy Preference

It is important to note that during the other data analysis methods we do not sample the choice probability for each decision, but use the following procedure, which also eliminates ordering effects: When determining an agent's preference between two lotteries L_1 and L_2, we present the agent with two choices $C(L_1, L_2)$ and $C(L_2, L_1)$ (see Equation 4.8). The agents' choices may not be consistent when the choice is presented in the two orders i.e. $O_{L_1,L_2} \neq 1 - O_{L_2,L_1}$. Thus, we define the concept of *fuzzy preference*, where the preference is determined as the decision made with higher probability. Specifically we have:

$$L_1 \succeq L_2 \iff O_{L_1,L_2} \geq O_{L_2,L_1} \ , \tag{5.3}$$

where $\succeq$ represents the fuzzy preference relation. We then interpret the choice probability of choosing L_1 as $\frac{O_{L_1,L_2}}{O_{L_1,L_2}+O_{L_2,L_1}}$.

5.2 Measuring Environmental Fitness

This method aims to investigate if there are large discrepancies in the performance of the agents trained in different environments. Specifically, we aim to study how the complexity of the environments and the GA settings affect the performance of the agents when making simpler decisions. To do so, we compare the performance of agents on two datasets, each data point being a binary decision task. Since these datasets represent the oversimplified laboratory setups, they contain lotteries with no ambiguity, i.e. the interval $[p_{min}, p_{max}]$ the outcome probability p lies in shrinks to a point, $p_{min} = p_{max}$. Consequently, the probability p is known with certainty. We randomly generate two datasets, one with choices where the difference between the expected values of the two lotteries are high (≥ 0.4) and one where the competing lotteries may have similar expected values (≥ 0.1). Thus, we can compare the performance between two levels of simplicity. The performance is computed using the payoff of the chosen lotteries i.e. we compute the average performance of an agent on a set of M choices $C = \{(L_1^1, L_2^1), ..., (L_1^M, L_2^M)\}$ as follows:

$$L_{\mathsf{chosen}}(L_1, L_2) = \begin{cases} L_1, & \text{if } L_1 \succeq L_2 \\ L_2, & \text{otherwise} \end{cases} \tag{5.4}$$

$$\mathsf{performance}(C) = \frac{1}{M} \sum_{i=1}^{M} \mathsf{payoff}(L_{\mathsf{chosen}}(L_1^i, L_2^i)) \tag{5.5}$$

where $\succeq$ refers to the fuzzy preference relation. The average stochasticity of the population is simply the average stochasticity of the agents in the population.

5.3 Detecting Expected Utility Theory Violations

To detect violations of EUT we consider the common ratio version of the Allais Paradox experiment, which involves only two-outcome lotteries. The experiment aims to determine how certain outcomes are evaluated relative to outcomes, which are merely probable. The experiment is set up as a pair of choices (A and B), each consisting of a choice between two binary lotteries. The first choice, A, is between a (near) certain prospect and a risky prospect. The second choice, B, consists of the same lotteries as in A, except that the probabilities for the higher outcomes are scaled by the same common ratio i.e. mixed with a common lottery. The general setup can be summarized as follows:

$$\begin{array}{llll} A & {}^*L_{1a} \text{ (a (near) certain prospect)} & vs & L_{2a} \text{ (a risky prospect)} \\ B & L_{1b} = \lambda L_{1a} + (1-\lambda)L_C & vs & {}^*L_{2b} = \lambda L_{2a} + (1-\lambda)L_C \end{array} \tag{5.6}$$

where the mixing factor λ determines to what extent the lotteries are scaled by the common lottery L_C. The choice pattern indicated by the asterisks is referred to as the common ratio effect [11]. Both the common ratio and the reverse common ratio effect have been observed in empirical studies [65] and violate of the independence axiom of EUT, which prescribes one of the following preference pairs: $[L_{1a} \succ L_{2a}$ and $L_{1b} \succ L_{2b}]$ or $[L_{1a} \prec L_{2a}$ and $L_{1b} \prec L_{2b}]$.

For concreteness, consider the following example:

A	$^*L_{1a} =$ 100% chance of 60	vs	$L_{2a} =$ 80% chance of 75	
			20% chance of 0	
B	$L_{1b} =$ 25% chance of 60	vs	$^*L_{2b} =$ 20% chance of 75	
	75% chance of 0		80% chance of 0	

where we have a mixing factor $\lambda = 0.25$ and can express L_{1b} and L_{2b} as follows: $L_{1b} = (25\%$ chance of $L_{1a}, 75\%$ chance of $0)$ and $L_{2b} = (25\%$ chance of $L_{2a}, 75\%$ chance of $0)$. The empirically observed choice pattern indicated by the asterisks, violates the independence axiom.

In Tables 5.1 and 5.2 we list each of the pairs of choices we use to investigate the common ratio effect. We test the robustness of the results in three dimensions: First, we vary how large the gains are, second we vary the mixing factor λ and third we measure the results for certain and near-certain lotteries.

Pairs (at certainty)		Option 1		Option 2
Large gains				
$\lambda = 0.6$	A	100% chance of 60	vs	80% chance of 75 20% chance of 10
	B	60% chance of 60 40% chance of 10	vs	48% chance of 75 52% chance of 10
$\lambda = 0.3$	A	100% chance of 60	vs	80% chance of 75 20% chance of 10
	B	30% chance of 60 70% chance of 10	vs	24% chance of 75 76% chance of 10
$\lambda = 0.1$	A	100% chance of 60	vs	80% chance of 75 20% chance of 10
	B	10% chance of 60 90% chance of 10	vs	8% chance of 75 92% chance of 10
Small gains				
$\lambda = 0.6$	A	100% chance of 20	vs	80% chance of 24 20% chance of 5
	B	60% chance of 20 40% chance of 5	vs	48% chance of 24 52% chance of 5
$\lambda = 0.3$	A	100% chance of 20	vs	80% chance of 24 20% chance of 5
	B	30% chance of 20 70% chance of 5	vs	24% chance of 24 76% chance of 5
$\lambda = 0.1$	A	100% chance of 20	vs	80% chance of 24 20% chance of 5
	B	10% chance of 20 90% chance of 5	vs	8% chance of 24 92% chance of 5

Table 5.1: Pairs of choices used to investigate the common ratio effect, comparing the evaluation of *certain* outcomes relative to outcomes which are merely probable. (See Table 5.2 for the choices dealing with near-certainty.) Experiment A is a choice between a *certain* lottery and a binary lottery with a higher EV. Experiment B is the same as A, except that the lotteries are mixed with a common lottery L_C, a certain prospect with outcome 5 and 10 in the small gain and large gain setting. λ refers to the mixing factor defined in Equation 5.6. *Gains were multiplied by 100 for readability.*

Pairs (near-certainty)		**Option 1**		**Option 2**
Large gains				
$\lambda = 0.6$	A	95% chance of 60 5% chance of 10	vs	80% chance of 69 20% chance of 10
	B	57% chance of 60 43% chance of 10	vs	48% chance of 69 52% chance of 10
$\lambda = 0.3$	A	95% chance of 60 5% chance of 10	vs	80% chance of 69 20% chance of 10
	B	28% chance of 60 72% chance of 10	vs	24% chance of 69 76% chance of 10
$\lambda = 0.1$	A	95% chance of 60 5% chance of 10	vs	80% chance of 69 20% chance of 10
	B	10% chance of 60 90% chance of 10	vs	8% chance of 69 92% chance of 10
Small gains				
$\lambda = 0.6$	A	95% chance of 20 5% chance of 5	vs	80% chance of 24 20% chance of 5
	B	57% chance of 20 43% chance of 5	vs	48% chance of 24 52% chance of 5
$\lambda = 0.3$	A	95% chance of 20 5% chance of 5	vs	80% chance of 24 20% chance of 5
	B	28% chance of 20 72% chance of 5	vs	24% chance of 24 76% chance of 5
$\lambda = 0.1$	A	95% chance of 20 5% chance of 5	vs	80% chance of 24 20% chance of 5
	B	10% chance of 20 90% chance of 5	vs	8% chance of 24 92% chance of 5

Table 5.2: Pairs of choices used to investigate the common ratio effect, comparing the evaluation of *near-certain* outcomes relative to outcomes which are merely probable. Experiment A is a choice between a *near-certain* lottery and a risky lottery with equal EV. Experiment B is the same as A, except that the lotteries are mixed with a common lottery L_C, a certain prospect with outcome 5 and 10 in the small gain and large gain setting. λ refers to the mixing factor defined in Equation 5.6. *Gains were multiplied by 100 for readability.*

5.4 Measuring Ambiguity Preferences

To provide intuitive representations of the ambiguity preferences of the agents trained by our model, we use the recently developed *ambiguity triangle* [4].

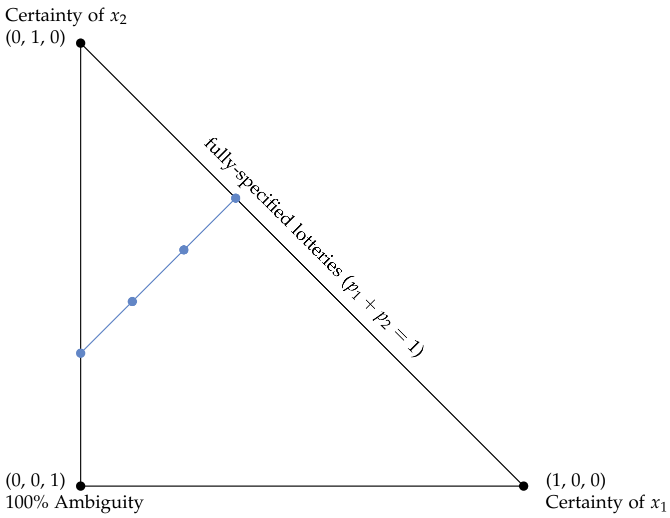

Figure 5.1: Ambiguity triangle depicting lower envelope lotteries that have two possible outcomes x_1 and x_2, where $x_1 > x_2$.

The ambiguity triangle is based on the concept of lower envelope lotteries, defined in Equation 2.2. Recall that such lower envelope lotteries specify the lower bounds on probabilities, $\underline{p}_1, \ldots, \underline{p}_k$, for a set of outcomes $x_1, \ldots, x_n$ and the ambiguity of the lottery is defined by the amount of "unassigned" probability mass $y = 1 - \sum_{i=1}^{n} \underline{p}_i$. The binary lotteries used in our model, which are defined by the lower and upper bound of the probability of x_1 (p_{min}, p_{max}), can easily be transformed into binary lower envelope lotteries by setting $\underline{p}_1 = p_{min}$ and $\underline{p}_2 = 1 - p_{max}$. When presenting the agents with such lotteries we set the lottery type t (see Section 4.1) to the neural 'uniform' distribution setting.

All possible lotteries in the set of binary lower envelope lotteries with outcomes x_1 and x_2 (with $x_1 > x_2$) can be defined by $(\underline{p}_1, \underline{p}_2, y)$ and represented graphically in one ambiguity triangle. Each point within the triangle represents one lottery, where the coordinates of the point are $(\underline{p}_1, \underline{p}_2)$. The vertices of the ambiguity triangle represent the extreme cases in the set of lotteries: (1,0,0) represents a certain lottery with the preferred outcome x_1, (0,1,0) represents a certain lottery with the non-preferred outcome x_2 and (0,0,1) represents a lottery where nothing is known about the probability distribution over the two outcomes. This corresponds to the probability of outcome x_1 (and therefore of outcome x_2) assuming any value between 0 and 1.

We examine choices between lower envelope lotteries that lie on lines with the same expected value, referred to as *EV-constant lines* (such as the blue line in Figure 5.1). Lotteries on one EV-constant line have the same expected value if one makes no further assumptions about the actual probabilities from their lower bounds (following the principle of insufficient reason [66]) . Specifically, this means that if one reduces the ambiguity of a lottery evenly i.e. when reducing the ambiguity of a lottery by x, the minimum probabilities of both outcomes increase equally by $x/2$ units, then the resulting lottery is on the same line. As an example we consider the lotteries $L_1 : (0, 0.3, 0.7)$ and $L_2 : (0.35, 0.65, 0)$, which are the endpoints of the EV-constant line depicted in Figure 5.1. When taking the lottery L_1 and eliminating all ambiguity without biased assumptions, one obtains $(0 + 0.35, 0.3 + 0.35, 0.7 - 0.7) = (0.35, 0.65, 0) = L_2$. Thus this fully-specified lottery (without ambiguity) is the endpoint of the EV-constant line on the hypotenuse of the ambiguity triangle.

Varying the position of these EV-constant lines allows us to examine the ambiguity preferences of the trained agents. Specifically we illustrate in the next subsection how we can use the ambiguity triangle to assess if the agents behave in line with the typical results of the Ellsberg experiment. Furthermore, we will be able to assess whether the trained agents exhibit constant ambiguity attitudes or show large variation in their ambiguity preferences (depending on the values for $x_1, x_2, \underline{p}_1$ and $\underline{p}_2$).

The agents' (fuzzy) choice-preference are probabilities of choosing a choice option. Because it is extremely unlikely that this choice probability will be exactly 0.5, the agent is never completely indifferent between two choice options. Thus, to determine if the ambiguity preference of an agent is neutral, we define a indifference range $[0.5 - \delta, 0.5 + \delta]$ for choice probabilities: If a choice probability p lies in this range we consider the agents to be indifferent in this choice. If $\delta = 0$, we do not allow indifference and agents will have a preference in every choice.

The Ellsberg Experiment in the Ambiguity Triangle

The Ellsberg paradox, introduced in Section 2.3.2, can be illustrated by preference patterns within the ambiguity triangle [67]. To do so the choices in the Ellsberg experiment must first be transformed into lower envelope lotteries. In the following paragraph we walk through one example of how to construct lower envelope lotteries from the Ellsberg lotteries. We also provide an overview of the Ellsberg experiment described with lower envelope lotteries in Table 5.3 and a graphical illustration using the ambiguity triangle (Figure 5.2).

Consider the first lottery in Experiment B of the Ellsberg experiment denoted as L_{1b}, where the outcome is \$100 if a red or yellow ball is drawn and \$0 otherwise. The number of red balls in the urn is known to be exactly 30 and the number of yellow balls is uncertain lying somewhere between 0 and 60. Thus, the minimum probability of receiving \$100 is $\frac{1}{3}$ while the maximum is 1. Because the maximum probability of receiving \$100 is 1, the minimum probability of receiving \$0 is 0. Thus, defining the lower lottery outcome x_1 as receiving \$100 and the outcome x_2 as receiving \$0, we can denote L_{1b} as a lower envelope lottery $(\frac{1}{3}, 0, \frac{2}{3})$.

Lotteries	**Payoffs of outcomes [$]**			**Lower envelope lotteries**
	red	black	yellow	$(\underline{p}_1, \underline{p}_2, y)$
L_{1a}	100	0	0	$(\frac{1}{3}, \frac{2}{3}, 0)$
L_{2a}	0	100	0	$(0, \frac{1}{3}, \frac{2}{3})$
L_{1b}	100	0	100	$(\frac{1}{3}, 0, \frac{2}{3})$
L_{2b}	0	100	100	$(\frac{2}{3}, \frac{1}{3}, 0)$

Table 5.3: The Ellsberg experiment described with lower envelope lotteries, where outcome x_1 and x_2 is receiving \$100 and \$0 respectively.

Typically in the Ellsberg experiment participants prefer the fully-specified lotteries (circled in Figure 5.2) in both choices. Thus, one can assess if agents behave according to the canonical choice pattern of the Ellsberg experiment by analysing if they prefer lotteries on EV-constant lines that are closer to the hypotenuse of the triangle. This corresponds to making ambiguity-averse choices.

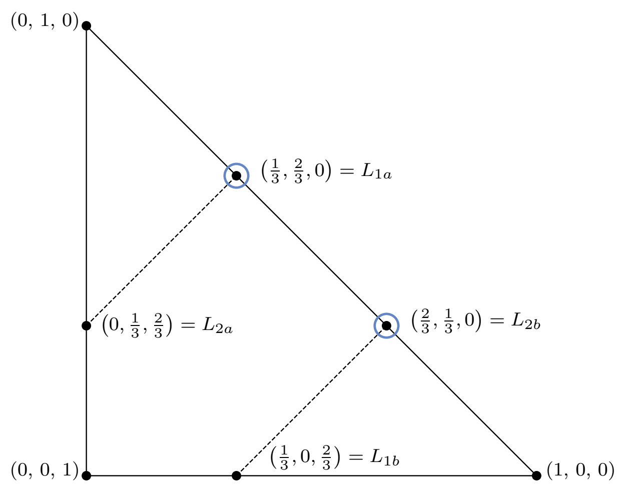

Figure 5.2: Choices in the Ellsberg experiment depicted in the ambiguity triangle. The lotteries lying on the same dashed line make up a choice. The blue circles indicate the empirically observed typical choices.

Chapter 6

Results

Due to the stochastic nature of the ABM, we need to simulate the model multiple times to investigate the expected decision-making behavior. Specifically, we ran 20 repetitions of the ABM with 500 iterations for each training environment. We report the standard deviation between runs when applicable. The model was implemented using Python and all simulations were run on a computing cluster provided by ETH Zürich (Euler).

We structure the results by data analysis method. In Sections 6.1 and 6.2, we present the general results on the stochasticity and performance of the decision-making behavior in simplified environments. Then we present the (detected) EUT violations and ambiguity preferences we found in the trained agent populations in Sections 6.3 and 6.4, respectively. Finally, we present our findings on the cooperation mechanism in Section 6.5.

6.1 Stochastic Behavior

We analyze the stochasticity of the agents preferences in two simplified environments and their respective training environment. Recall that we measure the stochasticity using the agents' choice probabilities, i.e. more certain choices (with choice-probabilities near 0 or 1) correspond to low stochasticity.

When looking at Table 6.1, where we report the results of the stochasticity analysis, it is immediately noticeable that the outcome variance $\mathcal{V}$ of the lotteries in the training environment has by far the greatest effect on the degree of choice stochasticity. Specifically, agents trained in environments with low outcome variance evolve to make highly stochastic choices, i.e. have weaker preferences. This may be seen as paradoxical, as one could argue that lotteries with low outcome variance are less noisy and thus should lead to less stochasticity. However, because the stochasticity is also high in

Training environment			Simplified environment I *(Large EV-diff)*	Simplified environment II *(Small EV-diff)*	Training environment
$\mathcal{V}$	$\mathcal{A}$	$\mathcal{D}$	Stochasticity	Stochasticity	Stochasticity
high	high	U-shaped	0.065	0.073	0.104
high	high,low	U-shaped	0.113	0.128	0.139
high	low	U-shaped	0.083	0.088	0.137
high	high	uniform	0.077	0.097	0.148
high	high,low	uniform	0.062	0.079	0.107
high	low	uniform	0.074	0.091	0.134
high	high	bell-shaped	0.084	0.107	0.130
high	high,low	bell-shaped	0.082	0.083	0.111
high	low	bell-shaped	0.078	0.126	0.186
high,low	high	U-shaped	0.096	0.130	0.128
high,low	high,low	U-shaped	0.094	0.120	0.134
high,low	low	U-shaped	0.090	0.117	0.128
high,low	high	uniform	0.109	0.125	0.140
high,low	high,low	uniform	0.100	0.126	0.135
high,low	low	uniform	0.105	0.130	0.147
high,low	high	bell-shaped	0.124	0.136	0.153
high,low	high,low	bell-shaped	0.098	0.123	0.136
high,low	low	bell-shaped	0.121	0.150	0.163
low	high	U-shaped	0.215	0.234	0.210
low	high,low	U-shaped	0.176	0.208	0.185
low	low	U-shaped	0.208	0.236	0.220
low	high	uniform	0.180	0.196	0.200
low	high,low	uniform	0.187	0.215	0.199
low	low	uniform	0.192	0.226	0.201
low	high	bell-shaped	0.200	0.231	0.207
low	high,low	bell-shaped	0.210	0.247	0.224
low	low	bell-shaped	0.198	0.236	0.205

Table 6.1: Average stochasticity of the agents' decisions in the simplified and training environments. Stochasticity measures the uncertainty of a choice, where higher values represent higher stochasticity i.e. for the choice probability p, we have stochasticity $\text{stoch}_p = 0.5 - |0.5 - p|$. See Equations 5.1 and 5.2 for details. The three characteristics that define the lottery types in the training environment are the outcome variance $\mathcal{V}$, probability range $\mathcal{A}$ and probability distribution $\mathcal{D}$.

the training environment (with low outcome variance), this indicates that the agents do not learn to differentiate strongly between lotteries. Perhaps lotteries with highly contrasted outcomes are necessary for the agents to learn the structure of the input and develop clear preferences.

For the rest of the analysis we look more closely at the 'variance-neutral' environments which include lotteries with both high and low outcome variance (i.e. the middle section of Table 6.1). When comparing the average stochasticity, we find that the simpler the testing environments, the less stochastic the agents' decisions are. This indicates that for easy decisions, such as choices between lotteries with large differences in expected value, the preferences are stronger. We also see that there is a consistent trend in each of the three testing environments: The degree of stochasticity increases, the more certain the probability distribution $\mathcal{D}$ in the training environment. Interestingly, the probability range $\mathcal{A}$ of the environment, which directly influences the ambiguity of the environment, did not have a large effect on the stochasticity of the agents' decision-making.

6.2 Performance in Simplified Environment

In this section, we report the results on the effect of the training environment on the agents' performance in simplified environments. The data shows a very high positive correlation (+0.9) between the performances in the two simplified environments. Thus, we only present the results of the 'Simplified environment I' here. The full data on the performance (Simplified environment I and II and the training environment) can be found in the appendix A.1. However, comparing the performance in the training environments is problematic; although the datasets have been generated to have the same total score for a decision-maker with a random strategy (choice probability 0.5 for each choice), this does not mean that the maximum score is the same. We also find a high negative correlation (-0.8) between the average performance in the simplified environments and the standard deviation between runs.

As in the stochasticity analysis, the outcome variance $\mathcal{V}$ is the environment characteristic that has the largest effect on the performance. On average, agents trained in environments with high outcome variance perform worse in the simplified environment, than agents trained in environments with low outcome variance: For training environments with high outcome variance, the average score (as defined in Equation 5.5) is only 279 on average, compared to 285 in training environments with low outcome variance. The best performances (288 on average) result from agents trained in environments with varied variance (high and low). Both of the other environmental characteristics, the probability range $\mathcal{A}$ and the probability distribution $\mathcal{D}$, which directly affect the ambiguity of the lottery, have a similar affect on the performance, which we summarize in Table 6.3. In both cases, the more ambiguous the training environment the worse the performance is in the simplified environment, which confirms our hypothesis that populations evolving in uncertain environments develop decision-making behavior that is sub-optimal in simplified laboratory setups.

While the measured performance differences are small ($\sim 2 - 10\%$), most of the differences between the performance in corresponding high and low outcome variance settings (other parameters equal), are considered statistically significant by conventional criteria (t-test, $p < 0.05$). However, the same performance score can result from many different choice combinations. To better understand to what extent the performance difference corresponds to differing decision-making behavior, one can consider the percentage of differing choices made by populations trained in differing environments. Using a few settings as random samples, we found that a small increase in the difference in performance (e.g. 2% and 4%), can correspond to a large increase in the percentage of differing choices (e.g. 6% and 25%). This shows that the performance metric does not provide enough insight on its own. In the following sections we look more closely at the choice patterns that affect the performance.

Training environment			Simplified environment I *(Large EV-diff)*
$\mathcal{V}$	$\mathcal{A}$	$\mathcal{D}$	Score (σ)
high	high	bell-shaped	280 (± 5.8)
high	high	uniform	278 (± 10.9)
high	high	U-shaped	270 (± 10.0)
high	high, low	bell-shaped	284 (± 7.3)
high	high, low	uniform	277 (± 9.6)
high	high, low	U-shaped	276 (± 9.7)
high	low	bell-shaped	292 (± 8.4)
high	low	uniform	286 (± 7.2)
high	low	U-shaped	274 (± 11.0)
high, low	high	bell-shaped	286 (± 4.0)
high, low	high	uniform	287 (± 4.2)
high, low	high	U-shaped	286 (± 3.5)
high, low	high, low	bell-shaped	288 (± 2.8)
high, low	high, low	uniform	286 (± 5.7)
high, low	high, low	U-shaped	286 (± 3.1)
high, low	low	bell-shaped	291 (± 2.3)
high, low	low	uniform	289 (± 3.4)
high, low	low	U-shaped	289 (± 3.0)
low	high	bell-shaped	285 (± 5.1)
low	high	uniform	285 (± 5.4)
low	high	U-shaped	283 (± 7.2)
low	high, low	bell-shaped	286 (± 5.9)
low	high, low	uniform	286 (± 4.1)
low	high, low	U-shaped	282 (± 6.6)
low	low	bell-shaped	288 (± 3.5)
low	low	uniform	285 (± 4.0)
low	low	U-shaped	284 (± 7.1)

Table 6.2: Total average score of the agents in the 'Simplified environment I' dataset depending on their training environment. The choices in the 'Simplified environment I' dataset consist of lotteries that (1) are not ambiguous (probabilities are fixed to a value and not range) and (2) differ in expected value (EV) by at least 0.4. The score of an agent with a random strategy is 200. The payoffs of each choice lie between 0 and 1. The reported standard deviation σ refers to the amount of variation between runs (not agents). The three characteristics that define the lottery types in the training environment are the outcome variance $\mathcal{V}$, probability range $\mathcal{A}$ and probability distribution $\mathcal{D}$.

Probability range $\mathcal{A}$			Probability distribution $\mathcal{D}$		
low	low,high	high	bell-shaped	uniform	U-shaped
286	283	282	287	284	281

$\longrightarrow$ Increasing Ambiguity $\qquad$ $\longrightarrow$ Increasing Ambiguity

Table 6.3: Performance (total average score) of agents in the 'Simplified environment I' dataset, depending on ambiguity characteristics of their training environment. See Table 6.2 caption for details on the dataset.

6.3 Expected Utility Theory Violations (Certainty Effect)

The amount and type of EUT violations detected in the common ratio experiment (described in Section 5.3), were very similar across all training environments. We thus summarize the data and present the percentage of runs in which the majority of agents violated independence, averaged over all 27 environments, in Table 6.4. For detailed data on the individual environments, see Appendix A.2.

We detect a significant prevalence of independence violations: In the true certainty setting, we find that the agents almost always evolve to independence violating decision-making behavior. The effect is slightly less robust for small gains than for large gains. In the near certainty setting we find that in most of the runs the agents did not evolve to independence violating strategies. Only in around 10% of the runs, the dominant strategy of the agents violated independence and thus EUT. We also find more frequent violations of independence as the mixing factor λ (defined in Equation 5.6) decreases. This could be because the lower λ the least similar the mixed lotteries are to the original lotteries.

Almost all of the independence violations occur due to the certain lottery L_{1a} being chosen in choice A and the less certain lottery L_{2b} being chosen in the mixed choice B. This choice pattern is known as the 'certainty effect'. The opposite effect, the reverse certainty effect was very rarely found ($< 1\%$ of runs).

	% of runs violating independence	
	large gains	small gains
certainty		
$\lambda = 0.6$	99.4	89.3
$\lambda = 0.3$	99.6	95.0
$\lambda = 0.1$	99.4	96.7
near-certainty		
$\lambda = 0.6$	0.7	7.0
$\lambda = 0.3$	1.3	11.9
$\lambda = 0.1$	22.0	13.0

Table 6.4: Percentage of runs where the dominant strategy (majority of agents) violates the independence axiom of EUT, averaged over environments. The mixing factor λ, near-certainty vs. certainty and large vs. small gain settings uniquely describe the choice-pairs used to detect the independence violations (See Tables 5.1 and 5.2 for details on the choice-pairs).

6.4 Ambiguity Preference

We assess the ambiguity preferences by presenting the agent populations with choices between lotteries on five EV-lines on the ambiguity triangle. We find that in all environments the ambiguity preferences are consistent for different outcome values and over EV-lines i.e. for example, ambiguity-seeking agents are ambiguity-seeking on each EV-line. In the cases where the population has a non-neutral ambiguity preference, the preferences lie on the endpoints of the EV-lines and not in the interior of the ambiguity triangle. We can, thus, group the ambiguity preferences into three types: ambiguity-seeking, ambiguity-averse and ambiguity-neutral, displayed in Figure 6.1. The fact that the agents choose the most ambiguous or the certain option means that they are maximally ambiguity-seeking or maximally ambiguity-averse and allows for their ambiguity preferences to be represented as *linear* indifference curves in the triangle.

In each environment, we found that ambiguity-seeking, ambiguity-neutral and ambiguity-averse populations could emerge. The agents within a population were generally homogeneous in regards to their ambiguity preferences, however, between runs we observed large discrepancies. Generally, we find that the dominant preference of the population evolves to ambiguity-seeking preferences slightly more often than to ambiguity-averse preferences (47% vs 39% for indifference range = 0.48-0.52). On average, less ambiguous training environments, evolve to ambiguity-seeking populations slightly more often (40% for high $\mathcal{A}$ and 49% for low $\mathcal{A}$, for indifference range = 0.48-0.52). The

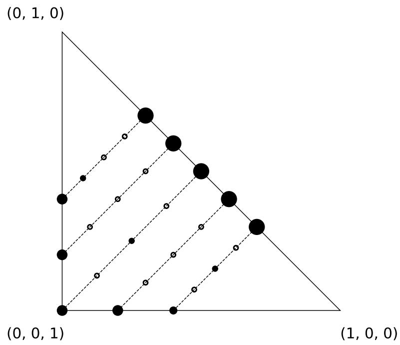

(a) Ambiguity-averse

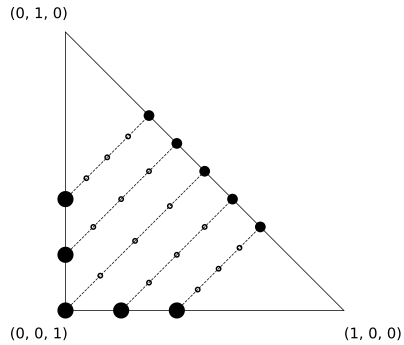

(b) Maximally ambiguity-seeking

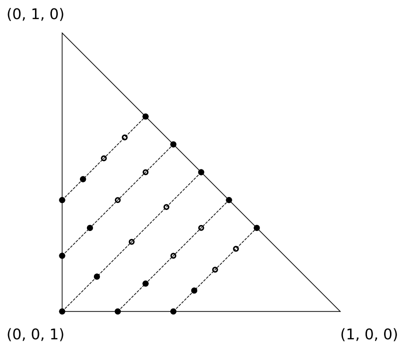

(c) Ambiguity-neutral

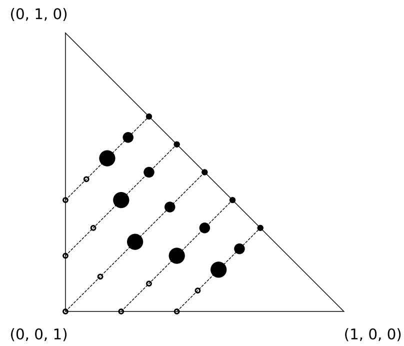

(d) Non-maximally ambiguity-seeking

Figure 6.1: Consistent ambiguity preferences of a population displayed on the ambiguity triangle. The size of the dots indicates the number of agents in the population that prefer this lottery over all others on the same EV-line. (a-c) are the three types of preference behavior we find in our simulations, (d) is a case that does *not* occur.

percentage of runs that evolve to ambiguity neutral preferences depend on the indifference range of the choice probability. In Appendix A.3 we present detailed results on the percentage of runs that evolved to ambiguity seeking, neutral and averse preferences for several indifference ranges and all environments.

6.5 Cooperation Dynamics

When running the simulations with multiple groups with different cooperation rates, we found that the group with the higher cooperation rate dies out. However, the developments (e.g. number of iterations until a group does out) depend strongly on the selection pressure of the genetic algorithm.

Chapter 7

Discussion

Our findings support our main hypothesis that very complex and ambiguous evolutionary environments may result in sub-optimal decision-making in simpler environments. We found that the outcome variance was the environmental characteristic with the largest effect on both the agents' performance and choice-stochasticity in simplified environments. Our findings suggest that strategies that are evolutionary beneficial in environments with high outcome variance do not perform best in simpler, less ambiguous environments. This is an important insight, because lab experiments consisting of lotteries with slightly differing monetary payoffs have very small outcome variance compared to some evolutionary prospects such as, for example, the prospect of going hunting, which could result in death or food. Also, as expected, evolutionary environments with higher ambiguity (reflected through the probability range and distribution of the lotteries' probabilities), resulted in poorer performance in the simple environment, suggesting that the agents were not trained to perform optimally in these environments. The stochasticity analysis showed that, as expected, more difficult choices were generally more stochastic. We also found that agents trained in more certain environments, made noisier choices. This is somewhat counter-intuitive, but may be linked to the fact that the agents are less well-adapted, i.e. show lower performance, in the environments with uncertain probability distributions.

Further compelling support for our hypothesis was given by the fact that the trained agents' decision-making behavior in the Allais' common ratio experiment was robust over all training environments. We detected violations of expected utility theory that are in line with general empirical findings [68, 69]. The choice patterns support the hypothesis of the certainty effect, where agents place disproportionate weight on outcomes when they are certain. Several alternative theories can accommodate the certainty effect, the most prominent such theory being prospect theory with its probability weighting mechanism [11]. The finding that the certainty effect is more pronounced with large gains is consistent with empirical studies [70, 71]. We found that the certainty effect did not occur in near certainty choices but required true certainty (probability one). This is at odds with both expected utility and probability weighting, which does not allow discontinuous preferences at certainty. Previous studies regarding this are not consistent, some report the same effects at and near certainty [72], while other evidence suggests that decisions at certainty and near certainty differ [73, 74].

A consistent finding across environments regarding ambiguity preferences was that populations either had no preference or predominately chose the most ambiguous or most certain option. That is to say, populations never preferred non-maximally ambiguous prospects. Experiments using the same ambiguity triangle setup report observe this as well [4]. Other aspects of the ambiguity preferences varied strongly across runs. In each training environment ambiguity-averse, ambiguity-neutral and ambiguity-seeking populations could emerge. This instability in the dominant ambiguity preference suggests that either the ambiguity preferences were not a crucial factor driving environmental fitness in the training environment or that there are multiple equilibria. This heterogeneity in ambiguity preference has been found in empirical studies as well, which find considerable heterogeneity in ambiguity preferences among individuals and depending on the size of gains [75, 76, 77]. We observed that less ambiguous training environments, evolved to ambiguity-seeking populations slightly more often. This finding supports our hypothesis that ambiguous environments lead to classically considered paradoxical behavior (i.e. ambiguity-averse behavior) in more simple environments. However, because of the variability in this metric, more simulations are necessary to assess if there is a significant affect of the training environment on the ambiguity behavior. Also, it needs to be further investigated if there are in fact multiple equilibria (ambiguity-averse, ambiguity-neutral and ambiguity-seeking) and how likely they are to occur. Looking more closely at the preferences of the populations in the early generations of the GA, could provide important insights on the development of the ambiguity preferences, which may be a strongly path-dependent process.

The findings regarding the cooperation mechanism show that group cooperation is does not emerge naturally, when wealth is simply redistributed. Similarly, previous studies show that in the equivalent public good games setting, where the multiplication factor is set to one, almost no contributions to the public good are made [78]. The reasoning behind the simple redistribution mechanism we implemented, was that this type of cooperation could provide valuable diversity leading to an evolutionary advantage over groups without such individuals. This was not observed, however.

Chapter 8

Conclusion

The results of our simulations clearly support the hypothesis that decision-making strategies emerging through evolutionary processes can result in expected utility theory violations, in a broad range of environmental settings. We also show that the characteristics of the evolutionary environment affect the performance and choice-stochasticity in simplified environments.

While our simulations thoroughly addressed the question of the effect of differing evolutionary environments, our findings are limited to the specific design choices we made for our input representation and genetic algorithm. To more precisely estimate the effect of the evolutionary environment, we suggest extending the simulations to a wider range of settings for the evolutionary algorithm and strategically analysing the influence of these design choices. Another interesting question that merits future investigation is the effect of more complex cooperation mechanisms, specifically introducing incentives for cooperation by, for example, using a multiplication factor > 1 in the public goods game.

Appendix A

Appendix

A.1 Performance in Simplified and Training Environments

Training environment			Simplified environment I *(Large EV-diff)*	Simplified environment II *(Small EV-diff)*	Training environment
$\mathcal{V}$	$\mathcal{A}$	$\mathcal{D}$	Score (σ)	Score (σ)	Score (σ)
high	high	bell-shaped	280 (±5.8)	249 (±4.5)	225 (±2.8)
high	high	uniform	278 (±10.9)	247 (±7.3)	216 (±3.4)
high	high	U-shaped	270 (±10.0)	242 (±7.2)	213 (±2.4)
high	high, low	bell-shaped	284 (±7.3)	252 (±5.3)	218 (±2.5)
high	high, low	uniform	277 (±9.6)	246 (±6.2)	220 (±3.1)
high	high, low	U-shaped	276 (±9.7)	247 (±6.8)	216 (±3.2)
high	low	bell-shaped	292 (±8.4)	257 (±6.5)	222 (±4.2)
high	low	uniform	286 (±7.2)	253 (±5.3)	224 (±2.8)
high	low	U-shaped	274 (±11.0)	245 (±8.1)	221 (±5.1)
high, low	high	bell-shaped	286 (±4.0)	254 (±3.3)	240 (±1.4)
high, low	high	uniform	287 (±4.2)	253 (±3.5)	236 (±2.2)
high, low	high	U-shaped	286 (±3.5)	253 (±3.4)	243 (±1.8)
high, low	high, low	bell-shaped	288 (±2.8)	255 (±2.9)	237 (±1.2)
high, low	high, low	uniform	286 (±5.7)	253 (±3.7)	237 (±1.6)
high, low	high, low	U-shaped	286 (±3.1)	254 (±3.7)	244 (±1.9)
high, low	low	bell-shaped	291 (±2.3)	257 (±2.4)	243 (±1.3)
high, low	low	uniform	289 (±3.4)	255 (±3.1)	234 (±1.2)
high, low	low	U-shaped	289 (±3.0)	255 (±3.1)	240 (±1.2)
low	high	bell-shaped	285 (±5.1)	252 (±5.0)	256 (±0.7)
low	high	uniform	285 (±5.4)	252 (±4.4)	253 (±0.8)
low	high	U-shaped	283 (±7.2)	251 (±5.6)	251 (±1.0)
low	high, low	bell-shaped	286 (±5.9)	253 (±4.8)	250 (±1.2)
low	high, low	uniform	286 (±4.1)	255 (±3.0)	254 (±0.8)
low	high, low	U-shaped	282 (±6.6)	251 (±4.5)	254 (±0.7)
low	low	bell-shaped	288 (±3.5)	255 (±3.1)	255 (±0.8)
low	low	uniform	285 (±4.0)	254 (±4.4)	250 (±1.0)
low	low	U-shaped	284 (±7.1)	250 (±5.9)	252 (±0.8)

Table A.1: Total average score in simplified and training environments, depending on the agents training environment. The score of an agent with a random strategy is 200. The reported standard deviation λ refers to the amount of variation between runs (not agents). The three characteristics that define the lottery types in the training environment are the outcome variance $\mathcal{V}$, probability range $\mathcal{A}$ and probability distribution $\mathcal{D}$.

A.2 Results on Allais' Common Ratio Experiment

Training environment			**% of runs where dominant strategy violates EUT (at certainty)**					
			Large gains			Small gains		
$\mathcal{V}$	$\mathcal{A}$	$\mathcal{D}$	$\lambda = 0.6$	$\lambda = 0.3$	$\lambda = 0.1$	$\lambda = 0.6$	$\lambda = 0.3$	$\lambda = 0.1$
high	high	bell-shaped	100	100	100	70	90	95
high	high	uniform	100	100	100	85	95	100
high	high	U-shaped	100	100	100	80	90	90
high	high, low	bell-shaped	100	100	100	100	100	100
high	high, low	uniform	95	95	95	80	90	90
high	high, low	U-shaped	100	100	100	75	85	100
high	low	bell-shaped	95	95	95	65	70	85
high	low	uniform	100	100	100	75	90	90
high	low	U-shaped	95	100	95	65	75	80
high, low	high	bell-shaped	100	100	100	100	100	100
high, low	high	uniform	100	100	100	80	95	95
high, low	high	U-shaped	100	100	100	95	95	95
high, low	high, low	bell-shaped	100	100	100	100	100	100
high, low	high, low	uniform	100	100	100	100	100	100
high, low	high, low	U-shaped	100	100	100	100	100	100
high, low	low	bell-shaped	100	100	100	100	100	100
high, low	low	uniform	100	100	100	100	100	100
high, low	low	U-shaped	100	100	100	100	100	100
low	high	bell-shaped	100	100	100	85	95	95
low	high	uniform	100	100	100	90	100	100
low	high	U-shaped	100	100	100	100	100	100
low	high, low	bell-shaped	100	100	100	100	100	100
low	high, low	uniform	100	100	100	95	100	100
low	high, low	U-shaped	100	100	100	80	100	100
low	low	bell-shaped	100	100	100	100	100	100
low	low	uniform	100	100	100	90	95	95
low	low	U-shaped	100	100	100	100	100	100
Average			99.4	99.6	99.4	89.3	95.0	96.7

Table A.2: Percentage of runs, where the dominant strategy (majority of agents) violates the independence axiom of EUT in the Allais common ratio experiment *at certainty*. The decision of the majority of agents in a populations is considered the dominant strategy.($\mathcal{V}$: outcome variance, $\mathcal{A}$: probability range and $\mathcal{D}$: probability distribution.)

Training environment			% of runs where dominant strategy violates EUT (near-certainty)					
			Large gains			Small gains		
$\mathcal{V}$	$\mathcal{A}$	$\mathcal{D}$	$\lambda = 0.6$	$\lambda = 0.3$	$\lambda = 0.1$	$\lambda = 0.6$	$\lambda = 0.3$	$\lambda = 0.1$
high	high	bell-shaped	5	5	55	5	30	30
high	high	uniform	5	10	25	10	15	20
high	high	U-shaped	0	0	25	5	15	15
high	high, low	bell-shaped	0	0	40	10	20	20
high	high, low	uniform	0	0	30	5	15	15
high	high, low	U-shaped	0	5	35	0	25	30
high	low	bell-shaped	5	5	50	15	25	30
high	low	uniform	0	0	40	10	20	20
high	low	U-shaped	0	5	20	10	15	20
high, low	high	bell-shaped	0	0	20	0	0	0
high, low	high	uniform	0	0	30	15	20	20
high, low	high	U-shaped	0	0	10	20	20	20
high, low	high, low	bell-shaped	0	0	5	5	5	5
high, low	high, low	uniform	0	0	20	10	10	10
high, low	high, low	U-shaped	0	0	5	0	0	0
high, low	low	bell-shaped	0	0	15	5	5	5
high, low	low	uniform	0	0	20	0	0	0
high, low	low	U-shaped	0	0	10	0	0	0
low	high	bell-shaped	0	0	10	10	10	10
low	high	uniform	5	5	25	0	10	10
low	high	U-shaped	0	0	25	15	15	15
low	high, low	bell-shaped	0	0	10	0	0	0
low	high, low	uniform	0	0	15	5	5	5
low	high, low	U-shaped	0	0	25	10	20	20
low	low	bell-shaped	0	0	15	5	5	5
low	low	uniform	0	0	10	10	10	10
low	low	U-shaped	0	0	5	5	5	5
Average			0.7	1.3	22.0	7.0	11.9	13

Table A.3: Percentage of runs, where the dominant strategy (majority of agents) violates the independence axiom of EUT per environment in the Allais common ratio experiment *near certainty*. ($\mathcal{V}$: outcome variance, $\mathcal{A}$: probability range and $\mathcal{D}$: probability distribution.)

A.3 Detailed Results on Ambiguity Preferences

Training environment			**Dominant ambiguity preferences [% of runs]**					
			Indifference range: 0.48 - 0.52			Indifference range: 0.45 - 0.55		
$\mathcal{V}$	$\mathcal{A}$	$\mathcal{D}$	seeking	neutral	averse	seeking	neutral	averse
high	high	bell-shaped	55	5	40	55	15	30
high	high	uniform	30	15	55	20	35	45
high	high	U-shaped	50	5	45	45	20	35
high	high, low	bell-shaped	65	5	30	45	25	30
high	high, low	uniform	45	10	45	30	25	45
high	high, low	U-shaped	50	15	35	45	30	25
high	low	bell-shaped	50	0	50	35	25	40
high	low	uniform	40	5	55	35	15	50
high	low	U-shaped	45	25	30	30	45	25
high, low	high	bell-shaped	45	10	45	35	35	30
high, low	high	uniform	35	15	50	30	20	50
high, low	high	U-shaped	35	5	60	25	25	50
high, low	high, low	bell-shaped	45	25	30	35	50	15
high, low	high, low	uniform	75	0	25	35	45	15
high, low	high, low	U-shaped	60	5	35	40	40	20
high, low	low	bell-shaped	55	0	45	40	35	25
high, low	low	uniform	50	5	45	40	25	35
high, low	low	U-shaped	60	10	30	45	40	15
low	high	bell-shaped	40	25	35	25	55	20
low	high	uniform	35	30	35	25	55	20
low	high	U-shaped	40	40	20	25	60	15
low	high, low	bell-shaped	20	25	55	5	50	45
low	high, low	uniform	50	20	30	40	35	25
low	high, low	U-shaped	55	25	20	25	65	10
low	low	bell-shaped	50	15	35	10	85	5
low	low	uniform	35	30	35	15	65	20
low	low	U-shaped	60	15	25	35	50	15
Average			47	14	39	32	40	28

Table A.4: Percentage of runs where the population's dominant strategy is ambiguity-seeking, ambiguity-averse and ambiguity-neutral per environment.The three characteristics that define the lottery types in the training environment are the outcome variance $\mathcal{V}$, probability range $\mathcal{A}$ and probability distribution $\mathcal{D}$.

Bibliography

[1] M. Allais, “Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’ecole americaine,” *Econometrica*, vol. 21, no. 4, pp. 503–546, 1953.

[2] D. Ellsberg, “Risk, ambiguity, and the savage axioms,” *The Quarterly Journal of Economics*, vol. 75, no. 4, pp. 643–669, 1961.

[3] V. Nanotti, “How has evolution shaped our decision making? A neural network agent based model for the development of heuristic.,” *Master’s thesis, Chair of Entrepreneurial Risks, ETH Zurich*, 2020.

[4] E. Fehr, D. Burghart, and T. Epper, “The uncertainty triangle – uncovering heterogeneity in attitudes towards uncertainty,” *Journal of Risk and Uncertainty*, vol. 60, 04 2020.

[5] D. Bernoulli, “Exposition of a new theory on the measurement of risk,” *Econometrica*, vol. 22, no. 1, pp. 23–36, 1954.

[6] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[7] L. J. Savage, *The Foundations of Statistics*. Wiley Publications in Statistics, 1954.

[8] S. Huck and W. Müller, “Allais for all: Revisiting the paradox in a large representative sample,” *Journal of Risk and Uncertainty*, vol. 44, pp. 261–293, June 2012.

[9] S. Trautmann and G. van de Kuilen, *Ambiguity attitudes*, pp. 89–116. John Wiley & Sons Ltd, 2015.

[10] V. Yukalov and D. Sornette, “Quantum decision theory as quantum theory of measurement,” *Physics Letters A*, vol. 372, no. 46, pp. 6867–6871, 2008.

[11] D. Kahneman and A. Tversky, “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, vol. 47, pp. 263–291, March 1979.

[12] A. Tversky and D. Kahneman, “Advances in prospect theory: Cumulative representation of uncertainty,” *Journal of Risk and Uncertainty*, vol. 5, no. 4, pp. 297–323, 1992.

[13] M. Abdellaoui, “Parameter-free elicitation of utility and probability weighting functions,” *Management Science*, vol. 46, no. 11, pp. 1497–1512, 2000.

[14] A. Bruhin, H. Fehr-Duda, and T. Epper, “Risk and rationality: Uncovering heterogeneity in probability distortion,” *Econometrica*, vol. 78, no. 4, pp. 1375–1412, 2010.

[15] M. O. Rieger, M. Wang, and T. Hens, “Estimating cumulative prospect theory parameters from an international survey,” *Theory and Decision*, vol. 82, no. 4, pp. 567–596, 2017.

[16] A. Oliver, “Your money and your life: Risk attitudes over gains and losses,” *Journal of Risk and Uncertainty*, vol. 57, pp. 29–50, August 2018.

[17] M. H. Birnbaum and W. R. McIntosh, “Violations of branch independence in choices between gambles,” *Organizational Behavior and Human Decision Processes*, vol. 67, no. 1, pp. 91–110, 1996.

[18] G. M. Ferro and D. Sornette, “Stochastic representation decision theory: How probabilities and values are entangled dual characteristics in cognitive processes,” *Plos One*, vol. 15, no. 12, 2020.

[19] N. C. Barberis, “Thirty years of prospect theory in economics: A review and assessment,” *Journal of Economic Perspectives*, vol. 27, pp. 173–96, February 2013.

[20] P. Bossaerts, N. Yadav, and C. Murawski, “Uncertainty and computational complexity,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 374, 2018.

[21] A. Tversky, “Elimination by aspects: A theory of choice,” *Psychological Review*, vol. 79, no. 4, pp. 281–299, 1972.

[22] J. W. Payne, J. R. Bettman, and E. J. Johnson, *The Adaptive Decision Maker*. Cambridge University Press, 1993.

[23] G. Gigerenzer and R. Selten, *Bounded Rationality: The adaptive toolbox*. Cambridge: The MIT Press, 2001.

[24] J. Busemeyer and J. Townsend, "Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment," *Psychological review*, vol. 100, pp. 432–59, 08 1993.

[25] J. Johnson and J. Busemeyer, "Decision making under risk and uncertainty," *Wiley Interdisciplinary Reviews: Cognitive Science*, vol. 1, pp. 736–749, 09 2010.

[26] J. M. Smith and G. R. Price, "The logic of animal conflict," *Nature*, vol. 246, no. 5427, pp. 15–18, 1973.

[27] J. R. Stevens, "The evolutionary biology of decision making," *Better than Conscious? Decision-Making, the Human Mind and Implications for Institutions (MIT Press)*, 05 2008.

[28] B. M. Knauft, T. S. Abler, L. Betzig, C. Boehm, R. K. Dentan, T. M. Kiefer, K. F. Otterbein, J. Paddock, and L. Rodseth, "Violence and sociality in human evolution [and comments and replies]," *Current Anthropology*, vol. 32, no. 4, pp. 391–428, 1991.

[29] C. L. Apicella and J. B. Silk, "The evolution of human cooperation," *Current Biology*, vol. 29, no. 11, pp. R447–R450, 2019.

[30] L. R. Santos and A. G. Rosati, "The evolutionary roots of human decision making," *Annual Review of Psychology*, vol. 66, no. 1, p. 321–347, 2015.

[31] D. Dunning, C. Heath, and J. Suls, "Flawed self-assessment implications for health, education, and the workplace," *Psychological Science in the Public Interest*, vol. 5, pp. 69–106, 12 2004.

[32] M. Alicke and O. Govorun, *The self in social judgement*, pp. 83–106. Psychology Press, 01 2005.

[33] S. E. Taylor, *Positive illusions: Creative self-deception and the healthy mind*. Basic Books, 1989.

[34] D. D. Johnson, D. T. Blumstein, J. H. Fowler, and M. G. Haselton, "The evolution of error: error management, cognitive constraints, and adaptive decision-making biases," *Trends in Ecology & Evolution*, vol. 28, no. 8, p. 474–481, 2013.

[35] A. Abbey, "Sex differences in attributions for friendly behavior: Do males misperceive females' friendliness," *Journal of Personality and Social Psychology*, pp. 830–838, 1982.

[36] R. L. Shotland and J. M. Craig, "Can men and women differentiate between friendly and sexually interested behavior?," *Social Psychology Quarterly*, vol. 51, no. 1, pp. 66–73, 1988.

[37] F. E. Saal, C. B. Johnson, and N. Weber, "Friendly or sexy?: It may depend on whom you ask," *Psychology of Women Quarterly*, vol. 13, no. 3, pp. 263–276, 1989.

[38] M. Haselton and D. Buss, "Error management theory: A new perspective on biases in cross-sex mind reading," *Journal of personality and social psychology*, vol. 78, pp. 81–91, 02 2000.

[39] A. Bernardo and I. Welch, "On the evolution of overconfidence and entrepreneurs," *Journal of Economics & Management Strategy*, vol. 10, pp. 301–330, 09 2001.

[40] A. Kacelnik, "Meanings of rationality," in *Rational Animals?*, pp. 87–106, Oxford University Press, 2006.

[41] D. W. Stephens and J. R. Krebs., *Foraging Theory*. Princeton Univ. Press, 1986.

[42] H. Arkes and P. Ayton, "The sunk cost and concorde effects: Are humans less rational than lower animals?," *Psychological Bulletin*, vol. 125, pp. 591–600, 1999.

[43] A. Houston and J. McNamara, *Models of Adaptive Behaviour: an approach based on state*. Cambridge University Press, 1999.

[44] A. I. Houston, J. M. Mcnamara, and M. D. Steer, "Do we expect natural selection to produce rational behaviour?," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 362, no. 1485, p. 1531–1543, 2007.

[45] A. Houston, "Natural selection and context-dependent values," *Proceedings of the Royal Society B: Biological Sciences*, vol. 264, pp. 1539 – 1541, 1997.

[46] N. G. Walworth, J. Hinners, P. A. Argyle, S. G. Leles, M. A. Doblin, S. Collins, and N. M. Levine, "The evolution of trait correlations constrains phenotypic adaptation to high co2 in a eukaryotic alga," *bioRxiv*, 2020.

[47] E. Firnberg and M. Ostermeier, "The genetic code constrains yet facilitates Darwinian evolution," *Nucleic Acids Research*, vol. 41, pp. 7420–7428, 06 2013.

[48] L. Fogel, A. Owens, and M. Walsh, *Artificial intelligence through simulated evolution*. Wiley, 1966.

[49] J. H. Holland, *Adaptation in Natural and Artificial Systems*. University of Michigan Press, 1975.

[50] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., 1st ed., 1989.

[51] T. Blickle and L. Thiele, "A comparison of selection schemes used in evolutionary algorithms," *Evolutionary Computation*, vol. 4, p. 361–394, Dec. 1996.

[52] M. de la Maza and B. Tidor, "An analysis of selection procedures with particular attention paid to proportional and boltzmann selection," in *Proceedings of the 5th International Conference on Genetic Algorithms* (S. Forrest, ed.), pp. 124–131, Morgan Kaufmann Publishers, 1993.

[53] J. H. Gillespie, "Natural selection for variances in offspring numbers: a new evolutionary principle," *The American Naturalist*, vol. 111, no. 981, pp. 1010–1014, 1977.

[54] T. Bäck, "Selective pressure in evolutionary algorithms: A characterization of selection mechanisms," in *In Proceedings of the First IEEE Conference on Evolutionary Computation*, pp. 57–62, IEEE Press, 1994.

[55] H. Mühlenbein and D. Schlierkamp-Voosen, "Predictive models for the breeder genetic algorithm," *Evolutionary Computation*, vol. 1, no. 1, 1993.

[56] J. E. Baker, "Adaptive selection methods for genetic algorithms," in *Proceedings of the 1st International Conference on Genetic Algorithms*, p. 101–111, L. Erlbaum Associates Inc., 1985.

[57] V. Grimm and S. F. Railsback, *Individual-based modeling and ecology*. Princeton university press, 2013.

[58] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, 1986.

[59] W. Jenkins, "Neural network weight training by mutation," *Computers & Structures*, vol. 84, no. 31, pp. 2107–2112, 2006.

[60] P. Palmes, T. Hayasaka, and S. Usui, "Mutation-based genetic neural network," *IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council*, vol. 16, pp. 587–600, 06 2005.

[61] P. A. Samuelson, "The pure theory of public expenditure," *The Review of Economics and Statistics*, vol. 36, no. 4, pp. 387–389, 1954.

[62] X. Shi, W. Long, Y. Li, and D. Deng, "Multi-population genetic algorithm with er network for solving flexible job shop scheduling problems," *PLOS ONE*, vol. 15, pp. 1–23, 05 2020.

[63] "Multi-subpopulation evolutionary algorithms for coverage deployment of uav-networks," *Ad Hoc Networks*, vol. 68, pp. 16–32, 2018.

[64] A. Hintze, R. S. Olson, C. Adami, and R. Hertwig, "Risk sensitivity as an evolutionary adaptation," *Scientific Reports*, vol. 5, Feb 2015.

[65] P. R. Blavatskyy, "Reverse common ratio effect," *Journal of Risk and Uncertainty*, vol. 40, no. 3, pp. 219–241, 2010.

[66] B. Eva, "Principles of indifference," *The Journal of Philosophy*, vol. 116, no. 7, pp. 390–411, 2019.

[67] D. Burghart, T. Epper, and E. Fehr, "The ambiguity triangle: Uncovering fundamental patterns of behavior under uncertainty," discussion paper, June 2015.

[68] A. Tversky and D. Kahneman, "Rational choice and the framing of decisions," *The Journal of Business*, vol. 59, no. 4, pp. S251–S278, 1986.

[69] M. Cohen and J.-Y. Jaffray, "Certainty effect versus probability distortion: An experimental analysis of decision making under risk," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 14, pp. 554–560, 11 1988.

[70] P. Blavatskyy, A. Ortmann, and V. Panchenko, "Now you see it, now you don't: How to make the Allais Paradox appear, disappear, or reverse," Discussion Papers 2015-14, School of Economics, The University of New South Wales, July 2015.

[71] C. F. Camerer, "An experimental test of several generalized utility theories," *Journal of Risk and Uncertainty*, vol. 2, no. 1, pp. 61–104, 1989.

[72] R. Jain and K. Nielsen, "A Systematic Test of the Independence Axiom Near Certainty," IEAS Working Paper : academic research 20-A001, Institute of Economics, Academia Sinica, Taipei, Taiwan, June 2020.

[73] J. Andreoni and C. Sprenger, "Uncertainty equivalents: Testing the limits of the independence axiom," Working Paper 17342, National Bureau of Economic Research, August 2011.

[74] J. Andreoni and C. Sprenger, "Certain and Uncertain Utility: The Allais Paradox and Five Decision Theory Phenomena," levine's working paper archive, David K. Levine, Feb. 2010.

[75] D. O. Stahl, "Heterogeneity of ambiguity preferences," *The Review of Economics and Statistics*, vol. 96, no. 4, pp. 609–617, 2014.

[76] G. van de Kuilen, R. Cubitt, and S. Mukerji, "The strength of sensitivity to ambiguity," *Theory and Decision*, vol. 85, pp. 275–302, Oct. 2018.

[77] S. T. Trautmann and G. van de Kuilen, *Ambiguity Attitudes*, ch. 3, pp. 89–116. John Wiley & Sons, Ltd, 2015.

[78] R. Isaac, J. M. Walker, and A. W. Williams, "Group size and the voluntary provision of public goods: Experimental evidence utilizing large groups," *Journal of Public Economics*, vol. 54, no. 1, pp. 1–36, 1994.

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every semester paper, Bachelor's thesis, Master's thesis and any other degree paper undertaken during the course of studies, including the respective electronic versions.

Lecturers may also require a declaration of originality for other written papers compiled for their courses.

I hereby confirm that I am the sole author of the written work here enclosed and that I have compiled it in my own words. Parts excepted are corrections of form and content by the supervisor.

Title of work (in block letters):

EMERGENCE OF HUMAN DECISION-MAKING BEHAVIOR USING EVOLUTIONARY SELECTION OF AN ECOLOGY OF LEARNING NEURAL NETWORKS

Authored by (in block letters):

For papers written by groups the names of all authors are required.

Name(s):	**First name(s):**
Schmidt-Rohr	Anna

With my signature I confirm that

- I have committed none of the forms of plagiarism described in the 'Citation etiquette' information sheet.
- I have documented all methods, data and processes truthfully.
- I have not manipulated any data.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for plagiarism.

Place, date	**Signature(s)**
Zürich, 11.06.2021	

For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.