

The Evolutionary Foundation of Probability Weighting and Hyperbolic Discounting

And Their Intimate Connection

Abstract

In this paper, we delve into the evolutionary foundation of both probability weighting and hyperbolic discounting. It is argued that it is evolution which selects between various kinds of people those who have the optimal preference profile. As such, utility functions and behavioural biases can be explained from an evolutionary perspective. It is also argued that there may be an intimate connection between hyperbolic discounting and probability weighting. This follows from the similarities between these two deviations from what normative economic theory would consider 'rational'. Following the suggestion of this intimate connection, we analyse a model which is capable of explaining why probability weighting may have been advantageous for our ancestors and we look whether we can reinterpret this model to also include hyperbolic discounting. This turns out to be difficult, however, as an analysis of the properties of the model makes clear.

Acknowledgements

The author wishes to thank Jerker Holm for the supervision during the writing of this paper and for helping make difficult parts more understandable. The author also wishes to thank Niels Kouwenhoven and Matt van der Gronde for the discussions which helped materialise the concepts within this paper.

Contents

Cover Page	1
Contents	2
1 Introduction	3
1.1 Behavioural Economics	3
1.2 Probability Weighting	4
1.3 Hyperbolic Discounting	5
1.4 An Intimate Connection?	7
2 Evolutionary Psychology	8
2.1 Fitness Equals Utility	8
2.2 Fitness Does not Equal Utility	9
3 Probability Weighting	9
3.1 Replication of the Model	9
3.2 Extension of the Model	12
4 Hyperbolic Discounting	14
4.1 Appropriate Properties	14
4.2 Inappropriate Properties	14
4.3 Direct Interpretation	15
4.4 Indirect Interpretation	16
5 Limitations	17
5.1 What Games Are Played in Nature	17
5.2 Multiple Self	18
5.3 Learning Quick Enough or Observing	18
5.4 Prospect Theory Is not Cumulative Prospect Theory	18
5.5 Inflection Point	18
6 Conclusion	19
6.1 Achievements of Evolutionary Psychology	19
Bibliography	20

1 Introduction

1.1 Behavioural Economics

Behavioural economics has in recent times caused a shift in the focus of economists from theoretically rational models to empirically correct models. The purpose of behavioural economics is not to create a normative theory of human behaviour, but rather to produce a descriptive theory adequate to deal with certain counter-intuitive behaviour seen in many situations [12]. For example, it seems people are willing to drive across town to save 5 euro on a 15 euro calculator, but not to save 5 euro on a 125 euro jacket [14]. Research is now being done on how this irrational economic behaviour can be explained by psychological tendencies. The clear divide between psychology and economics has subsequently become more blurry and many suggestions from psychology for descriptions of more realistic human behaviour have entered into the field of economics. Mental accounting describes how people categorise financial activities in their brain in a similar manner to how accountants in firms categorise such activities and can explain, for example, why cab drivers stop work early on a good day while continue working longer on rainy days [29]. The problem which normative economic theory has with this behaviour is that it would be wiser to work as long as possible on a good day in which there is a lot of money to be made. To maximise pay-offs, it is rational to work relatively more on such days than on bad ones. The availability heuristic says that people determine the likelihood of an event occurring by the ease at which it comes to mind and the representativeness heuristic describes how people judge the probability of an object A belonging to a class B by seeing how much A resembles B [26]. These tendencies can help explain why we see certain irrational behaviour; behaviour which does not directly lead to utility maximisation. However, it does not explain why humans show this type of behaviour in the first place. We have merely pushed the question of explaining the data one step further, from the field of economics to the field of psychology. We have found a proximate cause for such seemingly irrational behaviour in these psychological tendencies, but we still lack an ultimate cause. It is the purpose of evolutionary psychology and evolutionary game theory to provide such an ultimate cause. The idea is that behaviour which is advantageous to the survival of the individual will spread in the population and become dominant. A more precise definition of this view will be given later in the description of our model. One needs only to show how certain behaviour, which normative economics may consider irrational, can be advantageous to human ancestors, in order to explain why this behaviour persists today. Evidence that this view of behaviour being shaped in our evolutionary past is correct follows from experiments with animals. When examining loss-aversion in capuchin monkeys, Chen et al. found that the behaviour was strikingly similar to that of humans. They conclude that, because capuchin monkeys lack the ability of social learning, this behaviour must have had a early-evolutionary origin and be innate to human ancestors [1]. Furthermore, studies on rats have shown that they exhibit the same irrational, hyperbolic way of discounting the future as we humans do, but we will see more on this particular topic later [20].

Coming back to the examples mentioned earlier, we can posit that mental accounting simplifies decision making by creating clear divisions between activities. Furthermore, representativeness leads to accurate predictions in most common situations; situations in which the resemblance between object A and class B really does say something about their connection. In general, we can say that applying heuristics greatly reduces complex tasks; something evolution would favour. Only in very specific situations do these heuristics lead to strong biases, as illustrated by the example of the cab driver [26]. Furthermore, it has been shown that loss-aversion may be beneficial in situations with strategic interaction. By irrationally overvaluing your own property, loss-aversion leads to toughness in bargaining. This may cause the opponent in a Nash-bargaining game to adjust his behaviour in a way that benefits the biased bargainer [10, 9]. If we view this Nash-bargaining game as representative of situations in which human ancestors found themselves, we conclude that loss-aversion would survive the evolutionary process in the long-run.

We have seen some examples of irrational economic behaviour and their underlying psychological tendencies and have found a number of global reasons for why this irrational behaviour may be advantageous when viewed through the lens of natural selection. We will now turn to two other forms of behaviour found in humans: probability weighting, which is part of Prospect Theory, and hyperbolic discounting. It is the main purpose of this paper to provide an account of why these types of behaviour, as well, may be advantageous in an evolutionary setting.

1.2 Probability Weighting

When examining data on gambles, it seems people do not adhere to the normative theory of expected utility. As such, they do not maximise their utility. Moreover, people seem to violate the independence axiom and are unwilling to back away from a losing bet, as illustrated by their reluctance to sell a falling stock [29]. In 1979, Kahneman and Tversky published their seminal paper on Prospect Theory, in which they developed a theory of human behaviour under risk and uncertainty which was capable of explaining the empirical data [13]. The authors proposed two main alterations to the expected utility theory: the value function would replace the standard utility function and weighted probabilities would be used in gambles, or what they called prospects.

The key feature of the value function is its reference point. People do not look at their total wealth and total utility when making a financial decision, but frame every separate action using their present state as a null-point or reference point. From there, a loss will decrease utility temporarily while a gain will increase utility temporarily, until a new reference-point is established. Another part of Prospect Theory is the diminishing sensitivity for both gains and losses. Finally, the value function also includes loss aversion, by making the part in the negative somewhat steeper than the part in the positive.

From an evolutionary perspective, these three features all make sense. We had already discussed loss-aversion above, which can be advantageous in situations with strategic interaction. The fact that people view gains and losses as compared to a reference point also makes sense when one realises that evolution is an ever increasing arms race in which you are running to stand still. It does not matter how good you are, only that you are better than the rest. The research by Chen et al. on capuchin monkeys also revealed this reference dependency [1]. Finally, the diminishing sensitivity can be explained by examining the environment in which animals struggle for life. In a paper by McDermott et al., the authors build a model where animals need a certain number of resources to survive to the next period [17]. In a situation with some uncertainty, animals who are at the verge of dying may need to take high risks in order to have at least a small chance of surviving. For these animals, it's all or nothing. As such, the marginal effect of losses on survival diminishes as the potential losses get larger.

The second, and for the present discussion the most important part of Prospect Theory is the probability weights which people seem to use when judging gambles. Although the theory has been adjusted and updated in more recent papers, the general idea is that people overweight prospects with low probability and underweight prospects with almost certain probability. An example of the first case is the betting on long-shot horses in horse racing; even though there is only a small chance that these horses will win, people behave as if the probability is higher. In their 1992 paper, Kahneman and Tversky propose a weighting function which has the following form [27]:

$$\pi(p, \gamma) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{\frac{1}{\gamma}}} \quad (1.1)$$

The factor γ in this equation determines the level of weighting applied; for $\gamma = 1$, the function is linear and the subjective probability π is equal to the true probability p . For values lower than $\gamma = 1$, over- and underweighting is observed in accordance with Figure 1 on the following page.

Violations of the expected utility theory have not only been observed in humans, but also in other animals [11]. Similar to the discovery of loss-aversion in monkeys, this provides evidence to back up the claim that evolution is responsible for the psychological tendencies we have and ultimately for the irrational economic behaviour we observe. But is probability weighting also advantageous? Would it survive in an evolutionary setting? One of the three key points this paper wants to propose is an extension of a previous model in which probability weighting is evolutionarily advantageous. Before we move on to this model, we continue with our introduction into behavioural economics and discuss hyperbolic discounting.




Figure 1: Prospect Theory probability weighting for $\gamma = 0.6$.

1.3 Hyperbolic Discounting

When rewards are realised at a time later than the present, the benefit of such rewards is discounted. In the market place, this is often done using the interest rate such that the present value of the reward is determined by calculating what pay-off in the present, together with its compound interest over time, would equal the future reward in terms of benefit. In a more social setting, a similar behaviour is reasonable when there is an uncertainty involved with the realisation of the reward. If rewards further in the future are less certain, their present value should be less high. For a constant uncertainty rate through time, or a constant interest rate, the shape of the discount function, which describes by how much to discount a future reward, should be exponentially decaying.

$$f(t) = f_0 \times \delta^t \quad \text{where } 0 < \delta < 1 \quad (1.2)$$

Several experiments, however, have shown that the shape of this function according to actual human behaviour is more closely related to a hyperbolic function [6, 19]. In this case, the subjective value from the present to the near future makes a large drop, i.e. people discount the near future heavily. For example, when you are given the option between 1 euro now or 1 euro tomorrow, you have a very strong preference for the 1 euro now. Simultaneously, the difference in value associated with two rewards in the far future, separated by the same amount of time, is only small¹. This means people do not care all that much about either obtaining 1 euro in one year time or obtaining 1 euro in one year plus one day. Put more simply, the discount factor is not constant, such as the interest rate or the uncertainty rate is constant, but increases for periods further into the future.

One of the important implications of hyperbolic discount function is that two functions at different time periods may cross. This means that there is a possibility of preference reversal in which a person would first favour A over B, but as some time passes, changes his mind and favours B over A². Similar to the discovery of loss-aversion among capuchin monkeys, Richards et al. found hyperbolic discounting to also be present in rats, suggesting a similar evolutionary mechanism for this behaviour [20].

¹In mathematics, the term 'hyperbolic' refers to the particular functional form $u(t, k) = \frac{u_0}{1+kt}$. Initial experiments used this functional form to fit their data [20]. However, in recent times, economists use the term hyperbolic to merely mean the behaviour just described: large drops from the present to the near future and smaller drops in the far future.

²One may wonder whether such a reversal of preferences is not so detrimental to optimal decision making that it could never survive through natural selection. As a reply to this, Netzer posits that animals were never able to change their minds in the type of situations in our evolutionary past we are considering [18]. Dynamic inconsistencies could never arise back then because decisions were irreversible. Only in recent times, through the market place, decisions have increasingly become more reversible. As such, in the modern world, hyperbolic discounting truly does lead to dynamic inconsistencies and the changing of one's mind.

Let us examine one possible mechanism for the occurrence of hyperbolic discounting suggested by Sozou³. For simplicity, let us assume there is no uncertainty regarding the realisation of a reward in the future if one is still alive. The only factor, in that case, that would cause future rewards to be less beneficial to the biological fitness of an animal is the chance that such an animal has died by the time the reward could potentially be realised. The theory in evolutionary biology which deals with such situations is called Life History Theory. In this framework, it is analysed how an animal can best spread out resources and mating opportunities over its lifetime. Following the uncertainty about being alive at some moment in the future, created by environmental hazards, it is often better to have a bias towards the present. If the hazard rate is a known constant, discounting future rewards should be done in an exponential manner as in Equation 1.2. However, Sozou argues that, throughout evolutionary times, the hazard rate was not always the same and a newborn would not know into what kind of environment he was born. In that case, there is some uncertainty about the hazard rate and the optimal shape of the discount function may no longer be exponential. This approach suggests one possible explanation for hyperbolic discounting⁴. While the hazard rate is initially not known by newborn animals, with time, more information about the specific hazard rate is discovered as the animals learn in what kind of environment they actually live. With this additional information, they update their idea about the hazard rate and change the shape of the discount function for the remaining time. Following some basic assumptions, Sozou shows that, in this case, a hyperbolic discount function is optimal [24].

As was rightfully observed by Dasgupta, however, this does not account for the type of preference reversal seen in most empirical data [2]. If people apply a hyperbolic discount function to future rewards, they may at one moment prefer A over B, while at a later time reverse their preference and go for B instead. This type of dynamic inconsistency arises because two hyperbolic discount functions can cross at different points in time. To see why this cannot occur under the function proposed by Sozou, consider what happens to the discount function, which started out hyperbolic, when some time has passed. The 'present' moves along the curve to the right and the discount function itself rescales to unity, conditional on still being alive at that moment. That is because a resource which was initially expected at some moment in the future, and therefore discounted, is now in the present, and therefore should no longer be discounted. The part of the discount function which lies in the future is not the full hyperbolic function, but merely its flat tail. This follows from the fact that the discount rate is age-dependent according to Sozou, who suggested that the hyperbolic shape relies on an updating of the discount rate with age, as more information becomes available to the animal. As the discount function at a later age is merely the last part of the original discount function, it can never cross the original function and lead to dynamic inconsistencies. However, the account by Sozou may still be relevant within Life History Theory to provide an idea of what factors influence the shape of the discount function.

The previous analysis suggests that there is a special property of the discount function which we require for it to be realistic and independent of age⁵. As time progresses, we move along the discount function to the right. If we are still alive at that moment, we update the discount function by rescaling it to unity. As nothing has really changed in the environment, the shape of the updated discount function should be the same of the original one. If the discount function is truly age-independent, it should be the same shape now as it is in one year time. There is only one function which has this property and that is the exponential function. If we move any step to the right and rescale the function to unity, the original exponential function appears. In the following section, we will see how it is possible to combine this important property with a hyperbolic function and how this method is the only way to obtain the preference reversal behaviour seen in experiments.

³A similar account is provided by Robsons, who looks at the aggregate uncertainty concerning survival rates to explain the hyperbolic shape of the discount function [22].

⁴There are other ideas about the origin of the hyperbolic discount function as well. Notably, Lehmann and Rousset determine the best way to discount future rewards when genetic drift is included, i.e. when the geographical spread of animals and their kin becomes important. They find that the function departs from the standard exponential function, possibly towards a more hyperbolic shape [15].

⁵Some experiments suggest that there is a slight dependency on age. Discounting is highest for very young and very old people, while lowest for middle aged people. A model by Sozou and Seymour, who combine their idea of an uncertain hazard rate with the theory of ageing, is capable of explaining this behaviour [25].

1.4 An Intimate Connection?

As was just mentioned, one of the reasons people discount the future is due to uncertainty. From an evolutionary perspective, this also makes sense as animals live in uncertain times. Promises of future rewards cannot be guaranteed to be fulfilled; there is not even a guarantee of being alive by the time the rewards can benefit the individual. The combination of these two types of uncertainty explains why animals, including humans, discount the future, but it does not suggest any reason for why the method of discounting should be hyperbolic. Following the previous discussion about probability weighting, however, it does seem there is a connection between the two ideas. Discounting occurs when there is uncertainty about the future, i.e. when the benefit of the reward is not secure, but is actually a gamble. The probability of realisation decreases exponentially with time as in Equation 1.2. The further into the future the reward will be realised, the less certain one can be about its realisation and its benefit. The probability associated with such a reward may decay exponentially when there is a constant uncertainty rate, but the subjective probability associated with the gamble does not follow a linear trend, as Prospect Theory teaches us. Humans overweight small probabilities, while they underweight large ones. The weighting function which we defined earlier in Equation 1.1 is a mapping from the true, objective probabilities to the subjective probabilities. If we apply this same mapping to the exponentially decaying discount function, we obtain the graph below:




Figure 2: The exponential discount function $f(t)$ and the function $u(t)$ obtained by applying the mapping from equation 1.1 to $f(t)$ with $\delta = 0.9$ and $\gamma = 0.6$. It seems $u(t)$ follows the economist's definition of 'hyperbolic'.

As can clearly be seen, the function which arises from applying probability weights to the exponentially decaying discount function has the feature that it makes a large drop from the present to the near future, while only smaller drops between periods further into the future. The second of the three key points that this paper wishes to propose is that there may be an intimate connection between probability weighting and the particular shape of the discount function, i.e. that they may actually be two sides of the same coin. Potential mechanisms behind this connection will be analysed in order to give credit to the concept. Besides making argumentative sense, the benefit of this idea is that the underlying discount function is, in fact, exponential⁶. This follows from the constant hazard rate; the chance of not being alive by the time the future reward is realised. As such, when a certain period of time has passed, and you are still alive, you can rescale the objective discount function to unity, obtain the same exponential function as before and then apply the mapping of probability weights which results in a hyperbolic function. This discount function is age-independent. At any moment in time, the function is fully hyperbolic. We can say that as time passes, we slide down the exponential function, but look at it through a lens. This lens distorts the image of the smooth exponential function and turns it hyperbolic. Crossing discount functions, dynamic inconsistencies and preference reversals arise naturally.

⁶Rogers suggests the underlying exponential function discounts future rewards by about 2 percent per year [23]. This is based on the fact that future rewards often benefit one's offspring, who are genetically only partially related to you.

With this striking similarity between hyperbolic discounting and probability weighting, one immediately wonders whether this is a coincidence or whether a deeper force is at play. Part of the purpose of this paper is to investigate this idea. Assuming for a moment that there really is a connection between the two effects, and assuming, as argued above, that evolution is, in fact, responsible for shaping our preferences, there are only three possibilities capable of explaining the observed similarity.

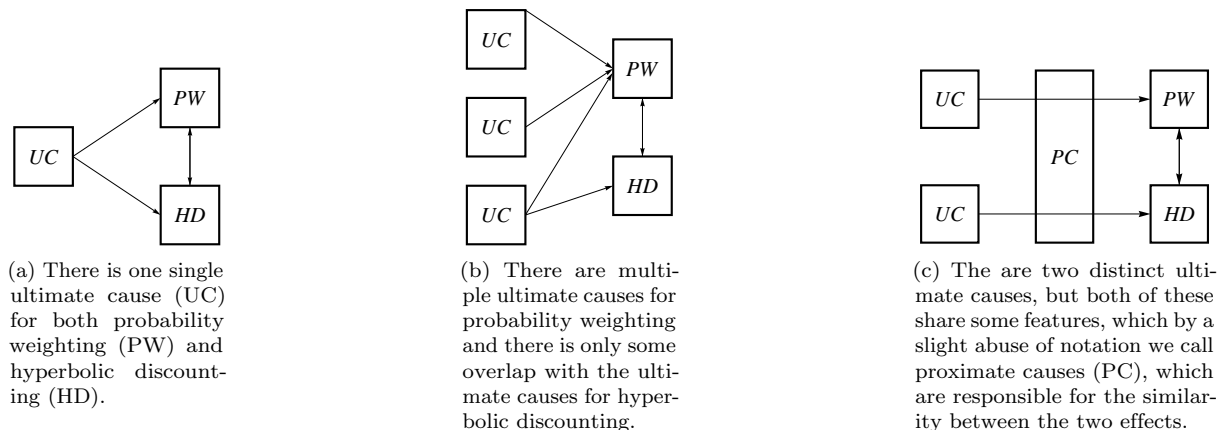


Figure 3: Graphical representation of the three possibilities concerning the similarity between probability weighting and hyperbolic discounting.

The first of these possibilities is that there is one single ultimate cause for why humans, and some animals, weight probabilities and for why they discount the future in a hyperbolic manner. The second possibility is that there are several ultimate causes for probability weighting, i.e. there are several reasons why it may be advantageous to weight probabilities in an evolutionary setting. It is, then, possible that only one of these causes is also responsible for the type of weighting which is applied to the exponentially decaying discount function, resulting in hyperbolic discounting. Finally, it is possible that there are two distinct ultimate causes for probability weighting and hyperbolic discounting. If this is the case, the similarities between the two effects may only be due to similarities in the features of the two ultimate causes, i.e. there may be a more proximate cause which is identical in both cases, but does not stem from the same underlying mechanism. The last of the three key points this paper wants to address is a model which tries to explain hyperbolic discounting by assuming there truly is an intimate connection between probability weighting and the hyperbolic shape of the discount function.

2 Evolutionary Psychology

2.1 Fitness Equals Utility

In the same way as that physics is needed to understand chemistry and chemistry is needed to understand biology, psychology is needed to understand economics. It is the purpose of evolutionary psychology to complete this list by finding the missing link between biology and psychology. Once this is established, much of economics may be explained from first principle, with less need for premises and assumptions.

The most fundamental insight, which lays at the heart of evolutionary psychology, is that happiness or utility is our body's way of controlling our actions. In fact, it is the long reach of the gene, which determines our behaviour in a way that is optimal for the genes. For example, people enjoy sex; they derive utility from it. This mechanism ensures we repeat the behaviour many times, increasing our genes' chance of survival. In this way, the behaviour which gives us utility also gives us a higher biological fitness. But utility is derived from many other things as well, such as eating or consuming goods and resources. If utility does, in fact, equal biological fitness, and we seek to maximise our utility, we ultimately end up maximising our fitness. It may seem surprising that evolution is capable of calculating which actions enhance our fitness, but explaining complexity from simplicity is the prime achievement of the theory of

evolution by natural selection. One example to show how good evolution can be at producing optimal behaviour comes from an experiment with woodpeckers. In this experiment, woodpeckers had to determine how many holes in a tree to try for food before accepting that that tree did not contain any food and move on to the next one. The average number of holes tried by the woodpeckers was remarkably close to the correct number calculated with advanced mathematics and statistics [29]. This shows that evolution is capable of solving difficult optimisation problems in its attempt to increase our chances of survival.

2.2 Fitness Does not Equal Utility

All is well with maximising utility when animals are in simple situations where no strategic interactions need to be taken into account. We can say that in these cases, the animal is playing a game on its own in which maximising utility equates to maximising fitness. Things change, however, when there are multiple players in the game and interdependencies and externalities arise. In those cases, it may be advantageous to back away from the most rational decision, from the strategy which directly maximises utility [4]. This is because there are some games in which choosing an irrational, sub-optimal strategy may induce a behaviour from the other players which eventually leads to a higher fitness for yourself. We have already encountered an example of such a situation. In the case of loss-aversion, Huck et al. had posited that in a simple Nash-bargaining game, players may obtain a higher pay-off if they value their property higher than would seem rational [10]. The reason for this was that such loss-aversion leads to toughness in bargaining. This toughness forces the opponent to accept a lower pay-off for himself and, therefore, leads to a higher fitness for the biased player. In such cases with strategic interaction, therefore, maximising a utility function which does not exactly equate to the fitness function, may well end up maximising fitness.

3 Probability Weighting

3.1 Replication of the Model

We have talked about evolutionary psychology, and its aim to fundamentally understand economic behaviour, in very general terms. Now it is time to investigate a specific model which tries to explain the probability weighting of Prospect Theory from an evolutionary perspective. The following model is based on Rieger’s paper ‘Evolutionary Stability of Prospect Theory Preferences’ which itself drew heavily from Goeree et al. [21, 5]. We combine the model with works by Guth and Von Widekind and some original ideas [8, 7, 28].

We will consider a 2 by 2 game matrix which resembles the matching pennies game. Arguments for why these particular pay-offs are reasonable and natural will be presented below, in Section 5.1. However, we can already build on the intuition behind the general dynamic of such a matching pennies game. Goeree et al. gives the example of a football goalie who would prefer to dive in the same direction as the ball is shot in, while the opponent prefers to shoot in a different part of the goal than where the goalie dives to. In nature, this type of situation may also have occurred regularly.

	B1	B2
A1	2	3
	4	-2
A2	2	0
	3	0

Table 1: Pay-off matrix of the generalised matching pennies game.

In our model, there are two levels, following what is known as the indirect evolutionary approach. The lowest level contains the actually pay-offs to the players in terms of biological fitness. The pay-offs are some reward such as resources or the opportunity to mate and enhance the survival of the player. The level above contains the preferences over the reward. As was explained, in settings of strategic interaction, evolution may favour agents whose utility function does not perfectly trace its fitness function. We allow

various types of players in our model to have all possible preferences over the rewards and let natural selection decide which type has the most advantageous preference profile. It is assumed that, given their utility function, the various types act 'rational' in the sense that they try to maximise their utility. This need not coincide with the direct maximisation of fitness. However, it is the fitness which in the end determines which type will become dominant in the population as the mechanism of evolution is such that those with the highest fitness produce the largest number of offspring and will, therefore, become relatively more numerous.

To make this model more precise for our analysis, take the weighting function in Equation 1.1. The distinction between the types of people in our model is determined by the weighting variable γ . In the setting which we will describe below, the type with the optimal level of probability weighting will emerge dominant in the population. To allow for a richer variety of weighting, let us go beyond Rieger's model and extend this equation to the following:

$$\pi(p, \gamma, \lambda, \mu) = \frac{p^\gamma}{(p^\lambda + (1-p)^\lambda)^{1/\mu}} \quad (3.1)$$

In this case, the distinction between the types of people is determined by the weighting variable γ , as well as the variables λ and μ . The reason why the two exponents on the probabilities p and $1-p$ are the same is the symmetry in the function⁷. The four possible behaviours in terms of weighting probabilities are shown in Figure 4. What we then hope to find through our model is that only one of these possible behaviours survives the evolutionary process.




Figure 4: The four possible behaviours in terms of weighting probabilities: General overweighting, general underweighting, underweighting low probabilities while overweighting high ones and overweighting low probabilities while underweighting high ones.

The game described above has no pure strategy equilibria. It does have a mixed equilibrium, where both players A and B choose their first strategy with probability $\frac{2}{3}$ and their second with probability $\frac{1}{3}$. This is exactly what we need if we want to consider the effects of weighting on probabilities. According to the strategy chosen by their opponent, the players know they will end up in either one of two outcomes. We assume that the players weight the probabilities with which they believe their opponent chooses his strategies. In this case, the players will not end up in the objective mixed equilibrium, but in a different equilibrium, based on their level of probability weighting. To see precisely how the players do this, look at the following equation which determines the subjective utility a player gets from a particular outcome.

$$\begin{aligned} u_A(\sigma) &= \sum_{s \in S} \sigma_A(s_A) \pi(\sigma_B(s_B)) f_A(s) \\ &= 4p_A \pi(p_B) - 2p_A \pi(1-p_B) + 3(1-p_A) \pi(p_B) \end{aligned} \quad (3.2)$$

⁷By including even more variables in our equation, we could model higher order deviations from the linear, objective probability function. However, this level of complexity is unnecessary, especially seeing the fact that we already know what behaviour should come out of the model. Empirical data suggests the deviation is of order two, with overweighting of low probabilities and underweighting of high ones.

In this equation, the utility of A is determined by his own mix of strategies σ_A as well as the mix of strategies by player B, σ_B . These latter probabilities are weighted using the weighting function $\pi(p)$. As such, the utility one feels from a particular strategy is not necessarily equal to the fitness one gets from that outcome. Whereas the fitness of an outcome is determined by the rational expected utility theory⁸, using objective probabilities, the utility deviates from this following the players beliefs due to probability weighting.

Just as there are two levels in terms of fitness and utility, there are also two levels in this model related to the speed at which the game is played and at which evolution selects the optimal preference profile. We assume that the game is played much faster than the speed at which evolution operates [21]. That means that, from the point of evolution, the game is always in its equilibrium. This is an adiabatic approximation. There are two ways to justify the idea that the game will always be in its equilibrium position. First, we can say that the players repeat this game several time over, such that they can learn quickly what type their opponent is. As such, they will know what mixed strategy is optimal for them. Particularly, the opponent will know how to adjust his mixing of strategies so as to compensate for any 'irrational' behaviour by the other player. Alternatively, we can follow Von Widekind and Dekel et al. by assuming that the type of the players is known to the other player, i.e. one can observe the type and play the optimal strategy immediately [28, 3]. For the first of these two alternatives, the dynamics of the strategy space can be analysed⁹. The replicator equation was originally conceived to describe the dynamics of populations under natural selection, however, it was soon realised that the same equation may also describe a learning process.




Figure 5: Strategy spaces with replicator dynamics as a learning process in the generalised matching pennies game. In the first figure, both players are perfectly rational. In the second figure, player A applies probability weighting.

In Figure 5a above, we see the strategy space of the two players and the typical matching pennies behaviour around the mixed equilibrium. In this objective equilibrium, both players obtain a fitness of 2. If player A, however, weights the probabilities with which he expects player B to choose between his strategies, the position of the mixed equilibrium changes. Goeree et al. applied this idea to experimental data and found that it fit the data well [5]. In Figure 5b we see the dynamics change slightly when player A weights¹⁰ with $\gamma = 0.5$. The mixed equilibrium has now shifted, specifically towards to outer sides of

⁸Or should we say, expected fitness theory?

⁹The following graphs, as well as all the calculations done in this paper, were produced using a *Mathematica* notebook.

¹⁰As we shall see later on, the values of λ and μ are irrelevant for this result.

the strategy space. The shift in the equilibrium is advantageous¹¹ for player A, who now obtains a fitness of 2.67. In fact, for the most general pay-off matrix¹², the subjective mixed equilibrium is positioned at:

$$p_A \rightarrow \frac{(b_3 - b_4)^{\frac{1}{\gamma}}}{(b_1 - b_2)^{\frac{1}{\gamma}} + (b_3 - b_4)^{\frac{1}{\gamma}}} \quad (3.3)$$

$$p_B \rightarrow \frac{(a_2 - a_4)^{\frac{1}{\kappa}}}{(a_1 - a_3)^{\frac{1}{\kappa}} + (a_2 - a_4)^{\frac{1}{\kappa}}} \quad (3.4)$$

Here, we call the level of probability weighting by player B κ and not γ to avoid confusion between the two different players. What is striking about this result, however, is that the variables λ and μ have literally dropped out of the equation¹³. This means we are able to see the effect of changing γ but still have no idea about optimal values for λ and μ . We see that a lower γ leads to a higher fitness. As such, there will be an evolutionary pressure towards a lower γ . We can cross out two of the general types of behaviour found in Figure 4, namely, those for which $\gamma > 1$. What remains possible in the model is a general overweighting of all probabilities or an overweighting of only low probabilities and underweighting of the high ones, depending on the precise values of λ and μ . The properties of this model which account for the increase in fitness will be analysed later, in the Section 4.1.

3.2 Extension of the Model

The original paper by Rieger shows that a population of probability weighters is evolutionarily stable, i.e. stable against the infiltration of a mutant who weights less [21]. We have shown the equivalent by demonstrating that a player with a higher level of probability weighting will obtain a higher fitness. Types which receive a higher fitness will grow in number in the population. By showing that the fitness will always increase for a higher degree of weighting, we have proven that no mutant with a lower level of weighting can infiltrate in the population and that the high weighting type will dominate in the long-run. Unlike Rieger, however, our argument is not based on the notion of group selection for which probability weighting is beneficial to the group as a whole¹⁴. We have shown that even the selfish gene would favour a higher level of weighting, because the shift in the position of the mixed equilibrium causes an increase in the fitness of the player. As such, one would expect the process of increasing probability weights to continue until $\gamma = 0$. Rieger argues, however, that such a jump away from the more rational, true probabilities would not be beneficial to an animal in the end as the class of games described above is not the only game he will encounter in Nature [21]. As mentioned, a utility function which does not perfectly trace the fitness function may be beneficial under strategic interactions, but will certainly not be beneficial in all situations, for example, where there is no strategic interaction. A realistic picture of evolution is a setting in which an animal plays several types of games, with and without strategic interactions, and applies a level of probability weighting which is optimal for such a combination of games. Figure 6 below sketches a simple version of this idea, in which a player either ends up playing the matching pennies game or ends up playing a simple game where applying correct probabilities is advantageous. An example of the latter game is a situation in which an animal has the choice between a nearby tree with average food and a further-away tree which may or may not have very good food. The application of correct probability estimates is very important here, as it will have an effect on the quality of the resources you can consume and, therefore, on the chances of survival.

¹¹In the more general model developed by Von Widekind, the author also finds that under the indirect evolutionary approach, the population becomes dominated by a type which always plays the strategy with the highest fitness level, and therefore the most efficient outcome, if the players can observe the opponent's type [28].

¹²One which has the following form: $\begin{pmatrix} a_1 & a_2 \\ a_3 & 0 \end{pmatrix}$ and $\begin{pmatrix} b_1 & b_2 \\ b_3 & 0 \end{pmatrix}$.

¹³The reason that λ and μ are irrelevant in the determination of the position of the mixed equilibrium is that the denominator is perfectly symmetric in p and $1-p$. What matters in the model is only the relative value of $\pi(p)$ and $\pi(1-p)$. As such, the entire denominator drops out.

¹⁴Rieger also applies his probability weights to the 'War of Attrition' game to show that it will lead to lower waiting times for both players, which is beneficial in terms of getting a higher pay-off [21]. However, this argument, as well, is based on group selection and does not hold. A mutant with a lower level of probability weighting would be able to wait a little bit longer than the other players and thereby obtain all the benefits. Although the pay-offs for the population as a whole is lower, this is the only type of population which is stable under true evolution at the gene-level.




Figure 6: Animals encounter various kinds of games. With density $[G_1]$, they will play the matching pennies sub-game, with density $[G_2]$, they play a simple, one-player game, in which it is important they judge the probability R accurately.

The extensive form above has particular pay-offs and a particular density of which game is played, but what is most important is the general notion that probability weighting may be advantageous in some situation while disadvantageous in others. As such, there is some level of probability weighting, which will be between $\gamma = 1$ and $\gamma = 0$, which is optimal for the environment as a whole. If we combine the two sub-games with 50-50 probability, the following fitness will be obtained for various levels of γ .




Figure 7: The two sub-games in our extension of Rieger's model, for which we have used weighting as in Equation 1.1, give a different response in terms of fitness to decreases in γ . Whereas in the matching pennies sub-game the fitness increases with lower levels of probability weighting, the fitness decreases for our simple, one-player game. A combination of the two, in this case with equal densities, shows what the total fitness will be and for what level of γ this reaches its maximum.

As can be seen, the pay-off of the matching pennies game increases when γ decreases, while for the other game, the benefit becomes lower. The total fitness obtained is highest for a γ somewhere in between 0 and 1. This gives us a notion of what kind of probability weighting is evolutionarily stable. We see that a population which does not weight at all may be infiltrated by a probability weighter with a lower γ because the person with the lower γ obtains a slightly higher fitness. Likewise, a population for which

$\gamma = 0$ is also not evolutionarily stable, as a lower degree of probability weighting leads to a higher fitness. The evolutionarily stable state of this model is found for values of γ somewhere in between these two extremes.

What is more important, however, is that this approach gives us a clear answer for the values of λ and μ . As these variable did not enter the equation of the position of the mixed equilibrium in our matching pennies game, they were, up until now, still undetermined. As such, we were not yet able to decide which of the four possible methods of weighting, described in Figure 4, would survive the evolutionary process. Using our extension, however, we obtain the optimal values of $\gamma = 0.13$, $\lambda = 0.31$ and $\mu = 0.80$. We can say that the only evolutionarily stable state of our model is given by these three values. The most important notion that this gives is that, like γ , λ and μ lie somewhere between 0 and 1. From the original four possibilities found within Equation 3.1, we had narrowed it down to two, for which $0 < \gamma < 1$ while λ and μ were undetermined, by applying Rieger's model. Now we are left with just one possibility following our extension. Only those types of people who overweight low probabilities and underweight high probabilities will dominate in the long-run.

4 Hyperbolic Discounting

4.1 Appropriate Properties

We have seen that there are some similarities between probability weighing and hyperbolic discounting by showing that someone who applies weights to the probability of the realisation of a reward in the future will end up with a hyperbolic discount function. We have also seen one particular model which can explain the existence of probability weighting from an evolutionary perspective. If it is true that the ultimate cause of this probability weighting is also responsible for the hyperbolic discount function, we should try to analyse what properties of the previous model exactly result in the increase in fitness for probability weighters.

The most important property which results in an increased fitness is one very much related to a credible threat. Imagine the two players start in the objective mixed equilibrium of the game. If player A is a probability weighter, he will misjudge the action of his opponent B and believe it to be appropriate to move downwards on his strategy space. He will feel that it is optimal to choose his strategy A2 more often. In reality, however, this is irrational as it would result in a lower fitness. To be more precise, the strategy A2 is not the only best-reply to the strategy of player B. However, as there is interdependency between the players, this move by A will cause B to also shift his strategy mix. He will respond to the threat by A by moving more closely to his first strategy, B1. This shifts the mixed equilibrium towards the outer side of the strategy space and results in a higher pay-off to player A. Take the example of the goalie. By choosing one corner more frequently, he induces a reaction from his opponent to shoot in the other corner more often. If the pay-offs are such that this response actually benefits the goalie, his initial 'irrational' behaviour ends up benefitting him.

At the position of this new mixed equilibrium, the equilibrium in the higher-level, subjective game, it would be rational for player A to choose his strategy A1 more often. However, for the same reason, player A misjudges his opponent's actions and decides not to move away from his mix of strategies. Irrationally, he decides to continue mixing his strategies A1 and A2 with probabilities $\frac{2}{3}$ and $\frac{1}{3}$. This is ultimately a good thing, because otherwise player B would also need to adjust his strategy, resulting in a return to the original equilibrium in which player A obtained a lower fitness.

4.2 Inappropriate Properties

There are also properties of the model which are essential for the type of probability weighting we tried to explain earlier, but which will prove inappropriate for explaining hyperbolic discounting. As we shall see in the next section, a direct interpretation of this model as one of future discounting turns out to be impossible, or at the very least, unreasonable, due to these properties.

When player A applies probability weighting, he does so on the strategy of his opponent and on both p and $1 - p$. He sees that player B will choose between one of two alternatives, B1 or B2. There is a perfect symmetry between these two strategies. The pay-off for players A and B may not be the same, but whatever B chooses, A will receive some pay-off and B will receive some pay-off. If B mixes between his two strategies, player A will obtain the reward associated with strategy B1 discounted by the probability of him choosing that strategy. At the same time, he will obtain the reward associated with strategy B2, discounted by the probability of him choosing B2. As player B moves his strategy from B1 more towards B2, A will receive the reward associated with strategy B2 with an increasingly higher discount factor, until he receives it with certainty when B move all the way to the right. The important thing to take from this is that A will always obtain a reward and that it does not matter which part we call 'B1' and which 'B2'. There is a symmetry in the names, just as there is a symmetry in the probabilities: it does not matter whether we redefine $p \rightarrow 1 - p$, the model will still be the same and both probabilities will still be weighted, resulting in the same subjective probabilities and the same mixed equilibrium.

4.3 Direct Interpretation

To see what is so important about the symmetry between the names of the strategies, and likewise the symmetry between the probabilities p and $1 - p$, let us consider the same game as above, but reinterpreted to include future discounting. As we saw in the introduction, there may be an intimate connection between probability weighting and hyperbolic discounting following the similarity between discounting future rewards due to uncertainty about their realisation and discounting gambles due to the inherent uncertainty within the gamble. If the previous model is capable of explaining probability weighting from an evolutionary perspective, it may also be able to explain hyperbolic discounting if we simply reinterpret the variables of the model and the strategies of the players.

Initially, the probabilities p and $1 - p$ associated with mixed strategies were viewed as the chance of the realisation of the two pay-offs, determined by the actions of the players. Now, these variables should be viewed as the probabilities of the realisation of a reward after some period of time which is uncertain due to the hazard rate. As such, the axes on the strategy space now no longer represent mixed strategies, but are time variables. Players A and B each choose a time t_A and t_B within this reinterpretation of the game.




Figure 8: The strategy space of the model by Rieger, but reinterpreted such that the probabilities of a strategy being chosen are now probabilities of still being alive at some time t in the future.

The further into the future the reward will be available, the less likely it is that you will obtain it. As such, the reward is discounted according to the probability of realisation. If player A would obtain a fitness of 4 at (A1, B1), he would get $4 \times \delta^t$ if player B moves t steps to the right within his strategy space. This is not very different from the original model. If player B moved to the right, away from his strategy B1, the probability of B1 occurring was less and player A had to discount the associated pay-off more. However, to follow the model above accurately, we need to acknowledge that while the chance of still being alive to see the first reward realised decreases with time, the chance of being dead increases. The original interpretation of this was that when player B moved to the right, the probability of strategy B1 occurring may have decreased, but the probability of B2 occurring would subsequently increase. This means that, when we move to the right on the strategy space, the benefit of the first reward becomes lower, while we would need to consider the benefit of the second reward, the one associated with (A1, B2), to be increasing. In some way, we would need to associate a benefit to being dead which has an increasing present value further into the future, as the probability of being dead is higher there. We see that the symmetry between the two strategies B1 and B2 of player B is no longer valid in this case, because there is no natural symmetry between being dead and alive. Biological advantages can only be included in the fitness while being alive. There is no advantage to being dead and, therefore, we cannot give an interpretation to the pay-off originally associated with strategy B2 in this variation on Rieger's model. Likewise, we could not simply redefine $p \rightarrow 1 - p$ and expect nothing to change. If p is the probability of still being alive after some period of time and $1 - p$ is the probability of being dead, redefining these two terms would give vastly different results. A direct interpretation of the model for probability weighting in terms of future discounting seems impossible due to this property. If it were not impossible, the equivalent result as before would follow in that players benefit from hyperbolic discounting by its effect on the position of the equilibrium.

4.4 Indirect Interpretation

We have discussed above in what way the similarity between probability weighting and hyperbolic discounting can be explained from an evolutionary perspective. There were three alternatives, of which the first suggested that the ultimate cause for both effects was the same, Figure 3a. We can now cross out this possibility. What remains is that there are different ultimate causes for the effects and that only some of them overlap. This is not ruled out, though we can say with certainty that the particular model developed by Rieger to explain probability weighting is not appropriate for explaining hyperbolic discounting. There may be other models which explain not only probability weighting, but also hyperbolic discounting. Further research into whether other models for probability weighting can be reinterpreted as models of future discounting still needs to be done, but with such a striking similarity between the two effects, it is a promising approach. Finally, there was another possibility which explained the similarity between probability weighting and hyperbolic discounting. This was based on the idea that there were two distinct ultimate causes, two distinct models, but that they shared some properties which may lead to the effects. We have already discussed such a property: the credible threat.

In the case of loss-aversion, we saw that a certain level of toughness in bargaining may be advantageous as it shifted the equilibrium solution in favour of the tough player. As Huck et al. point out, the individuals behave sincerely and truly follow their preferences; they do not lie nor commit a non-credible-threat, they simply develop an innate sense of loss-aversion [10]. In the case of probability weighting, we saw that by misjudging the probabilities with which the opponent chose his strategies, the player was led to engage in a threat. This threat was credible because the player really did feel his subjective probabilities were correct, he was not merely pretending to obtain a better outcome. Can we come up with a situation in which hyperbolic discount may lead a player to put out a threat towards his opponent? A threat which in first instance may seem irrational because it would lead to a lower pay-off, but ultimately proves beneficial because of the adjustments made by the opponent.

Consider a situation similar to the ultimatum game. Two players own a particular good which gives them a benefit equal to $\frac{1}{4}$ on every time period. The present value of this good can be evaluated by integrating over time, discounted for the probability of dying in the future. However, it turns out that when both these goods are combined, their total benefit is equal to 1. One example could be that one player owns a long stick while the other owns a sharp blade. Used together, these goods form a very useful spear. As

such, it would be beneficial to both players if they could decide on a sharing plan. Player A would obtain the two goods first, for a certain amount of time t , before giving them both to player B. Player B would enjoy the benefit of the goods for the remainder of the time.




Figure 9: Both players own a good which benefits them in every period, conditional on them still being alive. If they share the goods by letting A have them both for a length of time t and B for the remainder of the time, they may both benefit more. Player B makes an offer of this length of time t which A can either accept or reject.

It is up to player B to decide on the time period t for which player A is allowed to hold the two goods. Player A merely has to accept the offer, or reject it. If we assume, for simplicity, that $\delta = 0.9$, the benefit to both players when A rejects the offer is 9.49. As such, if under the sharing plan A can get a higher benefit, he should accept the offer. When one follows the calculations, it seems that A should accept the offer whenever $t > 2.73$. Similarly, B would benefit from the exchange whenever he obtains the two goods before $t = 13.16$. To maximise his profit in this ultimatum game, B should offer $t = 2.73$, which A should accept. However, if A is not perfectly rational, if A discounts the future in a hyperbolic fashion, the result may be different. Let us assume that A follows a hyperbolic function such as in Figure 2 with a $\gamma = 0.6$. A is now of the opinion that he will not benefit from an exchange in which he is only allowed to use the two good for a time period of $t = 2.73$, so he rejects that offer. A is making a threat towards player B, one which is credible because A truly believes he has acted rationally. If we redo the calculation, it now seems that A would only accept offers from B where $t > 3.66$. As this length of time is still below the maximal level B would accept, B is very willing to make this offer. As such, the fact that player A has engaged in hyperbolic discounting has led him to receive a higher pay-off and obtain a higher biological fitness. Players such as A, who show hyperbolic discounting, will then increase within the population.

5 Limitations

5.1 What Games Are Played in Nature

We have now seen that the model by Rieger, with our extension, is very capable of explaining probability weighting. We have seen that there is only one type of weighting, for a particular γ , λ and μ , which gives the optimal fitness level. However, all of this is based on the idea that the type of game presented by Rieger was common in the history of our ancestors.

The idea of the matching pennies game is simple. In general, it may be true that one player benefits from a match, while the other player benefits from a mismatch. We presented the example by Goeree et al. of the football goalie above, who would prefer to dive in the same direction as the ball is shot in,

while the opponent would prefer to shoot in a different part of the goal than where the goalie dives to. Rieger explains his generalised matching pennies games as one of social control. In this game, player B needs to cooperate with player A, but also has the option of defecting. Player A has the task to check on B, but this checking comes at a cost. Plugging in some values for these four strategies according to the description, one ends up with the pay-off matrix in Table 1. Does this story sound reasonable? Would such situations really occur in Nature?

We had already discussed that humans and other animals encounter various types of games, some with strategic interaction and some without. We created a model in which Nature chooses which game we play with a certain probability as in Figure 6. To be more realistic, we would need to extend this model by including a wide variety of sub-games; not only 2 by 2 games, but even some with many more strategies and many more players. The appropriate density on each of the sub-games would then allow us to determine the optimal level of probability weighting. This is, of course, a daunting task and our simple model may be very close to the reality of life in the prehistory. At the very least, it illustrates the principle. However, it remains a strong simplification which suggests we should take the final results of our analysis with a grain of salt.

5.2 Multiple Self

There have also been experiments conducted on future discounting in neuroscience [16]. These suggest that there are two separate systems in the brain which are responsible for short-term and long-term discounting. The notion that there are multiple selves which deal with various calculations of the present value of a future reward does not agree with our idea that the discount function is hyperbolic due to probability weighting. What the ultimate cause is of this division between the short-term and the long-term system is yet to be discovered and whether hyperbolic discounting follows from this division may be too soon to tell, but this discovery certainly casts doubts on the conclusions of our model.

5.3 Learning Quick Enough or Observing

Another aspect of our model which may be regarded as a weakness is its assumption that the two players in the matching pennies game will know what type their opponent is. We have suggested two mechanisms for this; namely that the two players repeat the game several times such that they learn how their opponent plays and what their own optimal strategy is, or that the two players simply immediately observe the type of their opponent. The first of these two mechanisms has the limitation that it is only applicable to a subset of games played in Nature which are indeed repeated several times over. This reduces the likeliness that the game we have been considering is truly a reasonable game to hold as representative of situations in our evolutionary past. The second of these two mechanisms has the limitation that it requires the two players to be able to observe and identify various types within the population. This is in line with some previous work [8, 7]. But would humans and animals really be capable of doing this? This is another caveat in our model.

5.4 Prospect Theory Is not Cumulative Prospect Theory

Since the development of Prospect Theory, many advances have been made. Most notably, the idea of probability weighting has shifted from a direct mapping of objective probabilities to subjective probabilities to an indirect mapping through the cumulative distribution function [27]. This new approach changes the calculations of our analysis somewhat. Although we have not delved into this matter here, Rieger did apply Cumulative Prospect Theory to his research and found that the discrepancy between the two approaches was only minor [21]. However, whether this discrepancy is also minor when we apply the new approach to future discounting remains to be seen in future work.

5.5 Inflection Point

One last limitation of our extension of the model by Rieger is that it does not explain the position of the inflection point of the subjective weighting function. Various experiments have shown that the point at which the weighting function $\pi(p, \gamma)$ crosses the true probability line is well below 0.5 and that the

area of underweighting high probabilities is much wider than the area of overweighting low probabilities. In our model, the only behaviour that comes out includes an inflection point which is very close¹⁵ to $\pi(p, \gamma) = p = 0.5$. Why it is evolutionarily advantageous to have a weighting function with an inflection point at a lower p remains to be investigated by further research.

6 Conclusion

6.1 Achievements of Evolutionary Psychology

In this paper, we set out to give an evolutionary explanation for both probability weighting and hyperbolic discounting. We have presented a model, mostly based on previous work by Rieger, which shows that people who weight probabilities in the manner suggested by Prospect Theory will achieve a greater biological fitness than people who weight less. As such, there is an evolutionary pressure towards greater levels of weighting. Through our extension, however, we show that at some moment, there will also be a counter-pressure which results in an evolutionarily stable equilibrium position in which a particular level of weighting is applied.

Initially, we had established that there was a striking similarity between probability weighting and hyperbolic discounting. This similarity both made intuitive sense and had the additional feature that it was capable of explaining the strange behaviour of dynamic inconsistent preference reversals. Following our model on probability weighting, we made a suggestion of how to reinterpret this model to become one of future discounting in which hyperbolic discounting became evolutionarily advantageous. However, particular properties of the model by Rieger did not allow for a direct interpretation and we were forced to create a new model which merely presented one aspect of hyperbolic discounting, namely its relation to a credible threat, which could cause it to indeed be evolutionarily advantageous.

To sum up, let us look at the most important points that this paper hopes to contribute, which come in threefold. Firstly, the similarity between hyperbolic discounting and probability weighting, together with the fact that the idea of an intimate connection can truly explain preference reversals, are key. Secondly, in cases where animals play games both with and without strategic interaction, it can be advantageous to weight probabilities in a particular manner, namely for which the three variables γ , λ and μ are between 0 and 1. Finally, a direct interpretation linking a model of probability weighting with a model for hyperbolic discounting may be possible, but cannot be found in the matching pennies game suggested by Rieger.

One of the predictions of the idea that probability weighting and hyperbolic discounting stem from the same ultimate cause is that the two effects should always be observed simultaneously. This means that we can falsify this idea by finding evidence in the biological world of animals who either weight probabilities or discount hyperbolically, but not both. Further research needs to be done to not only check this claim, but also to deal with the various limitations of our model.

Behavioural economics has spurred great enthusiasm among economists, not only because of its great successes in explaining economic data, but also because its accomplishments entail a grand sensation of truly understanding the essence of behaviour from the more fundamental psychological tendencies inside the human brain. Evolutionary psychology comes with a similar sensation as it uses the grandest, yet simplest theory ever conceived, evolution by means of natural selection, to understand the human brain from the most basic principles. From simplicity to complexity, evolutionary psychology establishes the link from the fields of physics, chemistry, biology and psychology down to the theory of economics and has the power to explain the rich economic behaviour we observe in the world today.

¹⁵In fact, for the values of λ and μ we obtained, the inflection point is located at $p = 0.51$.

Bibliography

- [1] Keith Chen, Venkat Lakshminarayanan, and Laurie Santos. How basic are behavioral biases? evidence from capuchin-monkey trading behavior. *Journal of Political Economy*, pages 517--537, 2006.
- [2] Partha Dasgupta and Eric Maskin. Uncertainty and hyperbolic discounting. *The American Economic Review*, pages 1290--1299, 2005.
- [3] Eddie Dekel, Jeffrey C. Ely, and Okan Yilankaya. Evolution of preferences. *Review Of Economic Studies*, pages 685--704, 2007.
- [4] Robert H. Frank. Cooperation through emotional commitment. In *Passions Within Reason - The Strategic Role Of The Emotions*. Russell Sage Foundation, 1988.
- [5] Jacob K. Goeree, Charles A. Holtb, and Thomas R. Palfrey. Risk averse behavior in generalized matching pennies games. *Games And Economic Behavior*, pages 97--113, 2003.
- [6] Leonard Green, Astrid F. Fry, and Joel Myerson. Discounting of delayed rewards: A life-span comparison. *American Psychological Society*, pages 33--36, 1994.
- [7] W. Guth. An evolutionary approach to explaining cooperative behavior by reciprocal incentives. *International Journal Of Game Theory*, pages 323--344, 1995.
- [8] W. Guth and M. E. Yaari. An evolutionary approach to explain reciprocal behavior in a simple strategic game. In *Explaining Process And Change: Approaches In Evolutionary Economics*. The University Of Michigan Press, 1992.
- [9] Aviad Heifetz and Ella Segev. The evolutionary role of toughness in bargaining. *Games And Economic Behavior*, pages 117--134, 2004.
- [10] Steffen Huck, Georg Kirchsteiger, and Jorg Oechssler. Learning to like what you have - explaining the endowment effect. *The Economic Journal*, pages 689--702, 2005.
- [11] John H. Kagel, Donn. Macdonald, and Raymond C. Batialio. Tests of fanning out of indifference curves: Results from animal and human experiments. *The American Economic Review*, pages 912--921, 1990.
- [12] Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Experimental tests of the endowment effect and the coase theorem. *Journal Of Political Economy*, pages 1325--1348, 1990.
- [13] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, pages 263--291, 1979.
- [14] Daniel Kahneman and Amos Tversky. Choices, values and frames. *American Psychologist*, pages 341--350, 1984.
- [15] Laurent Lehmann and Francois Rousset. The evolution of social discounting in hierarchically clustered populations. *Molecular Ecology*, pages 1--25, 2011.
- [16] Samuel M. McClure, David I. Laibson, George Loewenstein, and Jonathan D. Cohen. Separate neural systems value immediate and delayed monetary rewards. *Science*, pages 503--507, 2004.
- [17] Rose McDermott, James H. Fowler, and Oleg Smirnov. On the evolutionary origin of prospect theory preferences. *The Journal Of Politics*, pages 335--350, 2008.
- [18] Nick Netzer. Evolution of time preferences and attitudes toward risk. *The American Economic Review*, pages 937--955, 2009.
- [19] Howard Rachlin, Andres Raineri, and David Cross. Subjective probability and delay. *Journal Of The Experimental Analysis Of Behavior*, pages 233--244, 1991.

- [20] Jerry B. Richards, Suzanne H. Mitchell, Harriet De Wit, and Lewis S. Seiden. Determination of discount functions in rats with an adjusting-amount procedure. *Journal Of The Experimental Analysis Of Behavior*, pages 353--366, 1997.
- [21] Marc Oliver Rieger. Evolutionary stability of prospect theory preferences. 2009.
- [22] Arthur J. Robson and Larry Samuelson. The evolution of time preference with aggregate uncertainty. *American Economic Review*, pages 1925--1953, 2009.
- [23] Alan R. Rogers. Evolution of time preference by natural selection. *The American Economic Review*, pages 460--481, 1994.
- [24] Peter D. Sozou. On hyperbolic discounting and uncertain hazard rates. *Proceedings Of The Royal Society B*, pages 2015--2020, 1998.
- [25] Peter D. Sozou and Robert M. Seymour. Augmented discounting: Interaction between ageing and time-preference behaviour. *Proceedings Of The Royal Society B*, pages 1047--1053, 2003.
- [26] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, pages 1124--1131, 1974.
- [27] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal Of Risk And Uncertainty*, pages 297--323, 1992.
- [28] Sven Von Widekind. *Evolution Of Non-expected Utility Preferences*. Springer, 2008.
- [29] Nick Wilkinson. *An Introduction To Behavioral Economics*. Palgrave MacMillan, 2008.