

Why humans care about sunk costs while (lower) animals don't. An evolutionary explanation

Felix Höffler¹

March 31, 2008.

¹WHU - Otto Beisheim School of Management, Burgplatz 2, 56179 Vallendar, Germany. ++49(0)261 6509 220. felix.hoeffler@whu.edu. Part of the work was done while the author was at the Max Planck Institute for Collective Goods, Bonn. Financial support by Deutsche Forschungsgemeinschaft through SFB-TR 15 is gratefully acknowledged. I would like to thank Tilman Börgers, Werner Güth, Hendrik Hakenes, Martin Hellwig and Georg Nöldeke for stimulating discussions, as well as seminar participants at the Max Planck Institute in Jena, University of Munich, Tilburg University, and at the SFB-TR 15 Meeting, Fall 2005, in Tutzing. Francis Ratnieks introduced me to the biological literature on the Concorde fallacy. Angelika Praus provided excellent research assistance. All remaining errors are mine.

Abstract

While humans often care about sunk investment, animals are not subject to this sort of sunk cost behavior or "Concorde fallacy". This paper analyzes sunk cost behavior in a simple two stage decision problem under uncertainty. At the second stage, subjects can commit the Concorde fallacy by sticking to the first stage decision, independent of the state of nature revealed in-between. We investigate whether this can be beneficial in a standard payoff monotonic adaptation process. Committing the Concorde fallacy reduces the payoffs but accelerates the adaptation since it acts like "self-punishment". It will, however, not only reduce the population growth rate in the long run but also the population size at any point in time in a biological evolutionary process. In this sense, animals can never benefit from the Concorde fallacy. By contrast, risk averse humans using a simple learning rule can increase their expected utility by committing the Concorde fallacy. Risk aversion gives an extra benefit to a behavior that more rapidly learns to avoid bad outcomes. If the wrong initial decision leads occasionally, albeit very infrequently, to a very low payoff, then risk averse humans will be better off by committing the Concorde fallacy.

JEL-Classification: C73, D81, D83

Keywords: Concorde fallacy, Sunk costs, Evolutionary game theory, Replicator Dynamics, Risk aversion

1 Introduction

The "Concorde fallacy" and "sunk cost effects" refer to the same problem of suboptimally attributing attention to past investments while only future, marginal effects should be taken into account. While the first term has been used by biologists in the context of animal behavior, the second term is familiar to economists and psychologists alike, referring to human behavior.

Everyday experience tells us that humans are vulnerable to sunk cost behavior: When we buy a ticket for the opera and – on the evening of the performance – recognize that there is an interesting football match on TV, which we would actually prefer to watch, we feel somehow obliged to go to the opera in order to avoid the sense that we have "wasted" the money on the opera ticket. This sort of behavior is well documented in many psychological experiments.

Sunk cost behavior also plays a crucial role in many economic contexts, in particular with respect to follow-up investment decisions. The term "Concorde fallacy" refers to the project of building the "Concorde", the first civil supersonic aircraft. It became obvious quite early that it would not be a commercial success (Woolley (1972), Edwards (1973)). However, a strong argument to finalize the project was that otherwise the sunk investment would have been in vain.

In contrast to humans, animals seem less prone to such fallacies. Biologists' have investigated many species and have not found a clear incidence where animals commit the Concorde fallacy. This has given rise to the question: "Are humans less rational than lower animals?" the title of a review article by Arkes and Ayton (1999) comparing experimental evidence for humans and animals.

The current paper tries to provide an evolutionary interpretation of why animals are less likely to care about sunk cost than humans. We take up the psychological idea of "overgeneralization": Humans apply a certain rule (Don't give up too early!) – which often leads to optimal decisions – sometimes in situations where this is not adequate. Caring for sunk cost can be beneficial if it helps to get important decisions right at the cost of errors in less important situations.

This idea is formalized using a simple two stage decision problem. At stage one an individual has to choose between two projects. At stage two she can either carry on or stop the project. Which of the two is optimal depends on the state of the world which is revealed between stage one and stage two. The Concorde fallacy

consists of carrying on independent of the state of the world. Most of our analysis focuses on the simplest situation, in which one of the two alternatives always yields higher payoffs, provided an optimal stage two decision.

We investigate how the right initial project choice is "learned" in a standard adaptation process, the replicator dynamics. We discuss two different applications of this process: (i) A standard population game of biological evolution; (ii) individual decision-making of a boundedly rational risk averse agent using a reinforcement learning rule. Both can be discussed using a similar replicator dynamics formulation.

The key effect is that committing the Concorde fallacy can serve as a self-punishment mechanism for wrong initial decisions: this can speed up the adaptation process. Thus, there is a general trade-off between accelerating the adaptation (getting the initial decision right) and the average payoff in each of the underlying decision problems (lower payoffs whenever one sticks to the initial decision although this is not optimal). For linear objective functions (e.g. the number of offspring in biological evolution), the latter effect always dominates. Therefore, the results are in line with the observations that animals are not subject to sunk cost behavior. For concave objective functions (e.g. utility of a risk averse human decision maker), committing the Concorde fallacy is always beneficial if only the worst of all possible outcome is sufficiently bad. Risk aversion gives a particular weight to avoiding bad outcomes, even if they are very infrequent. Committing the Concorde fallacy can help here by smoothing payoffs. It reduces payoffs on average but accelerates learning such that drastic mistakes become rare more rapidly.

The remainder of the paper is organized as follows. Section two provides an overview on the empirical results on the Concorde fallacy. Section three introduces the decision problem. Section four discusses the first application to the problem of maximizing fitness in a biological population context. Section five investigates the learning application. The results are discussed in the last section.

2 Empirical evidence for Concorde fallacy and sunk cost effects

There is extensive empirical biological research on the "Concorde fallacy". Though it impossible to prove the non-existence of an empirical phenomenon like the Concorde fallacy, it is fair to state that it is highly unlikely that animals commit the Concorde

fallacy. There is at least no clear cut example in which animal behavior could not equally well be explained in terms of "adaptive" behavior, i.e. in line with traditional explanations of maximizing fitness (see Curio (1987) and Arkes and Ayton (1999) for summaries of the experimental literature). A good example for this stems from the early discussion about the question: "Do digger wasps commit the Concorde fallacy?" the title of a paper by Dawkins and Brockman (1980). Female digger wasps lay a single egg in a burrow, and before doing so provide the nest with food (paralyzed katydids). Occasionally, a second wasp moves in and the two jointly provide for the nest. When the two meet at the nest, they fight for the nest. Dawkins and Brockman report that a wasp will fight more intensely the more katydids it has already stored in the nest. This could accord with a Concorde fallacy, since the wasp should care for the total number of katydids, not only the "own" katydids. Dawkins and Brockman, however, convincingly argue that this behavior need not be "maladaptive" if the wasps cannot count the number of katydids. Then their own effort of supplying the nest is the best estimator available to the wasps for the total number of katydids in the nest. It is in this sense "rational" to condition the intensity of the fight only on their own past investment in the nest.

Empirical results are different for humans. For humans, the "sunk cost" behavior is a phenomenon well-established by many experiments. Arkes and Ayton (1999) provide a summary of the literature and report a typical experiment: Season tickets for the Ohio University Theater were sold at three different price levels: \$15, \$13, and \$8. The allocation to the different price classes happened randomly, thus, there is no correlation between the willingness to attend the plays and the price levels to be expected. However, those with a higher price attended more plays. One explanation is the sunk cost effect: because they had invested more in the past, they felt more obliged to use the tickets.

An explanation offered by psychologists is the *overgeneralization of rules* (Arkes and Ayton, 1999, 592). Rules like "Don't waste resources", or "Don't give up too early" are important rules humans have to learn. Children shall eat up their meals, students should finish their studies and so on. This is, however, sometimes not optimal. Therefore, the application of a rule, which originally was meant to prevent the waste of resources, occasionally leads to suboptimal decisions because decision-makers cannot determine when the application of the rule is adequate and when not. Arkes and Ayton argue that this is an important reason why humans commit

the Concorde fallacy. They also cite evidence that the only experimental subjects reported to behave more rationally are children below the age of 6 (Arkes and Ayton, 1999, 598) – thus, it seems that we teach our children to become irrational and to commit the Concorde fallacy.

3 A sunk cost problem

Sunk cost problems occur in dynamic decision problems only. Figure 3 depicts the minimum structure of such a problem.




Figure 1: Time Structure of a Sunk Cost Problem

There are at least two consecutive, interrelated decision problems. The initial decision involves options that in some sense are associated with options of the follow-up decision. One typical example is that the initial decision is a project choice, while the second regards whether to continue with the project or to abandon it. In-between the initial decision and the follow-up decision, there is some payoff-relevant signal about the state of the world. An optimal follow-up decision will be conditioned upon this signal. Committing the Concorde fallacy means ignoring the signal and selecting the "associated option" in any event.¹

We want to analyze the simplest possible decision problem in which a sunk cost problem can occur. There are two alternative projects, a good project \mathcal{A} and a bad project \mathcal{B} . There are two states of the world, a good state θ_A and a bad state θ_B . To simplify matters, assume that after choosing a good project \mathcal{A} and after the realization of the good state, the game is over and a large payoff of A obtains. In the opposite case, after choosing the bad project \mathcal{B} and after the realization of the bad state of the world, again the game is over and a bad payoff B is realized. Cases

¹Game theory has some difficulties in grasping this dimension of the Concorde fallacy. When confronted with a sunk cost problem, most people can identify what the "associated" action is, e.g. carrying on with a project. In an extensive form game, however, the Concorde fallacy has only one property: it yields suboptimal payoffs. However, there may be many other suboptimal options which we would not consider as committing the Concorde fallacy. In the game theoretic sense, "associated" is just nomenclature.

in between are more interesting, since they allow for an additional action. If after the choice of \mathfrak{A} the bad state θ_B occurs, the player can either stop the project (s), which yields C , or continue (choose c), which gives $C - d$, $d > 0$. With the bad project \mathfrak{B} , the good state of the world θ_A provides an opportunity of correction (s), i.e. stopping the project, yielding D , or continuing (c) with it, resulting in $D - d$. Continuation of the project (i.e. choosing c) is interpreted as committing the Concorde fallacy. The payoff structure is depicted in Figure 2.




Figure 2: Payoff Structure

Assuming:

$$A > C \geq D > B \tag{1}$$

implies that \mathfrak{A} yields a higher payoff in both states of the world, provided an optimum follow-up decision is taken. In the Appendix, we relax this assumption and show that the main insights of the paper remain unchanged if \mathfrak{A} is better only on average.

The assumption $d > 0$ reflects the sunk cost fallacy. To investigate the effect of committing the sunk cost fallacy, we can think of a "competition" between two populations, where the first does not commit the sunk cost fallacy, while the second does commit it. Depending on the project choice and the state of the world, the former faces the payoff matrix 1, while the latter faces the payoff matrix 2:

Matrix 1			Matrix 2		
	\mathfrak{A}	\mathfrak{B}		\mathfrak{A}	\mathfrak{B}
θ_A	A	D	θ_A	A	$D - d$
θ_B	C	B	θ_B	$C - d$	B

Investigating the effect of the sunk cost fallacy is therefore equivalent to asking: Can decreasing the payoff by some $d > 0$ in those situations where stopping the project would be optimal increase the success of the second population? In the psychological terminology, committing the sunk cost fallacy would correspond to an overgeneralization of the rule "finish what you have started!".

More formally, we can compare the performance of populations (or individuals) using either payoff matrix 1 or matrix 2, by starting from matrix 1, and the investigating what happens if we introduce a deduction d .

4 Biological Evolution

4.1 Adaptation Process

Consider a population of animals where each animal is programmed to either play \mathfrak{A} or \mathfrak{B} . The population size N is large but finite. Each animal inherits its strategy from its parents. The payoff equals the number of offspring in each period and the objective function is just the number of offspring (i.e. the expected payoff). Call $p(t)$ the fraction of animals in the population which choose \mathfrak{A} . Using standard techniques (see Weibull (1995), 124, and Samuelson (1997), 65, and the Appendix for details) the discrete dynamics can be approximated by a continuous time formulation and this process can be described by the standard replicator dynamics, where the rate of change is increasing in the expected payoff difference $x - y$:²

$$\frac{dp}{dt} = p(t)(1 - p(t))(x - y), \text{ where} \quad (2)$$

where x is the expected payoff from choosing \mathfrak{A} and y is the expected payoff from choosing \mathfrak{B} :

$$\begin{aligned} x &= \alpha A + (1 - \alpha)(C - \delta) \\ y &= \alpha(D - \delta) + (1 - \alpha)B, \end{aligned}$$

²Our results can be generalized to more general forms of the adaptation process, i.e. to all functions where dp/dt is an increasing, concave and non-converging function of $(x - y)$. See Höfler (2006).

where $\delta = 0$ for those using payoff matrix 1, and $\delta = d > 0$ for those using payoff matrix 2.

The growth rate of a strategy is a monotone function of its payoff difference to the population average. The growth rate of the overall population is given by:

$$\frac{dN_t/dt}{N_t} = \pi(t) = p(t)x + (1 - p(t))y. \quad (3)$$

The solution to the ordinary differential equation (2) equals:

$$p(t) = \frac{e^{zt}}{e^{zt} + A_0}, \quad (4)$$

where $z = (x - y)$, i.e.:

$$z = [\alpha(A - (D - \delta)) + (1 - \alpha)((C - \delta) - B)]. \quad (5)$$

A_0 reflects the initial condition:

$$p(0) = \frac{1}{A_0 + 1}.$$

The equation of motion (2) has two obvious characteristics: First, $dp/dt > 0$ for $0 < p < 1$, since due to (1) $(x - y) > 0$. Evolution will drive out the suboptimal choice of \mathfrak{B} in the long run, since \mathfrak{B} yields a lower payoff.

Lemma 1 $p \rightarrow 1$ for $t \rightarrow \infty$ for the replicator dynamics (2) for $p(0) > 0$.

A second observation concerns the effect of committing the Concorde fallacy on the speed of adaptation. The payoff difference $(x - y)$ increases in δ only for $\alpha > \frac{1}{2}$. Thus, only if the "good" state θ_A is sufficiently likely, will committing the Concorde fallacy (i.e. increasing δ from zero to some positive level) increase the payoff difference and thereby it will also accelerate the adaptation process.³

Lemma 2 *Given assumption (1), the Concorde fallacy accelerates the evolutionary process for a given initial condition if and only if $\alpha > \frac{1}{2}$.*

³Note that "accelerates" does not mean that dp/dt increases by increasing d . This cannot be true for all values of t since p converges to 1 for all values of d . Thus, if the initial "speed" is larger for $\delta > 0$ in the beginning, it must be slower at some other point in time.

Proof: See Appendix.

This highlights a key feature of the underlying structure. As argued above, any sunk cost problem must involve at least two interrelated decision problems. Committing the Concorde fallacy means doing wrong with respect to the second problem, or generally, in some situations. The gain may consist in learning to get things right in other, potentially more important situations. The Concorde fallacy, however, does not always teach the right lesson. In our game, in state θ_A of the world it does: it reduces the payoff whenever the wrong first period decision had been taken, and thus subjects adapt faster to what to do in θ_A . In the state of the world θ_B , the sunk cost fallacy decreases the payoff of the correct first period choice \mathfrak{A} , therefore hindering the process towards the optimal decision. With $\alpha < \frac{1}{2}$, the latter occurs more often, thus, $\alpha > \frac{1}{2}$ is a necessary condition for a positive evolutionary effect of committing the sunk cost fallacy. Throughout the paper, $\alpha > \frac{1}{2}$ is assumed.

4.2 Biological evolution and the sunk cost problem

From the perspective of evolutionary biology, a typical question to ask in this context is: Can committing the Concorde fallacy increase the growth rate π , defined in (3), of a population in time T , in particular for large values of T ? The behavior with the highest growth rate will drive out any other behavior in the long run. Figure 3 illustrates that in the long run the growth rate π will be larger when a population does not commit the Concorde fallacy. In this example $A = 8 > C = D = 5 > B = 2$, $\delta = 3.5$ and $\alpha = 0.96$. The initial condition is $A_0 = 1$, i.e. at $t = 0$ half of the population chooses \mathfrak{A} , and the other half chooses \mathfrak{B} .

The thin line reflects the population growth rates of a behavior that avoids the Concorde fallacy. The bold line reflects the population growth rates of a population committing the Concorde fallacy, namely, earning $C - \delta = D - \delta = 1.5$ instead of $C = D = 5$. The population that commits the Concorde fallacy starts with a lower growth rate at $t = 0$. It then "overtakes" for some intermediate period, i.e. there is a time interval in which the population committing the Concorde fallacy grows faster. In the long run, the growth rates of both populations converge, but the growth rate of the Concorde fallacy population converges to a lower level.




Figure 3: Biological Evolution and the Concorde Fallacy

The properties of this example can be generalized as stated in the following proposition, where $\pi(t, \delta)$ denotes the growth rate at time t for a given level of δ .

Proposition 1 *Compared to $\delta = 0$, a deviation $\delta > 0$ from the correct second stage decision, will have the following effect on the growth rate $\pi(t)$: (i) The growth rate will be lower in the beginning, i.e. for t close to zero $\pi(t, \delta = 0) > \pi(t, \delta > 0)$. (ii) The growth rate will be lower in the long run: For $t \rightarrow \infty$ again $\pi(t, \delta = 0) > \pi(t, \delta > 0)$. (iii) For intermediate values of t , the growth rate will be higher for α sufficiently close to 1 and $(x - y)$ sufficiently large.*

Proof: See Appendix.

For extreme values of t , the intuition for the result is straightforward. In either extreme, the population states p are almost the same for both kinds of behavior. They start at the same point, and for large values of t both will be close to 1. If the probability of avoiding the wrong first decision is the same, committing the sunk cost fallacy must yield a lower payoff due to the second stage mistake. In-between, committing the Concorde fallacy might result in a sufficient advantage in the speed of learning so that committing the Concorde fallacy can provide a lead for some time interval. This can happen only if the good state of the world is highly likely, i.e. α is close to 1. As mentioned already, only in the good state of the world does the punishment effect of the Concorde fallacy drive evolution in the right direction. What the Concorde fallacy delivers is a faster transition to a high p . The prize to be gained from that is exactly the difference between x and y , the difference between the payoffs of selecting the right initial option \mathfrak{A} and the wrong initial option \mathfrak{B} .

Proposition 1 provides a clear answer to the biologist's question: In the long run, avoiding the Concorde fallacy yields higher growth rates. Thus, for evolutionary

processes where the relevant time horizon is sufficiently large, evolution will drive out the Concorde fallacy.

What about the higher intermediate growth rates? Not all biological processes might have enough time to reach the long-run state in which committing the Concorde fallacy again yields lower growth rates. Proposition 2 states that, although the *growth rate* might be higher at some points in time, the *population size* of the Concorde fallacy population will never be higher.

Proposition 2 *The population size of a population committing the Concorde fallacy (i.e. for all $\delta > 0$) will never exceed the population size of a population not committing the Concorde fallacy.*

Proof: See Appendix.

The higher evolutionary speed is not able to compensate for the lower reproduction rates in the beginning. Thus, in this framework, committing the Concorde fallacy is driven out by the biological evolutionary process, a result in conformity with the biological evidence that animals do not commit the Concorde fallacy.

5 Individual (reinforcement) learning of a risk averse individual

5.1 Adaptation Process

Arguably, the line between humans and "lower" animals is not always easy to draw. However, it might be fair to state that human behavior is to a larger extent driven by learning, compared to, e.g., digger wasps. A part of the difference in the observed behavior might therefore well stem from the differences of the adaptation process (that is, differences between learning and biological selection processes). However, we want to abstract from such differences here, which we can easily do by employing the replicator dynamics to describe also human learning processes.

As it is well known, many learning processes can be described by the replicator dynamics. One such process is reinforcement learning. Consider a boundedly rational player who repeatedly plays the game. She employs a simple learning rule which is made conditional on her own past payoff. She uses a mixed strategy, where $p(t)$ denotes the probability to select \mathfrak{A} at time t . If she played \mathfrak{A} and received a

payoff π , then she would reduce the deviation from selecting \mathfrak{A} with certainty, which is given by $1 - p(t)$, by the factor π/A . For \mathfrak{B} , vice versa. This learning process, up to a constant factor $1/A$, yields the same adaptation process as described by the replicator equation (2) (details can be found in the Appendix).

We want to consider human decision making as "boundedly rational" in the sense of Simon (Simon, 1997, 291). A boundedly rational agent may be well aware of his or her cognitive limitations in deciding in every situation which alternative is the best. The agent does not optimize but uses "Simple heuristics that make us smart" (Gigerenzer, Todd, and the ABC Research Group (1999)). To stick to the initial decision is a simple heuristic; but does the Concorde fallacy make us "smart", i.e. is it in some sense beneficial? In our context, this can be rephrased in the following form: Will the utility of a subject over its finite lifetime T increase when it ex-ante commits to the Concorde fallacy in order to speed up learning? Formally, this means that we want to compare the expected utility for the time period until T for the cases in which the Concorde fallacy is committed to those in which it is not.

Since payoffs in this context are not meant as offspring but as some sort of consumption (or monetary payoffs that allow for consumption), it is interesting to allow for risk aversion. We will therefore use a utility function $u(\cdot)$, with $u' > 0$ and $u'' < 0$. Since risk aversion will play an important role, comparative statics with respect to risk aversion are of particular interest. We therefore specify the utility function to exhibit constant absolute risk aversion γ for consumption of the quantity ζ :

$$u(\xi) = -e^{-\gamma\xi}, \quad \gamma > 0. \quad (6)$$

5.2 Learning of risk averse individuals and the sunk cost problem

Again, we first derive the result for not committing the Concorde fallacy and then ask whether switching to a behavior which reduces the payoff sometimes by $\delta > 0$ can increase the expected utility. The expected utility from not committing the concord fallacy is given by:

$$EU(\pi(t)) = \int_0^T \left(\frac{e^{zt}}{e^{zt} + A_0} u(x) + \frac{A_0}{e^{zt} + A_0} u(y) \right) dt, \quad (7)$$

This can be rewritten (see the proof of Proposition 2 for details) to get:

$$EU(\pi(p)) = \frac{u(x) \ln(e^{zT} + A_0) - u(y) \ln(A_0 e^{-zT} + 1) + (u(y) - u(x)) \ln(1 + A_0)}{z}. \quad (8)$$

We are interested in the question whether the expected utility can be increased by $\delta > 0$. Figure 4 shows a case where this actually happens for intermediate values of T . The parameters of the example are: $A = 1.2$, $C = D = 1$, $B = -1.8$, $d = 0.3$, i.e. in the case where switching in the second stage is optimal, committing the Concorde fallacy yields a lower payoff of 0.7 instead of 1. The probability of the state of the world θ_A is $\alpha = 0.96$, the individual starts with a 50-50 probability of selecting \mathfrak{A} or \mathfrak{B} , and it is risk averse with $\gamma = 1$. The properties of the functions are the same as in the biological application. Committing the Concorde fallacy (dashed line) yields lower payoffs in the beginning and in the end, while it might yield higher payoffs for an intermediate time period.




Figure 4: Expected utilities over time. Solid = avoiding Concorde fallacy; dashed = committing Concorde fallacy

However, analyzing the payoff difference yields a surprising result. Figure 5 depicts the difference between the two payoff functions of Figure 4 (the utility of not committing the Concorde fallacy minus the utility of committing the Concorde fallacy). Since the integral of this function equals the expected utility difference over time, it is obvious that for some T , committing the Concorde fallacy is worthwhile.




Figure 5: Difference in Expected Utilities

If the subject lives for $T = 20$, it can expect a utility of $EU^{NC} = -6.867$ when avoiding the Concorde fallacy, while it can expect $EU^C = -6.768$ if it commits the Concorde fallacy. Thus, the fact that occasionally committing the Concorde fallacy speeds up the learning process, eventually leading to fewer mistakes in the first decision, overcompensates for the lower utility whenever the second decision is taken incorrectly by 0.1. The following proposition states sufficient conditions for the occurrence of this beneficial effect from committing the Concorde fallacy.

Proposition 3 *Assume that, due to B being sufficiently small, the payoff difference $(x - y)$ is sufficiently large. Then, for any time interval T a decision-maker with constant absolute risk aversion γ can increase its expected utility by committing the Concorde fallacy if γ is sufficiently large, or if α is sufficiently large:*

$$\forall \gamma > \hat{\gamma} > 0 : \frac{\partial EU}{\partial \delta} > 0 \text{ for } B \rightarrow -\infty; \text{ and} \quad (9)$$

$$\exists \hat{\alpha} > \frac{1}{2} \text{ such that } \forall \alpha \in (\hat{\alpha}, 1) : \frac{\partial EU}{\partial \delta} > 0 \text{ for } B \rightarrow -\infty. \quad (10)$$

Proof: See Appendix.

Proposition 3 stands in strong contrast to Proposition 2. While for a biological application, the Concorde fallacy can never be regarded as superior, with humans we can always find conditions such that this is the case. Condition (9) states that the Concorde fallacy increase expected utility when there is a very low payoff possible, if only the decision maker is sufficiently risk averse (for all values of $\alpha > \frac{1}{2}$). Condition

(10) states that this will also happen for all positive levels of risk aversion, if only α is sufficiently close to unity, i.e. the bad outcome B is very unlikely.⁴

The key for explaining the difference in the results of Proposition 2 and Proposition 3 lies in the assumption of risk aversion. In both contexts committing the Concorde fallacy increases the evolutionary speed but reduces the average return. However, what is important for Proposition 3 is that learning to take the optimal initial decision more quickly also produces a more steady stream of payoffs, since it more rapidly leads to the avoidance of the bad payoff B . Risk aversion attributes an extra benefit to lower variance of payoffs which can compensate for the loss from a lower average return.

It can be shown (see the Appendix for details) that a similar logic applies if the other payoff of the wrong decision \mathfrak{B} is very small, i.e. $D \rightarrow -\infty$. Again, if only the decision maker is sufficiently risk averse, he or she can increase expected utility by committing the Concorde fallacy.

It is worth mentioning that the benefit from the Concorde fallacy stems from *avoiding very low payoffs* – not from increasing the probability of getting a far better alternative in general. Consider the case where the payoff difference $x - y$ becomes very large not due to the fact the y becomes very low but x becomes very large. Then, the Concorde fallacy can not be beneficial in the same way. The reason is that risk aversion treats very low payoffs differently from very high payoffs. The marginal utility from avoiding a low payoff increases the lower this payoff is. The marginal utility from increasing the probability to achieve a high payoff becomes smaller the larger this payoff gets, and might even go to zero. Thus, the marginal gain from the Concorde fallacy is too small to compensate for the reduction in expected payoffs.

5.3 Social (imitation) learning

A boundedly rational decision maker, who is aware of the fact that she uses reinforcement learning, can find it optimal to commit the sunk cost fallacy. An alternative interpretation would be the following. Imagine that individuals observe the "well being" (i.e. the utility level) of other individuals and tend to imitate the behavior of those better off. (This would reflect the situation that I imitate my neighbor if he appears to be happy, rather than that I imitate him if he has a higher salary).

⁴ $B \rightarrow -\infty$ implies that $y \rightarrow -\infty$, and therefore $z = (x - y) \rightarrow -\infty$. However, when inspecting the effect of $B \rightarrow -\infty$ on (7), one needs to take into account that z as well as $u(y)$ become large.

All individuals in the population use the reinforcement learning rule, while some commit the sunk cost fallacy, while others do not. Then those with the sunk cost fallacy would, under the conditions described in the last section, be better off and would therefore be imitated and would gain a majority in the population. This would require that the individual adaptation process is fast, compared to the social learning process.

6 Discussion

Whenever biological evolution aims at maximizing the expected payoff, e.g. the expected number of offspring, committing the Concorde fallacy can never be optimal. For boundedly rational, risk averse humans who want to maximize their expected utility it will be beneficial if only there is a bad outcome which is sufficiently infrequent or the level of risk aversion is sufficiently high. By choosing a simple rule like "stick to your initial decision", humans learn faster to avoid the bad outcome. This reduces the variance of payoffs at the cost of a lower expected payoff, which risk averse individuals find worthwhile.

In recent experimental studies on imitation learning, there is an interesting parallel to the sunk cost "puzzle" that humans seem to behave less rational than animals. Similar observations are made by Horner and Whiten (2005), who report from experiments with children and chimpanzees. Both had to open a box containing some reward. A human demonstrator opened the box by a series of actions, where some were necessary to open the box, while others were not. In a first treatment, the experimental subjects could not distinguish what was necessary and what was unnecessary. Both, children and chimpanzees, imitated all the demonstrator's actions to open the box. In an alternative treatment, the subjects could infer which were unnecessary actions. Chimpanzees copied only the necessary actions, while children still copied all of the demonstrator's actions, including the unnecessary ones. One possible explanation might be that humans are confronted with much more complex problems than animals, where causal relations are very difficult to discern. The more complex a problem the less likely it is to be successful by experimenting and trying to find the right solution alone. Thus, imitation learning is much more important for humans. This implies also a form of overgeneralization: Although sometimes it would be better to experiment on one's own, humans follow the general rule "imitate

successful behavior".

In some sense, our discussion fits well into this interpretation. Whenever the Concorde fallacy increases expected utility, the decision maker is confronted with a "difficult" problem: There is a very bad outcome, but it happens only very rarely. Since one sees the bad outcome rarely it might be difficult to grasp the causal relationship. Therefore, imitating might have an advantage over experimenting in this context. This advantage is, however, limited, since those experimenting with the wrong alternative (\mathfrak{B}) are not that bad off most of the time. Hence, the Concorde fallacy is very beneficial here by punishing the wrong choices of \mathfrak{B} even if it did not yield the very low outcome B .

Avoiding very bad outcomes is of particular importance if this could lead to an extinction of the characteristic, e.g. because it kills (all) hosts of a gene. A "selfish gene" (Dawkins (1989)) will want to avoid this risk and might therefore favor risk averse behavior. So far we have just assumed that humans are risk averse while lower animals are not. However, risk aversion among animals has long been discussed among biologists', see Philippi and Seger (1989) and Grafen (1999) on the issue of "bet hedging". Ever since Gillespie (1974) evolutionary biology has noted that in stochastic environments it is not good to "put all eggs in one basket". This suggests that animals in populations with a large number of individuals with the similar genes, e.g. social insects like ants or bees, are almost risk neutral in their individual behavior (since extinction of a gene is a very unlikely event), while others, where individuals differ more with respect to their genetic endowment, behave more risk averse. This could, however, offer the chance to test whether risk aversion plays a role. At first sight, this thinking harmonizes with our hypothesis: Humans differ more with respect to their genes than insects, birds, or fish do. And it is among these sorts of "lower animals" that the absence of the Concorde fallacy has been found.

Although our model does not capture formally the possibility of the extinction of a gene or strategy, it is worthwhile pointing to the close connection between the circumstance, in which it makes evolutionary sense to be risk averse and the circumstances in which committing the Concorde fallacy is beneficial. If in our model the outcome B is very bad, such that realizing B , say, three times in a row, leads to an extinction, the gene wants the agent to avoid this. Thus, it wants the agent to be risk averse. Committing the Concorde fallacy is an additional means

towards the same ends. Given that the agent is risk averse, he is "self interested" in committing the Concorde fallacy (since it provides a higher expected utility) and prevents the agent more quickly from earning the low payoffs B (which is in the interest of the "selfish gene").

Such a reasoning might apply not only to genetically programmed characteristics but also to behavior adapted by socialization or learning. If a certain behavior - although it yields high expected returns - can lead to the "extinction" of the one exhibiting it (e.g. because it implies a higher probability of bankruptcy), it can not be imitated in the future. Thus, only those who avoid such extreme outcomes quickly, can be imitated in the long run. Again, the Concorde fallacy works into this direction.

The Concorde fallacy shows that "throwing away" payoff can sometimes be beneficial in adaptation processes. However, the Concorde fallacy is only one very specific way of throwing away payoff. Further research would be of interest which tries to characterize other situations where giving up payoffs is beneficial - and how to do so optimally.

Appendix A: Proofs

Proof of Lemma 2: Compare two different situations with the same initial condition p_0 : In situation 1, in the underlying decision problem, subjects have $\delta = 0$; in situation 2 subjects receive less, due to $\delta = d > 0$. Both situations define evolutionary processes, $p(t, \delta = 0)$ and $p(t, \delta > 0)$, which cannot intersect: At $p(0)$, $p(t, \delta = 0)$ will have the larger growth rate if and only if $\alpha > \frac{1}{2}$. Now assume that at some $\hat{t} > 0$ the two functions intersect, which would imply that $p(\hat{t}, \delta = 0) = p(\hat{t}, \delta > 0)$ and $\frac{\partial p(\hat{t}, \delta = 0)}{\partial t} > \frac{\partial p(\hat{t}, \delta > 0)}{\partial t}$, which cannot be the case for $\alpha > \frac{1}{2}$. Thus, $p(t, \delta = 0) \leq p(t, \delta > 0) \forall t > 0$ for $\alpha > \frac{1}{2}$. ■

Proof of Proposition 1: The starting point is the growth rate π defined already in (3):

$$\pi = p(t)x + (1 - p(t))y,$$

where

$$\frac{\partial p}{\partial t} = p(1 - p)(x - y)$$

The term z is defined in (5), thus $\partial z/\partial \delta = 2\alpha - 1 > 0$. We know that

$$p(t) = \frac{e^{zt}}{e^{zt} + A_0} \text{ with } p(0) = \frac{1}{1 + A_0}.$$

Thus (x' and y' denote partial derivatives with respect to δ):

$$\begin{aligned} \frac{\partial \pi}{\partial \delta} &= \frac{\partial}{\partial \delta} \left(\frac{x e^{zt} + y A_0}{e^{zt} + A_0} \right) \\ &= \frac{x' e^{zt} + y' A_0}{e^{zt} + A_0} + \frac{A_0 e^{zt} t z' (x - y)}{(e^{zt} + A_0)^2}. \end{aligned} \quad (11)$$

In (11) the first term is negative while the second is positive, implying that $\partial \pi/\partial \delta > 0$ if and only if:

$$-(e^{zt} + A_0) (x' e^{zt} + y' A_0) < A_0 e^{zt} t z' (x - y). \quad (12)$$

For $t \rightarrow 0$, the right hand side of (12) goes to zero while the left hand side goes to $(A_0 + 1) ((1 - \alpha) + \alpha A_0) > 0$, proving claim (i) of the Proposition.

Rewrite (12) to get:

$$(1 - \alpha) \frac{e^{zt}}{t} \frac{1}{z' z} + \frac{A_0 + \frac{A_0^2}{e^{zt}}}{t z' z} < A_0, \quad (13)$$

and note that for $t \rightarrow \infty$, the first term on the left hand side goes to $+\infty$ and the second to zero, implying again a contradiction and proving claim (ii).

Finally, for $\alpha \rightarrow 1$ we find that (13) becomes:

$$\begin{aligned} -y' - y' \frac{A_0}{e^{zt}} &< t z' (x - y), \text{ where } y' = -\alpha = -1 \text{ for } \alpha \rightarrow 1, \\ \frac{e^{zt} + A_0}{e^{zt} t z'} &< x - y, \end{aligned}$$

which will be satisfied for $z = (x - y)$ sufficiently large, proving claim (iii). ■

Proof of Proposition 2: First note that:

$$\int_0^T \pi(t) dt = \int_0^T \frac{dN_t}{dt} \frac{1}{N_t} dt = [\ln N_t]_0^T = \ln N_T - \ln N_0.$$

Thus, the larger the integral from 0 to T , the larger the population size N_T in T , and

therefore, if increasing $\delta = 0$ to some $d > 0$ increases the integral, it also increases the population size:

$$\int_0^T \pi(\delta = 0, t) dt < \int_0^T \pi(\delta > 0, t) dt \Leftrightarrow N_T(\delta = 0) < N_T(\delta > 0).$$

Call

$$EU(T, \delta) = \int_0^T \pi(t, \delta) dt.$$

To prove the proposition, it is sufficient to show that

$$\frac{\partial EU}{\partial \delta} < 0.$$

The proof is by contradiction. Using (3) for π and the solution to the ordinary difference equation for $p(t)$ in (4) we can write:

$$\begin{aligned} EU(\pi(t)) &= \int_0^T \left(\frac{e^{zt}}{e^{zt} + A_0} x + \frac{A_0}{e^{zt} + A_0} y \right) dt \\ &= \int_0^T \left(\frac{e^{zt}}{e^{zt} + A_0} x + \frac{A_0 e^{-zt}}{1 + A_0 e^{-zt}} y \right) dt. \end{aligned}$$

In order to use that for $\gamma = \ln f(x)$ we have $\gamma' = \frac{f'(x)}{f(x)}$, this expression needs to be expanded by z and $-z$, respectively, to get:

$$EU = \frac{x}{z} \int_0^T \frac{ze^{zt}}{e^{zt} + A_0} dt + \frac{y}{-z} \int_0^T \frac{-zA_0 e^{-zt}}{1 + A_0 e^{-zt}} dt, \quad (14)$$

since

$$\begin{aligned} \frac{\partial}{\partial t} (e^{zt} + A_0) &= ze^{zt} \text{ and} \\ \frac{\partial}{\partial t} (1 + A_0 e^{-zt}) &= -zA_0 e^{-zt}, \end{aligned}$$

(14) becomes:

$$\begin{aligned}
EU &= \frac{x}{z} [\ln(e^{zt} + A_0)]_0^T - \frac{y}{z} [\ln(1 + A_0 e^{-zt})]_0^T \\
&= \frac{x}{z} [\ln(e^{zT} + A_0) - \ln(1 + A_0)] \\
&\quad - \frac{y}{z} [\ln(1 + A_0 e^{-zT}) - \ln(1 + A_0)] \\
&= \frac{x \ln(e^{zT} + A_0) - y \ln(A_0 e^{-zT} + 1) + (y - x) \ln(1 + A_0)}{z}. \tag{15}
\end{aligned}$$

Taking the derivative of (15) with respect to δ yields (where x', y', z' denote partial derivatives with respect to δ):

$$\begin{aligned}
\frac{\partial EU}{\partial C} &= -\frac{z'}{z^2} [x \ln(e^{zT} + A_0) - y \ln(A_0 e^{-zT} + 1) + (y - x) \ln(1 + A_0)] \\
&\quad + \frac{1}{z} \left[x' \ln(e^{zT} + A_0) + x \frac{z'T e^{zT}}{e^{zT} + A_0} - y' \ln(A_0 + 1) + y \frac{z'T A_0 e^{-zT}}{A_0 e^{-zT} + 1} + \ln(1 + A_0) (y' - x') \right] \\
&= \frac{1}{z^2} [\ln(e^{zT} + A_0) (x'z - z'x) + \ln(1 + A_0 e^{-zT}) (y'z - z'y) \\
&\quad + \ln(1 + A_0) (y'z - yz' + xz' - x'z) \\
&\quad + z'zT \frac{e^{zT}x + A_0y}{A_0 + e^{zT}}]. \tag{♠}
\end{aligned}$$

Straightforward calculations – using the definitions for x, y , and z – show that the term in line (♠) is zero. $\frac{\partial EU}{\partial \delta} > 0$ therefore requires:

$$\ln(e^{zT} + A_0) (x'z - z'x) + \ln(1 + A_0 e^{-zT}) (y'z - z'y) > -z'zT \frac{e^{zT}x + A_0y}{A_0 + e^{zT}}. \tag{16}$$

Since $z' = 2\alpha - 1 > 0$, both sides of the inequality are negative. And since $\ln(e^{zT} + A_0) \geq zT$, we obtain an upper bound for the left hand side in (16):

$$zT (x'z - z'x) > \ln(e^{zT} + A_0) (x'z - z'x) + \ln(1 + A_0 e^{-zT}) (y'z - z'y).$$

Given this upper bound, a necessary condition for (16) to hold is:

$$\begin{aligned}
zT(x'z - z'x) &> -z'zT \frac{e^{zT}x + A_0y}{A_0 + e^{zT}} \\
z'x \left(\frac{e^{zT}}{e^{zT} + A_0} - 1 \right) + x'z &> -z'y \frac{A_0}{e^{zT} + A_0} \\
z' \frac{A_0}{e^{zT} + A_0} (y - x) + x'z &> 0,
\end{aligned}$$

a contradiction, since $z' > 0$ and $y < x$ and $x' < 0$. ■

Proof of Proposition 3: Differentiating (7) with respect to δ yields:

$$\begin{aligned}
\frac{\partial EU}{\partial \delta} &= \int_0^T \frac{z'te^{zt}(e^{zt} + A_0) - z'te^{zt}e^{zt}}{(e^{zt} + A_0)^2} + u'(x) \frac{e^{zt}}{e^{zt} + A_0} - \frac{z'te^{zt}}{(e^{zt} + A_0)^2} u(y) + \frac{A_0}{e^{zt} + A_0} u'(y) dt \\
&= \int_0^T z' \frac{e^{zt}}{(e^{zt} + A_0)^2} t (A_0 u(x) - u(y)) + \frac{e^{zt} u'(x) + A_0 u'(y)}{e^{zt} + A_0} dt.
\end{aligned} \tag{17}$$

The second term is negative, while the first can have either sign. We want to show that (17) is positive for B sufficiently low for every $T > 0$. We show that the integral converges pointwise to infinity. Note that $B \rightarrow -\infty$ implies $z = x - y \rightarrow +\infty$. Therefore, the second term in the integral goes to a (negative) constant $u'(x) = -(1 - \alpha) \gamma e^{-\gamma(C-\delta)}$. The limit of the first term in the integral is not straightforward:

$$\begin{aligned}
&tz' \lim_{B \rightarrow -\infty} \frac{e^{zt}}{(e^{zt} + A_0)^2} (A_0 u(x) - u(y)), \\
&= tz' \left[\lim_{B \rightarrow -\infty} \frac{e^{zt} A_0 u(x)}{(e^{zt} + A_0)^2} + \lim_{B \rightarrow -\infty} \frac{-e^{zt} u(y)}{(e^{zt} + A_0)^2} \right].
\end{aligned}$$

The first term in the square brackets is negative (since $u(x) < 0$), but goes to zero for $B \rightarrow -\infty$. The second term is positive (since $u(y) < 0$), and both, nominator and denominator go to infinity for $B \rightarrow -\infty$. Applying l'Hospital's rule twice for

the second term yields:

$$\begin{aligned}
-tz' \lim_{B \rightarrow -\infty} \frac{e^{zt} u(y)}{(e^{zt} + A_0)^2} &= -tz' \lim_{B \rightarrow -\infty} \frac{e^{zt} (-\alpha e^{-\gamma(D-\delta)} - (1-\alpha) e^{-\gamma B})}{(e^{zt} + A_0)^2} \\
&= -tz' \lim_{B \rightarrow -\infty} \frac{t\alpha e^{-\gamma C} + [(1-\alpha)t + \gamma] e^{-\gamma B}}{2te^{zt} + 2tA_0} \\
&= tz' \lim_{B \rightarrow -\infty} \frac{\gamma [(1-\alpha)t + \gamma]}{2t^2(1-\alpha)} e^{-\gamma B - zt}.
\end{aligned}$$

Note that:

$$-\gamma B - zt = B(t(1-\alpha) - \gamma) - t[\alpha(A - (D-\delta)) + (1-\alpha)(C-\delta)],$$

which will go to $+\infty$ for $B \rightarrow -\infty$ if:

$$t(1-\alpha) < \gamma,$$

which holds for any t if α is close to 1 and/or γ is sufficiently large, as stated in the proposition. ■

References

- ARKES, H. R., AND P. AYTON (1999): “The Sunk Cost and Concorde Effects: Are Humans Less Rational Than Lower Animals?,” *Psychological Bulletin*, 125(5), 591–600.
- CURIO, E. (1987): “Animal Decision-Making and the Concorde Fallacy,” *Trends in ecology and evolution*, 2(6), 148–152.
- DAWKINS, R. (1989): *The Selfish Gene*. Oxford University Press, 2. edn.
- DAWKINS, R., AND H. J. BROCKMAN (1980): “Do Digger Wasps Commit the Concorde Fallacy?,” *Animal Behavior*, 28, 892–896.
- EDWARDS, C. (1973): “A Cost-Benefit Analysis of the Concorde Project: A Comment,” *Journal of Transport Economics and Policy*, 7(3), 300–301.
- GIGERENZER, G., P. M. TODD, AND THE ABC RESEARCH GROUP (1999): *Simple Heuristics That Make Us Smart*. Oxford University Press, New York.

- GILLESPIE, J. H. (1974): “Natural Selection for Within-Generation Variance in Offspring Number,” *Genetics*, 76, 601–606.
- GRAFEN, A. (1999): “Formal Darwinism, the Individual-as-Maximizing-Agent Analogy and Bet-Hedging,” *Proceedings of the Royal Society, London*, 266, 799–803.
- HÖFFLER, F. (2006): “Why Humans Care About Sunk Cost While (Lower) Animals Don’t. An Evolutionary Explanation.” mimeo.
- HORNER, V., AND A. WHITEN (2005): “Causal Knowledge and Imitation/Emulation Switching in Chimpanzees (*Pan Troglodytes*) and Children (*Homo Sapiens*),” *Animal Cognition*, 8, 164–181.
- PHILIPPI, T., AND J. SEGER (1989): “Hedging One’s Evolutionary Bets, Revisited,” *Trends in ecology and evolution*, 4(2), 41–44.
- SAMUELSON, L. (1997): *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge, MA.
- SIMON, H. A. (1997): *Models of Bounded Rationality*. MIT Press, Cambridge, MA.
- WEIBULL, J. W. (1995): *Evolutionary Game Theory*. MIT Press, Cambridge, MA.
- WOOLLEY, P. K. (1972): “A Cost-Benefit Analysis of the Concorde Project,” *Journal of Transport Economics and Policy*, 6(3), 224–239.

Appendix B: Derivation of replicator dynamics

This section derives explicit functional forms for the replicator dynamics for the two interpretations in the main body of the paper. As pointed out, the formulations are identical up to a constant factor $1/A$, which does not change any of the dynamic properties of the differential equation.

Biological evolution Consider a population of animals where each animal is programmed to either play \mathfrak{A} or \mathfrak{B} . The population size is denoted by N and is large but finite. Each animal inherits its strategy from its parents. The payoff equals the number of offspring in each period. We will refer to the payoff from playing \mathfrak{A} as $x = \alpha A + (1 - \alpha) C$, and to the payoff from playing \mathfrak{B} as $y = \alpha D + (1 - \alpha) B$. Call $p(t)$ the fraction of animals in the population which choose \mathfrak{A} . The fraction of \mathfrak{A} -types in $t + 1$ will be equal to:

$$p(t+1) = \frac{p(t)(1+x)}{p(t)(1+x) + (1-p(t))(1+y)}. \quad (18)$$

The total number of offspring at time t is given by:

$$N(t) \cdot [p(t)x + (1-p(t))y]. \quad (19)$$

Using standard techniques (see Weibull (1995), 124, and Samuelson (1997), 65) the discrete dynamics can be approximated by a continuous time formulation in which the growth rate of the population is given by:⁵

$$\frac{dN_t/dt}{N_t} = \pi(t) = p(t)x + (1-p(t))y, \quad (20)$$

⁵We have in mind an overlapping generations model in which, in each instant in time, reproduction and death takes place for a fraction r of animals per unit of time, where a unit of time is $\tau = 1/r$. For details, see e.g. (Weibull, 1995, 124). Thus:

$$\begin{aligned} \frac{dN}{dt} &= \lim_{\tau \rightarrow 0} \frac{N_{t+\tau} - N_t}{\tau} = \lim_{\tau \rightarrow 0} \frac{N_t + \tau px + \tau(1-p)y - N_t}{\tau} \\ &= N_t(px + (1-p)y), \end{aligned}$$

implying

$$\frac{dN_t/dt}{N_t} = (px + (1-p)y).$$

and the population share $p(t)$ can be approximated by the following expression:⁶

$$\begin{aligned}\frac{dp}{dt} &= p(t)(1-p(t))x - p(t)(1-p(t))y \\ &= p(t)(1-p(t))(x-y).\end{aligned}\tag{21}$$

Equation (21) reflects the standard replicator dynamics, where the growth rate of a strategy equals the difference of the strategy's payoff to the average population payoff:

$$\begin{aligned}\frac{dp}{dt} &= x - (px + (1-p)y) \\ p &= (1-p)(x-y) \Leftrightarrow (21)\end{aligned}$$

Individual (reinforcement) learning Consider a bounded rational player who repeatedly plays the game. She employs a simple learning rule which is made conditional on her own past payoff. She uses a mixed strategy, where $p(t)$ denotes the probability to select \mathfrak{A} at time t . If she played \mathfrak{A} and received a payoff π , then she would reduce the deviation from selecting \mathfrak{A} with certainty, which is given by $1-p(t)$, by the factor π/A . For \mathfrak{B} , vice versa. This payoff monotonic learning rules yields the following as an expected probability for playing \mathfrak{A} at time $t+1$:

$$\begin{aligned}p(t+1) &= p(t) + p(t)(1-p(t))\frac{\alpha A + (1-\alpha)C}{A} \\ &\quad - (1-p(t))p(t)\frac{\alpha C + (1-\alpha)B}{A}.\end{aligned}$$

When analyzing the payoffs in expected terms, the learning processes can thus be described by the same formalism, namely the replicator dynamics (21).

6

$$\begin{aligned}p(t+1) &= \frac{p(t) + p(t)x}{p(t)(1+x) + (1-p(t))(1+y)} \\ &= \frac{p(t) + p(t)x}{1 + p(t)x + (1-p(t))y} \\ p(t+\tau) &= \frac{p(t) + \tau p(t)x}{1 + \tau p(t)x + \tau(1-p(t))y} \\ \lim_{\tau \rightarrow 0} \frac{p(t+\tau) - p(t)}{\tau} &= \lim_{\tau \rightarrow 0} \frac{1}{\tau} \left(\frac{p(t) + \tau p(t)x}{1 + \tau p(t)x + \tau(1-p(t))y} - p(t) \right) \\ &= (21)\end{aligned}$$

Appendix C: Alternative Payoff Structures

In the main body of the paper we assumed $A > D$ and $C > B$. Alternatively, one might assume that $A > D$ but $C < B$, while still $(x - y) > 0$ and $u(x) - u(y) > 0$, i.e. \mathfrak{A} is still the better choice in terms of payoffs and of utility from payoffs, but only in expected terms. This would fit to an investment decision to sell either umbrellas (\mathfrak{B}) or ice cream (\mathfrak{A}). Umbrellas sell better under θ_B (bad weather), ice cream better under θ_A (sunshine). It might still be that ice cream is on average the better business, e.g. because sales of umbrellas are very bad when the sun is shining (D is very low).

First note that Proposition 1 and Proposition 2 still hold, since the proof uses only the expected payoff difference $(x - y)$. Proposition 3, however, requires some qualifications, since we discussed only the case where $(x - y)$ goes to infinity due to $B \rightarrow -\infty$. We also want to investigate the interesting cases where $(x - y)$ becomes large due to either $D \rightarrow -\infty$ or $A \rightarrow +\infty$.

Proposition 4 *Assume that, due to D being sufficiently small, the payoff difference $(x - y)$ is sufficiently large. Then, for any time interval T , a decision-maker with constant absolute risk aversion γ , can increase its expected utility by committing the Concorde fallacy if γ is sufficiently large:*

$$\forall \gamma > \hat{\gamma} > 0 : \frac{\partial EU}{\partial \delta} > 0 \text{ for } D \rightarrow -\infty.$$

Proof. The proof is along the lines of the proof of Proposition 3. We want to show the (17) is positive for D sufficiently small. We do so by showing that the integral converges pointwise to infinity for $D \rightarrow -\infty$ (z', x', y' denote partial derivatives with respect to δ):

$$\lim_{D \rightarrow -\infty} z' \frac{e^{zt}}{(e^{zt} + A_0)^2} t (A_0 u(x) - u(y)) + \frac{e^{zt} u'(x) + A_0 u'(y)}{e^{zt} + A_0}. \quad (22)$$

The second term in (22) goes to a (negative) constant $u'(x)$. The first term can be written as:

$$\lim_{D \rightarrow -\infty} \frac{e^{zt} (\alpha e^{-\gamma(D-\delta)} + (1-\alpha) e^{-\gamma B})}{e^{2zt} + 2A_0 e^{zt} + A_0^2},$$

applying l'Hospital's rule twice yields:

$$\begin{aligned}
& \lim_{D \rightarrow -\infty} \frac{e^{zt} (\alpha e^{-\gamma(D-\delta)} + (1-\alpha) e^{-\gamma B})}{e^{2zt} + 2A_0 e^{zt} + A_0^2} \\
= & \lim_{D \rightarrow -\infty} \frac{\alpha t (\alpha e^{-\gamma(D-\delta)} + (1-\alpha) e^{-\gamma B}) + \alpha \gamma e^{-\gamma(D-\delta)}}{2\alpha t e^{zt} + 2\alpha t A_0} \\
= & \lim_{D \rightarrow -\infty} \frac{\alpha t \gamma + \gamma^2}{2\alpha t^2} e^{-\gamma(D-d)-zt}. \tag{23}
\end{aligned}$$

Note that

$$-\gamma(D-d) - zt = -D(\gamma - t\alpha) + \gamma\delta - t(\alpha A + \alpha\delta + (1-\alpha)((C-\delta) - B)),$$

implying that (23) goes to $+\infty$ for $D \rightarrow -\infty$ for $\gamma > t\alpha$. ■

A similar logic as for Proposition 3 applies. If an individual is sufficiently risk averse it will benefit from committing the Concorde fallacy if there is a very low outcome which the individual will learn to avoid faster due to the Concorde fallacy. The condition is, however, more restrictive. It is no longer true that for any level of risk aversion committing the Concorde fallacy will be beneficial if only the occurrence of the bad outcome is sufficiently unlikely. The reason is that if the bad outcome D is realized, this will already have a strong learning impact. Reducing the payoff to $D - d$ has only a relatively small additional learning effect. Only very risk averse individuals will find this learning gain worthwhile compared to the loss of expected payoff.

Finally, the payoff difference might be large not due to the danger of very low payoffs but also due to the chance of very high payoffs. Proposition 5 shows that the effect that causes the Concorde fallacy to be beneficial is not symmetric in this sense.

Proposition 5 *If the payoff difference becomes infinitely large due to $A \rightarrow +\infty$, the expected utility EU can not be increased by committing the Concorde fallacy,*

$$\frac{\partial EU}{\partial \delta} < 0 \text{ for } (x - y) \rightarrow \infty \text{ due to } A \rightarrow \infty.$$

Proof. The proof is again along the lines of the proof of Proposition 3. For $A \rightarrow \infty$, the integral (17) converges pointwise to some negative constant. The second term in the integral of (17) goes to a negative constant $u'(x) = \partial u(x) / \partial \delta =$

$-\gamma(1 - \alpha) e^{-\gamma(C-\delta)} < 0$. Therefore, investigate:

$$\lim_{A \rightarrow \infty} tz' \left[\frac{e^{zt} A_0 u(x)}{(e^{zt} + A_0)^2} - \frac{e^{zt} u(y)}{(e^{zt} + A_0)^2} \right].$$

The second term in the square brackets goes to zero ($u(y)$ does not depend on A). The first term goes to zero as well; it can be written as:

$$\frac{A_0 u(x)}{e^{zt} + 2A_0 + \frac{A_0^2}{e^{zt}}},$$

where the numerator goes to zero while the denominator tends to infinity for $A \rightarrow \infty$.

■

The reason for this asymmetry stems from the two opposing effects from high payoff differences. High payoff differences imply a very fast learning. In the replicator dynamics the effect is exponential. Any further increase in learning speed must therefore have a very high marginal valuation to be worthwhile giving up some payoff d . If the payoff structure is such that there is a very bad payoff to be avoided, any risk averse individual attributes a very high utility to marginally decreasing the risk of this very bad outcome. These are the cases of Proposition 3 and 4. However, risk aversion is not "symmetric" in that it attributes the same marginal utility to increasing the probability of getting very high payoffs. The marginal utility from very high payoffs is bounded above for any concave utility function. Therefore, the Concorde fallacy cannot increase expected utility in this case.