Working Paper , pages 1–14
(draft)

Base rate neglect for the wealth of populations

Diemo Urbig

Humboldt-Universität zu Berlin

School of Business and Economics and Department of Computer Science

Abstract. Base rate neglect has been shown to be a very robust bias in human information processing. It has also been show to be ecologically rational, more precisely, the bias may be present since it does not make individuals significantly worse off in their environments. However, when arguing about base rate neglect usually isolated individuals are considered. I complement these results by showing that in many scenarios of social learning a base rate neglect can increase a population's wealth. I thereby strengthen the argument that the presence of base rate neglect can be evolutionary stable. I pick up a model of social learning that has been used to demonstrate the potential benefits of overconfidence. Individuals are confronted with a safe and a risky option with a high and low payoff. They receive a private signal about the risky option's outcome, they decide in an exogenously given sequence, and they observe decisions of preceding individuals. I first deviate from the original model by incorporating base rates that differ from fifty-fifty and show that under weighting this base rate is for the wealth of a population. In a second step I analyze how the optimal base rate neglect reacts to changes in payoffs that are associated with the safe and risky options. I show that for large set of settings under weighting the base rate is still positive, but for a smaller subset it decreases wealth instead.

Keywords: cognitive biases, base rate neglect, social learning, information cascades, ecological rationality

JEL codes: D83, D7

Address for corespondance: Diemo Urbig, Institute for Entrepreneurial Studies and Innovation Management, School of Business and Economics, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany, urbig@diemo.de

1. Introduction

Individuals placed in uncertain environments experience a large set of cognitive biases (Kahneman & Tversky, 2000a). Often individuals have some a priori information about the environment and during their interaction with the environment they receive additional cues and update their beliefs about some parts of the environments. Often this updating process is described by Bayesian updating, which combines a base rate with some additional signal to form a new updated probability judgment. In such scenarios many individuals experience a cognitive bias that is called base rate neglect or base rate fallacy. This cognitive biases causes people to underweight significantly the information coming across by the base rate (Kahneman & Tversky, 2000b). If people consider the base rate as having a causal nature or as being very specific and strongly related to the specific event under consideration, then the effect is not that strong (Kahneman & Tversky, 2000b). Also repetition of the same task and the incentive structure seem to be a moderating factor (Grether, 1980). Since base rate neglect causes individuals to deviate from the individually rational behavior, these individuals are usually worse off. However, Todd and Goodie (2002) show that in dynamic environments where base rates may change, which is very similar to being confronted with different choices with different base rates, base rate neglect does not reduce individually expected payoffs very much. This effect gives some explanation why individuals biased by a base rate neglect can survive in an environment where evolution favors the survival of the fittest. This fit of cognitive reasoning process and environment is called ecological rationality.

In this article I will complement the work on the ecological rationality of base rate neglect by looking at scenarios of social learning, where individuals not only deal with base rates and private information but also observe others' decisions. Being part of a population, where one can – at least partially – observe others is a type of environments that can reasonable be assumed to dominate human decision making. For our first approach to this issue I use the economic model of social learning as introduced by Bikhchandani et al. (1992). In this model another cognitive bias has already been investigated. Bernardo and Welch (2001) and Kariv (2004) show that overconfidence of all but also of some individuals can be, but does not need to be beneficial for the whole population. Kariv concludes that it is not clear in what types of environments overconfidence is beneficial. The positive effect of overconfidence for the population is mainly based on an information externality caused by biased individuals. However, Bernardo and Welch (2001) and Kariv (2004) only consider the case of fifty-fifty base rates. I will extend their analysis by considering different base rates as well as a different cognitive bias, namely the base rate neglect. I will also discuss briefly the relations that exist between overconfidence in own knowledge as modeled by the mentioned authors on the one side and base rate neglect on the other side.

The next section introduces our basic model. The third section will then show the welfare effects of social learning without any cognitive bias for different base rates. The fourth section identifies scenarios where base rate neglect is beneficial from a population's point of view. For a specific payoff structure I derive the optimal level of base rate neglect. The fifth section then describes how this optimal base rate neglect reacts to changes in the payoffs.

2. A social learning model (BHW model)

Consider individuals being confronted with a choice between a safe and a risky option. The risky option's payoff depends on the some hidden state of the environment V, which causes either high payoffs $V = H$ or low payoffs $V = L$. The individual has some a priori information $p_v = \mathbb{P}[V = H]$, but receives an additional signal X, which either signals a high payoff, $X = H$, or a low payoff, $X = L$, respectively. The signal is symmetric and it is correct only with probability $p_x = \mathbb{P}[X = V|V]$. Without loss of generality we assume $p_s > 0.5$. Individuals decide in an exogenously given sequence. Individual i can observe previous decisions of $i-1$ individuals about adopting (a) the risky options or rejecting (r) the risky option. These observations are summarized in the history $h_i \in \{a, r\}^{i-1}$. The safe payoff is given by S with $L < S < H$; without this assumption one of the two options would dominate the other without regard to the probabilities associated with the risky option's payoffs. The rational decision for this choice is given by condition 1. Since the mathematics behind this result are standard we just refer to Chamley (2004) for a more detailed introduction Bayesian reasoning in models of social learning and their basic properties.

$$0 \; < \; \underbrace{ln\left(\frac{\mathbb{P}\left[h_i|V=H\right]}{\mathbb{P}\left[h_i|V=L\right]}\right)}_{\text{history information } q_H} + \underbrace{ln\left(\frac{\mathbb{P}\left[X|V=H\right]}{\mathbb{P}\left[X|V=L\right]}\right)}_{\text{private information } q_X} + \underbrace{ln\left(\frac{\mathbb{P}\left[V=H\right]}{\mathbb{P}\left[V=L\right]}\right)}_{\text{a priori probability } q_A} + \underbrace{ln\left(\frac{H-S}{S-L}\right)}_{\text{payoff structure } s} \quad (1)$$

The information, q_h, coming across by the history we calculate based on our sequential model. For this we need to further specify our model, since the indifference case is not unambiguously specified. As indifference rule we apply a random choice between safe and risky option as Bernardo and Welch did as well. This rule is in fact equivalent to a random choice between following the own signal and not following the own signal. However, Figure 1 visualizes the effects of different indifference rules, which are discussed in the next subsection. Given that the private information makes a difference, the history information is given by equation 2 with d representing the difference of meaningful observations of adoptions and rejections without the possibility of indifference and d^- representing the same difference for observations where indifference is possible and following the history information is observed. The second term related to d^- we need because we applied another indifference rule than Bernardo and Welch. This term describes the situation that although an individual received a private signal that contradicts the history information the individual is indifferent. That possibility of such indifference is predictable by observers since they are confronted with the same history. In this situation following the herd reveals information because a different behavior was possible. However this information is less significant. If in a world with $q_A = s = 0$ an individual observes an adoption and receives a signal for low payoffs then she is indifferent. Rejecting in this situation reveals perfectly the negative private information. However, adopting does not reveal completely the information, i.e. $\frac{\mathbb{P}[adopt|V=H,h=\{a\}]}{\mathbb{P}[adopt|V=L,h=\{a\}]} = \frac{\mathbb{P}[X=H|V=H]+0.5\mathbb{P}[X=L|V=H]}{\mathbb{P}[X=H|V=L]+0.5\mathbb{P}[X=L|V=L]}$.

$$q_H = log\left(\left(\frac{p_x}{1-p_x}\right)^d \left(\frac{1+p_x}{2-p_x}\right)^{d^-}\right) \tag{2}$$

Most models of social learning make strong simplifications, e.g. $q_A = s = 0$, see for instance Bikhchandani et al. (1992), Anderson and Holt (1997), Hung and Plott (2001). They exclude the base rate from the model. Since we want to focus on a cognitive bias related to base rates we must keep it in the model. Nevertheless, we start our analysis with assuming $s = 0$. In this case our model is equivalent to models where individuals do not have a choice between a safe and a risky option depending on a hidden state of the environment (adoption game), but must only predict if the environment's state is high or low (prediction game). Examples for the first can be found in Bikhchandani et al. (1992), while examples for the second can be found in Anderson and Holt (1997). In a subsequent section we then relax the assumption of $s = 0$ and analyze the impact of non-neutral payoff structures.

3. Information revelation, social learning, and welfare effects

To analyze welfare effects we have to define welfare for a population of cognitively biased individuals. As welfare we define the sum of expected payoffs over all individuals in a population and follow thereby the work by Bernardo and Welch (2001). For this we assume that the biases considered do not affect the utility of individuals but only the quality of their decisions. For instance, if we assumed risk-avers individuals, which is a cognitive bias as well, then we should assume a risk utility function, which then does not allow to compare easily the welfare of two populations that differ in the risk-aversion of their members. The problem is close to the analysis of welfare consequences of preference changes (see Weizsäcker, 2005). However, the assumption simplifies our analysis to a significant extend; probably it is currently a necessary assumption to get interpretable results. Because we also assume an exogenous random sequence of individuals we can use the individually expected payoff as indicator for the wealth of a population. Similar to Bernardo and Welch (2001) we iteratively calculate the probabilities of specific histories and signals as well as expected payoffs associated with adoption and rejection in these situations. These numerical calculations are implemented in Java. Due to the problem of limited precision of numbers we do not apply equality operators but consider all numbers as equal that are within a range of $5 \cdot 10^{-15}$.

For a detailed analysis of the basic model without reference to welfare aspects we refer to Bikhchandani et al. (1992). An increasing quality of the private signal, p_x, increases the probability of correct decisions. If the history contains two more observations of adoptions or rejection then the private signal is ignored and all follow the herd, i.e. a cascade has emerged. We add some results related to relaxing the assumption of a fifty-fifty base rate, i.e. $p_v = 0.5 \rightarrow q_A = 0$. Figure 1 plots the expected payoffs for an individual against different base rates and as already mentioned for three different decision rules for indifference cases. For this analysis as well as for all following results we consider a population of $n = 50$. Most results derived for the BHW model and its derivatives do not change significantly for larger populations. Just to mention, our first numerical analyses we did for a smaller n and got qualitatively the same results.

Without any private signals individuals can only base their decisions on the base rate. Below $p_v = 0.5$ they go for the safe option while above $p_v = 0.5$ they choose the risky option. If individuals have private signals (in the example with $p_x = 0.7$) but can not observe others, then they base their decision only on the private signal and the base rate. In scenarios, where the base rate is not too extreme such that the signal can counter-balance the base rate, i.e. $1 - p_x < p_v < p_x$,

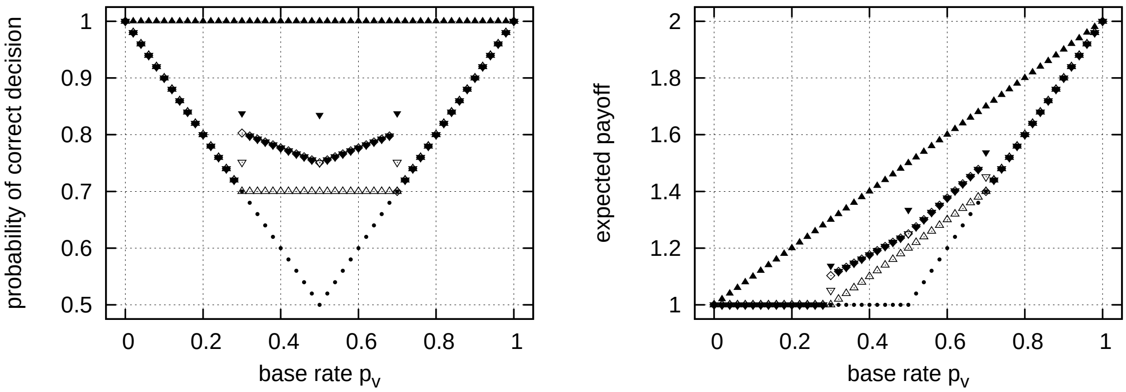

Figure 1. Probability of correct decision (left) and average payoff (right) for no information (•), for private information only (△), for perfect private information (▲), for private information and social learning with AR(0.5) (▽), for private information and social learning with XP(1.0) (▼), for private information and social learning with AR(1.0) (◊)

the individuals' expected payoffs increase due to better information. Only in these scenarios, where signals may influence decisions, observing others' decision can reveal information to followers. This information revelation finally let diffuse the private signals through a population, which reduces the risks of following individuals. Figure 1 also visualizes different indifference rules. We see that the rule "follow the own signal" (as applied in the work by Banerjee, 1992, Anderson & Holt, 1997) has some interesting discontinuities, which implied that increasing the base rate from 0.5 only by a very small amount reduces the expected payoff very much. The rule "adopt the risky option" obviously introduces an asymmetry, which favors adoption. Indifference rule "adopt with probability fifty-fifty" seems reasonably plausible. Nevertheless there are also two points where an increasing base rate decreases expected payoffs. However, these point describe limits, beyond which social learning does not occur and are thus not only artifacts. Bernardo and Welch (2001) allow indifferent individuals to abstain from a decision between risky and safe option, and thus give indifferent individuals the opportunity to perfectly reveal their private information. Due to the perfect information revelation for indifference cases and the same expected payoff for the individual, this specific indifference rule is equivalent to the rule ”follow the own signal”; and it has therefore the same discontinuities.

There is an interesting observation for the basic BHW model. In fact, for those scenarios that are sensible to the private signal, changes in the base rate have no affect on the decisions, i.e. the base rate is within the interval $[1 - p_x, p_x]$. Such changes only affect the outcome. The intuition for this is easy to get. If for the first individual the private signal is decisive then obviously the base rate does not matter. Also the second individual’s decision does not change if the base rate changes. This robustness is an artifact of the model. If the observations are made noisy, such that observations have less impact, then base rate changes may have an effect. However, we stick to the simple model, since the noise does not change our basic results. On the other hand, if including noise we had to separate the effects of noise from the effects of base rate neglect, which would blow up the analysis without gaining more insights into our basic question.

4. The effect of base rate neglect

We now introduce a cognitive bias into our decision model. Base rate neglect is an under weighting of the base rate. The specific functional dependency is to my knowledge not completely clear. However, we follow Grether (1980) and operationalize a base rate neglect as a linear under weighting of the history information q_A. Then, the adoption condition can be written as in (3) with $\beta = 1$ representing the rational reasoning process and $\beta < 0$ and under weighing and $\beta = 0$ representing a complete base rate neglect.

$$0 \ < \ q_H + q_X + \beta \cdot q_A + s \tag{3}$$

As we have seen above, extreme base rates prevent social learning. If a base rate neglect weakens the impact of base rate then we may expect that social learning processes occur for a bigger set of scenarios, i.e. for more extreme base rates. Since previous analysis has shown that social learning processes can be good for a population we can expect increased expected payoffs for scenarios where social learning did not take place. Additionally, since the basic model is robust to some changes in the base rate if people follow the private information, we also do not expect a loss in expected payoffs for scenarios where social learning took already place in the original model. This intuition is supported by Figure 2. We see that an increasing base rate neglect, i.e. a decreasing β, extends the range of base rates where a social learning process is started. Without any base rate neglect ($\beta = 1.0$) the social learning process is restricted to the range $[1 - p_x, p_x]$. With a base rate weight of $\beta = 0.8$ the range expands while keeping the expected payoff for all base rates above the values reached for no base rate neglect. The loss in expected payoffs of the individuals that starts the process are counter-balanced by the risk reduction for following individuals due to the information externality. However, if the bias is too big (i.e. $\beta = 0.3$), then individuals may perform worse for some extreme base rates. The range of social learning is expanded too much. If the base rate is completely neglected, then the probability of correct decision does not depend on the base rate and is completely determined by the signal quality p_x. Only for cases of certainty for the risky option, $p_v = 0$ or $p_v = 1$, the individual does consider the signal. Altogether, there seems to be an optimal level of base rate neglect, which we now determine.

The role of base rate neglect is reducing the impact of the base rate such that a social learning process is triggered. The social learning process in our model can at most double the information that an individual considers for its own decision. If the social learning process accumulates, i.e. the history provides, more information than the private signal provides then the private signal can not counter-balance the history information, which in turn causes that the decision does not reveal any of the private information. The social learning process stops. Therefore, a social learning process is useful if the amount of information accumulated though a social learning process can counter-balance the base rate. For deriving the optimal base rate we additionally as-

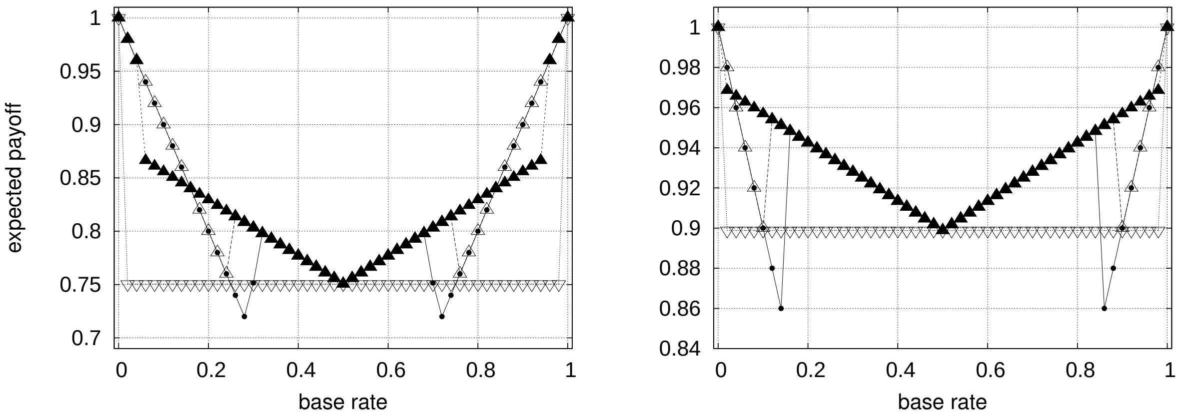

Figure 2. The expected probability of correct choice for different base rates (from 0 to 1), for different qualities of private signals (left: $p_x = 0.7$ and right: $p_x = 0.85$), and for different levels of weights for the base rate ($\beta = 1.0$ (●), $\beta = 0.8$ (△), $\beta = 0.3$ (▲), $\beta = 0.0$ (▽))

sume that the social learning process accumulates positive and negative information in the same way.[1] This gives us a rule for the optimal weight for the base rate: $\beta\frac{\mathbb{P}[V=H]}{\mathbb{P}[V=L]} = \max_{h_i}(\frac{\mathbb{P}[h_i|V=H]}{\mathbb{P}[h_i|V=L]})$. In the BHW model the maximal amount of information that is accumulated before the social learning model stops is sharply below twice the information provided by the private signal. Related to the insights by Bikhchandani et al. (1992) it is the amount of information provided by a history that triggers an informational cascade. This implies that in the very artificial scenario of the BHW model the optimal base rate neglect is 0.5. It is independent of the quality of the private signal, p_x.

$$\beta^* = 0.5 \tag{4}$$

[1]However, when relating the model back to real life one should be aware of empirical work that shows that in general processes of social learning as for instance in word-of-mouth recommendation negative information seems to be overweighted, which may violate the assumption of a symmetric social learning process.

5. The base rate neglect for favorable and unfavorable payoff structures

Our analysis is so far restricted to scenarios with $s = 0$, which means that relative gains are of the same amount as relative losses. This is a very artificial assumption, which we now relax. From a mathematical point of view our previous analysis can be extended to arbitrary initial decision structures. Every setting with a payoff structure s different from zero can be transformed into a setting with $s = 0$ without affecting the adoption condition. The transformation basically adapts p_v such that $\frac{p'_v}{1-p'_v}$ equals $\frac{p_v}{1-p_v}\frac{H-S}{S-L}$. This implies that in our extended model with arbitrary payoff structure under weighting the sum $q_A + s$ by one half leads to an increased expected payoff due to triggering social learning processes. However, the base rate neglect does only weight the base rate. Nevertheless, this idea gives us easily the optimal weight for the base rate for our extended model with arbitrary payoff structures s. The weight for the base rate in the extended model, β^*_s, must yield the same effect as if the whole sum of $q_A + s$ is multiplied by $\beta^* = 0.5$. From $\beta^*_s \cdot q_A + s = \beta^*(q_A + s)$ we then get our optimal base rate neglect for the extended model as in (5). Figure 3 visualizes all combinations of base rates and payoff structures with optimal weights for the base rate below 1, which means that a base rate neglect is appropriate for a population.

$$\beta^*_s = \frac{q_A - s}{2q_A} \tag{5}$$

Equation 5 and figure 3 give very interesting insights, because the positive benefit of a base rate neglect is not anymore given for every environment. If s and q_A are in the same direction, then a base rate neglect is optimal. The bigger s is compared to q_A the more the base rate should be neglected. At some point even an inversion of the base rate information would be beneficial, i.e. $\beta^* < 0$. If both are inversely directed and the base rate is of smaller absolute value than the payoff structure, then the base rate should be over weighted. In such scenarios reducing

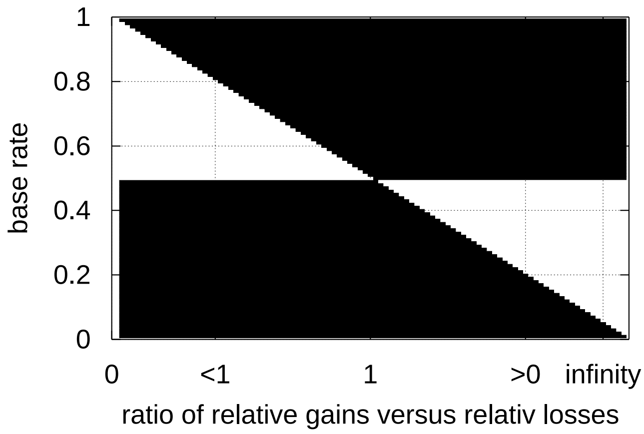

Figure 3. Black areas mark settings with base rate neglect being beneficial for a population: If the base rate is unfavorable, $pv < 0.5$, but the relative gains are higher than the relative losses, $\frac{H-S}{S-L} > 1$, then a base rate neglect is not appropriate. The same holds for settings where the base rate is favorable, $p_v > 0.5$, but the relative gains are lower than the relative losses, $\frac{H-S}{S-L} < 1$

the weight for the base rate would even strengthen the individuals' predisposition to favor one or the other option and thereby weakens the chances for a social learning process. We see that when relaxing the payoff assumption, $s = 0$, the result regarding the beneficial nature of a base rate neglect changes for some environments.

6. Conclusion

The paper demonstrates that a base rate neglect can be beneficial for a population. The primary positive effect of base rate neglect is that it can trigger social learning in domains where such processes would usually not start. We have characterized different settings where a base rate neglect is beneficial and those where it is not beneficial. If populations are situated in the first type of environments or if this type of environments is very dominant then we can expect that a populations with biased individuals outperform populations with individually rational members.

Our results supplement the insights by Bernardo and Welch (2001) and Kariv (2004), who showed similar positive effects for another cognitive bias, i.e. overconfidence in own information. However, their analysis is restricted to the case $q_A = s = 0$. If we relax this assumption then a combination of overconfidence and base rate neglect might be appropriate. If we consider overconfidence as over weighting private signals, then it is equivalent to under weighting history information, under weighting the base rate, and under weighting the payoff structure. Hence, overconfidence has already incorporated a base rate neglect. However, it is not clear if history information as well as base rate need to be under weighted by the same factor. If not, then both biases may need to be combined. In a first trial for a setting with relative losses equal to relative gains and with a positive base rate of 0.7 we have seen that overconfidence together with a base rate neglect are optimal for the population. However, the effect of the base rate neglect is only extremely small. Although our work as well as the work by Bernardo and Welch (2001) and Kariv (2004) focus only on a single cognitive bias I believe that we should not forget that even if single biases are not beneficial for a population they may be beneficial if present as a set of biases.

References

Anderson, L. A., & Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, *87*(5), 847–862.

Banerjee, A. V. (1992, August). A simple model of herd behavior. *The Quarterly Journal of Economics*, *107*(3), 797–817.

Bernardo, A. E., & Welch, I. (2001). On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy*, *10*(2), 301–330.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *The Journal of Political Economy*, *100*(5), 992–1026.

Chamley, C. (2004). *Rational herds: Economic models of social learning*. Cambridge University Press.

Grether, D. M. (1980, Nov). Bayes rule as descriptive model: The representative heuristic. *The Quarterly Journal of Economics*, *95*(3), 537–557.

Hung, A. A., & Plott, C. T. (2001, Dec). Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions. *The American Economic Review*, *91*(5), 1508–1520.

Kahneman, D., & Tversky, A. (2000a). *Choices, values, and frames* (D. Kahneman & A. Tversky, Eds.). Cambridge University Press.

Kahneman, D., & Tversky, A. (2000b). Conflict resolution: A cognitive perspective. In D. Kahneman & A. Tversky (Eds.), *Choices, values, and frames* (pp. 473–487). Cambridge University Press.

Kariv, S. (2004, July). *Overconfidence and informational cascades.*

Todd, P. M., & Goodie, A. S. (2002). Testing the ecological rationality of base rate neglect. In B. Hallam, D. Floreno, J. Hallam, G. Hayes, & J. A. Meyer (Eds.), *From animals to animats 7 : Proceedings of the seventh international conference on simulation of adaptive behavior.* MIT Press / Bradford Books.

Weizsäcker, C. C. von. (2005, June). *The welfare economics of adaptive preferences* (Tech. Rep. No. 2005/11). Bonn: Max Planck Institute for Research on Collective Goods.