Advances in Complex Systems, Vol. 13, No. 5 (2010) 607–619

DOI: 10.1142/S0219525910002761

BASE RATE NEGLECT FOR THE WEALTH OF INTERACTING PEOPLE

DIEMO URBIG

Department of Management, University of Antwerp,
Prinsstraat, 13 (Z.107), 2000 Antwerp, Belgium
diemo.urbig@ua.ac.be

Received 4 December 2009
Revised 31 March 2010

Previous research investigating base rate neglect as a bias in human information processing has focused on isolated individuals. This study complements this research by showing that in settings of interacting individuals, especially in settings of social learning, where individuals can learn from one another, base rate neglect can increase a population's welfare. This study further supports the research arguing that a population with members biased by neglecting base rates does not need to perform worse than a population with unbiased members. Adapting the model of social learning suggested by Bikhchandani, Hirshleifer and Welch (*The Journal of Political Economy* **100** (1992) 992–1026) and including base rates that differ from generic cases such as 50–50, conditions are identified that make underweighting base rate information increasing the population's welfare. The base rate neglect can start a social learning process that otherwise had not been started and thus base rate neglect can generate positive externalities improving a population's welfare.

Keywords: Cognitive biases; base rate neglect; social learning; information cascades; ecological rationality.

1. Introduction

Individuals placed in uncertain environments experience a large set of cognitive biases [1]. Often individuals have some prior information about the uncertain environment, and during their interaction with the environment they receive additional cues, and update their beliefs about the environment. This updating process has frequently been modeled as Bayesian updating, which combines base rate information with additional signals to form a new updated judgment [2]. In such scenarios individuals can experience a cognitive bias that is called *base rate neglect* or base rate fallacy. This cognitive bias causes people to significantly underweigh the information coming across as base rate [3]. If people do not consider the base rate as having a causal nature or as being specific and strongly related to the specific event under consideration, then the tendency to ignore it becomes stronger [3]. Also repetition of the same task and the incentive structure seem to be moderating factors [4]. Wherever the effect comes from, a base rate neglect causes individuals to deviate from the individually rational behavior that maximizes the unbiased individually expected payoff. Everything else kept equal, individuals suffering from this bias are worse off in terms of objectively expected payoffs.

Once individuals have been found to suffer from cognitive biases, researchers usually try to reduce the bias, i.e., to debias them [5]. While researchers might do this for testing the robustness of the bias, there is also research deriving normative conclusions. Some researchers argue that public administration and policy makers should, for instance, discourage unrealistically optimistic entrepreneurs from becoming entrepreneurs [6]. This type of argument refers to normative recommendations based on behavioral findings, which is subject to various problems, and are discussed in depth by Berg [7]. More specifically, for dealing with biases of individuals, it is important to understand whether such individual biases translate into a loss of welfare for the population the individual is embedded in. Individual rationality in terms of individual maximization of profits (or utility) does not need to translate into a maximization of the aggregate profits (or welfare) at a population level. Especially if externalities are present and individual decisions positively or negatively affect other people's outcomes, individual maximization can fail in maximizing a population's overall welfare [8]. As a consequence, detecting individually irrational behavior does not imply that a population is not maximizing its welfare at an overall level. Thus, rational maximization of welfare and individual rationality can diverge. For instance, Bernardo and Welch [9] and Berg and Lien [10] show that overconfidence, which is a systematic deviation in subjective beliefs from rational expectations, can increase a population's welfare by providing a social benefit that more than offsets individual costs of that deviation. For base rate neglect, there are analyses showing that depending on the environment, the bias does not hurt that much [11]. This fit of cognitive reasoning processes with specific environments is called ecological rationality [12, 13].

This article complements the work on the ecological rationality of base rate neglect by looking at scenarios of social learning, where individuals not only deal with (public) base rate information and additional (private) signals but also observe others' decisions. These observations generate interdependencies between members of a group or a population. Being part of a population, where one can — at least partially — observe others is a type of environment that can reasonably be considered relevant in human decision-making (see [2] for an overview). Such settings of interacting individuals frequently raise the questions for understanding the dependencies between individuals' behavior and the overall population's performance. The analysis presented in this paper builds on the model of social learning as introduced by Bikhchandani, Hirshleifer and Welch [14]. In this model the beneficial effects of another cognitive bias has already been investigated. Bernardo and Welch [9] show that overconfidence of all but also of some individuals can sometimes be beneficial for the whole population. Following an analysis in a different but related model, Kariv [15] concludes that it is not clear in what types of environments overconfidence is beneficial. In both models, however, the potentially positive effect of overconfidence for the population is based on an information externality that is caused by biased individuals. This externality prolongates a social learning process that otherwise had stopped. Bernardo and Welch [9] and Kariv [15] only consider the case of 50–50 base rates. Their analyses are complemented here by considering different base rates as well as a different cognitive bias, specifically the base rate neglect.

The next section introduces the basic model. The third section will then show the welfare effects of social learning for different base rates and unbiased individuals. The fourth section introduces a weight for base rate information into the model and gives a first impression of the beneficial effects of individuals underweighting the base rate. In the fifth section the optimal weight is determined, while the sixth section concludes the paper with a brief discussion.

2. The Basic Social Learning Model (BHW Model)

Following Ref. 14, the basic model is built around individuals who are confronted with a discrete choice between a safe and a risky option. The risky option's payoff depends on the some hidden state of the environment, $V \in \{H, L\}$, which causes either high payoffs ($V = H$) or low payoffs ($V = L$). Individuals have some *a priori* information about the probabilities of the environment's state, i.e., $p_v = \mathbb{P}[V = H]$ with $0 \leq p_v \leq 1$, but each receive independently an additional imperfect private signal $X_i \in \{H, L\}$, which either signals a high payoff, $X_i = H$, or a low payoff, $X_i = L$, respectively. The signal is symmetric and it is correct only with probability $p_x = \mathbb{P}[X = H|V = H] = \mathbb{P}[X = L|V = L]$ with $1 > p_x > 0.5$.[a] Individuals decide in an exogenously given sequence. Individual i with $i > 1$ can observe previous decisions of $i - 1$ individuals about adopting (a) or rejecting (r) the risky option and choosing the safe option. These observations are summarized in the history $h_i \in \{a, r\}^{i-1}$. The history is thus a sequence of letters "a" and "r". The safe option's payoff is given by S with $L < S < H$.

Assuming no perceptual distortions, perfect information about the decision rules, information aggregation by Bayesian updating, and maximization of expected individual payoffs, the rational decision for the choice between the safe and risky option is given by condition (1), written in terms of log-odd ratios. Since the algebra behind this result is standard, the reader is referred to Ref. 2 for a more detailed introduction to Bayesian reasoning in models of social learning. In contrast to other adaptations of the BHW model, e.g., [14, 16, 17], the model developed in this

[a]For $p_x = 1$, the private information is perfect and there would not be any uncertainty anymore. For $p_x < 0.5$, the signal would systematically signal the wrong thing and it would be better interpreted inversely, which leads to a signal with $p_x > 0.5$. For $p_x = 0.5$, the private signal would not be informative at all, such that the individual decision could not reveal any private information, which is the basic mechanisms social learning is concerned with, i.e., the diffusion of private beliefs [2].

paper deviates from the previous models by relaxing the assumption of a 50–50 base rate.

$$0 < \underbrace{\log\left(\frac{\mathbb{P}[X|V=H]}{\mathbb{P}[X|V=L]}\right)}_{\text{private information } q_x} + \underbrace{\log\left(\frac{p_v}{1-p_v}\right)}_{\text{base rate information } q_A} + \underbrace{\log\left(\frac{H-S}{S-L}\right)}_{\text{payoff structure } s}$$
$$+ \underbrace{\log\left(\frac{\mathbb{P}[h_i|V=H]}{\mathbb{P}[h_i|V=L]}\right)}_{\text{history information } q_h}. \tag{1}$$

For the case of indifference between the risky and safe option, a random choice between safe and risky option is assumed [14]. Figure 1, which is discussed later on, visualizes the effects of an alternative indifference rule (individuals following the private signal). As will be shown below, the random choice leads to a more convincing model.

The *base rate information* is included as in previous models but is not fixed at a value where it would disappear, i.e., 0.5. For simplification and consistent with Refs. 14 and 9 we assume that $H - S = S - L$, which makes the payoff structure disappearing from (1), i.e., $s = 0$. The assumption implies that the potential gain of the risky option relative to the safe option is as large as the potential loss relative to the safe option.

The *private information* q_x is given by (2); it is equivalent to the analysis in [14]. The private information q_x has only two possible instantiations and these two values depend on p_x and the signal X_i. For simplification we assume that the values q_x^+ and q_x^- with $q_x^+ > 0 > q_x^-$ represent the values for cases $X = H$ and $X = L$, respectively. Due to assuming a symmetric signal the absolute values of q_x^+ and q_x^- are the same, i.e., $|q_x^+| = |q_x^-| = \log\left(\frac{p_x}{1-p_x}\right)$, where $|\cdot|$ gives the absolute value.

$$q_x = x\log\left(\frac{p_x}{1-p_x}\right) \quad \text{with } x = \begin{cases} -1 & \text{if } X = L \\ +1 & \text{if } X = H \end{cases}. \tag{2}$$

The history information, q_h is calculated based on the sequential decision model and it follows the analysis by Bikhchandani *et al.* [14]. Note that the model assumes almost perfect information such that only the private information is hidden. Therefore, the only information that individuals might learn from observing others' decisions, is about others' private information. This information is, however, not always revealed. From the analysis in Ref. 14 it is obvious that private information is only revealed if it affects the decision and thus if the remaining parts of condition (1) are small enough such that the decision changes depending on the value of the private information. Since everything in condition (1) except the private information is public knowledge, individuals observing other individuals can identify whether or not the others will consider the private information. Therefore, individuals can identify whether or not the observed decision reveals something about the private information.

In fact, others' decisions can be classified into three classes: cascades, full revelation, and incomplete revelation. If the private information does not affect the decision, i.e., the decision does not differ for the cases of a signal indicating a high payoff and of a signal indicating a low payoff, then the observation does not reveal any information and it can be ignored.

If the decision completely depends on the private information, i.e., it adopts if the private signal indicates a high payoff and reject otherwise, then the private information is fully revealed.

Individuals can get into a state of indifference, which happens if the private information is as large as the remaining parts of the adoption condition, i.e., $|q_h + q_A| = \log\big(\frac{p_x}{1-p_x}\big)$. In such cases the private information might only be incompletely revealed. For instance, if $q_h + q_A$ equals the value of q_x for a negative private signal, then receiving a negative signal makes the individual reject the risky option as it further strengthens the negative evaluation. However, for a positive signal the individual becomes indifferent and chooses randomly such that adoption occurs with a probability of 50%. Note that observing the adoption in this case perfectly reveals the private information, while observing a rejection incompletely reveals information. The rejection could be caused by a negative signal or by chance for a positive signal. Let d_f be the difference between adoptions and rejections with full revelation. The information revealed by observations with full revelation can be written as $d_f \log\big(\frac{p_x}{1-p_x}\big)$. Further, let d_l be the difference between adoptions and rejections with incomplete revelation. The information revealed by observations with incomplete revelation can be written as $d_l \log\big(\frac{1+p_x}{2-p_x}\big)$. The complete term for the history information is then given by (3).

$$q_h = d_f \log\left(\frac{p_x}{1-p_x}\right) + d_l\left(\frac{1+p_x}{2-p_x}\right). \tag{3}$$

For a detailed analysis of the basic model without reference to welfare aspects, the reader might consider Ref. 14.

3. Limits to Social Learning

A central feature of the BHW model is that social learning is restricted; the more extreme the information derived from observing other's decisions, the higher is the tendency that the private information does not matter and individuals just follow these observations. Depending on whether they adopt or reject regardless of their private signal, individuals are considered to be in an adoption cascade or in a rejection cascade, respectively. In the more general model developed here, an adoption cascade occurs if $q_h > -\big(-\log\big(\frac{p_x}{1-p_x}\big)x + q_A\big)$ and a rejection cascade happens if $q_h < -\big(\log\big(\frac{p_x}{1-p_x}\big) + q_A\big)$, social learning is prevented if $\log\big(\frac{p_x}{1-p_x}\big) < |q_A + q_h|$. Under this condition, decisions are not affected by the private signal, because either the history information q_h or the base rate information q_A or both together are as extreme as the private information does not counter it. The proofs in Ref. 14 show that the probability of reaching a cascade approaches 1.0 for an infinite number of individuals, this can be generalized to our model for those cases where social learning occurs.

In previous models such as Refs. 14 and 9, base rates are assumed to be 50–50. In these cases, the adoption condition only contains the term for history information besides private information. Therefore, the only reason for individuals to end up in a cascade in these models is information derived from observing other individuals' decisions. The model developed here allows non-trivial base rates. Thus, the base rate might be as extreme such that already the first individual's decision (for the first individual q_h is zero) is not affected by the private information, i.e., for all possible instances of q_x the individual decides the same way. In this case, no social learning occurs and all decide as if they had no private information. There are therefore two types of information cascades, those caused by information derived from observing others and those cascades triggered by extreme base rates. To distinguish one case from the other, we will refer to the former type as *endogenous cascade*, because the observations are endogenously created in the model, and to the latter as *exogenous cascade*, because the base rate is not affected by individuals' decisions.

Most parts of the analysis of the single decision can either be found directly in Ref. 14 or the arguments can easily be extended to the model developed here. Instead of analyzing the individual decision, this study focuses on welfare implications. More precisely, it asks the question of how do changes in the decision rule affect the social welfare.

In order to analyze welfare effects, we have to define welfare for a population of cognitively biased individuals. This study follows [9, 10] and defines welfare as the sum of objectively, i.e., unbiased, expected payoffs over all individuals in a population. Note that the expected payoff used for calculating the actual welfare is based on true probabilities and not on distorted perceptions. Further, due to the exogenous random sequencing of individuals, one can use the objectively expected payoff of an individual as proxy for the welfare of a population.[b] By referring to the unbiased expected payoff, it is assumed that the base rate neglect does not affect the utility of individuals but only the quality of their decisions. The base rate neglect is therefore considered as a mistake that people would actually avoid if they were capable of doing it, it is a perceptual bias only. If instead one assumes biases that change utilities such as risk attitudes, then comparing the welfare of two populations that differ in the bias of their members gets difficult to analyze. The problem is close to the analysis of welfare consequences of preference changes (see Ref. 18). However, the assumption that a perceptual bias does not change the utility function simplifies the analysis.

[b]Similar to Ref. 9 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.

Figure 1 plots the expected payoffs for an individual against different base rates and as already mentioned for two different rules for decisions under indifference. This visualization as well as all the following are based on a population size of $n = 50$.[c] Since the interesting effects are a bit difficult to see in Fig. 1, Fig. 2 plots the same expected payoffs but relative to the case of a population with private information but without social learning, i.e., individuals cannot observe each other.

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 receiving a payoff independent of the base rate, while above $p_v = 0.5$, they choose the risky option which gives a payoff proportional to the base rate. If individuals have private signals (in the example

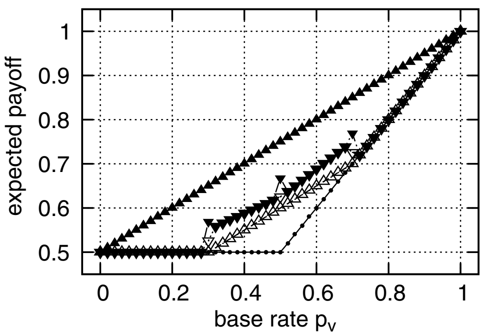

Fig. 1. Expected payoffs for no information (●), for private information only (△), for perfect private information (▲), for private information and social learning with random choice when indifferent (▽), for private information and social learning with following private information when indifferent (▼). Private signal quality p_x equals 0.7.

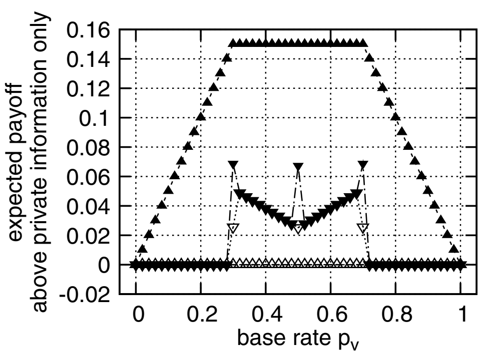

Fig. 2. Expected payoffs relative to expected payoffs for private information only (△), for perfect private information (▲), for private information and social learning with random choice when indifferent (▽), for private information and social learning with following private information when indifferent (▼). Private signal quality p_x equals 0.7.

[c]Most results derived for the BHW model and its derivatives do not change substantially for larger populations.

where $p_x = 0.7$ is assumed), but cannot observe others, then they base their decision only on the private signal and the base rate. Their decision is based on the private signal if the base rate is not too extreme such that the signal can counter-balance the base rate. Here the individuals' expected payoffs conditional on the base rate increase due to better information. Only in these scenarios, where private signals may influence decisions, decisions can reveal information to following individuals. In these settings, we also see an additional benefit due to observing others' decisions.

Figures 1 and 2 also demonstrate an alternative indifference rule, "follow the own signal," which has been employed in Refs. 19 and 16. Bernardo and Welch [9] 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, their specific indifference rule is equivalent to the rule "follow the own signal." Figures 1 and 2 show that this rule has some discontinuities, which implies that increasing the base rate from 0.5 only by a very small amount reduces the expected payoff very much. The indifference rule "random choice" seems to generate more plausible dynamics. Nevertheless, also for the latter indifference rule there is a point where an increasing base rate decreases expected payoffs. However, this point describes a reasonable effect, i.e., reaching a limit beyond which social learning does not occur; it is thus not only an artifact.

Figure 3 visualizes the welfare of populations across different base rates for three levels of the private signal quality p_x. The better the private signal the wider the range where the private information matters and, thus, where social learning occurs, i.e., if $\log\left(\frac{p_x}{1-p_x}\right) \geq |q_A|$. As discussed above with respect to Figs. 1 and 2, the range of base rates that do not prevent social learning is bounded from the left side by a jump from expected payoff of 0.5 to a value above 0.5 and bounded from the right side by a non-monotonic behavior, where the expected payoff drops slightly. This drop is due to the fact that the increase in the base rate increases

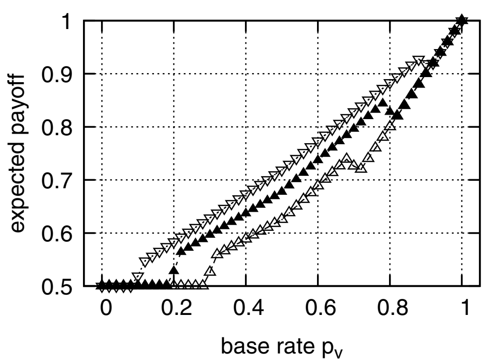

Fig. 3. Expected payoff versus base rates (safe payoffs) for different qualities of the private signal: $p_x = 0.7$ (△), $p_x = 0.8$ (▲), $p_x = 0.9$ (▽).

the expected profit just by a very small amount, but social learning stops, which decreases expected profits by a substantially larger amount. Social learning stops because the first individual already follows the base rate information and does not reveal the private information.

Before getting into the analysis of base rate neglect, let us briefly consider two observations for the BHW model which becomes relevant later on.

The first observation refers to endogenous cascades, i.e., those that are triggered by observations of others' decisions. Note that whenever the history information is more informative than the private information, individuals do not care about the private information and thus do not reveal this private information. This implies that whenever the history information is stronger than a single private signal, social learning stops. Therefore the benefit of social learning in this model is limited to the provision of information equivalent to a single private signal.

The second observation refers to the influence of base rate information. For those scenarios where social learning starts, i.e., base rates are sufficiently balanced ($\log\left(\frac{p_x}{1-p_x}\right) > |q_A|$), changes in the base rate that keep the base rate in this range do not affect the decisions. In other words, besides affecting the expected payoff, the base rate is irrelevant with respect to decisions that are taken.[d]

4. Base Rate Neglect: Expanding the Limits of Social Learning

After developing and solving the model for unbiased decision makers, we can now consider the base rate neglect as a deviation from unbiased decision-making. Following Ref. 4, the base rate neglect is operationalized as a linear underweighting of the base rate information q_A. The adoption condition can be re-written as in (4) with $\beta = 1$ representing the unbiased reasoning process, $0 < \beta < 1$ representing an underweighting, and $\beta = 0$ representing a complete base rate neglect.

$$0 < q_h + q_x + \beta q_A \quad \text{with } 0 \leq \beta \leq 1 \tag{4}$$

As shown above, extreme base rates can prevent social learning. If a base rate neglect weakens the impact of base rates, then one might expect social learning processes occurring for a bigger set of scenarios, i.e., for more extreme base rates. Since the previous analysis has shown that social learning processes can be good for a population, we might expect an increased expected payoff for scenarios where social learning does not take place without introducing the base rate neglect. Further, since the basic model is robust to some changes in the base rate if people follow the private information, we may also do not expect a loss in expected payoffs for scenarios where social learning took already place in the original model.

[d]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 the basic results, but if including the noise, one had to separate the effects of noise from the effects of base rate neglect, which blew up the analysis without gaining more insights into the basic question.

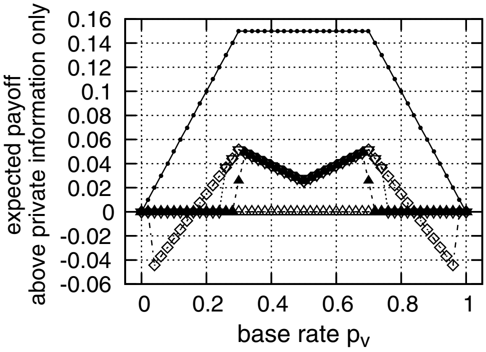

Fig. 4. Expected payoffs relative to expected payoffs for private information only (△), social learning with $\beta = 1.0$ (▲), with $\beta = 0.75$ (▽), with $\beta = 0.25$ (◇). Graph on the right-hand side plots the expected payoff relative to the case of private information only. Quality of private information p_x equals 0.7.

The intuition regarding the potentially beneficial effects of base rate neglect is supported by Fig. 4, which plots expected payoffs for different β and for different base rates. One observes 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 levels reached in a setting without a base rate neglect. The loss in expected payoffs of the individuals that starts the process is 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 more extreme base rates. The range of social learning is expanded too much. Altogether, there seems to be an optimal level of base rate neglect.

Running analyses equivalent to Fig. 4 but for many values of β between zero and one (and assuming a signal quality of $p_x = 0.7$) finds that for $\beta = 0.5$ the population maximizes its expected payoff over all possible base rates. This means that the range of social learning is expanded to the interval between 15.5% and 84.5% but not into the area where it would lead to a loss in expected payoffs. These values reflect the points where the zero line is crossed. Note that the range where social learning occurs is between 30% and 70% if there is no base rate neglect and in this range the base rate neglect does not change the expected payoff too. However, in scenarios where social learning starts but has not been started without base rate neglect, the whole population is better off.

The intuition behind this result is the idea that the base rate neglect should not suppress more information than is generated through social learning triggered by the base rate neglect. Recalling that social learning provides at most as much information as a single private signal (see discussion above), it should not suppress more information than a single private signal. This is in fact the case for a base rate neglect of 0.5. If the base rate is less than a private information, i.e., $|\log\left(\frac{p_v}{1-p_v}\right)| < \log\left(\frac{p_x}{1-p_x}\right)$, then an underweighting of this base rate does not change the decision (see discussion above). If the base rate is more extreme than two private signals, i.e., $|\log\left(\frac{p_v}{1-p_v}\right)| > 2\log\left(\frac{p_x}{1-p_x}\right)$, then unbiased decision makers would not care about the private signal. However, an underweighting of the base rate by 0.5 still leaves a base rate which is larger than a single private signal and lets a decision maker to still ignore the private signal. The base rate neglect only changes behavior for settings where the base rate is in between and thus for settings where the base rate carries information of a strength between a single private signal and two private signals. In this range, an individual that would rationally not care about the private information is made caring about it by underweighting the base rate.

There is a caveat to the reasoning, which is the size of the population. In this model, assuming a single individual, base rate neglect is welfare decreasing because nobody can profit from the information externality generated through the base rate neglect. In other models, however, such as such as Berg and Hoffrage [20], ignoring relevant information can be absolutely cost-free to the individual in terms of expected payoffs. The negative effect in the model presented in this paper exists for all first individuals independent of the size of a population. This first individual will generate an externality but does not benefit from it. However, the larger the population becomes, the more individuals can benefit and the less the impact of the first individual. For a sufficiently small population, a base rate neglect becomes suboptimal. These arguments together with the finding in this paper for larger populations guarantee the existence of such a threshold.

5. 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 settings where such processes would usually not start. This study shows that for the model developed by Ref. 14 where one relaxes the assumption of a 50–50 base rate the expected payoffs of a population is maximized with a base rate neglect that underweighs the base rate by a factor of one half.

Together with the work by Bernardo and Welch [9], Berg and Lien [10], and Berg and Hoffrage [20], this study illustrates the fact that a cognitive bias does not maximize the objectively expected payoffs of an individual (keeping everything else equal) but does not generalize to the fact that this bias destroys the welfare of a population. The interaction between members of a population can even lead to the observation that a cognitive bias enhances the population's welfare. This is exactly what this study shows for the base rate neglect and what Bernardo and Welch [9, 10] show for overconfidence. Policymakers should therefore be careful when judging the value of debiasing based on the analysis of a single individual. It is not unlikely that the link between social welfare maximization and maximization of individually expected payoffs is not straightforward. While this study focuses on the base rate neglect, Berg and Gigerenzer [21] discuss the issue at a more general level.

While this study discourages policymakers from funding projects that debias people with respect to base rate neglect, there is still the opportunity that free markets debias people. If for instance, biased individuals recognize that their expected payoffs are below the expected payoff of unbiased individuals, they might have an incentive to hire consultants or trainers who help them avoid these biases. While this might increase the payoffs in the short run, the population-wide employment of debiasing will decrease the overall payoffs. Therefore, policymakers might even have an incentive to restrict the free market forces to avoid the spread of debiasing practices.

Acknowledgments

The author would like to thank two anonymous reviewers for valuable ideas and comments. An earlier version of the article was presented at the 12th International Conference on Computing in Economics and Finance (CEF), 2006. All remaining errors are mine. Most of the work were done while the author was at the Department of Computer Science at Humboldt-Universität zu Berlin. The author gratefully acknowledges the financial support through the Odysseus program of the Flemish Science Foundation (FWO).

References

[1] Kahneman, D. and Tversky, A. (eds.), *Choices, Values, and Frames* (Cambridge University Press, 2000).
[2] Chamley, C., *Rational Herds: Economic Models of Social Learning* (Cambridge University Press, 2004).
[3] Kahneman, D. and Tversky, A., Conflict resolution: A cognitive perspective, in *Choices, Values, and Frames*, Kahneman, D. and Tversky, A. (eds.) (Cambridge University Press, 2000), pp. 473–487.
[4] Grether, D. M., Bayes rule as descriptive model: The representative heuristic, *The Quarterly Journal of Economics* **95** (1980) 537–557.
[5] Fischhoff, B., Debiasing, in *Judgement Under Uncertainty: Heuristics and Biases*, Kahneman, D., Slovic, P. and Tversky, A. (eds.), Chapter 31 (Cambridge University Press, 1982), pp. 422–444.
[6] Parker, S. C., Policymakers beware!, in *Handbook of Research on Entrepreneurship Policy*, Audretsch, D. B., Grilo, I. and Thurik, A. R. (eds.), Chapter 4 (Edward Elgar, 2007), pp. 54–63.
[7] Berg, N., Normative behavioral economics, *Journal of Socio-Economics* **32** (2003) 411–427.
[8] Dahlman, C. J., The problem of externality, *Journal of Law and Economics* **22** (1979) 141–162.
[9] Bernardo, A. E. and Welch, I., On the evolution of overconfidence and entrepreneurs, *Journal of Economics & Management Strategy* **10** (2001) 301–330.

[10] Berg, N. and Lien, D., Does society benefit from investor overconfidence in the ability of financial market experts? *Journal of Economic Behavior and Organization* **58** (2005) 95–116.

[11] Todd, P. M. and Goodie, A. S., Testing the ecological rationality of base rate neglect, in *From Animals to Animats 7: Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior*, Hallam, B., Floreno, D., Hallam, J., Hayes, G. and Meyer, J. A. (eds.), From animals to animats (MIT Press/Bradford Books, 2002), pp. 215–223.

[12] Todd, P. M., The ecological rationality of mechanisms evolved to make up minds, *American Behavioral Scientist* **43** (2000) 940–956.

[13] Gigerenzer, G., Todd, P. M. and the ABC Research Group, *Simple Heuristics that Make us Smart* (Oxford University Press, 1999).

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

[15] Kariv, S., Overconfidence and informational cascades (2005).

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

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

[18] von Weizsäcker, C. C., The welfare economics of adaptive preferences, Technical Report 2005/11, Max Planck Institute for Research on Collective Goods, Bonn (2005).

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

[20] Berg, N. and Hoffrage, U., Does society benefit from investor overconfidence in the ability of financial market experts? *Journal of Economic Psychology* **29** (2008) 792–809.

[21] Berg, N. and Gigerenzer, G., Psychology implies paternalism? Bounded rationality may reduce the rationale to regulate risk-taking, *Social Choice and Welfare* **28** (2007) 337–359.