The evolution of optimism: A multi-agent based model of adaptive bias in human judgement

Dylan Evans[*]

[*] Biomimetics Group
Department of Mechanical Engineering
University of Bath
Bath BA2 7AY
UK
D.Evans@bath.ac.uk

Annerieke Heuvelink[*~]

[~] Department of Philosophy
University of Utrecht
Heidelberglaan 8
3584 CS Utrecht
The Netherlands
A.Heuvelink@phil.uu.nl

Daniel Nettle[°]

[°] Departments of Psychology and Biological Sciences
The Open University
Walton Hall
Milton Keynes MK7 6AA
UK
D.Nettle@open.ac.uk

Abstract

Human judgement under uncertainty has been shown to involve consistent departures from normative rationality, but it is hard to tell whether these biases are adaptive or not. We have attempted to throw light on this question by constructing a multi-agent based computer model. In the model, a variety of agents with different decision rules are allowed to compete in a variety of environments. We find that under certain environmental conditions biased agents who behave in ways similar to those observed in humans outperform the classically rational agent, who acts purely to maximise expected utility. These conditions are similar to those in which humans find themselves.

1 Introduction

Human judgement under uncertainty has been shown to involve consistent departures from normative rationality. In particular, people show 'motivational biases' in judgements of probability, over-estimating the probability of events with a positive return to the self and under-estimating the probability of events with a negative return (Miller & Ross, 1975; Zuckerman, 1979).

These biases disappear when people are depressed ('depressive realism') and when people are asked to estimate the probability of the same events happening to *others*. Apart from these situations, however, the optimistic bias is remarkably resilient, persisting despite repeated disappointment and evidence to the contrary.

From the standpoint of rational choice theory, these biases are clearly maladaptive. Some psychologists, however, have argued that they are adaptive (Taylor & Brown, 1988). We have attempted to adjudicate between these two possibilities by constructing a multi-agent based computer model.

In the model, a variety of agents with different decision rules are allowed to compete in a variety of environments. We find that under certain environmental conditions biased agents who behave in ways similar to those observed in humans outperform the classically rational agent, who acts purely to maximise expected utility. Moreover, it is plausible that these conditions are similar to those that humans find themselves in.

Our findings therefore add support to the view that motivational biases are in fact adaptive.

2 Methods

In order to test our hypothesis we designed a multi-agent based simulation using NetLogo. NetLogo is an agent-based parallel modelling and simulation environment produced by the Center for Connected Learning and Computer-Based Modelling at Northwestern University.

In our model, patches represent 'opportunities'. Each opportunity has a 'probability of success' (p, ranging from 0-1), a benefit for success (b, ranging from 0.0001 to 10 energy points) and a cost of failure (c, ranging from 0.0001 to 10 energy points). The colour of the patch is determined by the probability of success, with darker patches representing more difficult opportunities.

Agents have only one goal - to maximise their energy points. In other words, their utility function is a linear function of their energy level. Agents have some knowledge of the cost of failure (c), the benefit for success (b), and the probability of success (p), for each opportunity they face.

The values of c, b and p are properties of the patch that the agents find themselves on at any given moment. The level of noise affecting the agents' knowledge of these values can be set by means of the error-sliders on the interface (see Figure 1).

The error-sliders can vary from 0 (perfect information) to 10 (great uncertainty). This error determines the standard deviation used for the normal distribution of which the mean is the true value of c, b or p of the patch. A random number drawn from this distribution determines the agents' guesses about the values of c, b and p. There are two error sliders; one affects the agents' knowledge of p, the other affects the agents' knowledge of c and b.

Agents do not move, but since each patch is updated each turn, each agent is presented with a different opportunity at every time step. Each turn, every agent must decide whether or not to 'play' that opportunity or not. This decision is made according to the agent's 'decision rule'. There are three types of agent, each with a different decision rule:

1. The RATIONAL agent uses the principle of expected utility; i.e. it only plays when the expected utility of playing is greater than that of not playing.
2. The OPTIMISTIC agent also uses the principle of expected utility, but uses a biased estimate of p (its estimate of p is multiplied by its estimate of b divided by its estimate of c).
3. The EMOTIONAL agent always plays if it estimates b to be more than twice c, and never plays when it estimates b to be less than half c. When it estimates that b and c are between these limits, its chance of playing is proportional to its estimate of p.

If an agent decides to play an opportunity, its chance of success is determined by the probability of success associated with that opportunity. If it plays and succeeds, its energy level is increased by the benefit for success associated with the opportunity. If it plays and fails, its energy level is decreased by the cost of failure associated with that opportunity. If an agent does not play, its energy level remains the same for that turn. Agents start with zero energy. Agents never die, and there is no reproduction. Figure 1 shows the interface of the program.[2]

Figure 1: Interface Program

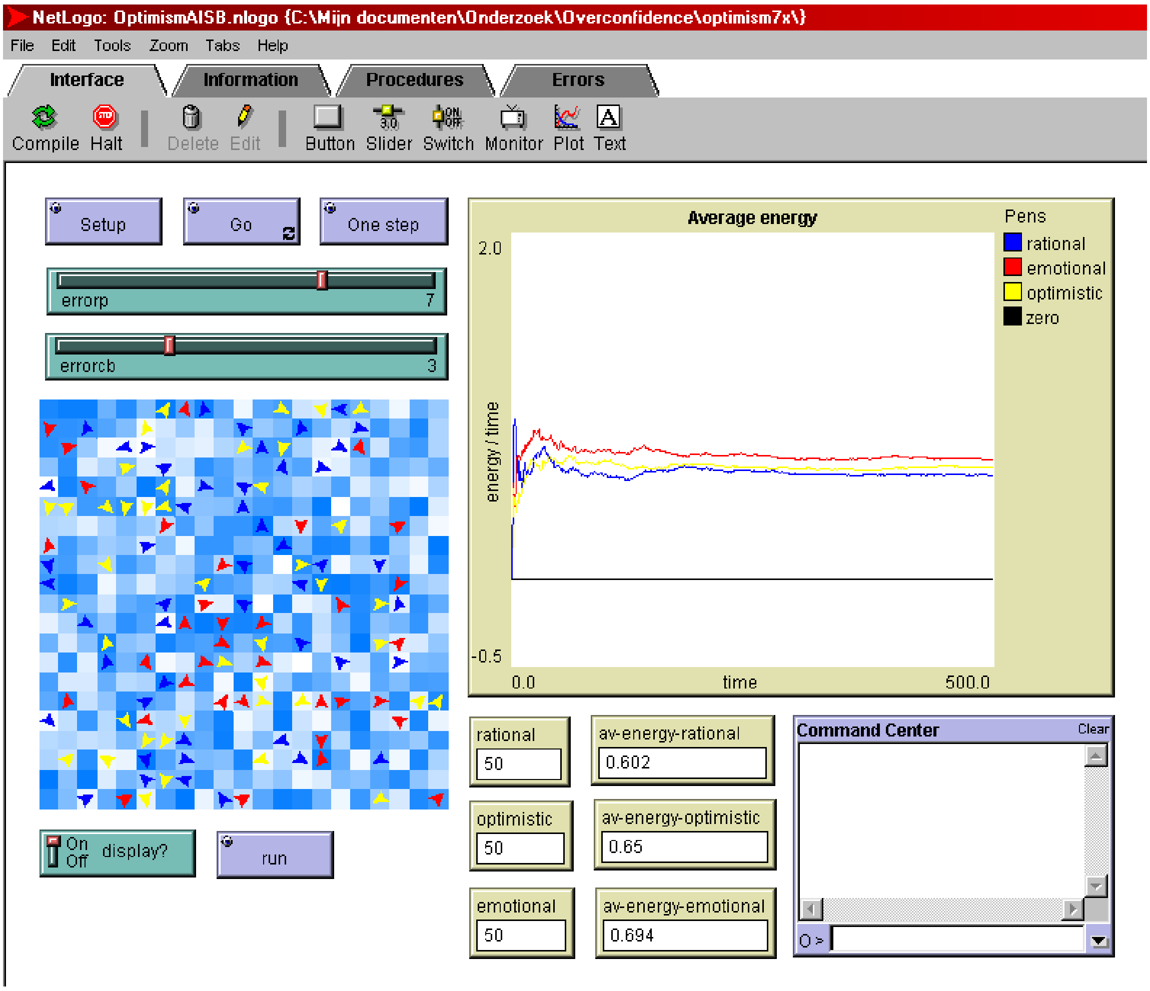

3 Results

We let the program run 10 times for every possible combination of the error of p and the error of c and b. One run lasts for 500 time steps. For each run we recorded the final average energy of each type of agent. We used this data to calculate the means and 95% confidence interval for all those runs for each kind of agents.

As can be seen from the graphs (Figures 2 – 5) there are significant differences (the error bars do not overlap). We show four specific 2D graphs, and an overall 3D graph seen from two different angles (Figures 6 and 7). The 3D graph does not show error bars but gives an overall view of our results.

Unsurprisingly, the rational agents do better than all other agents under most conditions. More interesting is the fact that there are conditions under which the rational agent is outperformed. This occurs when the error affecting agents' knowledge of c and b is quite small but the error affecting agents' knowledge of p is high. Both the emotional and optimistic agents do better then.

When the error affecting agents' knowledge of all three variables (c, b, and p) is high, all the agents perform at similar levels. Under these conditions there is great uncertainty and so they all perform quite badly.

Figure 2:

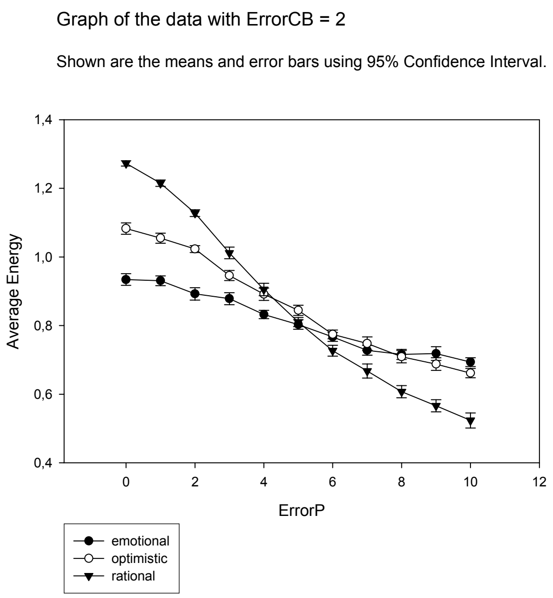

Figure 3:

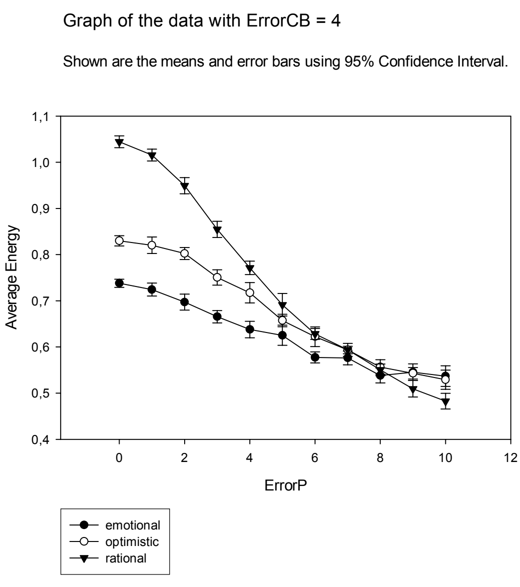

Figure 4:

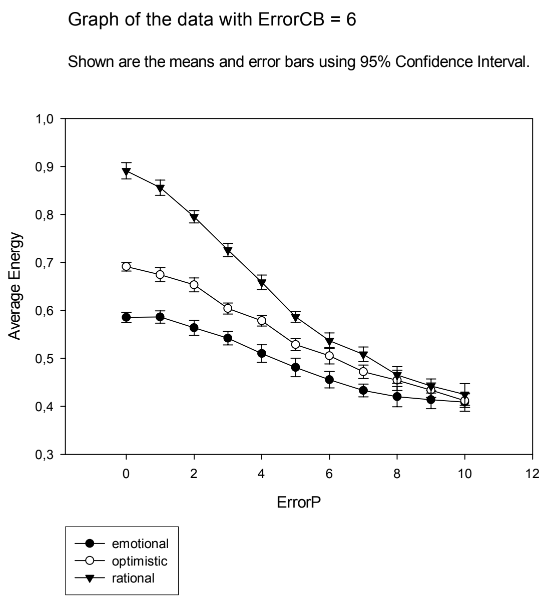

Figure 5:

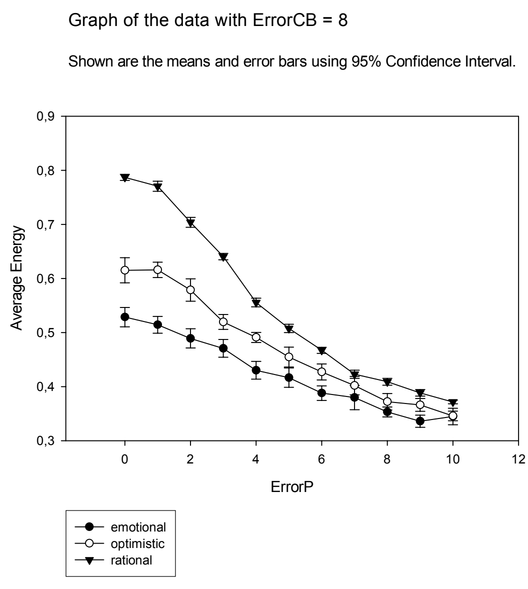

Figure 6:

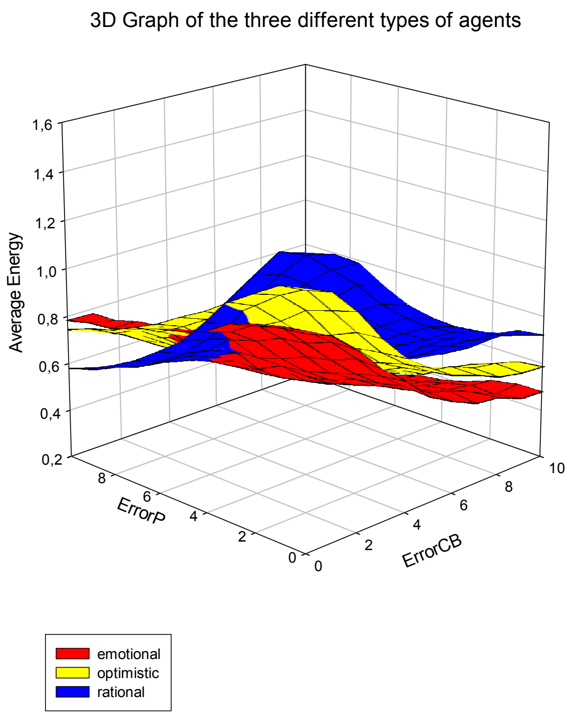

Figure 7:

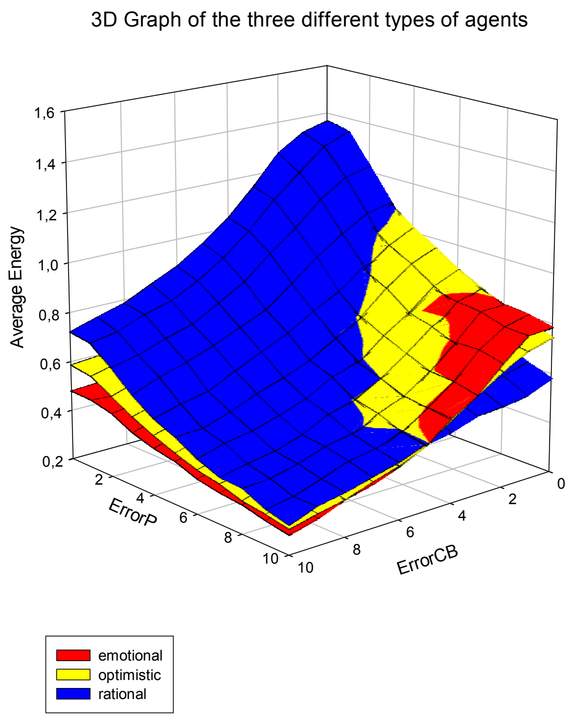

4 Discussion

Dealing with uncertainty is a common problem for agent systems. The ability to reason with uncertain information is an indispensable requirement for modeling intelligent behavior in a complex and dynamic environment. This is why uncertain reasoning has become a major research topic of AI with many important applications. These circumstances of uncertainty are the ones we investigated.

Psychologists have shown that human judgement under uncertainty involves consistent departures from normative rationality (Nettle, 2003; Kahneman & Tversky, 1973; Kahneman et al, 1982). In particular, people show motivational biases in judgements of probability, over-estimating the probability of events with a positive return to the self and under-estimating the probability of events with a negative return (Harris & Middleton, 1994; Larwood & Whitaker, 1977; Weinstein, 1980;1982).

These biases disappear when people are depressed ('depressive realism') (Alloy & Ahrens, 1987; Alloy & Abramson, 1979) and when people are asked to estimate the probability of the same events happening to *others* (for example Mirels, 1980). Apart from these situations, however, the optimistic bias is remarkably resilient, persisting despite repeated disappointment and evidence to the contrary.

From the standpoint of rational choice theory, these biases are clearly maladaptive. Some psychologists, however, have argued that they are adaptive (Taylor & Brown, 1988).

Our results show that there are environmental conditions under which biased agents outperform the classical rational agent, who acts purely to maximise expected utility. Our findings therefore add support to the view that motivational biases are in fact adaptive.

It is interesting to see that the biased agents, which we have designed to have biases similar to those shown by humans, do best when the error for b and c is low, but the error for p is high.[1] This is arguably the situation that people mostly encounter in the real world. It is plausible to think that people can estimate costs and benefits of an opportunity quite accurately, by observing other people faced with similar opportunities and by memories of past experiences. However, the chances of success of any specific opportunity depend on the interaction with other human beings and many other imponderable factors. Hence our ability to estimate probably is much poorer than our ability to estimate the cost of failure or the benefit of success.

[1] We have based the decision rules used by our biased agents necessarily on our interpretation of the empirical literature, but this literature is complex and clearly open to different interpretations.

Our emotional agents are based on the observation that when the benefits of a certain action are high people tend to play, independent of what the probability of success is.

The lottery is a case in point. The chance of success here is so small that nobody with a rational view will buy a lottery ticket. However, the benefit of success is so much high and the cost of failure so low that people do play. On the other hand, if there is a certain decision to make and there are high costs associated with failure but low benefits associated with success, then it is not likely that people will play, even if the probability of success is quite high. Only when the difference between costs and benefits is not that big do people seem to attend to the probability of success (Rottenstreich & Hsee, 2001).

We translated this empirical data to the rule for the emotional agents, such that they always play if the benefit of success is equal to or greater than twice the cost of failure, and never play if the cost of failure is equal to or greater than twice the benefit of success. Between those two extremes they only pay attention to the probability of success.

There is a certain degree of arbitrariness associated with the cut-off points in this rule, but the point here is not to provide a detailed model of the psychological process in real humans. Rather, we are interested merely in showing that biased agents that resemble humans to some first approximation can do just as well, or better than the purely rational agent.

Our optimistic agents are inspired by the empirical data about the situations in which people's judgement is biased.

A recurrent finding in the empirical literature is that when the benefits of an opportunity are higher than the costs, people tend to be more optimistic about their chances of success. They estimate the probability of success as higher than it really is and hence increase their change of playing. Conversely, when the costs of failure are higher than the benefits of success, people are pessimistic in the sense that they estimate their changes of success to be smaller then they actually are, so decreasing the change of playing.

Thus while the normative model of expected utility theory suggests that behavioural decisions should be based on the simple product of the probability of success and the benefit, people's weightings of probabilities in fact follow an inverse S-shaped function (Kahneman & Tversky, 1984).

We modelled this by giving our optimistic agents a biased estimate of p; in their case, the estimate of p is multiplied by the estimate of b divided by the estimate of c. The result is that the agents play more than would be rational if the benefits are higher than the costs, and play less if it is the other way around.

Our rational agents are clearly inspired by the classical decision theory (Von Neuman & Morgenstern, 1944). They only play if the expected utility of playing is greater than that of not playing.

In this program we did not investigate the possibility of giving different computational costs to different decision rules. This would be interesting, since then we could introduce an evaluation function that favoured cheaper rules.

The problem with such an experiment is that we do not have any empirical data about the computational costs associated with different rules, at least insofar as they are implemented in the human brain.

It would, however, be possible to calculate the computational cost of implementing the various decision rules in a specific robot. This could be done by measuring how long it takes the robot to make a certain decision following one rule and compare that to the other rules. In such circumstances we might find that a cheaper rule, that performs a little worse than the more expensive rule, will perform better in total.

5 Conclusion

From the standpoint of classical decision theory, motivational biases are clearly irrational. We have found that under certain environmental conditions biased agents, who behave in ways similar to those observed in humans, outperform the classically rational agent, who acts purely to maximise expected utility. Moreover, it is plausible that these conditions are similar to those that humans find themselves in.

Our findings therefore add support to the view that motivational biases are in fact adaptive.[2]

Acknowledgements

A Platform Grant from the Engineering and Physical Sciences Research Council, EPSRC GR/M97503 supported Dylan Evans' work.

[2] The code of our program can be found on: http://www.dylan.org.uk/optimismAISB.nlogo

References

Alloy, L. B. & Abramson, L. Y. (1979) Judgement of contingency in depressed and non-depressed subjects: Sadder but wiser? *Journal of Experimental Psychology: General*, **108**, 443-479.

Alloy, L. B. & Ahrens, A. H. (1987) Depression and pessimism for the future: Biased use of statistically relevant information in predictions for self versus others. *Journal of Personality and Social Psychology*, **41**, 1129-1140.

The Center for Connected Learning and Computer-Based Modelling
Northwestern University
Annenberg Hall 311
2120 Campus Drive
Evanston, IL 60208
USA
Homepage: http://ccl.northwestern.edu/

Harris, P. & Middleton, W. (1994) The illusion of control and optimism about health: On being less at risk but no more in control than others. *British Journal of Social Psychology*, **33**, 369-386.

Kahneman, D., Slovic, P. & Tversky, A. (eds) (1982) *Judgement under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.

Kahneman, D. & Tversky, A. (1973) On the psychology of prediction. *Psychological Review*, **80**, 237-251.

Kahneman, D. & Tversky, A. (1984) Choices, values and frames. *American Psychologist,* **39**, 341-350.

Larwood, L & Whitaker, W. (1977) Managerial Myopia: Self-serving biases in organisational planning. *Journal of Applied Psychology*, **62**, 194-198.

Miller, D. T. & Ross, M. (1975) Self-serving biases in attribution of causality: Fact or fiction? *Psychological Bulletin*, **82**, 213-225.

Mirels, H. L. (1980) The avowal of responsibility for good and bad outcomes: The effects of generalized self-serving biases. *Personality and Social Psychology Bulletin*, **6**, 299-306.

Von Neuman, J., Morgenstern, O. (1944) *Theory of Games and Economic Behaviour*. Princeton: Princeton University Press.

Nettle, D. (forthcoming) Adaptive illusions: optimism, control and human rationality. In *Emotion, Evolution and Rationality*, eds. Dylan Evans & Pierre Cruse, Oxford: Oxford University Press.

Rottenstreich, Y. & Hsee, K. (2001) Money, kisses and electric shocks: On the affective psychology of risk. *Psychological Science*, **12**, 185-190.

Taylor, S. E. & Brown, J. D. (1988) Illusion and well-being: a social psychological perspective on mental health. *Psychological Bulletin*, **103**, 193-201.

Weinstein, N. D. (1980) Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, **39**, 806-820.

Weinstein, N. D. (1982) Unrealistic optimism about susceptibility to health problems. *Journal of Behavioural Medicine*, **5**, 441-460.

Zuckerman, M. (1979) Attribution of success and failure revisited: The motivational bias is alive and well in attribution theory. *Journal of Personality*, **47**, 245-287