Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://www.econ.upenn.edu/pier

PIER Working Paper 01-057

“ Confidence Enhanced Performance”

by

Olivier Compte and Andrew Postlewaite

Confidence Enhanced Performance*

Olivier Compte
CERAS

Andrew Postlewaite
University of Pennsylvania

May, 2001
This version: December, 2001

Abstract

There is ample evidence that confidence can affect performance: confidence can improve performance, while a lack of confidence may diminish it. For example, the fears induced by the possibility of failure or of negative evaluations have physiological consequences (shaking, loss of concentration) that may impair performance in sports, on stage or at school.

There is also ample evidence that individuals have distorted recollection of past events, distorted attributions of the causes of successes of failures. Recollection of good events or successes is typically easier than recollection of bad ones or failures. Successes tend to be attributed to intrinsic aptitudes or own effort, while failures are attributed to bad luck. In addition, these attributions are often reversed when judging the performance of others.

The objective of this paper is to reconcile the two phenomena described above, and show that in a world where performance depends on confidence biases in information processing can increase welfare.

1. Introduction

There is ample evidence that people are overly optimistic. Among the many studies that show that in comparison to others, individuals evaluate themselves more highly than others do is Guthrie, Rachlinski and Wistrich (2001). Guthrie et al. distributed a questionnaire to 168 federal magistrate judges as part of the Federal Judicial Center's Workshop for Magistrate Judges II in New Orleans in November 1999. The respondents were assured that they could not be identified from their questionnaires, and were told that they could indicate on their questionnaire if they preferred their answers not to

*Much of this work was done while Postlewaite was a visitor at CERAS. Their support is gratefully acknowledged, as is the support of the National Science Foundation.

be included in the research project. One of the 168 judges chose not to have his or her answers included.

To test the relationship of the judges' estimates of their ability relative to other judges, the judges were asked the following question to estimate their reversal rates on appeal: "United States magistrate judges are rarely overturned on appeal, but it does occur. If we were to rank all of the magistrate judges currently in this room according to the rate at which their decisions have been overturned during their careers, [what] would your rate be?" The judges were then asked to place themselves into the quartile corresponding to their respective reversal rates: highest ($> 75\%$), second highest ($> 50\%$), third highest ($> 25\%$), or lowest.

The judges answers are very interesting: 56.1% put themselves into the lowest quartile, 31.6% into the second lowest quartile, 7.7% in the second highest quartile and 4.5% in the highest quartile. In other words, nearly 90% thought themselves above average.

These judges are not alone in overestimating their abilities relative to others. People routinely overestimate themselves relative to others in driving, (Svenson (1981)), the likelihood their marriage will succeed (Baker and Emery (1993)), and the likelihood that they will have higher salaries and fewer health problems than others in the future (Weinstein (1980)).

People do not only overestimate their likelihood of success relative to other people; they overestimate the likelihood of success in situations not involving others as well. When subjects are asked questions, and asked the likelihood that they are correct along with their answers their "hit rate" is typically 60% when they are 90% certain (see, e.g., Fischoff, Slovic and Lichtenstein (1977) and Lichtenstein, Fischoff and Phillips (1982)).

Biases in judging the relative likelihood of particular events are frequent, and many scholars have attempted to trace the source of these biases. One often mentioned source is the over-utilization of simple heuristics, such as the *availability heuristic* (Tversky and Kahneman (1973)): in assessing the likelihood of particular events, people are often influenced by the availability of similar events in their past experience. For example, a worker who is often in contact with jobless individuals for example because he is jobless himself, would typically overestimate the rate of unemployment; similarly, an employed worker would typically underestimate the rate of unemployment. (See Nisbett and Ross (1980), page 19). In the same vein, biases in self-evaluations may then be the result of some particular past experiences being more readily available than others: if good experiences are more easily remembered than bad ones, or if failures tend to be disregarded or attributed to atypical circumstances, people will tend to have overly optimistic self-evaluations (see Seligman 1990 for more details on the connection between attributional styles and optimism).

These biases are often seen as "shortcomings" of judgement or pathologies that should be corrected. Certainly, classical decision theory views such biases as defects that can only lower the welfare of an individual afflicted by them in decision making.

And a natural question is then: why don't people learn to correct these biases? Psychology has been more open to possible positive benefits of biased perceptions. Taylor and Brown (1988) survey a considerable research that suggests that overly positive self-evaluation, exaggerated perceptions of mastery and unrealistic optimism are characteristic of normal individuals, and that moreover, these illusions appear to promote productive and creative work.

There are a number of ways in which an individual's psychological state might affect his or her performance. Steele and Aronson (1995) provide evidence that stress can impair performance. Blacks and whites were given a series of difficult verbal GRE problems to solve under two conditions: the first was described as diagnostic, to be interpreted as evaluating the individual taking the test, while the purpose of the second test was described as determining how individuals solved problems. Steele and Aronson interpret the diagnostic condition as putting Black subjects at risk of fulfilling the racial stereotype about their intellectual ability, causing self-doubt, anxiety, etc., about living up to this stereotype. Blacks performed more poorly when stress is induced (the diagnostic condition) than under the neutral condition.[1]

Aronson, Lustina, Good, and Keough (1999) demonstrate a similar effect in white males. Math-proficient white males did measurably worse on a math test when they were explicitly told prior to the test that there was a stereotype that Asian students outperform Caucasian students in mathematical areas than did similar students to whom this statement was not made.

Stress is not the only psychological condition that can affect performance. Ellis et al. (1997) showed that mood affected subjects' ability to detect contradictory statements in written passages. They induced either a neutral or a depressed mood by having participants read aloud a sequence of twenty-five self referent statements. An example of a depressed statement was "I feel a little down today," and an example of a neutral statement was "Sante Fe is the capital of New Mexico."[2] They found that those individuals in whom a depressed mood was induced were consistently impaired in detecting contradictions in prose passages. McKenna and Lewis (1994) similarly demonstrated that induced mood affected articulation. Depressed or elated moods were induced in participants, who were then asked to count aloud to 50 as quickly as possible. Performance was retarded in the depressed group.

Physical reaction time was shown to be affected by induced mood by Brand, Verspui and Oving (1997). Subjects were randomly assigned to two groups with positive mood induced in one and negative mood in the other. Positive mood was induced by showing the (Dutch) subjects a seven minute video consisting of fragments of the 1988 European soccer championship in which the Dutch team won the title, followed by a two minute

[1] See also Steele and Aronson (1998) for a closely related experiment.

[2] See Teasdale and Russell (1983) for a fuller description and discussion of this type of mood induction.

fragment of the movie "Beethoven" including a few scenes of a puppy dog. Negative mood was induced by showing subjects an eight minute fragment of the film "Faces of Death" consisting of a live recorded execution (by electric chair) of a criminal.[3] Subjects with positive induced mood showed faster response times than did subjects with negative induced mood.

There is a large number of studies that demonstrate that an individual's psychological state affects performance. We will mention one last study before moving on. Baker et al. (1997) also induce elated and depressed moods in subjects,[4] and show that induced mood affects subjects' performance on a verbal fluency task. In addition, Baker et al. measure regional cerebral blood flow using Positron Emission Tomography (PET). They find that induced mood is associated with activation of areas of the brain associated with the experience of emotion. This last finding is of particular interest in that it points to demonstrable physiological effects of mood.

Our aim in this paper is to model choice in a framework in which psychological state or beliefs can affect performance. We will show that when psychological state affects performance, biases in information processing are a natural response. Stated another way, far from being a liability, biases increase welfare.

Finally let us mention one application of our result to the Economics and Psychology literature that studies psychological games.[5] In these games, beliefs affect utilities and our analysis questions whether the standard assumption that beliefs should be correct in equilibrium is a valid one. If agents have some (possibly limited) control over their attributional styles, they should learn to bias their beliefs in a way that increases welfare, and in the long-run, we should not expect beliefs to be correct.

We present the basic model, including the possibility that an individual's perceptions can be biased, in the next section.

2. Basic model

We consider an agent who faces a sequence of decisions of whether or not to undertake a risky activity. This activity may either be a success or a failure. We have in mind a situation as follows. Consider a lawyer who is faced with a sequence of cases that he may accept or decline. Accepting a case is a risky prospect as he may win or lose the

[3] Subjects in this group were explicitly told in advance that they were free to stop the video if the emotional impact of the fragments would become too strong. Seven of the twenty subjects made use of this possibility about halfway through the video.

[4] Depressed, neutral and elated states are induced by having subjects listen to extracts of Prokofiev's "Russia under the Mongolian Yoke" at half speed in the depressed state, "Stressbusters," a recording of popular classics in the neutral condition, and Delibes' "Coppelia" in the elated condition.

[5] See Geanakoplos, Pearce and Stacchetti (1989), and Rabin (1993) for applications of psychological games.

case. The lawyer will receive a payoff of 1 in case of success, and a payoff normalized to 0 in case of failure.

Our primary departure from conventional decision theory is that we assume that the lawyer's probability of winning the case depends on his confidence: if he is unsure about his abilities, or anxious, his arguments will have less chance of convincing the jury. To capture the idea that confidence affects performance in the simplest way, we identify these feeling of anxiety or self assurance with the lawyer's belief about his chance of success.[6] A lawyer who is more confident about his chances of success will be assumed to have greater chances of succeeding than a lawyer who is not confident.

In what follows, we present two main building blocks of our model: First we describe the risky activity and the true prospects of the agent if he undertakes it. Second, we formalize the process by which an agent recollects past outcomes and describe how past outcomes translate into perceived chances of success. Finally, we describe the decision rule followed by the agent.

2.1. The risky activity.

Formally, we denote by p the *belief* of the agent about his chance of success, and by ρ the *actual* probability of success. In standard models, the probability ρ is independent of the agent's belief p; we assume that the probability of success is a differentiable function of the agent's beliefs, $\boldsymbol{\rho} : [0,1] \to [0,1]$, and that $\boldsymbol{\rho}'(p) \in [0,1)$ (see figure 1 below). The positive slope captures the idea that high (respectively low) perceptions affect performance positively (respectively negatively).

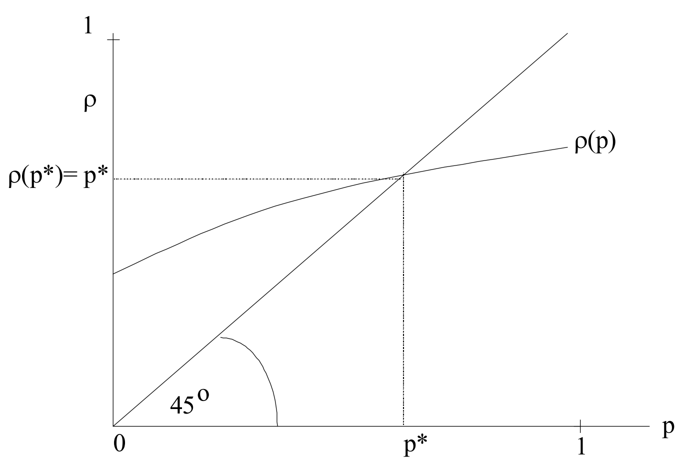

Figure 1

[6] We will discuss alternative formalizations in later sections.

Among the possible beliefs that an agent may have about the probability of success, there is one that will play a focal role. Consider the belief[7] p^* such that

$$p^* = \boldsymbol{\rho}(p^*).$$

Whenever the agent holds a belief p that is below p^*, the actual probability of success exceeds his belief. Similarly whenever the agent holds a belief p that lies above p^*, the actual probability of success is below his belief. The probability p^* corresponds to correct belief, that is, the belief at which the agent's perception is statistically accurate. An agent who holds this belief would not find his average experience at variance with his perception.

We will say that an agent has *currently optimistic beliefs* if $p > \boldsymbol{\rho}(p)$ and that he has *currently pessimistic beliefs* if $p < \boldsymbol{\rho}(p)$. (We use the term "currently" because the agent's beliefs may change over time and the comparison is between his current belief and the probability of winning given that belief.)

2.2. The agent's perceived chance of success.

We assume that the agent's *perception of the probability of success of the risky activity* (or belief) is based on his recollection of past outcomes. An agent who mostly recollects successes should believe that the probability of success is high, while an agent who mostly recollects failure should believe the probability of success is low. Formally, we denote by f the number of failures that the agent recalls, and by s the number of successes that the agent recalls. Assume that

$$p = \beta(s, f),$$

where β satisfies the following properties:

(i) $\forall s, f,\ 0 < \beta(s, f) < 1$
(ii) $\beta_1' > 0,\ \beta_2' < 0$
(iii) for any $a > 0$, $\lim_{f \uparrow \infty} \beta(af, f) = \frac{a}{a+1}$

The third property is key. It says that when the number of recalled past outcomes is large, the belief of the agent approaches the ratio $\frac{s}{s+f}$, that is, it approaches the *(perceived) empirical frequency of successes.*

These assumptions can be interpreted as a formalization of the *availability heuristic*: the numbers s and f describe the data that is available to the agent, and he forms his belief exclusively based on this data; in particular, he ignores the fact that the process by which this data is generated might be biased, and lead him to a biased estimate of the frequency of success.

[7] There is a unique p^*satisfying this equation due to our assumption that $\boldsymbol{\rho}' \in [0, 1)$.

The critical issue is thus what the agent recollects. Obviously, there should be some close relationship between true past outcomes and what the agent recalls. In many contexts however, the psychology literature has identified various factors such as generation of attributions or inferences, memorability, perceptual salience, vividness, that may affect what the agent perceives, records and recalls, hence the data that is available to the agent.[8]

To illustrate the role of attributions, for example, consider the case of our lawyer. When he loses, he may think that this was due to the fact that the defendant was a member of a minority group while the jury was all white: the activity was a failure, but the reasons for the failure are seen as atypical and not likely to arise in the future. He may then disregard the event when evaluating his overall success. In extreme cases, he may even convince himself that had the circumstances been "normal", he would have won the case and record the event as a success in his mind.[9]

This basic notion is not new; there is evidence in the psychology literature that people tend to attribute positive experiences to things that are permanent and to attribute negative experiences to transient effects.[10] If I get a paper accepted by a journal, I attribute this to my being a good economist, while rejections are attributed to the referee not understanding the obviously important point I was making. With such attribution, successes are likely to be perceived as predicting further successes, while failures have no predictive content.

We now formalize these ideas.

The perception matrix We assume that when he undertakes the activity, the agent receives a signal $y \in Y$ that may affect either his perception of the outcome, or the way he will later recollect the current experience. We will refer to y as the agent's *perception*. Formally, the set of outcomes is denoted X, and for any $x \in X$ and $y \in Y$, we let $\pi_{x,y}$ denote the probability that the agent's perception is y when x is the true outcome, and let $\pi = (\pi_{x,y})_{x,y \in X \times Y}$ denote the corresponding matrix.

In our basic model, there are only two possible outcomes:

$$X = \{S, F\}.$$

For example, in the lawyer problem, we would set

$$Y = \{\text{normal(n), atypical(a)}\} \tag{2.1}$$

[8] An example of biased recall, presumably induced by better memorability of own actions, is given in Ross and Sicoly (1979), who report that when married couples are asked to estimate the percentage of household tasks they perform, their estimates typically add up to far more than 100%.

[9] We mentioned above the paper by Svenson (1981) in which a majority of people felt they were safer drivers than average. This is consistent with a situation in which drivers attribute accidents or near accidents in which they are involved to atypical circumstances.

[10] See, e.g., Seligman (1990).

where y =n indicates normal circumstances, and y =a indicates an attribution of the outcome to atypical circumstances; and we would set

$$\pi = \begin{pmatrix} & n & a \\ S & 1 & 0 \\ F & 1-\gamma & \gamma \end{pmatrix}, \tag{2.2}$$

where γ is the probability that the lawyer attributes a failure to the particular circumstances surrounding the case.

The recall technology Perceptions are assumed to play a key role in the recollection process. Consider an agent who wishes to measure the number of occurrences of the outcome $x_0 \in X$. If the agent could perfectly recollect past outcomes, he would count the number of experiences for which $x = x_0$. However, we assume that both the true outcome x and the perception y may affect whether a particular experience xy (that is, an outcome x and perception y) is included in the total.

Formally, let Ω denote the set of past experiences $\omega = xy$, that is, $\Omega = X \times Y$. We denote by $\mathbf{1}_{\omega,x_0}$ the random variable that takes a value equal to 1 when the experience $\omega = xy$ is included in the agent's enumeration of the number of past occurrences of x_0, and 0 otherwise. The *recall technology* is characterized by the vector of probabilities $\mathbf{q} = \{q_{\omega,x_0}\}_{\omega,x_0}$, where

$$q_{\omega,x_0} = \Pr\{\mathbf{1}_{\omega,x_0} = 1\},$$

and we define the random variable:

$$\nu(x_0) = \sum_{\omega\in\Omega} \mathbf{1}_{\omega,x_0}.$$

For any given set of past experiences, the numbers of recalled successes and failures, s and f, are thus random variables, respectively equal to

$$s = \nu(S)$$

and

$$f = \nu(F).$$

To illustrate with the lawyer example, we define the recall technology q in which the agent only recalls experiences for which circumstances were viewed as normal:

$$\begin{aligned} q_{xy,x_0} &= 1 \text{ if } x = x_0 \text{ and } y = \text{n, and} \\ q_{xy,x_0} &= 0 \text{ otherwise.} \end{aligned} \tag{2.3}$$

Long-run perceived frequency of success. We are now equipped to define the long-run perceived frequency of success, as a function of the true outcome frequencies. Let $\alpha = (\alpha_x)_{x\in X}$ be the true frequencies over outcomes. The expected number of events when tallying the occurrences of x_0 is equal to

$$E_{(\pi,\mathbf{q},\alpha)}\nu(x_0) = \sum_{x\in X, y\in Y} q_{xy,x_0}\pi_{xy}\alpha_x.$$

At any date t, the perceived frequency of outcome x_0 is a random variable. As experience accumulates however, its distribution will be concentrated on values close to the ratio

$$\frac{E_{(\pi,\mathbf{q},\alpha)}\nu(x_0)}{\sum_{x\in X} E_{(\pi,\mathbf{q},\alpha)}\nu(x)}.$$

2.3. Decision rule.

As described in the outline of our model above, the agent faces a sequence of decisions of whether or not to undertake a risky activity. We assume that undertaking the activity entails a cost c, possibly an opportunity cost. For example, in the lawyer example above, there might always be available some riskless alternative to taking on a new case, for instance drawing up a will. We assume the cost c is stochastic and takes values in $[0, 1]$. We further assume that the random variables $\{c_t\}_{t=1}^{\infty}$, the costs at each time t, are independent, that the support of the random variables is $[0, 1]$, and that at the time the agent chooses whether to undertake the activity at t, he knows the realization of c_t.

At any date, the agent has a belief p_t about his chance of success if he undertakes the activity. The agent compares the expected payoff from undertaking the activity to the cost of undertaking it. That is, the agent undertakes the activity if and only if

$$p_t \geq c_t.$$

Finally, we define the expected payoff $v(p)$ to the agent who has a belief p and who does not yet know the realization of the cost of the activity:[11]

$$v(p) = \Pr\{p \geq c\} E[\boldsymbol{\rho}(p) - c \mid p \geq c].$$

[11] Note that this expected utility is from the perspective of an outside observer, since it is calculated with the true probability of success, $\rho(p)$. From the agent's point of view, the expected value is

$$\Pr\{p \geq c\} E[p - c \mid p \geq c]$$

.

This completes the description of the model. The key parameters of the model are the technology $\boldsymbol{\rho}(\cdot)$, the belief function β, the perception matrix π, and the recall technology $\mathbf{q}$. Each quadruple $(\boldsymbol{\rho}, \beta, \pi, \mathbf{q})$ induces a probability distribution over beliefs, decisions and outcomes. We are interested in the limit distribution and the agent's expected gain under that limit distribution when the agent's perceptions are given by the matrix π. Formally, define

$$V_t(\pi, \mathbf{q}) = E_{(\boldsymbol{\rho}, \beta, \pi, \mathbf{q})} v(p_t).$$

We are interested in the long-run payoff defined by $V_\infty(\pi, \mathbf{q}) = \lim_{t \to \infty} V_t(\pi, \mathbf{q})$. We will suppress the dependence on π and $\mathbf{q}$ when there is no confusion about the agent's perceptions and simply write V_∞.

3. A benchmark: the cost of biased perceptions.

We begin with the standard case in which confidence does not affect performance, that is, the case in which $\boldsymbol{\rho}(p) = \rho_0$, independent of p. Our objective here is to highlight the potential cost associated with biased perceptions. To simplify presentation, we now restrict our attention to the simple perception matrix π and recollection technology $\mathbf{q}$ defined in (2.1-2.3) for which the agent perceives some failures as atypical, and recalls only experiences for which circumstances were viewed as normal. This case can thus be parameterized by the single parameter γ, the probability that experiences were seen as atypical. It will be convenient to refer to the case where $\gamma = 0$ as the correct perception case.

We let ρ_0 denote the true probability of success. As the number of instances where the activity is undertaken increases, the frequency of success gets close to ρ_0 (with probability close to 1). When perceptions are correct, the true and perceived frequencies of success coincide. Hence, given our assumption (iii) concerning β, the agent's belief must converge to ρ_0 as well. It follows that

$$V_\infty = \int_{\rho_0 \geq c} (\rho_0 - c) g(c) dc$$

where $g(\cdot)$ is the density function for the random cost.

When perceptions are not correct, the true and perceived frequencies of success generally do not coincide. Let α denote the true frequency of success, and define for any π,$\mathbf{q}$, the function $\psi^{\pi,\mathbf{q}} : R \to R$ as

$$\psi^{\pi,\mathbf{q}}(\alpha) = \frac{E_{(\pi,\mathbf{q},\alpha)} \nu(S)}{E_{(\pi,\mathbf{q},\alpha)} \nu(S) + E_{(\pi,\mathbf{q},\alpha)} \nu(F)}.$$

$\psi^{\pi,\mathbf{q}}(\alpha)$ is essentially the proportion of recorded events that are recorded as success when the true frequency of success is α and π is the perception matrix. For the particular perception matrix π and recollection technology $\mathbf{q}$ defined in (2.1-2.3), we have

$$\psi^{\pi,\mathbf{q}}(\alpha) = \frac{\alpha}{\alpha + (1-\gamma)(1-\alpha)}.$$

The perceived frequency of success at date t is a random variable. Yet for t large, its distribution is concentrated on values close to $\psi^{\pi,\mathbf{q}}(\alpha)$. And therefore, given our assumption (iii) concerning β, beliefs are also close to $\psi^{\pi,\mathbf{q}}(\alpha)$. Figure 2 illustrates how true frequencies of success may translate into beliefs. $(\psi^{\pi,\mathbf{q}})^{-1}(p)$ is the probability of success that would be necessary for the agent to perceive his success rate as p.

We shall say that perceptions are optimistic when $\psi^{\pi,\mathbf{q}}(\alpha) > \alpha$ for all α, and pessimistic when $\psi^{\pi,\mathbf{q}}(\alpha) < \alpha$ for all α.

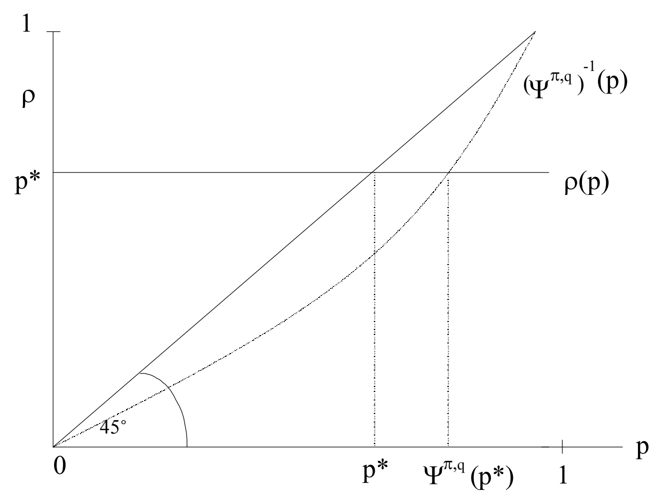

Figure 2

How do biased perceptions affect payoffs? By the law of large numbers, the true frequency of success must be close to ρ_0 (with probability close to 1). Yet the agent decides to undertake the project whenever $\psi_\pi(\rho_0) \geq c$. It follows that

$$V_\infty = \int_{\psi^{\pi,\mathbf{q}}(\rho_0)\geq c} (\rho_0 - c)g(c)dc.$$

The only effect of the perception bias is to (possibly)[12] change the circumstances under which the activity is undertaken. For biased perceptions that lead to optimistic

[12] While it is generally true that biased perceptions lead to incorrect beliefs, this may not be the case in specific circumstances. For example, if an agent records half his successes and half his failures as "no activity", his beliefs will be accurate in the long run.

beliefs, for example, there are events for which

$$\psi^{\pi,\mathbf{q}}(\rho_0) > c > \rho_0.$$

In these events, the agent undertakes the activity while he should not do so (since they have negative expected payoff with respect to the true probability of success), and in the other events, he is taking the correct decision. So the agent's (true) welfare would be higher if he had correct perceptions. This is essentially the argument in classical decision theory that biasing perceptions can only harm agents.

4. Confidence enhanced performance

When confidence affects performance, it is no longer true that correct perceptions maximize long-term payoffs. It will still be the case that agents with optimistic perceptions will be induced to undertake the activity in events where they should not have, but on those projects they undertake, their optimism leads to higher performance, that is, they have higher probability of success. We will compare these two effects and show that having some degree of optimism is preferable to correct perceptions.

The key observation is that in the long run, beliefs tend to be concentrated around a single value, say p_∞, and this value p_∞ is therefore easy to characterize. The realized frequency of successes $\rho_t = \boldsymbol{\rho}(p_t)$ is with high probability near $\rho_\infty = \boldsymbol{\rho}(p_\infty)$. True and perceived frequencies of success are thus concentrated around ρ_∞ and $\psi^{\pi,\mathbf{q}}(\rho_\infty)$ respectively. The only possible candidate for p_∞ must therefore satisfy

$$\psi^{\pi,\mathbf{q}}(\boldsymbol{\rho}(p_\infty)) = p_\infty.$$ [13]

Figure 3 illustrates geometrically p_∞ for an agent with optimistic beliefs.

[13] We restrict attention to the case that there is a unique such p_∞.

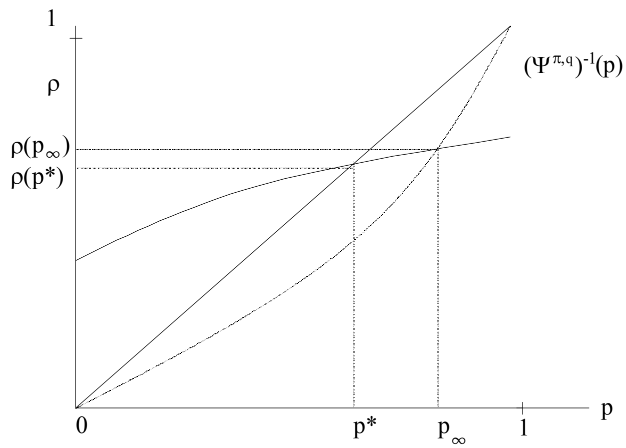

Figure 3

Note that in the case that perceptions are correct, p_∞ coincides with the correct belief p^* defined earlier, but that for an agent for whom $\boldsymbol{\rho}'(\cdot) > 0$,

$$p_\infty > \boldsymbol{\rho}(p_\infty) > \boldsymbol{\rho}(p^*) = p^*.$$

Definition: We say that an agent has *optimistic beliefs* (is *optimistic)* if $p_\infty > p^*$.

Turning to the long-run payoff, we obtain:

$$V_\infty = \int_{p_\infty \geq c} (\boldsymbol{\rho}(p_\infty) - c) g(c) dc.$$

As a benchmark, let V_∞^* denote the long-run payoff when perceptions are correct, that is,

$$V_\infty^* = \int_{p^* \geq c} (p^* - c) g(c) dc.$$

To understand how V_∞ compares to V_∞^*, we write $V_\infty - V_\infty^*$ as the sum of three terms:

$$\begin{aligned} V_\infty - V_\infty^* \quad = \quad & \int_{c \leq p^*} (\boldsymbol{\rho}(p_\infty) - \boldsymbol{\rho}(p^*)) g(c) dc + \int_{p^*}^{\boldsymbol{\rho}(p_\infty)} (\boldsymbol{\rho}(p_\infty) - c) g(c) dc \qquad (4.1) \\ & + \int_{\boldsymbol{\rho}(p_\infty)}^{p_\infty} (\boldsymbol{\rho}(p_\infty) - c) g(c) dc \end{aligned}$$

The first term is positive and corresponds to the increase in performance that arises due optimism for the activities that would have been undertaken even if perceptions had been correct. The second term is positive: it corresponds to the realizations of costs for which the activity is profitable to the agent only because he is optimistic. Finally, the third term is negative: it corresponds to the realizations of costs for which the agent should not have undertaken the activity and undertakes it because he is optimistic. The shaded regions in figure 4 represent these three terms when c is uniformly distributed.

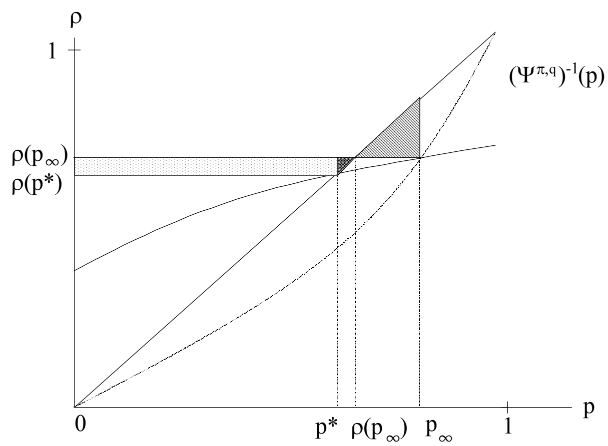

Figure 4

One implication of equation (4.1) is that when confidence positively affects performance ($\boldsymbol{\rho}' > 0$), some degree of optimism always generates higher long-run payoff. The reason is the following. As p_∞ rises above p^*, the first term generates a first order increase in long-run payoff, while the second and third term only give rise to second order changes: the positive effect on performance induced by optimism applies to many realizations of costs, while the possible negative effect applies to few of them. Formally,

$$\left.\frac{dV_\infty}{dp_\infty}\right|_{p_\infty=p^*} = \boldsymbol{\rho}'(p^*)\,\Pr\{p^* \geq c\} > 0.$$

To formalize this argument, we introduce the following definition:

Definition: We say that $(\pi, \mathbf{q})$ *dominates* $(\pi', \mathbf{q}')$ if $V_\infty(\pi, \mathbf{q}) > V_\infty(\pi', \mathbf{q}')$.

That is, if $(\pi, \mathbf{q})$ *dominates* $(\pi', \mathbf{q}')$, an agent's long run payoff will be higher under $(\pi, \mathbf{q})$ than under $(\pi', \mathbf{q}')$. The following proposition summarizes the argument above.

Proposition: If $\boldsymbol{\rho}' > 0$, there exists a pair $(\pi, \mathbf{q})$ that induces optimistic beliefs and that dominates correct perceptions and recall.

5. Sources of optimism and attributional styles

We discussed in the previous sections our simple model that incorporated biased perceptions, and showed that in the context of our model the biasedness of perceptions increased welfare. In incorporating biased perceptions, we chose a particularly simple form of bias: individuals selectively remembered past experience, recalling a greater proportion of their past successes than their past failures. This is a particularly salient means by which biased perceptions can lead to optimism, but by no means the only one.

The objective of this section is to explain how our insight extends to more complex environments, and discuss several additional ways in which biased perceptions or attributional styles can lead to optimism.

In the model analyzed above, the agent sometimes attributes the cause of failure to events that are not likely to occur again in the future, while he never makes such attributions in case of success (non-permanence of failure/permanence of success). An alternative way in which an agent's perceptions may be biased is that rather than attributing failures to transient circumstances, he attributes successes to a broader set of circumstances than is appropriate; in the psychology literature, this is referred to as pervasiveness of success bias (see Seligman 1990).

We modify the environment in which the agent is making decisions in the following way. There are two possible environments, and the probability of success may differ in the two environments. Consider again the lawyer example. Suppose that cases that are available to the lawyer can be of two types: high profile cases that will attract much popular attention and low profile cases that will not. Suppose that in high profile cases, there will be many more people in court, including reporters and photographers, and consequently, much more stress, which we assume impairs performance. If we denote high and low profile cases by H and L respectively, $\boldsymbol{\rho}^H(p)$ represents the probability of success in a high profile case when the lawyer believes the probability of success is p, and $\boldsymbol{\rho}^L(p)$ is the probability of success in a low profile case. Analogous to what we have done above, we assume $\frac{d\rho^H}{dp} > 0$, and, to simplify exposition, $\frac{d\rho^L}{dp} = 0$. We let $\lambda = \Pr\{L\}$.

The set of (state contingent) outcomes (x, k) is

$$X = \{S, F\} \times \{H, L\}.$$

Our interest in this section is to capture an aspect of optimism that we call "pervasiveness". We can illustrate the idea of pervasiveness with our lawyer example. Assume that it is more difficult to perform with public attention than when there is not the attention, and suppose that one day the lawyer wins a case when there was no attention. Although the lawyer realizes that success would have been less likely had there been attention, he thinks "The arguments I made were so good, I would have won even if this had been a high profile case." In thinking this way, the lawyer is effectively recording the experience as a success whether it had been high profile or not. We model the possibility of this type of biased recording by assuming that the set of perceptions consists of

$$Y = \{specific(s), non\text{-}specific(n)\}$$

with the interpretation that $y = s$ means that the outcome is specific to the particular current environment, while $y = n$ means that it is not.

The recall technology (see Section 2) specifies for each possible outcome $(x_0, k_0) \in X$, the probability $q_{(x,k)y,(x_0,k_0)}$ that a past experience $\omega = (x,k)y$ (an experience here consists of an outcome $(x,k) \in X$ and a perception $y \in Y$) is counted. We shall consider the recall technology defined by

$$\begin{aligned} q_{(x,k)y,(x_0,k_0)} &= 1 \text{ if } \left\{ \begin{array}{c} (x,k) = (x_0,k_0) \text{ and } y = s \\ \text{or} \\ x = x_0 \text{ and } y = n \end{array} \right. , \text{ and} \\ &= 0 \text{ otherwise.} \end{aligned} \tag{5.1}$$

and the perception matrix defined by

$$\pi^\gamma = \begin{pmatrix} & s & n \\ SH & 1 & 0 \\ SL & 1-\gamma & \gamma \\ FH & 1 & 0 \\ FL & 1 & 0 \end{pmatrix}. \tag{5.2}$$

This perception matrix captures the intuition described above, in which sometimes the agent attributes the cause of success to circumstances that are broader than appropriate, while the recall technology above captures the intuition that the inferences made at the time of the activity affect later recollection of past experiences. This inability to distinguish between outcomes and inferences when one attempts to recall past events is well documented.[14]

The next proposition provides sufficient conditions under which unbiased perceptions will necessarily be dominated by some type of biased perceptions.

[14] A particularly compelling example of this was demonstrated by Kahneman, Fredrickson, Schreiber and Redelmeier (1993). Subjects were asked to rate two painful experimental procedures that involved submersion of the hand in ice cold water. Later they were given a choice to repeat one or the other of the two painful procedures. The subjects could be induced to choose the more painful procedure (as measured by self reports at the time of the experience) if the more painful procedure happened to include a decrease in the pain at the end. Redelmeier and Kahneman (1996) find a similar dissociation between contemporaneous reporting of pain and later recall of that pain in an actual medical procedure.

Proposition:
$(\pi^\gamma, \mathbf{q})$ dominates perfect recall for some $\gamma > 0$.

Proof: It is easy to check that in the long run that beliefs about probability of success in state k tend to be concentrated around a single value, denoted p^k_∞.[15] Since the agent undertakes the activity in circumstance k when $p^k_\infty \geq c$, we have

$$V^\infty = \lambda \int_{p^L_\infty \geq c} (\boldsymbol{\rho}^L(p^L_\infty) - c)g(c)dc + (1-\lambda) \int_{p^H_\infty \geq c} (\boldsymbol{\rho}^H(p^H_\infty) - c)g(c)dc$$

For the pair $\pi, \mathbf{q}$ defined in (5.1)-(5.2), we will check that the values p^L_∞ and p^H_∞ depend solely on γ, and are equal to the correct beliefs defined by:

$$p^*_k = \boldsymbol{\rho}^k(p^*_k)$$

when $\gamma = 0$. Hence, since $\frac{d\boldsymbol{\rho}^L}{dp} = 0$, we have:

$$\left.\frac{dV^\infty}{d\gamma}\right|_{\gamma=0} = (1-\lambda) \left.\frac{dp^H_\infty}{d\gamma}\right|_{\gamma=0} \left.\frac{d\boldsymbol{\rho}^H}{dp}\right|_{p=p^*_H} G(p^H),$$

where G denotes the cumulative distribution associated with the density g. We will thus have shown our result if we can show that $\left.\frac{dp^H_\infty}{d\gamma}\right|_{\gamma=0} > 0$.

We now characterize the values p^L_∞ and p^H_∞. Consider a pair of beliefs (p^L, p^H), assumed to remain the same over time. In the long-run, these beliefs induce a frequency over outcomes, denoted $\alpha(p^L, p^H)$.[16] Similarly to the simple case analyzed in the previous section, we define the functions $\psi^k_{\pi,\mathbf{q}}$ as

$$\psi^k_{\pi,\mathbf{q}}(p^L, p^H) = \frac{E_{(\pi;\mathbf{q};\alpha(p^L,p^H))}\nu(S,k)}{E_{(\pi;\mathbf{q};\alpha(p^L,p^H))}\nu(S,k) + E_{(\pi;\mathbf{q};\alpha(p^L,p^H))}\nu(F,k)}$$

The perceived frequencies of successes (contingent on k) at date t are random variables. Yet for t large, their distributions are concentrated on values close to $\psi^k_\pi(p^L, p^H)$, $k =$

[15] This is obviously the case for beliefs in state L, because in this state performance does not depend on beliefs and perceptions are correct (hence beliefs in state L must converge to the true probability of success in state L). For beliefs in state H, the analysis is then similar to that of the basic model.

[16] These long-run frequencies satisfy:

$$\begin{aligned}\alpha(p^L, p^H) \quad &= \quad \frac{1}{\lambda G(p^L) + (1-\lambda)G(p^H)}(\lambda\rho^L(p^L)G(p^L), \lambda(1-\rho^L(p^L))G(p^L), ... \\ &\quad ...(1-\lambda)\rho^H(p^H)G(p^H), (1-\lambda)(1-\rho^H(p^H))G(p^H)) \qquad (5.3)\end{aligned}$$

0, 1. The values p^L_∞ and p^H_∞ therefore solve:

$$\begin{aligned} p^L_\infty &= \psi^L_{\pi,\mathbf{q}}(p^L_\infty, p^H_\infty) \\ p^H_\infty &= \psi^H_{\pi,\mathbf{q}}(p^L_\infty, p^H_\infty). \end{aligned} \tag{5.4}$$

For the pair $\pi, \mathbf{q}$ defined in (5.1)-(5.2), we have:

$$\psi^H_\pi(p^L, p^H) = \frac{\boldsymbol{\rho}(p^H) + \gamma\boldsymbol{\rho}(p^L)\mu(p^H, p^L)}{1 + \gamma\boldsymbol{\rho}(p^L)\mu(p^H, p^L)}(\geq \boldsymbol{\rho}(p^H)),$$

where

$$\mu(p^H, p^L) \equiv \frac{\lambda G(p^L)}{(1-\lambda)G(p^H)}.$$

Equation (5.2) implies

$$\boldsymbol{\rho}(p^H_\infty) - p^H_\infty = \gamma\boldsymbol{\rho}(p^L_\infty)(1 - p^H_\infty)\mu(p^H_\infty, p^L_\infty).$$

Thus for $\gamma = 0$, p^H_∞ must coincide with the correct belief p^*_H, and for $\gamma > 0$, we have $p^H_\infty > p^*_H$, hence $\left.\frac{dp^H_\infty}{d\gamma}\right|_{\gamma=0} > 0$, as desired.

Comments:

(1) One possible interpretation of environment H and L is that they each correspond to a particular activity. It is easy to check that our analysis could be adapted to deal with the case where at each date, the agent has to decide whether to undertake activity H, activity L, or undertake no activity.[17] Optimistic beliefs lead to a sub-optimal decision between these three alternatives, however the effect on welfare is small compared to the gain on performance that the agents derive from having beliefs more optimistic than correct ones.

(2) Attributions or inferences is only one channel by which beliefs about chances of success may become biased. Another possibility is that the agent has a biased memory technology, and for example, tends to remember more easily past successes than past failures.[18] Optimism may also stem from *self-serving bias*. Such bias occurs in settings where there is not a compelling unique way to measure success. When I give a lecture, I might naturally feel it was a success if the audience enthusiastically applauds when I'm finished. If that, unfortunately, doesn't occur, but an individual comes up to me following the lecture and tells me that it was very stimulating, I can choose that as a signal that the lecture was a "success".

Formally, consider an activity for which success may be measured in different ways: $x \in S, F$ or $y \in s, f$. Assume that x is what the agent cares about and that he decides

[17] The agent would choose the decision that maximizes $p^H - c^H$, $p^L - c^L$, 0.

[18] See Morris (1999) for a survey of the evidence that an individual's mood affects his or her recollection process, e.g.

to undertake the activity when $p_x > c$. Optimism may arise when the experiences $xy = Fs$ have a positive probability of being recalled when measuring the number of past successes (that is, the occurrences of past $x = S$).

6. Discussion

We have identified circumstances in which having correct beliefs is not welfare maximizing. Hence, to the extent that an agent has some (possibly limited) control over the way he forms his beliefs, we should not expect correct beliefs to emerge in the long run: rational agents should learn to bias their beliefs.

This conclusion seems to be at odds with the general view that having correct beliefs is welfare maximizing, and that rational agents should have correct beliefs. We discuss below why our model conflicts with these standard views.

6.1. Biased beliefs and welfare.

Our model belongs to a general class of problems in which correct beliefs are not welfare maximizing. Beliefs are correct when they coincide with the true probability of success, hence when $p = \boldsymbol{\rho}(p)$; beliefs are biased when, for example,

$$p = \lambda \boldsymbol{\rho}(p)$$

with $\lambda \neq 1$. The scalar λ parametrizes the bias. Let $\bar{p}(\lambda)$ denote the resulting belief. We may write the agent's expected welfare as a function of the bias λ:

$$w(\lambda) = \int_{\bar{p}(\lambda) > c} (\boldsymbol{\rho}(\bar{p}(\lambda)) - c) g(c) dc. \tag{6.1}$$

Differentiating (6.1), we get:

$$\begin{aligned} w'(\lambda) &= \bar{p}'(\lambda)[(\boldsymbol{\rho}(\bar{p}(\lambda)) - \bar{p}(\lambda))g(\bar{p}(\lambda)) + \boldsymbol{\rho}'(\bar{p}(\lambda))G(\bar{p}(\lambda))] \\ &= \bar{p}'(\lambda)[-(1 - \frac{1}{\lambda})\bar{p}(\lambda)g(\bar{p}(\lambda)) + \boldsymbol{\rho}'(\bar{p}(\lambda))G(\bar{p}(\lambda))] \end{aligned}$$

The two terms on the right hand side have the following interpretation: (i) biased beliefs (say, $\lambda > 1$) lead to suboptimal choices, and they are even more suboptimal when λ increases; (ii) increasing λ has a positive effect on performance in any event where the project is undertaken.

When beliefs are correct (that is, when $\lambda = 1$),

$$w'(1) = \bar{p}'(1)\boldsymbol{\rho}'(\bar{p}(1))G(\bar{p}(1)) > 0,$$

which illustrates why having correct beliefs is not welfare maximizing.

In standard models, beliefs only affect welfare through the decisions that the agent makes. In these models, having correct beliefs maximizes welfare. In our model, as in any other model in which beliefs would directly affect preferences or technology, beliefs affect welfare in two ways: indirectly through their effect on decisions, directly through their effect on preferences or technology. Our conclusion holds because (by the envelope theorem) the welfare loss induced by suboptimal decisions is a second order effect, while the direct effect on welfare is first order.[19]

Our model is not unique in breaking the link between correct beliefs and welfare maximization. In any model where agents do not have a correct criterion for taking their decisions, having incorrect beliefs may increase welfare. To illustrate, assume that in our model, performance is independent of beliefs, but the agent only undertakes the activity when

$$p > \mu c \text{ with } \mu > 1. \tag{6.2}$$

An interpretation might be that the agent lacks the will to undertake the risky activity, and that he only undertakes it when its expected return exceeds one by a sufficient amount. This formulation can be viewed as a reduced form of the agent's decision problem in Benabou and Tirole (2001) with μ corresponding to $\frac{1}{\beta}$ and β being the salience for the present.

Given this criterion for making the decision as to whether to undertake the risky activity, we may write expected welfare as a function of the bias λ. We obtain:

$$w(\lambda) = \int_{\lambda\rho > \mu c} (\rho - c) g(c) dc$$

which implies

$$w'(\lambda) = \frac{\rho}{\mu}(1 - \frac{\lambda}{\mu})\rho g(\frac{\lambda\rho}{\mu})$$

Hence $w'(\lambda)$ has the sign of $\mu - \lambda$, and it is positive at $\lambda = 1$ (correct beliefs) since $\mu > 1$.[20,21]

[19] Our insight would also apply to psychological games. In these games, beliefs directly affect preferences, and individuals would benefit from having biased perceptions, say of their opponent's behavior. We may thus question whether in these games, the standard assumption that players hold correct beliefs is compelling.

[20] See Samuelson (2001) for another example of a model in which an incorrectly specified decision criterion leads to welfare enhancing biases.

[21] In a strategic context, recent papers by Fang (2001) and Heifetz and Spiegel (2001) find that biased beliefs may increase welfare.

At first glance, these two papers do not seem to fall in either of the categories described above: preferences seemingly do not depend on beliefs, and players choose true best responses, given their beliefs.

Both papers, however, assume that each player (say player 1) becomes perfectly aware of his op-

6.2. Biased beliefs and rational agents

Our model conflicts with the standard view that rational agents should have correct beliefs, or that they should not be able to consistently fool themselves about their chances of success. In the standard view, it is impossible that rational agents should not eventually learn true probability distributions, while in our model, the process by which the agent recollects data is biased, and he ends up with incorrect beliefs.

A skeptic might argue that through introspection, an agent should have some awareness of the fact that the data he recollects is biased. If through introspection, he realizes that there is a good chance that his sample of past experiences over-represents successes, he should disregard selectively some of these past successful experiences.

Our basic model can be easily amended to allow for introspection. Formally, we would now set

$$\nu(x_0) = \sum 1_{x_0}.1_{xy,x_0}$$

where 1_{x_0} is a random variable that takes value 1 with probability ζ_{x_0}, and we would choose

$$\zeta_S < 1 \text{ and } \zeta_F = 1$$

with the interpretation that $\zeta_S < 1$ because, due to introspection, the agent is aware that the data he recollects is biased toward success, and consequently, he sometimes disregards experiences that he recalls as successes.

Belief formation would then be driven by two key parameters: γ, which captures the extent to which the agent attributes failures to atypical circumstances; ζ_S which captures the extent to which he believes his sample is biased toward success. In the long run, we should expect beliefs to be correct if and only if ζ_S exactly offsets γ.

Our view however is that there is no reason to believe ζ_S would exactly offset γ. First, we demonstrated above that if introspection were to exactly offset attribution biases, this would be welfare-detrimental in the types of settings we analyzed, hence, one should expect evolutionary pressures *against* such introspection. Second, such introspection would require that the agent *learns* the way he biases his perception of past outcomes. Learning this however would require that he (perfectly) recalls the precise way in which past experiences have been biased. Yet nothing prevents this recollection process to be biased as well.

6.3. Belief formation and control

How do agents form beliefs? What control do agents have over the way they form their beliefs?

---

ponent's (player 2's) belief (for example through "body language"), and of how this belief maps into actions. So in effect, these papers study a particular type of psychological game, in which, say, player 2's belief directly affects player 1's payoff function.

Throughout the paper, we have held the view (common in psychology) that agents form beliefs based on simple heuristics or routines. In our formal model for example, among the many routines and heuristics that may affect belief formation, we chose to focus on the two following aspects of belief formation:

(i) agents use the availability heuristic;

(ii) attributions affect later recollection of past experiences.

This has permitted us to highlight the role of attribution in biasing beliefs and identify two types of attribution that have received particular attention in the psychology literature.

Obviously, there is a rich array of heuristics, psychological routines, associations, and analogies that agents use to form beliefs. Our main purpose is not to catalog exhaustively the ways an agent might end up with optimistic beliefs. Our view is that agents have some (possibly limited) control over the type of attributions they perform, and more generally, the types of associations and analogies they construct. Given that they have such control, for the types of environment we have analyzed, agents should ultimately learn what is best for them and hold optimistic beliefs: for these environments the welfare that they derive when they are optimistic is greater than when their beliefs are correct.

In other words, we open the possibility that the agent learns which learning rule is best, possibly among a limited set of learning rules, and we find that even if the unbiased learning rule belongs to that set, agents will not use that rule even asymptotically.

6.4. Extensions and further research.

As stressed in the introduction, our main interest is in studying decision making when psychological state directly affects performance and explain how it departs from classical decision theory. To do so, we have made various simplifying assumptions. We now review some of these assumptions and propose extensions of the basic model.

(1) Psychological state.

We chose a particularly simple psychological state, namely, the agent's perceived chance of success. There are other obvious candidates that might also serve as foundation for belief-based performance.

(i) Performance could depend on one's doubt about whether the best decision was taken (especially when the risky project is undertaken – the other project gives c regardless); performance is then a function of $p - c$.

(ii) Performance could depend (positively) on the agent's belief about his own ability (even if ability is not actually key for good performance). In such settings, attributing successes to high ability and failures to lack of luck will generate optimistic beliefs about abilities.

(iii) Performance may also depend on the agent's belief about how he compares with others. If his confidence increases when he believes he fares better than others, we should expect him to attribute failures of others to permanent causes, and success to temporary causes: other's failures are pervasive.

(2) Timing.

We have not distinguished between the time at which the agent decides to undertake the activity, and the time during which he undertakes the activity. In general, these two dates need not coincide and therefore the agent's belief at these two dates need not coincide. This opens the possibility that an agent has optimistic beliefs while he undertakes the activity, and yet not so optimistic beliefs when deciding to undertake the activity. The confirmatory bias, according to which agents would tend to generate or recollect data that supports their decisions would quite naturally induce such an outcome: it boosts confidence, without overly impairing decisions.

(3) Direct evidence of success and failure.

We have assumed that the agent could observe whether the activity was a success or not. In many social interactions, the agent also (or only) gets indirect evidence, through the feedback that is given to him. The agent's perception of whether others have a vested interest in his performance is key, and optimistic beliefs would presumably arise when the agent underestimates the extent to which others have a vested interest in his performance, and as a consequence attach "too much" weight to their compliments.

Bibliography

Aronson, J., M. Lustina, C. Good, and K. Keough (1999), "When White Men Can't Do Math: Necessary and Sufficient Factors in Stereotype Threat," *Journal of Experimental Social Psychology* 35, 29-46.

Baker, L, and R. Emery (1993), "When Every Relationship is Above Average: Perceptions and Expectations of Divorce at the Time of Marriage," *Law and Human Behavior.*

Baker, S., C. Firth, and R. Dolan (1997), "The Interaction between Mood and Cognitive Function Studied with PET," *Psychological Medicine*, 27, 565-578.

Benabou, R. and J. Tirole (2000), "Self-Confidence: Intrapersonal Strategies," mimeo.

Block, R. Q. and D. R. Harper (1991), "Overconfidence in Estimation: Testing the Anchoring-and-Adjustment Hypothesis," *Organizational Behavior and Human Decision Processes,* 49(2), 186-207.

Brand, N., L. Verspui and A. Oving (1997), "Induced Mood and Selective Attention," *Perceptual and Motor Skills*, 84, 455-463.

Ellis, H. S. Ottaway, L. Varner, A. Becker and B. Moore (1997), "Emotion, Motivation, and Text Comprehension: The Detection of Contradictions in Passages," *Journal of Experimental Psychology: General*, 126(2), 131-46.

Fang, H. (2001), "Confidence Management and the Free-Riding Problem," mimeo, Yale University.

Fischoff, B., Slovic, P., and Lichtenstein, S. (1977), "Knowing with Certainty: The Appropriateness of Extreme Confidence," *Journal of Experimental Psychology: Human Perception and Performance,* 3(4), 552-564.

Geanakoplos, J., D. Pearce and E. Stacchetti (1989), "Psychological Games and Sequential Rationality," *Games and Economic Behavior* 1, 60-79.

Guthrie, C., J. Rachlinski and A. Wistrich (2001), "Inside the Judicial Mind: Heuristics and Biases," *Cornell Law Review* 86, 777-830.

Heifetz, A. and Spiegel, Y. (2001), "The Evolution of Biased Perceptions", mimeo.

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., and Redelmeier, D. A. (1993), "When More Pain is Preferred to Less: Adding a Better End," *Psychological Science*, 4, 401-5.

Lichtenstein, S. B. Fischoff, and L. Phillips (1982), "Calibration of Probabilities: The State of the Art to 1980," in D. Kahneman, P. Slovic, and A. Tversky, (eds.), *Judgement under Uncertainty: Heuristics and Biases*, pp. 306-334; Cambridge: Cambridge University Press.

McKenna, F. and C. Lewis (1994), "A Speech Rate Measure of Laboratory Induced Affect: The Role of Demand Characteristics Revisited," *British Journal of Clinical Psychology,* 33, 345-51.

Metcalfe, J. (1999), "Cognitive Optimism: Self-Deception or memory-Based Processing Heuristics?" *Personality and Social Psychology Review,* 2(2), 100-110.

Morris, W. N. (1999), "The Mood System," in *Well-Being: The Foundations of Hedonic Psychology*, D. Kahneman, E. Diener and N. Schwartz, eds., New York, NY: Russell Sage Foundation.

Nisbett, R. and L. Ross (1980), *Human Inference: Strategies and Shortcomings of Social Judgment,* Englewood Cliffs, NJ: Prentice-Hall, Inc.

Redelmeier, D. A. and Kahneman, D. (1996), "Patients' Memories of Painful Medical Treatments: Real-Time and Retrospective Evaluations of Two Minimally Invasive Procedures," *Pain*, 66, 3-8.

Rabin, M. (1993), "Incorporating Fairness into Game Theory and Economics", *American Economic Review*, 83, 1281–1302

Rabin, M. (1998), "Psychology and Economics," *Journal of Economic Literature* 36, 11-46.

Ross, M. and F. Sicoly (1979), "Egocentric Biases in Availability and Attribution," *Journal of Personality and Social Psychology* 322.

Samuelson, L. (2001), "Information-Based Relative Consumption Effects," mimeo.

Seligman, M. (1990), *Learned Optimism,* New York: Knopf.

Steele, C. M., & Aronson, J. (1995), "Stereotype Threat and the Intellectual Test Performance of African-Americans," *Journal of Personality and Social Psychology*, 69, 797-811.

Steele, C. M., & Aronson, J. (1998), "How Stereotypes Influence the Standardized Test Performance of Talented African American Students," in C. Jencks & M. Phillips (Eds.), *The Black-White Test Score Gap.* Washington, D. C.: Brookings Institution, 401-427.

Svenson, O. (1981), "Are We All Less Risky and More Skillful Than Our Fellow Drivers?", *Acta Psychologica* 143, 145-46.

Taylor, S. and J. Brown (1988), "Illusion and Well-Being: A Social Psychological Perspective on Mental Health," *Psychological Bulletin* 103, 193-210.

Teasdale, J. and M. Russell (1983), "Differential Effects of Induced Mood on the Recall of Positive, Negative and Neutral Moods," *British Journal of Clinical Psychology*, 22, 163-171.

Tversky, A. and D. Kahneman (1973), "Availability: A Heuristic for Judging Frequency and Probability," *Cognitive Science*, 5, 207-232.

Weinstein, N. (1980), “Unrealistic Optimism about Future Life Events,” *Journal of Personality and Social Psychology* 39, 806-820.