Available online at www.sciencedirect.com

Review of Economic Dynamics 7 (2004) 198–218

Review of Economic Dynamics

www.elsevier.com/locate/red

Optimism and overconfidence in search

Juan Dubra

Universidad de Montevideo, Montevideo, Uruguay

Received 1 November 2002; revised 21 March 2003

Abstract

In a standard search model I relax the assumption that agents know the distribution of offers and characterize the behavioral and welfare consequences of overconfidence. Optimistic individuals search longer than pessimists if they are equally "stubborn" and high offers are good news. Otherwise, the pessimists search longer. The welfare of unbiased individuals is larger than that of overconfident decision makers if the latter's biases are large and searchers stubborn. Otherwise, the overconfident may be better off. Finally, I give a testable implication of overconfidence and discuss some applications and policy issues.

1. Introduction and motivation

"Dozens of studies show that people generally overrate the chance of good events, underrate the chance of bad events and are generally overconfident about their relative skill or prospects. For example, 90 percent of American drivers in one study thought they ranked in the top half of their demographic group in driving skill" (Camerer, 1997).[1]

Despite the substantial evidence that overconfidence is pervasive, it has not received much attention in economic modeling. Given the wide applicability of search models, I study the implications of overconfidence in the search behavior of rational agents. To do so, I relax the usual assumption that the searchers know the true distribution of wage offers and suppose only that agents' beliefs are derived from a prior over a set of possible distributions.[2]

E-mail address: dubraj@um.edu.uy.

[1] I will not discuss this evidence here. See Camerer (1997) for experimental and psychological references.

[2] Several other authors have also studied search behavior when the distribution is not known. See Kohn and Shavell (1974), Rothschild (1974), Burdett and Vishwanath (1988) and Bikhchandani and Sharma (1996).

1094-2025/$ – see front matter
doi:10.1016/S1094-2025(03)00036-X

This paper has four objectives. The first is to establish the behavioral implications of optimism in a simple search model. Bikhchandani and Sharma (1996) have shown that when there is learning, the order of static optimism of two individuals may be reversed after observing the same information. Thus, they say that one individual is more optimistic than another if he assigns higher probabilities to high offers after all sequences of observations. I show that this definition of optimism fails to predict optimistic behavior. That is, even if one searcher is more optimistic, in their definition, the former may accept offers that the latter would reject. I then provide a new definition of optimism that guarantees that if one individual is more optimistic than another, the latter stops searching first. The reason why Bikhchandani and Sharma's definition of optimism fails to predict optimistic behavior is that in this context offers have informational value: as search evolves, individuals learn about the unknown distribution. Suppose then that a low offer implies that offers in the future will be high. In that case, optimism about today's offers may lead to a lower expected value of searching than pessimism. This, in turn, yields shorter search times for the more optimistic agent. The main result on behavior is that a searcher who is more optimistic (about the next offer) than another after all sequences of draws samples longer whenever there is a searcher who believes that high offers today mean high offers in the future, and whose priors lie "between" the more and the less optimistic priors.

The second objective is to study the welfare implications of overconfidence. I find conditions under which overconfident agents are worse off than unbiased searchers when welfare is computed using the true wage offer distribution. In this paper, an individual is overconfident if he believes that the distribution that generates the offers is better than it really is. In other words, an individual is overconfident if his prior is such that his beliefs about the first offer first-order-stochastically dominate the true distribution of offers. An individual is unbiased if his beliefs about the first offer are correct. I show that when searchers are not too patient, there are some overconfident individuals who obtain higher expected payoffs than some unbiased searchers. If agents have a degenerate prior, being unbiased means knowing the true distribution. In that case, unbiased searchers *must* be weakly better off than overconfident decision makers. However, if priors are non-degenerate, the comparison is not between an overconfident individual and a searcher who knows the truth, but between two searchers who are uncertain about the true distribution, one of whom happens to be unbiased. Thus, at least in principle, there is the possibility that an unbiased individual is worse off than an overconfident searcher. In fact, there *should* be an unbiased decision maker who is worse off than an overconfident individual. Along the search process high offers are accepted, so sampling continues only if offers have been low. Consequently, because priors are updated in each period, there is a tendency for beliefs to become pessimistic. Therefore, searchers who were initially unbiased and continue sampling today are likely to wrongfully accept a low offer tomorrow. Slightly overconfident searchers are more immune to this kind of mistake. Since they are not too biased, they do not mistakenly reject offers and, because they were originally optimistic, downward updating is not so harmful.

My third objective is to study the conditions under which the behavior and welfare consequences of overconfidence diminish over time. Since behavior and welfare are derived from beliefs, this amounts to finding conditions under which the overconfident individual's true average posterior approaches the true distribution. I show that, while unbiased priors remain unbiased on average, overconfident individuals may become pessimistic. If the true distribution allows only offers that are "too" low according to the overconfident decision maker's beliefs, he may become pessimistic after all offers. This cannot happen with unbiased priors. To insure that overconfident beliefs diminish over time and never become pessimistic, it suffices to assume that there is an unbiased belief that is (dynamically) more pessimistic than the overconfident. The condition is not trivial because it requires that the overconfident prior remains more optimistic than the unbiased after all sequences of draws. Then, the result follows because unbiased priors are a martingale and a lower bound for more optimistic beliefs.

Finally, I derive a testable implication of overconfidence. Using the results on behavior and on evolution of beliefs, I show that overconfident searchers tend to have shorter unemployment spells, whereas unbiased searchers tend to have constant spell lengths. Using structural estimation methods and the National Longitudinal Survey of Youth (NLSY) data, one can compute the proportion of overconfident, unbiased and underconfident individuals in an economy. This is very important, since it has been argued that most "evidence" for overconfidence comes from experimental and psychological evidence, and not from actual economic behavior and data.

I conclude with a discussion of some applications and policy issues.

There are three kinds of theoretical works related to the notion of overconfidence studied in this paper. The first class analyzes the effects of trader's overconfidence in financial markets in a static context. For instance, Benos (1998), Kyle and Wang (1997), and Odean (1998) show that increased overconfidence leads to greater expected trading volume and greater price volatility. The second class studies the emergence of trader's overconfidence in financial markets. For instance, Gervais and Odean (1997) study, in a dynamic setting, how biases in learning generate overconfidence. In their model, individuals attribute good trades to their ability and bad trades to chance. Thus, although overconfidence reduces expected payoffs, rich traders tend to be overconfident. A third class studies the consequences of entrepreneurs' overconfidence. For example, Manove (1995) shows that increased optimism leads to lower expected utility and inefficient allocation of resources in a growth model. Manove and Padilla (1999) show that the coexistence of optimistic and realistic entrepreneurs generates a screening problem for banks and leads to inefficient allocation of credit.

There are models that study optimism and other notions of overconfidence, but they are unrelated to my work. One notion of optimism is that in Beaudry and Portier (1998). In their model, agents observe a signal about an unknown productivity parameter and, if the signal is high, the individual is optimistic. However, he knows the distribution of the signal. In my model, the searcher is biased about the distribution. The second notion is that of self-fulfilling optimism, as in Kiyotaki (1988). In his model, if firms are optimistic about demand and investment, demand is high in equilibrium, so there is no over-optimism. In my model, the searcher is overly optimistic about the distribution. Another (seemingly unrelated) notion of overconfidence that has been studied can be defined as underestimation of volatility. For example, Alpert and Raiffa (1982) document how people systematically construct too narrow confidence intervals for random variables.

2. The model

For any compact topological space (X, τ) let $C(X)$ denote the set of all bounded continuous functions from X to $\mathbf{R}$ endowed with the sup norm. Also, let $P(X)$ represent the set of all probability measures on the Borel sets of X, endowed with the topology of weak convergence. Let $W \equiv \{w_1, w_2, \ldots, w_n\} \subset \mathbf{R}_+$, with $0 < w_1 < w_2 < \cdots < w_n$, and define $P^2(W) = P(P(W))$. I will represent any $g \in P(W)$ by $(g_1, \ldots, g_n)$, where $g_i = g(w_i)$.

At each date t, the individual receives independent and identically distributed wage offers from W and must decide whether to accept the current proposal or continue sampling. His objective is to maximize the expected discounted value of the offer he accepts. Thus, his decision depends on what he believes about future proposals. In most search models, it is assumed that the searcher knows the exact distribution from which offers are drawn. In this paper, I relax this assumption and assume only that the individual has beliefs over the set of possible distributions. Consequently, his beliefs are a distribution over probability measures, which can be represented by a prior $\pi \in P^2(W)$.

It is worth noting here that the model presented here encompasses the search model in which once an offer is accepted, the searcher receives the same wage in every period. For example, if a searcher accepts an offer of \$2, and will receive that salary in every period, in terms of my model he is accepting a one-time offer of $2/(1-\delta)$.

As offers arrive, the individual updates his priors according to Bayes' rule. Let $\Omega = W^\infty$ be the set of infinite sequences of offers. Also, for any *offer path* $\omega \in \Omega$ let ω^t stand for the first t elements of ω and ω_t for its tth element. Starting with beliefs π and after a history ω^t, the probability of any measurable set $C \subset P(W)$ is

$$B\big(\omega^t, \pi\big)(C) = \int\limits_C \frac{\prod_{i \leqslant t} g(\omega_i)}{\int \prod_{i \leqslant t} g(\omega_i) \pi(\mathrm{d}g)} \pi(\mathrm{d}g).$$

If ω^t is a zero π-probability event, $B(\omega^t, \pi)$ is arbitrary.

2.1. Optimal search behavior

In this section, I find the optimal policy for the searcher's maximization problem. In order to use dynamic programming to find the optimal rule, I need to specify a state space and the transition probabilities. In usual search models, the state space is the set of wage offers and the transition is given by the known distribution. Here, the state space must be extended to account for varying beliefs, and the transition function will depend on the history of draws.

At each date in which search continues, the searcher has some beliefs, belonging to $P^2(W)$, and is faced with an offer in W. If he has accepted a proposal, he receives offers of 0 thereafter. Thus, let $S \equiv P^2(W) \times \{W \cup \{0\}\}$ be the state space of the searcher's problem.

Any prior π induces a measure m_π over W, through

$$m_\pi(w) \equiv \int\limits_{P(W)} g(w)\pi(\mathrm{d}g).$$

Since π is a probability over distributions, m_π, usually called the *marginal* of π, is the average distribution that an agent with beliefs π expects to face. If beliefs are π and search continues, the only conceivable states tomorrow are of the form $(B(w,\pi),w)$, with $w \in W \subset \mathbf{R}_{++}$, and their probabilities are given by $m_\pi(w)$. Analogously, if an offer has been accepted, the only possible state is $(\pi,0)$. Then, the following measures over S describe the transitions:

$$C_\pi[s] = \begin{cases} m_\pi(w) & \text{for } s=(B(w,\pi),w), \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad D_\pi[s] = \begin{cases} 1 & \text{for } s=(\pi,0), \\ 0 & \text{otherwise.} \end{cases}$$

The measure C_π gives the subjective probability of each state tomorrow, given that beliefs today are π and search continues; D_π gives the probabilities if an offer has been accepted. Let $A=\{a,r\}$ be the action space, where a indicates that an offer is accepted, and r that it is rejected. For any state (π,w) and action c, define the transition $q(\cdot \mid (\pi,w),c)$ by

$$q\big(\cdot \mid (\pi,w),c\big) = \begin{cases} C_\pi & \text{if } w \in W, \\ D_\pi & \text{if } w=0 \text{ or } c=a. \end{cases}$$

Given state (π,w), if the searcher chooses an action c, $q(s \mid (\pi,w),c)$ gives the subjective probability of state s in the following date. In the next period, an offer is drawn, beliefs are updated, the searcher chooses an action, and the process is repeated.

Define $\mathbf{H}_t=(S\times A)^{t-1}\times S$. A policy is a sequence $p=\{p_t\}_1^\infty$ of functions such that $p_t:\mathbf{H}_t\to A$. For each policy p and $\omega\in\Omega$, let $\tau(p,\omega)$ stand for the date when an offer is accepted if p is followed. Then, for a discount factor $\delta\in(0,1)$ and beliefs π, the payoff of policy p is $E_\pi[\delta^{\tau(p,\omega)}\omega_{\tau(p,\omega)}]$ and the value function is $v:S\to\mathbf{R}$ is $v(s)=\sup_p E_\pi[\delta^{\tau(p,\omega)}\omega_{\tau(p,\omega)}]$. The following lemma states that $Ky(\pi,w)=\max\{w,\delta\int_S y(s')q(\mathrm{d}s' \mid (\pi,w),r)\}$, $y\in C(S)$, is a well defined function $K:C(S)\to C(S)$. All proofs can be found in the Appendix.

Lemma 1. *For any $y\in C(S)$, $Ky\in C(S)$.*

Since K is a contraction, it has a unique fixed point in $C(S)$. Moreover, the fixed point is the value function v.[3] Define $V(\pi)\equiv\int v[B(w,\pi),w]m_\pi(\mathrm{d}w)$, the maximum value of searching when beliefs are π. Then, in any state $(\pi,w)\in S$, accepting an offer if and only if

$$w \geqslant \delta V(\pi) \qquad \text{(optimal policy)}$$

is optimal. The optimal rule states that offers greater than the maximum expected continuation value of searching should be accepted. To see that the policy is in fact optimal, recall from Denardo (1967, Corollary 2) that an optimal policy exists. Then, let $x(s)$ be the expected return of following the above policy for one period and then following an optimal policy, when starting in an arbitrary state s. Since $x(s)=v(s)$ and s was arbitrary, following the rule in every period is optimal.

[3] See in Denardo (1967, Theorem 3). The result is for bounded functions, but his proof, as well as the one of Corollary 2 to be used later, applies to bounded and continuous maps.

Note that this rule does not imply a reservation value rule. Assume, as in Kohn and Shavell (1974), that a searcher believes that there are only two possible distributions. One that assigns probability 1 to \$1 and another with prob(\$2) $= 1 -$ prob(\$3) $= 0.01$. If the first draw is $w = 1$, the individual is certain that he will receive no higher offers and accepts the proposal. On the other hand, if he is patient and the first draw is $w = 2$, he will reject the offer and wait for a draw of \$3.

3. Dynamically consistent optimism and behavior

In static problems, a belief g about a parameter, or about wage offers, is "higher" or "better" than belief h if it can be ranked by first-order stochastic dominance. For $g, h \in P(W)$, g *first-order-stochastically dominates* h, denoted $g \trianglerighteq h$, if and only if $\int u(w)g(\mathrm{d}w) \geqslant \int u(w)h(\mathrm{d}w)$ for all non-decreasing functions u (Dubins and Savage, 1965). Of course, $\trianglerighteq$ is a partial order on the space of priors. For static decision problems, $\trianglerighteq$ captures the idea of optimism: for all utility functions u (above), g yields a higher expected utility than h. In this section, I will introduce a partial order in the space of priors about the sequences of wage offers which will characterize optimism in dynamic contexts. That is, I will provide a definition of optimism that will insure that if a searcher with a prior π is more optimistic than a searcher with a prior υ, then the former obtains a higher expected utility. Since a higher expected utility of searching implies that more offers are rejected, the definition of optimism that I will provide also insures longer search times.

Say that a prior $\pi \in P^2(W)$ is *monotonic* if and only if, for all $\omega, \kappa \in \Omega$ and t, $\omega^{t-1} = \kappa^{t-1}$ and $\omega_t \geqslant \kappa_t$ imply $m_{B(\omega^t,\pi)} \trianglerighteq m_{B(\kappa^t,\pi)}$. That is, posteriors after high offers first-order-stochastically dominate posteriors after low offers: high offers make a searcher more optimistic about the next draw. Monotonicity insures that the informational value of offers is ordered in the same manner as their monetary value. Dirichlet priors over multinomial distributions and arbitrary priors over binomial distributions satisfy monotonicity.[4] This condition is similar to those used by Bikhchandani and Sharma (1996), Burdett and Vishwanath (1988) and Milgrom (1981).

I now define optimism for the current context and then show that this partial order in the space of priors insures that the more optimistic searcher samples longer in *all* offer paths. For any priors π and υ, we will say that π is *more optimistic* than υ, written $\pi \geqslant \upsilon$, if there exists a monotonic ρ such that $m_\pi \trianglerighteq m_\rho \trianglerighteq m_\upsilon$ and $m_{B(\omega^t,\pi)} \trianglerighteq m_{B(\omega^t,\rho)} \trianglerighteq m_{B(\omega^t,\upsilon)}$ for all t and all $\omega \in \Omega$. In words, a prior π is more optimistic than υ if the marginals of π dominate those of υ after all sequences of draws, and there is a prior ρ between them which is "well-behaved."

For any prior π and $\omega \in \Omega$, let $\tau_\pi(\omega)$ be the *acceptance time*, the date when an offer is accepted if the optimal policy is followed.

Proposition 2. *Assume that $\pi \geqslant \upsilon$. Then, for all $\omega \in \Omega$, $\tau_\pi(\omega) \geqslant \tau_\upsilon(\omega)$.*

[4] A Dirichlet with parameter $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$, with $\pi_i > 0$ for all i, is a probability measure over $P(W)$. Let $S(\pi) = \sum_1^n \pi_i$, and $\mu_i = \pi_i / S(\pi)$. Then $m_\pi(w_i) = \mu_i$, and $B(w_i, \pi)$ is a Dirichlet with parameter $\pi + e_i$ (where e_i is the ith canonical vector). See De Groot (1970).

Remark 1. Notice that the proposition is not about average acceptance times, but about what happens along all offer paths. The idea behind this result is that the optimistic searcher believes that the future is good and thus rejects offers that the pessimist does not. Monotonicity guarantees that, in terms of information content, high offers are better than low ones. Then, the requirement that the posteriors are ordered by first-order stochastic dominance after all draws insures that the more optimistic individual searches longer.

Remark 2. Corollary 1 in Bikhchandani and Sharma (1996) proves that, if the marginals of the posteriors of π first-order-stochastically dominate those of υ after all sequences of draws, and either one of the priors is monotonic and searchers follow a reservation wage policy, optimistic searchers sample longer. Therefore, they have more and stronger assumptions, than in Proposition 2. However, since they concentrate on problems with no discounting, Proposition 2 is not a generalization of their result. However, my method of proof can be adapted to their context to avoid the assumptions of a reservation wage and the requirement that either prior is monotonic.

Remark 3. Proposition 2 can be proved as a corollary to the following result, which can be proved using the same steps as in Proposition 2. Let $\Omega = W^T$, for $T \in N \cup \{\infty\}$. If $\pi \geqslant \upsilon$, then $E_\pi(u) \geqslant E_\upsilon(u)$ for all increasing $u : \Omega \to \mathbf{R}$. See Müller and Stoyan (2002, Theorem 3.3.4) for another condition on priors that insures higher utilities.[5]

I now discuss what is being assumed and what is not with the postulate that $\pi \geqslant \upsilon$. First, the assumption that π is more optimistic that υ does not insure that the searchers follow reservation wage rules. That is, Proposition 2 does not assume, neither implicitly nor otherwise, that searchers follow reservation wage rules. To illustrate, consider again the example of Section 2.1. Recall that in that example searchers believed that there were only two possible distributions, one that assigned probability 1 to \$1 and another with prob(\$2) $= 1 -$ prob(\$3) $= 0.01$. For any $p > q$, a searcher who assigns probability p to the latter urn, call this prior π, is more optimistic than a searcher who assigns probability q to the same urn, call this prior υ. That is, $\pi \geqslant \upsilon$, and still none follows a reservation wage rule.

A second issue that is worth noting is that $\pi \geqslant \upsilon$ applies to the following model which was used by Burdett and Vishwanath (1988): the urn from which offers are drawn is a *normal* with known variance and unknown mean, and searchers believe that the mean of the distribution is also *normally* distributed, only that π's mean is higher than υ's.

Third, the condition $\pi \geqslant \upsilon$ requires m_π more than just first-order-stochastically dominate m_υ. As was shown by Bikhchandani and Sharma (1996), such a static condition is not sufficient to insure longer search times for the searcher with prior π. Because individuals whose prior about the first draw is low in first-order stochastic sense may also have priors that are less affected by updating than searchers with high beliefs, downward updating can lead the individual with the initially high prior to stop sampling before the

[5] Their condition requires checking that π assigns higher probability than υ to all increasing sets. Unfortunately, this condition is very hard to check, whereas the definition of optimism given here only involves comparisons of marginals.

searcher with the low prior. The following example is similar to Bikhchandani and Sharma (1996, Example 1). It illustrates how different propensities to update may lead a searcher with prior π to search less than a searcher with prior υ even though m_π, the belief about the first draw, first-order-stochastically dominates m_υ.

Example 1. Let $W = \{1, 2\}$ and $1/2 \geqslant q \geqslant 0$. Define $f, g, h \in P(W)$ by $h = (1, 0)$, $f = (1/2, 1/2)$ and $g = (0, 1)$. Also, define priors π, υ, by $\pi(g) = 1/2 + q$, $\pi(h) = 1/2 - q$ and $\upsilon(f) = 1$. Note that m_π first-order-stochastically dominates m_υ, so that π is "higher" or "better" than υ.

The posterior of π is degenerate in h after receiving a draw of 1. Thus, the searcher π accepts the offer of 1 in the first period. Since he also accepts a draw of 2 in any date, he stops sampling in the first period in every offer path. On the other hand, the searcher with prior υ never revises his priors and, for $\delta > 2/3$, continues sampling until a high draw occurs. Since the size of q indexes how "high" beliefs are, for *all* levels q of the prior and *all* offer paths, the individual with prior π never samples longer than the searcher with prior υ, and sometimes samples less, despite the fact that m_π first-order-stochastically dominates m_υ.

This result is driven by the fact that π is affected by updating and υ is not. This feature insures that even though m_π first-order-stochastically dominates m_υ, the reverse is true after an offer of 1 arrives. In a sense, υ is more "stubborn" in the face of new information. The above example shows that in order to say that one prior is more optimistic than another, we need the condition that posteriors are also ordered in first-order stochastic sense. That is in part what $\pi \geqslant \upsilon$ requires, and it is the only requirement in Bikhchandani and Sharma's definition of optimism (which, as we now show, does not capture optimistic behavior).

Fourth, even requiring that marginals are always ordered by first-order stochastic dominance does not necessarily yield longer search times. Since individuals learn about the true distribution as offers arrive, proposals have informational value. If the total value, monetary plus informational, of a low offer exceeds that of a high offer, assigning high probabilities to high proposals may lead to a low value of searching.

The next example shows that requiring that posteriors are ordered by first order stochastic dominance does not insure that π will search longer than υ.

Example 2. Let $W = \{2, 4, 5, 6\}$ and $1/4 \geqslant \epsilon \geqslant 0$. Define $g_\epsilon, j_\epsilon, h_\epsilon \in P(W)$ by: $g_\epsilon = (3/4 - \epsilon, 0, \epsilon, 1/4)$, $j_\epsilon = (1/4 - \epsilon, 0, \epsilon, 3/4)$, and $h_\epsilon = (0, 1 - \epsilon, \epsilon, 0)$.

Fix $\delta = 0.99$ and suppose $\epsilon = 0$. Assume also that a searcher is certain that the distribution is $j_0 = (1/4, 0, 0, 3/4)$. Because δ is close to 1, the searcher samples until $w = 6$ is drawn and obtains an expected payoff of approximately 6. The same is true for $g_0 = (3/4, 0, 0, 1/4)$. If the distribution is $h_0 = (0, 1, 0, 0)$ however, the searcher accepts the first offer of 4 and obtains an expected value of 4.

Define $\pi^\epsilon, \upsilon^\epsilon \in P^2(W)$ by $\pi^\epsilon(j_\epsilon) = 1 - \pi^\epsilon(h_\epsilon) = 1/2$ and $\upsilon^\epsilon(g_\epsilon) = 1 - \upsilon^\epsilon(h_\epsilon) = 4/5$. If priors are π^0, whenever $w = 2$ or $w = 6$ occur the searcher knows that the distribution is j_0. If $w = 4$ is drawn, the distribution is h_0. Thus, the value of searching when priors are π^0 is $V(\pi^0) \approx (6 + 4)/2 = 5$. Analogously, $V(\upsilon^0) \approx (4/5)6 + (1/5)4 > 5 \approx V(\pi^0)$. That is, the posteriors of π^0 first-order-stochastically dominate those of υ^0 after all sequences of draws, but yield a smaller value of searching. This result is driven by the fact that a draw of $w = 4$ signals a distribution with a value of 4, whereas $w = 2$ informs the individual that the value of searching is close to 6.

Note that because $5 \in (\delta V(\pi^0), \delta V(\upsilon^0))$, if $w = 5$ is drawn and priors are not updated, the searcher with prior π^0 accepts the offer and the one with υ^0 does not. However, $w = 5$ is a zero-probability event for both priors, so I will slightly modify them to insure that searchers can use Bayes' rule. For any ϵ, when $w = 5$ is drawn, updating does not change π^ϵ or υ^ϵ. Then $5 \in (\delta V(\pi^0), \delta V(\upsilon^0))$ guarantees that $5 \in (\delta V(\pi^\epsilon), \delta V(\upsilon^\epsilon)) = (\delta V(B(5, \pi^\epsilon)), \delta V(B(5, \upsilon^\epsilon)))$ for small enough ϵ. Therefore, when $w = 5$ is drawn, the searcher with "high" priors π^ϵ will accept the offer and the individual with "low" beliefs υ^ϵ will reject it. Moreover, since $V(\pi^\epsilon) < V(\upsilon^\epsilon)$, a searcher π^ϵ whose beliefs about draws dominate in each period the beliefs of another individual υ^ϵ, obtains a lower subjective expected utility.

In this example, a searcher with beliefs that are high in first-order stochastic sense, stops sampling before the searcher with low beliefs because a low offer has high informational value. That is, it is not true that high offers are good news.

In Example 1, searchers have monotonic beliefs. Nevertheless, the different propensities to update allowed the searcher with beliefs which were high in first-order stochastic sense to stop sampling before the searcher with low beliefs. Example 2 shows that if monotonicity fails, the searcher whose beliefs are high in first-order stochastic sense in every period may stop sampling before the individual with low beliefs.

4. Welfare implications

In this model, the optimal search rule calls for accepting high offers, so sampling continues only if offers are bad. Since searchers have non-degenerate priors and they update their beliefs in each period, this feature of the model makes them become more pessimistic over time. In this section, I analyze the welfare consequences of this fact.

Throughout, let $f \in P(W)$ be the true measure that generates the offers. I will say that prior π is *unbiased* if $m_\pi = f$. That is, a prior is unbiased if the expected value of the urn, according to beliefs and before any information has been received, is the true urn. In turn, I will say that a prior π is *overconfident* if $m_\pi \trianglerighteq f$. That is, a prior is overconfident if its initial estimate of the urn from which offers are drawn is larger, in first-order stochastic sense, than the truth.

Four comments about this definition are in order.

1. This definition captures the idea that the searcher thinks that he is better off than he really is, since he believes that the urn from which offers are drawn is better than it really is.
2. This definition encompasses other definitions of overconfidence that have been put forward in most of the literature, and is "better" in the following sense. An overconfident decision maker is often defined as one who thinks that some parameter is larger than it really is. Formally, it is assumed that the individual's beliefs are

degenerate and wrong. An individual with those beliefs would never learn the true distribution. My definition of overconfidence allows for those beliefs, but also for non-degenerate beliefs which, with enough learning, could converge to the truth.

3. The findings in Alpert and Raiffa (1982) that people systematically construct too narrow confidence intervals for random variables has been interpreted as a different notion of overconfidence than the one studied here. I believe, however, that the findings in Alpert and Raiffa (1982) are about how people think that their information is better than it really is, so that "point estimates" of random variables are thought to be very close to the true values. Therefore, those findings are also about overconfidence in the sense of this paper.
4. In this paper, optimism is a relative notion, about one searcher being more optimistic than another, whereas overconfidence is about how a searcher thinks he is, relative to the truth.

Let $\wp$ be the measure on Ω obtained by extending the probabilities that f induces on W^T for all T.[6] For the welfare criterion, I use the true expected value of searching,

$$V^{\wp}(\pi) = \int_{\Omega} \delta^{\tau_\pi(\omega)} \omega_{\tau_\pi(\omega)} \wp(\mathrm{d}\omega), \quad \pi \in P^2(W).$$

As in standard search models, the rule that maximizes the true expected value of searching is a stationary reservation wage policy. In addition, if a searcher is going to deviate from the optimal policy just once, the longer he follows the optimal policy, the higher is his expected payoff. Since there is a tendency for searchers to become pessimistic (and pessimistic searchers accept low offers), slightly overconfident individuals follow the truly optimal policy longer than unbiased searchers. As a consequence, in the following example an overconfident individual obtains a higher payoff than an unbiased searcher who is more pessimistic than the overconfident.

Example 3. Let $W = \{1, 2\}$, $8/11 > \delta > 2/3$, and $f = (1/2, 1/2)$. Then, the policy that maximizes the true expected value of searching is to reject offers of 1 and accept the first offer of 2.

Define $g, h \in P(W)$ by $g = (3/4, 1/4)$ and $h = (1/4, 3/4)$. Also, let $\pi(h) = 1 - \pi(g) = 3/4$ and $\upsilon(h) = 1 - \upsilon(g) = 1/2$. I will now show that the optimal search rule in this case calls for rejecting offers of 1 until the expected value of the next draw falls below $1/\delta$, and then accepting any offer. If the continuation value of searching falls below $1/\delta$, the agent accepts the current offer, so it suffices to show that whenever the expected value of the next draw falls below $1/\delta$, it is equal to the continuation value of searching. Suppose that the continuation value of searching after an offer of 1 is less than or equal to $1/\delta$. If the offer is rejected and $w = 1$ is drawn in the next period, the continuation value will be weakly smaller than it is today, which implies that the optimal strategy calls for accepting any offer tomorrow. Therefore, if the continuation value today is below $1/\delta$, it is equal to

[6] See Shiryayev (1984, Kolmogorov Extension Theorem, p. 161).

the expected value of the next draw. Also, for enough draws of 1, the continuation value is close to that of a prior which assigns probability 1 to $w = 1$. Thus, the continuation value eventually falls below $1/\delta$. Finally, since the expected value of the next draw is decreasing over time, whenever it falls below $1/\delta$ it must be the continuation value.

Since $m_{B(1,\upsilon)}(1) = 5/8$, the expected value of the next draw after observing $w = 1$ is $11/8 < 1/\delta$. Thus, the unbiased searcher stops sampling in the first period in any $\omega \in \Omega$ and obtains a true value of searching, $3/2$. Since $m_{B(1,\pi)}(1) = 1/2$, the expected value of the next draw after the first bad draw is $3/2 > 1/\delta$. Therefore, the overconfident searcher rejects the first low offer, but because $m_{B((1,1),\pi)}(1) = 5/8$ he accepts any offer in the second period. This yields a true value of searching, $1 + \delta(3/4) > 3/2$.

In this example, the overconfident searcher uses the true optimal strategy in period 1 whereas the unbiased one does not. As a consequence, the overconfident searcher is better off when the true distribution is used to compute welfare. Note that it is *not* the case that for some states of the world the overconfidence is better off (i.e., that he rejects a high offer and by chance he gets a higher offer in the next period). His *expected* payoff is larger than that of the unbiased searcher.

4.1. The benefits of overconfidence and costs of underconfidence

This section provides a generalization of the last example. Consider the following three features of the search model. First, the truly optimal search rule is a constant reservation wage policy. Second, searchers become pessimistic as search evolves, so there is a tendency for reservation wages to decrease. Third, overconfident searchers tend to have higher reservation wages than unbiased individuals. These features insure that one can always find overconfident searchers whose initial reservation wage is optimal and that, as search evolves, make fewer mistakes (relative to the truly optimal search rule) than unbiased searchers who are more pessimistic than the overconfident. These features are what drive Example 3, and thus suggest that, when these conditions hold, overconfident searchers are better off than unbiased ones. However, the following example shows that in general this is false.

Example 4. Let $W = \{1, 2, 100\}$, $f = (1/3, 1/3, 1/3)$ and $\delta = 58/1000$. The truly optimal strategy is to accept only offers of 100. A searcher who follows the optimal strategy in period 1 and then accepts any offer, obtains an expected payoff of $100/3 + (206/9)\delta$. Rejecting only offers of 1 in the first period and accepting any offer in the second yields a payoff of $34 + (103/9)\delta > 100/3 + (206/9)\delta$. Therefore, following the optimal strategy in the first period is harmful. In the reminder of the example I show how the above behavior can be derived from overconfident and unbiased priors.

Let $j = (1/1000, 1/1000, 998/1000)$ and $g = (999/2000, 999/2000, 1/1000)$. Define the overconfident prior π by $\pi(j) = 1 - \pi(g) = 0.98569$ and the unbiased prior υ by $\upsilon(j) = 1 - \upsilon(g) = 0.33266$. If $w = 100$ has not occurred in periods 1 or 2, the expected value of the next draw is lower than $1/\delta$ for both searchers, so they stop sampling. Therefore, the searchers know in period 1 that offers that yield a value smaller than the discounted expected value of the next draw must be accepted. Since $\delta E_{B(2,\pi)}[w] > 2$, the overconfident searcher only accepts offers of 100 in the first period. Since $2 > \delta E_{B(2,\upsilon)}[w] = \delta E_{B(1,\upsilon)}[w] > 1$, the unbiased searcher only rejects offers of 1 in the first period.

The example illustrates the point that if the optimal policy is not going to be followed tomorrow, it may not be optimal to follow it today. Therefore, although overconfident searchers may follow the optimal strategy more often than unbiased searchers, they are not always better off. To insure that overconfident searchers will be better off, it suffices to assume that searchers are not too patient. If they are impatient, the truly optimal policy is to reject all offers but the lowest. Then, because the individual receives in each period only worse news than he could imagine, reservation wages are decreasing. This, in turn, insures that the only possible deviation (for a searcher who starts off with the optimal reservation wage) is to accept any offer. Consequently, when searchers are not too patient and start off with the optimal reservation wage, they deviate from the optimal policy just once. This guarantees that overconfident searchers make exactly the same mistake as the unbiased individuals, but in a later period, in which case overconfident searchers are better off.

To formalize these arguments I first show that, if an individual would accept the next-to-the-lowest offer to which he assigns positive probability, his reservation wage is decreasing.[7]

Lemma 3. *Suppose that a searcher with monotonic prior π follows a reservation wage policy and that $w_2 > \delta V(\pi)$. Then, $\delta V(B(\omega^{t-1}, \pi)) \geqslant \delta V(B(\omega^t, \pi))$ for all $\omega \in \Omega$ and all t.*

Suppose that the optimal search rule calls for accepting w_2 and rejecting offers below that. Assume also, that π in the previous lemma is overconfident. Then, whenever π's search rule differs from the optimal one, he is accepting offers that he should not. Consider an unbiased searcher υ with $\pi \geqslant \upsilon$. By Proposition 2, the unbiased searcher makes a mistake before the overconfident does, and this makes him worse off. A similar reasoning applies to show that underconfident individuals are still worse off. This is summarized in Proposition 4. For any $\upsilon \in P^2(W)$, any $\epsilon > 0$ and metric d, define $N_\epsilon(\upsilon) \equiv \{\pi \in P^2(W)\colon d[\pi, \upsilon] < \epsilon\}$.

Proposition 4. *Define the prior υ^0 by $\upsilon^0(f) = 1$. Then, there exists $\bar{\delta}$ such that if $\bar{\delta} > \delta$,*

(i) *for any $\epsilon > 0$ there is an unbiased $\upsilon \in N_\epsilon(\upsilon^0)$ and an overconfident π such that $\pi \geqslant \upsilon$ and $V^{\wp}(\pi) \geqslant V^{\wp}(\upsilon)$. Moreover, if $f([0, \delta V(\upsilon^0))) > 0$, $V^{\wp}(\pi) > V^{\wp}(\upsilon)$;*

(ii) *there exists $\gamma > 0$ such that for all unbiased $\nu \in N_\gamma(\upsilon^0)$ that follow a reservation wage policy, if φ is an underconfident prior with $\nu \geqslant \varphi$, then $V^{\wp}(\nu) \geqslant V^{\wp}(\varphi)$.*

Proposition 4 says that there exists an unbiased searcher who is almost certain about the truth and a more optimistic overconfident searcher who is better off. The proof also shows

[7] Bikhchandani and Sharma (1996) provide sufficient conditions on priors to insure that searchers follow reservation wage rules.

that the unbiased prior is closer to the truth, in the metric on $P^2(W)$, than the overconfident prior. Therefore, the result is not about the searcher being overconfident but being "closer" to the truth than an unbiased decision maker. Second, it is overconfidence, and not an arbitrary bias, that makes the overconfident searchers better off. Underconfident searchers are still worse off. The reason driving this last result is that $\upsilon \geqslant \varphi$ and Proposition 2 insure that whenever υ accepts an offer, φ also does. Then, since the most likely mistake of υ, relative to the optimal policy, is to accept an offer that should be rejected, whenever υ makes a mistake, φ also does. Again, the reason why the most likely mistake is to accept an offer that should be rejected is that in search problems there is a tendency for beliefs to decrease over time, and a searcher whose prior is close to the truth starts off with the correct reservation wage.

Of course, these results may seem puzzling. In particular, should an individual, knowing that he will become more pessimistic over time, choose to behave as if he was not updating? No. But one could try to avoid the bias that I point to by trying to use all the draws that other individuals in a "peer group" receive on each date. In that way, no bias is introduced. But of course, after a few bad draws the individual could start updating about the probability that his peer group is really his peer group, and the problem mentioned in this paper arises again.

4.2. The costs of overconfidence

Proposition 4 shows that overconfident searchers are sometimes better off than unbiased decision makers. In this section, I examine the reasons why the converse may hold. The first reason why overconfidence can be harmful is the one illustrated in Example 4: following the optimal policy more often than not is not always beneficial. The second one is obvious: overconfident searchers may reject high offers that they should accept. However, since it is easy to construct examples where overconfident searchers with large biases are better off than unbiased searchers, the condition when searchers are stubborn (and keep making their original mistakes) needs to be added.

Consider an individual with priors υ close to the degenerate υ^0. By continuity of V (see Corollary 10 in the Appendix) one can make sure that, for almost any discount factor, the search rule of υ resembles that of υ^0 for a long period of time. Therefore, discounting insures that $V^{\wp}$, the true value of searching, is continuous at υ^0. Then, for υ and π close to the degenerate υ^0 and π^0, respectively, $V^{\wp}(\upsilon^0) > V^{\wp}(\pi^0)$ guarantees $V^{\wp}(\upsilon) > V^{\wp}(\pi)$. The result is summarized in the following proposition, which is just a statement about continuity of the true value of searching.

Proposition 5. *Fix any degenerate priors π^0 and υ^0. Assume that $f(\delta V(\upsilon^0)) = f(\delta V(\pi^0)) = 0$ and $V^{\wp}(\upsilon^0) > V^{\wp}(\pi^0)$. Then, there exists ϵ such that for all $\pi \in N_\epsilon(\pi^0)$ and $\upsilon \in N_\epsilon(\upsilon^0)$, $V^{\wp}(\upsilon) > V^{\wp}(\pi)$.*

The third reason why overconfident searchers may obtain lower payoffs than unbiased searchers is that reservation wages may be increasing for some offer paths. When reservation wages increase, even slightly overconfident decision makers will reject offers that they should accept. Although in general reservation wages do not increase, the following example shows that for some offer paths reservation wages may be increasing.

Example 5. Let $W = \{1, 2, 3, 4\}$, $\delta = 0.99$ and define $g, j \in P(W)$ by $g = (0, 5/12, 1/2, 1/12)$ and $j = (1/3, 1/3, 1/3, 0)$. Also, for $1 \geqslant \epsilon \geqslant 0$, define priors π^ϵ by $\pi^\epsilon(g) = 1 - \pi^\epsilon(j) = \epsilon$.

Since δ is close to 1, a searcher with beliefs π^0 accepts only offers of 3 and 4. Thus, for ϵ sufficiently small, the same is true for a searcher with beliefs π^ϵ. Suppose that some offer path ω starts with t draws of $w = 2$. Because $g(2) > j(2)$, for sufficiently large t, $B(\omega^t, \pi^\epsilon)$ assigns probability close to 1 to g. Consequently, for sufficiently large t, the searcher accepts only offers of 4.

Therefore, the searcher with priors π^ϵ accepts offers of 3 at the beginning of the search process, but after enough draws of 2, he only accepts proposals of $w = 4$. The reservation wage increases because an offer that is rejected at the start of the search process (i.e., a low offer) signals a good distribution.[8]

Adding an appropriate true distribution to this example, it is easy to show that overconfident searchers may be worse off than some unbiased individuals.

5. Evolution of beliefs

In this section, I give conditions that guarantee that true average posteriors diminish over time for overconfident priors. I first show that, although unbiased priors remain unbiased, overconfident beliefs may become underconfident. Then I show that, if there is an unbiased belief that is more pessimistic than the overconfident prior, the bias diminishes over time and the overconfident does not become pessimistic on average.

Suppose that υ is unbiased. Then, by the law of iterated expectation, $E_f[m_{B(w,\upsilon)}] = f$. That is, on average, unbiased searchers remain unbiased. The following example shows, however, that an overconfident prior may become pessimistic on average.

Example 6. Let $W = \{1, 2, 3\}$, define $g, f, j \in P(W)$ by $g = (3/4, 1/4, 0)$, $f = (1/2, 1/2, 0)$, $j = (0, 0, 1)$. Define priors π by $\pi(g) = \pi(j) = 1/2$. Since only offers of 1 and 2 will occur, the posterior of π is always $g = (3/4, 1/4, 0)$. Thus, although π is overconfident, he becomes underconfident with probability 1. But this implies that $f \rhd E_f[m_{B(w,\pi)}]$.

In the example, an unbiased belief that is more pessimistic than π does not exist. For any unbiased belief υ, $E_{m_\upsilon}[m_{B(w,\upsilon)}] = m_\upsilon$ implies that $m_{B(2,\upsilon)} \trianglerighteq m_\upsilon = (1/2, 1/2, 0) \rhd (3/4, 1/4, 0) = m_{B(2,\pi)}$, so first-order stochastic dominance of the marginals is not preserved, violating the definition of optimism. That is, while 2 is good news for υ, it

[8] In Burdett and Vishwanath (1988) this possibility is ruled out assuming that the cost of search is large, which insures that only "very" low offers are rejected.

is “very” bad news for π, and that causes their order of optimism to be reversed. If there existed an unbiased prior which was more pessimistic than π, the overconfident would remain overconfident on average. The reason is that the average posterior of the unbiased belief is a lower bound for the average posterior of π. Since the unbiased prior remains unbiased on average, the overconfident remains overconfident. The following proposition is a generalization of the previous argument.

Proposition 6. *Fix any prior π such that there exists an unbiased belief which is either more optimistic or more pessimistic than π. If π is overconfident,*

$$m_\pi \trianglerighteq E_\wp[m_{B(\omega^{t-1},\pi)}] \trianglerighteq E_\wp[m_{B(\omega^t,\pi)}] \trianglerighteq f$$

for all t, and conversely if π is underconfident.

Proposition 6 states that, for overconfident priors for which there is a more pessimistic unbiased belief, the true average posteriors decrease but never fall below the truth. They decrease because updating is, essentially, averaging priors and the information received and offers are generated by a distribution that is lower than beliefs in first-order stochastic sense. The overconfident beliefs do not fall below the truth because they are bounded from below by the unbiased priors which are martingale.

A corollary of Proposition 6 is that beliefs are martingale for unbiased priors. That is, the agent’s true average beliefs about the distribution that generates the offers do not change over time. This is not the usual “beliefs are martingale” claim of the literature on learning as, for example, in Kalai and Lehrer (1993). In that literature, the relevant distribution with respect to which the expectation is taken is m_π. Hence, in that context, “beliefs are martingale” means that one cannot expect any change in his beliefs. Here, the distribution with respect to which the expectation is taken, is the true measure f. Thus, the result is a statement about the true, and not subjective, evolution of beliefs.

6. Concluding remarks

One can apply the results on behavior and evolution of beliefs to obtain a testable implication of overconfidence. Suppose that the search problem was to be repeated a number of times, called spells, and that each problem were solved myopically. In addition, following Proposition 6, suppose that the overconfident’s prior at the beginning of today’s spell is dominated by his beliefs at the start of the last spell. Then, Proposition 2 insures that the expected search times decrease from spell to spell. On the other hand, an analogous construction for unbiased searchers yields constant spell lengths. Hence, one may be able to test whether people are overconfident through the analysis of search behavior of unemployed workers. This is very important, since it has been argued that most “evidence” for overconfidence comes from experimental and psychological evidence, and not from actual economic behavior and data.

There are, at least, four criticisms that have been put forward against this test of overconfidence. The first, and most obvious, is that if one observes search spells, they tend to get longer over time, so that one could conclude that what is pervasive is underconfidence. The problem with this criticism is that it fails to take into account the fact that the observation of search spells does not control for age. To be clear, suppose that search spells tend to get longer with age, conditional on the number of search spells. If this were the case, one would observe longer search spells only because of the age factor. Using, for example, the NLSY data, one can control for age, and get the "pure" belief effect. A similar comment applies to the criticism that spells might be getting shorter because older searchers have larger families to support, and hence hit their budget constraints earlier. In this case too, the NLSY data has information on family sizes, so family size can also be controlled for, just as age can.

A second criticism that has been put forward is that by just observing search spells one cannot rule out that spells are getting longer due to the fact that employers could be learning about the quality of workers (see Berkovitch, 1990). According to this theory, a firm who hires a worker gets to know his quality, and bad workers are fired more often. Also, a worker who is likely to be of bad quality, receives low offers, and thus has shorter search spells. Therefore, one more search spell in somebody's vita would indicate lower quality, and thus induce a shorter search spell. So this theory also predicts shortening search spells. This factor can be controlled for by carrying out the proposed test only among blue-collar workers where the complete past search history cannot always be observed by employers.

A third criticism that might arise is that searchers could be learning how to search over time, thus leading to shorter search spells. This hypothesis should then lead to the observation of constant, or increasing, accepted wages over time. If individuals are overconfident, however, accepted offers should tend to decline, so that one can tell the theories apart.

A fourth and last objection to the proposed test is that it is possible that search spells are getting shorter spell by spell, not because searchers are learning about the mean of the wage offer distribution, but about its variance. If more dispersed beliefs lead to longer search spells, as is usually the case, learning about the distribution would lower the search spells. At present I do not know how would one control for this last factor.

The results on welfare also suggest that one can build models where the pervasiveness of overconfidence is the consequence of evolutionary selection. In pre-agricultural societies subsistence depended on search activities, such as hunting and gathering. Thus, if overconfident searchers were better off than unbiased searchers, and that favored their reproduction, their progeny should tend to be overconfident.

An important technical open problem is whether the definition of optimism given in this paper is also necessary in dynamic problems of any kind. It can be easily shown that if $\pi \geqslant \upsilon$, then $E_\pi(u(w_1,\dots,w_t)) \geqslant E_\upsilon(u(w_1,\dots,w_t))$ for any u that is increasing in w_i (the wage received in period i) for all i. The question is then whether $\pi \geqslant \upsilon$ is also a necessary condition. Müller and Stoyan (2002, Theorem 3.3.4) prove that $E_\pi(u(w_1,\dots,w_t)) \geqslant E_\upsilon(u(w_1,\dots,w_t))$ for all increasing u if and only if π assigns a higher probability to all sets such that if w is in the set, every $z \geqslant w$ is also in the set. The problem with this condition is that it is very hard to check, so a characterization of optimism in terms of the marginals of the beliefs, such as the one given in this paper, would be very useful.

In closing, I note that all the results of this paper, except those on welfare, can be easily extended to the case of arbitrary $W \subset \mathbf{R}$. Also, all of this paper can be extended to the case where the individual can choose a search effort that affects his chances of receiving an offer. To understand why this is so, one must only note that the model in this paper can be thought of as the model with search effort, when a particular sequence of search efforts has been chosen. One then needs to optimize with respect to the sequence of search efforts. The results on behavior of this paper extend to that setting because, as was argued in the previous section, the definition of optimism given in this paper yields the correct comparative statics in a wide variety of contexts. The results about the welfare implications of overconfidence also extend to the model with search effort because the basic force behind those results is also present in the more general model: beliefs tend to decrease over time so that unbiased searchers stop sampling before they should, whereas slightly overconfident individuals make this kind of mistake less often.

Acknowledgments

I am indebted to Charles Wilson for his guidance. I am grateful to an anonymous referee in this journal, to Jean Pierre Benoît, Alberto Bisin, Federico Echenique, Néstor Gandelman, Efe Ok, and Rob Shimer for helpful comments. Support from the Sloan Foundation and the C.V. Starr Center for Applied Economics is gratefully acknowledged.

Appendix

Measures μ_n over X *converge weakly* to μ, denoted $\mu_n \Rightarrow \mu$, iff $\int y(x)\mu_n(\mathrm{d}x) \to \int y(x)\mu(\mathrm{d}x)$ for all $y \in C(X)$. For $\mu \in P(X)$ and measurable $h : X \to \mathbf{R}$, define $\mu h^{-1}(C) = \mu(h^{-1}(C))$ for all measurable C. The following is a corollary to Billingsley (1968, Theorem 5.5).

Lemma 7. *Let $\{\mu_n\}$, $\mu \in P(X)$, $h : X \to \mathbf{R}$ be continuous and $h_n : X \to \mathbf{R}$ converge uniformly to h as $n \to \infty$. Then, $\mu_n \Rightarrow \mu$ implies $\mu_n h_n^{-1} \Rightarrow \mu h^{-1}$.*

Lemma 8. *$B(w,\cdot) : P^2(W) \to P^2(W)$ is continuous.*

Proof. Fix any $y \in C(P(W))$. I have to show that $\pi_n \Rightarrow \pi$ implies

$$\int\limits_{P(W)} h_n(g)\pi_n(\mathrm{d}g) \equiv \int\limits_{P(W)} \frac{y(g)g(w)}{m_{\pi_n}(w)}\pi_n(\mathrm{d}g) \to \int\limits_{P(W)} \frac{y(g)g(w)}{m_\pi(w)}\pi(\mathrm{d}g)$$
$$\equiv \int\limits_{P(W)} h(g)\pi(\mathrm{d}g).$$

The range r_n of each h_n is bounded. Then, since $\int_{P(W)} h_n(g)\pi_n(\mathrm{d}g) = \int_{\mathbf{r}_n} y\pi_n h_n^{-1}(\mathrm{d}y)$, it suffices to show that $\pi_n h_n^{-1} \Rightarrow \pi h^{-1}$. Since h_n converges uniformly to h, continuity of h and Lemma 7 will complete the proof.

By finiteness of W, for arbitrary w_i, $g_n \Rightarrow g$ implies $g_n(w_i) \to g(w_i)$. This, and continuity of y guarantee that $|y(g_n)g_n(w) - y(g)g(w)| \to 0$. Noting that $|h(g_n) - h(g)| = |y(g_n)g_n(w) - y(g)g(w)|[m_\pi(w)]^{-1}$ completes the proof. □

Lemma 9. *For* $\{\pi_n\}_1^\infty$, $\pi \in P^2(W)$,

$$\int_W y\big(B(w, \pi_n), w\big) m_{\pi_n}[\mathrm{d}w] \to \int_W y\big(B(w, \pi), w\big) m_\pi[\mathrm{d}w]$$

if $y \in C(S)$ *and* $\pi_n \Rightarrow \pi$.

Proof. Lemma 8 and finiteness of W guarantee that $h_n(w) \equiv y(B(w, \pi_n), w)$ converges uniformly in w to $h(w) \equiv y(B(w, \pi), w)$. In addition, $m_{\pi_n} \Rightarrow m_\pi$, so Lemma 7 completes the proof. □

Proof of Lemma 1. Proofs of continuity when search has stopped and of boundedness are trivial and will be omitted. Assume $\pi_n \Rightarrow \pi$. Since $\int_S y(s) C_{\pi_n}[\mathrm{d}s] = \int_W y(B(w, \pi_n), w) m_{\pi_n}[\mathrm{d}w]$, Lemma 9 completes the proof. □

Using continuity of v, we obtain the following trivial corollary.

Corollary 10. $V : P^2(W) \to P^2(W)$ *is continuous.*

Lemma 11. *Assume that* $\pi \geqslant \upsilon$ *and that either* π *or* υ *is monotonic. Then, for all* t *and* $\omega \in \Omega$, $V(B(\omega^t, \pi)) \geqslant V(B(\omega^t, \upsilon))$.

Proof. I will say that $y \in C(S)$ is non-decreasing if $y(\pi, w) \geqslant y(\upsilon, w)$ whenever $\pi \geqslant \upsilon$ and either π or υ is monotonic. Let $N(S) \subset C(S)$ be the set of non-decreasing functions on S. Since K maps $N(S)$ into itself and $N(S)$ is closed, the value function v is non-decreasing. If υ is monotonic, $V(\pi) \geqslant \int v[B(w, \upsilon), w] m_\pi(\mathrm{d}w) \geqslant V_T(\upsilon)$. The first inequality follows from non-decreasingness of v and $\pi \geqslant \upsilon$. The second, because $v[B(w, \upsilon), w]$ is non-decreasing in w for monotonic priors. The result follows because monotonicity and $\geqslant$ are preserved by updating. For monotonic π the proof is symmetric. □

Proof of Proposition 2. Let ρ be as in the definition of optimism. Given the optimal policy, $V(B(\omega^t, \pi)) \geqslant V(B(\omega^t, \rho)) \geqslant V(B(\omega^t, \upsilon))$ for all t and $\omega \in \Omega$ will complete the proof. The result then follows from Lemma 11. □

Proof of Lemma 3. Given that $w_2 > \delta V(\pi)$, in the first period, search continues only if the first draw is w_1. Since priors are monotonic and $m_\pi = E_{m_\pi}[m_{B(w,\pi)}]$, $\pi \geqslant B(w_1, \pi)$. Since $\pi \geqslant B(w_1, \pi)$ and both are monotonic, Lemma 11 insures that $\delta V(\pi) \geqslant \delta V(B(w_1, \pi))$. Then, $w_2 > \delta V(\pi) \geqslant \delta V(B(w_1, \pi))$. Hence, in period 2, search continues only if w_1 occurs. Again, $B(w_1, \pi) \geqslant B(w_1, w_1, \pi)$, so $\delta V(B(w_1, \pi)) \geqslant \delta V(B(w_1, w_1, \pi))$. Continuing in this manner, the result follows. □

Lemma 12. *For any Dirichlet* $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$, $w_1 = \delta V(\pi)$ *implies* $V(\pi) = \int w m_\pi(\mathrm{d}w)$.

Proof. Since $w_1 > \delta V(B(w_1, \pi))$ implies that $V(\pi) = \int w m_\pi(\mathrm{d}w)$, it will suffice to show that $V(\pi) > V(B(w_1, \pi))$. Then, $V(\pi) \geqslant \int \max\{w, \delta V(B((w_1, w_1), \pi))\} m_\pi(\mathrm{d}w) > V(B(w_1, \pi))$. The first inequality follows from $B(w, \pi) \geqslant B((w_1, w_1), \pi)$ and Lemma 11. The second, since $\max\{w, \delta V(B((w_1, w_1), \pi))\}$ is strictly increasing and $\pi > B(w_1, \pi)$. □

Proof of Proposition 4. Trivially, for any f there exists $\bar{\delta}$ such that for all $\bar{\delta} > \delta$, $w_2 > \delta V(\upsilon^0)$. Then, (ii) follows directly from Theorem 2, Lemma 13, and Lemma 3.

To prove (i), I will find Dirichlet priors for unbiased and overconfident searchers. For $s > 0$, let υ^s be a Dirichlet prior with parameter $(f_1/s, f_2/s, \ldots, f_n/s)$. For small $\gamma > 0$, let $f_\gamma \in P(W)$ be defined by $f_\gamma = (f_1 - \gamma, f_2 + \gamma, \ldots, f_n)$ and let α^γ be degenerate in f_γ. For all $\gamma > 0$, $V(\alpha^\gamma) > V(\upsilon^0)$, so continuity of V guarantees that for small $\bar{\gamma}$, $\delta V(\alpha^{\bar{\gamma}}) \in (\delta V(\upsilon^0), w_2)$. Define $\pi^0 \equiv \alpha^{\bar{\gamma}}$ and, for $s > 0$, let π^s be a Dirichlet prior with parameter $((f_1 - \bar{\gamma})/s, (f_2 + \bar{\gamma})/s, \ldots, f_n/s)$. Then, continuity of B and of V guarantees that there exists an S such that for all $s < S$, $\delta V(B(w_2, \pi^s)) \in (\delta V(\upsilon^0), w_2)$. Hence, for all $s < S$, π^s satisfies the conditions of Lemma 3. This implies that π^s will never reject an offer that he should not. Thus, for all t, all ω and $s < S$, $\delta V(B(\omega^t, \pi^s)) \geqslant \delta V(B(\omega^t, \upsilon^s))$. To show that being overconfident is strictly better off than being unbiased, it suffices to prove that for some history with positive probability, the searcher with priors υ^s accepts w_1 and with π^s rejects it.

Let w_1^t denote a sequence of t draws of w_1. Then, for all t, $\wp\{\omega\colon \omega^t = w_1^t\} > 0$. It will suffice to show that for some s and some t, $\delta V(B(w_1^t, \upsilon^s)) \leqslant w_1 < \delta V(B(w_1^t, \pi^s))$.

Fix $s_1 < S$. Since, for t large enough, $B(w_1^t, \upsilon^{s_1})$ is close to a degenerate belief in a distribution that is degenerate in w_1, continuity of V implies $\delta V(B(w_1^t, \upsilon^{s_1})) < w_1$. Then, $\delta V(B(w_1^t, \upsilon^0)) > w_1$ and continuity of V and B guarantees that for some $s_2 < s_1$, $\delta V(B(w_1^t, \upsilon^{s_2})) = w_1$. Then, by Lemma 12, $w_1/\delta = \int w m_{B(w_1^t, \upsilon^{s_2})}(\mathrm{d}w) < \int w m_{B(w_1^t, \pi^{s_2})}(\mathrm{d}w) \leqslant V(B(w_1^t, \pi^{s_2}))$. Letting $\pi = \pi^{s_2}$ and $\upsilon = \upsilon^{s_2}$ completes the proof. □

For each $r \in [0, 1]$, let r_x denote the true value following the policy "in time t, if in the dyadic expansion of r the tth element is 1, accept w iff $w \geqslant x$; if the tth element is 0, accept w iff $w > x$." If r has two expansions, the choice between them is irrelevant.

Lemma 13. *For degenerate* $\pi^0 \in P^2(W)$, $V^\wp$ *is continuous at* π^0 *iff* $0_{\delta V(\pi^0)} = 1_{\delta V(\pi^0)}$.

Proof. I will first show sufficiency. Assume that $1_{\delta V(\pi^0)} = 0_{\delta V(\pi^0)}$. It is easy to see, by induction, that for all T, that if $q, r \in [0, 1]$ have a constant string of 0 or 1's after T, $q_{\delta V(\pi^0)} = r_{\delta V(\pi^0)}$. By continuity of V, for fixed $\gamma > 0$ and $T < \infty$, I can choose $\epsilon > 0$ so that for all $\pi \in N_\epsilon(\pi^0)$, all $t \leqslant T$ and $\omega \in \Omega$, $V(B(\omega^t, \pi)) \in N_\gamma(V(\pi^0))$. Then, for every $\omega \in \Omega$ there exists some $r(\omega) \in [0, 1]$ with $1_{\delta V(\pi^0)} = r(\omega)_{\delta V(\pi^0)}$, such that the choices made by a searcher with prior π who follows the optimal strategy are the same as those dictated by $r(\omega)$ for $t \leqslant T$. Note that for all ω and ω', the r's chosen are such that $r(\omega)_{\delta V(\pi^0)} = r(\omega')_{\delta V(\pi^0)} = 1_{\delta V(\pi^0)} \equiv r_{\delta V(\pi^0)}$. Then, I get $|V^{\wp}(\pi^0) - V^{\wp}(\pi)| = |r_{\delta V(\pi^0)} - V^{\wp}(\pi)| \leqslant \delta^T w_n$. Noting that T was arbitrary completes the proof of sufficiency.

Assume $1_{\delta V(\pi^0)} \neq 0_{\delta V(\pi^0)}$ and let π^0 be degenerate in $(q_1, q_2, \ldots, q_n) \in P(W)$. Since $q_i = \delta V(\pi^0)$ for some $i < n$, let π^s be degenerate in $(q_1, \ldots, q_i - \epsilon_s, \ldots, q_n + \epsilon_s)$ for $\epsilon_s \downarrow 0$. Then, for all s, $V(\pi^s) > V(\pi^0)$ and for large s, $|V^{\wp}(\pi^0) - V^{\wp}(\pi^s)| = |1_{\delta V(\pi^0)} - 0_{\delta V(\pi^0)}| \neq 0$. □

Proof of Proposition 5. $\{V(\upsilon^0), V(\pi^0)\} \cap \{w\colon f(w) > 0\} = \phi$ insures that the condition for Lemma 13 is met, so $V^{\wp}$ is continuous both at π^0 and υ^0. □

Proof of Proposition 6. The part of overconfidence will be proved by induction. The other is analogous and will be omitted. Monotonicity and $m_\pi \trianglerighteq f$ guarantee that $\int_W \int_{-\infty}^x m_{B(w,\pi)}(\mathrm{d}t) m_\pi(\mathrm{d}w) \leqslant \int_W \int_{-\infty}^x m_{B(w,\pi)}(\mathrm{d}t) f(\mathrm{d}w)$ and thus, $m_\pi \trianglerighteq E_f[m_{B(w,\pi)}]$. By assumption, there exists υ such that $\int m_{B(w,\pi)} f(\mathrm{d}w) \trianglerighteq \int m_{B(w,\upsilon)} f(\mathrm{d}w) = \int m_{B(w,\upsilon)} m_\upsilon(\mathrm{d}w) = m_\upsilon = f$.

Assuming $E_{\wp}[m_{B(\omega^{t-1},\pi)}] \trianglerighteq f$, $E_{E_{\wp}[m_{B(\omega^{t-1},\pi)}]}[m_{B(\omega^t,\pi)}] = E_{\wp}[m_{B(\omega^{t-1},\pi)}]$ and monotonicity, guarantees that $E_{\wp}[m_{B(\omega^{t-1},\pi)}] \trianglerighteq E_{\wp}[m_{B(\omega^t,\pi)}]$. Finally, for unbiased υ with $\pi \geqslant \upsilon$, $E_{\wp}[m_{B(\omega^t,\pi)}] \trianglerighteq E_{\wp}[m_{B(\omega^t,\upsilon)}] = f$. □

References

Alpert, M., Raiffa, H., 1982. A progress report on the training of probability assessors. In: Kahneman, D., Slovic, P., Tversky, A. (Eds.), Judgement Under Uncertainty: Heuristics and Biases. Cambridge Univ. Press, Cambridge.

Beaudry, P., Portier, F., 1998. An exploration into Pigou's theory of cycles. Mimeo.

Benos, A., 1998. Aggressiveness and survival of overconfident traders. Journal of Financial Markets. In press.

Berkovitch, E., 1990. A stigma theory of unemployment duration. In: Weiss, Y., Fishelson, G. (Eds.), Advances in the Theory and Measurement of Unemployment. MacMillan.

Bikhchandani, S., Sharma, S., 1996. Optimal search with learning. Journal of Economic Dynamics and Control 20, 333–359.

Billingsley, P., 1968. Convergence of Probability Measures. Wiley, New York.

Burdett, K., Vishwanath, T., 1988. Declining reservation wages. Review of Economic Studies 55, 655–665.

Camerer, C., 1997. Progress in behavioral game theory. Journal of Economic Perspectives 11 (4).

De Groot, M.H., 1970. Optimal Statistical Decisions. McGraw–Hill, New York.

Denardo, E.V., 1967. Contraction mappings in the theory underlying dynamic porgramming. SIAM Review 9 (2).

Dubins, L., Savage, L., 1965. How to Gamble if You Must. McGraw–Hill, New York.

Gervais, S., Odean, T., 1997. Learning to be overconfident. Working paper. Davis, Philadelphia.

Kalai, E., Lehrer, E., 1993. Rational learning leads to Nash equilibrium. Econometrica 61 (5).

Kiyotaki, N., 1988. Multiple expectational equilibria under monopolistic competition. Quaterly Journal of Economics 103, 695–714.

Kyle, A., Wang, F.A., 1997. Speculation duopoly with agreement to disagree: Can overconfidence survive the market test? Journal of Finance 52, 2073–2090.

Kohn, M.G., Shavell, S., 1974. The theory of search. Journal of Economic Theory 9 (2).

Manove, M., 1995. Enterpreneurs, optimism and the competitive edge. Universitat Autonoma WP 296.95, de Barcelona.

Manove, M., Padilla, A.J., 1999. Banking (conservatively) with optimists. RAND Journal of Economics 30 (2), 324–350.

Milgrom, P., 1981. Good new and bad news: representation theorems and applications. Bell Journal of Economics 12, 380–391.

Müller, A., Stoyan, D., 2002. Comparison Methods for Stochastic Models and Risks. Wiley, New York.

Odean, T., 1998. Volume, volatility, price, and profit when all traders are above average. Journal of Finance 53, 6.

Rothschild, M., 1974. Searching for the lowest price when the distribution of prices is unknown. Journal of Political Economy 82 (4).

Shiryayev, A.N., 1984. Probability. Springer–Verlag, Berlin.