

ECONOMIC RESEARCH REPORTS

Overconfidence in Search


By

Juan Dubra

RR# 99-10

September 1999

**C.V. STARR CENTER
FOR APPLIED ECONOMICS**



NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, NY 10003-6687

Overconfidence in Search*

By Juan Dubra[†]

Department of Economics, New York University

September 14, 1999

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 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 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

*I am indebted to Charles Wilson for his guidance. I am grateful to Jean Pierre Benoît, Alberto Bisin, Federico Echenique, Néstor Gandelman and Efe Ok for helpful comments.

[†]jmd228@nyu.edu

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

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. Several other authors have also studied search behavior when the distribution is not known. For example, Kohn and Shavell (1974) proves the existence of an optimal policy for a very general class of beliefs. Rothschild (1974) shows that, for a limited class of beliefs, the optimal policy has the same properties as that in standard search models. Burdett and Vishwanath (1988) proves that, if search costs are large, the reservation wage of workers decreases over time. Bikhchandani and Sharma (1996) give sufficient conditions for searchers to follow reservation wage policies. I address a different set of issues.

This paper has three objectives. The first is to establish the behavioral implications of optimism. 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 adopt this definition of optimism and show that it fails to predict optimistic behavior. That is, optimistic searchers may accept offers that pessimists reject. Thus, I find conditions that guarantee that optimistic individuals search longer than pessimists. Offers have informational value because, as search evolves, individuals learn about the unknown distribution. Suppose then that a low offer has a higher value than a high offer, violating what I will call Monotonicity. In that case, optimism may lead to a lower expected value of searching than pessimism. This, in turn, yields shorter search times for the optimistic agent. The main result on behavior is that an optimistic searcher samples longer than a pessimist, whenever one of them has monotonic 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. 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 ensure that overconfident beliefs diminish over time and never become pessimistic, it suffices to assume that there is an unbiased belief that is 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.

I conclude with a discussion of the testable implications of this model and 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 (1998) 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 invest, demand is high in equilibrium, so there is no over-optimism. In my model, the searcher is overly optimistic about the distribution. Another 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 topological space (X, \mathcal{T}) 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 $W \equiv \{w_1, w_2, \dots, w_n\} \subset \mathbf{R}_+$, with $0 < w_1 < w_2 < \dots < w_n$, and define $P^2(W) = P(P(W))$. I will represent any $g \in P(W)$ by (g_1, \dots, 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 $\pi \in P^2(W)$.

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 t th element. Starting with beliefs π and after a history ω^t , the probability of any measurable set $C \subset P(W)$ is

$$B(\omega^t, \pi)(C) = \int_C \frac{\prod_{i \leq t} g(\omega_i)}{\int \prod_{i \leq t} g(\omega_i) \pi(dg)} \pi(dg)$$

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

$$m_\pi(w) \equiv \int_{P(W)} g(w) \pi(dg)$$

Since, π is a probability over distributions, m_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:

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

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 r means that an offer is rejected, and a indicates that an offer is accepted. For any state (π, w) and action c , define the transition $q(\cdot | (\pi, w), c)$ by

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

Given state (π, w) , if the searcher chooses an action c , $q(s | (\pi, w), c)$ 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 $p_t : \mathbf{H}_t \rightarrow A$. 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 $E_\pi \left[\delta^{\tau(p, \omega)} \omega_{\tau(p, \omega)} \right]$ and the value function $v : S \rightarrow \mathbf{R}$ is $v(s) = \sup_p E_\pi \left[\delta^{\tau(p, \omega)} \omega_{\tau(p, \omega)} \right]$. The

following lemma states that $Ky(\pi, w) = \max \{w, \delta \int_S y(s') q(ds' | (\pi, w), r)\}$, $y \in C(S)$, is a well defined function $K : C(S) \rightarrow C(S)$. 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 .² Define $V(\pi) \equiv \int v[B(w, \pi), w] m_\pi(dw)$, the maximum value of searching when beliefs are π . Then, in any state $(\pi, w) \in S$, accepting an offer if and only if

$$w \geq \delta V(\pi) \quad (\text{Optimal Policy})$$

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 Corollary 2 in Denardo (1967) 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.

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 one to \$1 and another with $prob(\$2) = 1 - prob(\$3) = .01$. 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 this section, I derive conditions that guarantee that optimistic searchers obtain a higher subjective expected value of searching than pessimists after any history of draws. This ensures that the optimal strategy is to search longer. As was shown by Bikhchandani and Sharma (1996), a static definition of optimism is not sufficient to ensure longer search times for the optimist. Because pessimists may have priors that are less affected by updating than optimists, downward updating can lead the initially optimistic searcher to stop sampling before a pessimistic decision maker. I adopt their definition of optimism which ensures that the optimistic individual assigns higher probabilities (than the pessimist) to high offers after

²See Theorem 3 in Denardo (1967). 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.

any sequence of draws. Next I show that even this restriction 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 of a low offer exceeds that of a high offer, assigning high probabilities to high proposals may lead to a low value of searching. To rule out this possibility, I define a property called Monotonicity which ensures that the informational value of offers is ordered in the same way as their monetary value. Finally, I show that an optimistic searcher samples longer than a pessimist, whenever one of them has monotonic priors.

For $g, h \in P(W)$, g *first order stochastically dominates* h , denoted $g \succeq h$, if and only if $\int u(w)g(dw) \geq \int u(w)h(dw)$, for all non decreasing functions u (Dubins and Savage, 1965). For static decision problems, \succeq captures the idea of optimism. The following example, which is similar to Example 1 in Bikhchandani and Sharma (1996), illustrates how different propensities to update may lead an optimistic individual to stop sampling before a pessimistic searcher.

Example 1: Let $W = \{1, 2\}$ and $\frac{1}{2} \geq q \geq 0$. Define $f, g, h \in P(W)$ by $h = (1, 0)$, $f = (\frac{1}{2}, \frac{1}{2})$ and $g = (0, 1)$. Also, define priors π, v by $v(f) = 1$ and $\pi(g) = \frac{1}{2} + q$, $\pi(h) = \frac{1}{2} - q$.

The posterior of the optimistic prior π is degenerate in h after receiving a draw of 1. Thus, the optimistic searcher 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 pessimistic searcher (with prior v) never revises his priors and, for $\delta > \frac{2}{3}$, continues sampling until a high draw occurs. Since the size of q indexes the degree of optimism, for *all* levels of optimism and *all* offer paths, the optimistic individual never samples longer than the pessimistic searcher and sometimes samples less. ■

This result is driven by the fact that π is affected by updating and v is not, which causes the order of optimism to be reversed with the arrival of information. In a sense, π is more “stubborn” in the face of new information.

Equal Stubbornness: $\pi, v \in P^2(W)$ are *equally stubborn* if and only if, $m_\pi \succeq m_v$ implies $m_{B(\omega^t, \pi)} \succeq m_{B(\omega^t, v)}$ for all t and all $\omega \in \Omega$.

Equal Stubbornness states that if one prior is statically more optimistic than another, the relationship is maintained after receiving the same information. For example, two Dirichlet priors³ (π_1, \dots, π_n) and (v_1, \dots, v_n) are equally stubborn if $\sum_1^n \pi_i = \sum_1^n v_i$.

³A Dirichlet with parameter $\pi = (\pi_1, \pi_2, \dots, \pi_n)$, with $\pi_i > 0$ for all i , is a probability measure over $P(W)$. Let $S(\pi) = \sum_1^n \pi_i$, and $\mu_i = \frac{\pi_i}{S(\pi)}$. 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 i th canonical vector). See De Groot (1970).

To say that one prior is dynamically more optimistic than another, we need to restrict attention to priors for which static optimism is preserved when the same information is observed. The following definition is essentially the same as that in Bikhchandani and Sharma (1996). Define the partial order \succeq on $P^2(W)$ by $\pi \succeq v$ if and only if $m_\pi \supseteq m_v$ for equally stubborn $\pi, v \in P^2(W)$. If $\pi \succeq v$, I will say that π is *more optimistic* than v .

The next example shows that even Equal Stubbornness is not sufficient to ensure that π will search longer than v when $\pi \succeq v$.

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

Fix $\delta = .99$ and suppose $\epsilon = 0$. Assume also that a searcher is certain that the distribution is $j_0 = (\frac{1}{4}, 0, 0, \frac{3}{4})$. Because δ is close to 1, he samples until $w = 6$ is drawn and obtains an expected payoff of approximately 6. The same is true for $g_0 = (\frac{3}{4}, 0, 0, \frac{1}{4})$. 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 $\pi^\epsilon, v^\epsilon \in P^2(W)$ by $v^\epsilon(g_\epsilon) = 1 - v^\epsilon(h_\epsilon) = \frac{4}{5}$ and $\pi^\epsilon(j_\epsilon) = 1 - \pi^\epsilon(h_\epsilon) = \frac{1}{2}$. 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 $V(\pi_0) \approx \frac{6+4}{2} = 5$. Analogously, $V(v_0) \approx \frac{4}{5}6 + \frac{1}{5}4 > 5 \approx V(\pi_0)$. That is, π_0 is more optimistic than v_0 and yields a smaller value of searching. This result is driven by the fact that a draw of 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 $5 \in (\delta V(\pi^0), \delta V(v^0))$, if $w = 5$ is drawn and priors are not updated, the searcher with prior π_0 accepts the offer and the one with v_0 does not. However, $w = 5$ is a zero probability event for both priors, so I will slightly modify them to ensure that searchers can use Bayes' rule. For any ϵ , when $w = 5$ is drawn, updating does not change π^ϵ or v^ϵ . Then $5 \in (\delta V(\pi^0), \delta V(v^0))$ guarantees that $5 \in (\delta V(\pi^\epsilon), \delta V(v^\epsilon)) = (\delta V(B(5, \pi^\epsilon)), \delta V(B(5, v^\epsilon)))$ for small enough ϵ . Therefore, when $w = 5$ is drawn the optimistic searcher (with priors π^ϵ) will accept the offer and the pessimistic individual (with beliefs v^ϵ) will reject it. Moreover, since $V(\pi^\epsilon) < V(v^\epsilon)$, an optimistic searcher obtains a lower subjective expected value of searching than a pessimist. ■

In this example, an optimistic searcher stops sampling before a pessimist because a low offer has high informational value. That is, it is not true that high offers are good news. I now define a concept that formalizes this property.

Monotonicity: $\pi \in P^2(W)$ is *monotonic* if and only if, for all $\omega, \kappa \in \Omega$ and t , $\omega^{t-1} = \kappa^{t-1}$

and $\omega_t \geq \kappa_t$ imply $m_{B(\omega^t, \pi)} \succeq m_{B(\kappa^t, \pi)}$.

That is, after observing high offers, posteriors are statically more optimistic than after receiving low offers. Monotonicity ensures 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. This condition is similar to those used by Bikhchandani and Sharma (1996), Burdett and Vishwanath (1988) and Milgrom (1981). Unless otherwise stated, I will restrict attention to monotonic priors.

In Example 1 searchers have monotonic beliefs. Nevertheless, the failure of equal stubbornness allowed the statically optimistic agent to stop sampling before the pessimist. Example 2 shows that if Monotonicity fails, the optimistic searcher may stop sampling before the pessimist, even if they are equally stubborn. However, the following theorem shows that Monotonicity and Equal Stubbornness ensure that the statically more optimistic searcher samples longer in *all* offer paths.

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.

Theorem 2 *Assume that $\pi \succeq v$ and that either π or v are monotonic. Then, for all $\omega \in \Omega$, $\tau_\pi(\omega) \geq \tau_v(\omega)$.*

Notice, again, that the theorem 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 reject offers that the pessimist does not. Monotonicity guarantees that high offers are better than low offers and Equal Stubbornness ensures that the order of optimism is not reversed.

Corollary 1 in Bikhchandani and Sharma (1996) proves that, if priors satisfy the same assumptions as in Theorem 2 and searchers follow a reservation wage policy, optimistic searchers sample longer. However, since they concentrate on problems with no discounting, Theorem 2 is not a generalization of their result.

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 and only if $m_\pi = f$ and *overconfident* if and only if $m_\pi \succeq f$. Let \wp be the measure on Ω obtained by extending the probabilities that f induces on W^T for all T .⁴ For the welfare criterion I use the true expected value of searching,

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

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 equally stubborn unbiased searcher.

Example 3: Let $W = \{1, 2\}$, $\frac{8}{11} > \delta > \frac{2}{3}$ and $f = (\frac{1}{2}, \frac{1}{2})$. 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 $g = (\frac{3}{4}, \frac{1}{4})$ and $h = (\frac{1}{4}, \frac{3}{4})$. Also, let $\pi(h) = 1 - \pi(g) = \frac{3}{4}$ and $v(h) = 1 - v(g) = \frac{1}{2}$. 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 $\frac{1}{\delta}$ and then accepting any offer. If the continuation value of searching falls below $\frac{1}{\delta}$ the agent accepts the current offer, so it suffices to show that whenever the expected value of the next draw falls below $\frac{1}{\delta}$, 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 $\frac{1}{\delta}$. 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 $\frac{1}{\delta}$, it is equal to 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 one to $w = 1$. Thus, the continuation value eventually falls below $\frac{1}{\delta}$. Finally, since the expected value of the next draw is decreasing over time, whenever it falls below $\frac{1}{\delta}$ it must be the continuation value.

Since $m_{B(1,v)}(1) = \frac{5}{8}$, the expected value of the next draw after observing $w = 1$ is $\frac{11}{8} < \frac{1}{\delta}$. Thus, the unbiased searcher stops sampling in the first period in any $\omega \in \Omega$ and obtains a true value of searching of $\frac{3}{2}$. Since $m_{B(1,\pi)}(1) = \frac{1}{2}$ the expected value of the next draw after the first bad draw is $\frac{3}{2} > \frac{1}{\delta}$. Therefore the overconfident searcher rejects the first low

⁴See the Kolmogorov Extension Theorem in Shiryaev (1984) p. 161.

offer, but because $m_{B((1,1),\pi)}(1) = \frac{5}{8}$ he accepts it in the second. This yields a true value of searching of $1 + \delta\frac{3}{4} > \frac{3}{2}$. ■

In this example, the overconfident searcher uses the true optimal strategy in period 1 whereas the unbiased 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 overconfident 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 ensure 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 equally stubborn unbiased searchers. This suggests that, when these conditions hold, overconfident searchers are better off than unbiased individuals. However, the following example shows that this is false in general.

Example 4: Let $W = \{1, 2, 100\}$, $f = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and $\delta = \frac{58}{1000}$. 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 $\frac{100}{3} + \frac{206}{9}\delta$. Rejecting only offers of 1 in the first period and accepting any offer in the second yields a payoff of $34 + \frac{103}{9}\delta > \frac{100}{3} + \frac{206}{9}\delta$. Therefore, following the optimal strategy in the first period is harmful. In the remainder of the example I show how the above behavior can be derived from overconfident and unbiased priors.

Let $j = (\frac{1}{1000}, \frac{1}{1000}, \frac{998}{1000})$ and $g = (\frac{999}{2000}, \frac{999}{2000}, \frac{1}{1000})$. Define the overconfident prior π by $\pi(k) = 1 - \pi(g) = .98569$ and the unbiased prior v by $v(j) = 1 - v(g) = .33266$. If $w = 100$ has not occurred in periods 1 or 2, the expected value of the next draw is lower than $\frac{1}{\delta}$ 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 $2 > \delta E_{B(2,v)}[w] = \delta E_{B(1,v)}[w] > 1$ 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 ensure 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 the worse news he could imagine, reservation wages are decreasing. This, in turn, ensures 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 lowest offer to which he assigns positive probability, his reservation wage is decreasing.⁵

Lemma 3 *Suppose that a searcher with prior π follows a reservation wage policy and that $w_2 > \delta V(\pi)$. Then, $\delta V(B(\omega^{t-1}, \pi)) \geq \delta V(B(\omega^t, \pi))$ 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 with equally stubborn priors. By Theorem 2, the unbiased searcher makes a mistake before the overconfident and this makes him worse off. A similar reasoning applies to show that underconfident individuals are still worse off. This is summarized in the following theorem. For any $v \in P^2(W)$, any $\epsilon > 0$ and metric d , define $N_\epsilon(v) \equiv \{\pi \in P^2(W) : d[\pi, v] < \epsilon\}$.

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

i) for any $\epsilon > 0$ there is an unbiased $v \in N_\epsilon(v^0)$ and an equally stubborn overconfident π for which $V^\varphi(\pi) \geq V^\varphi(v)$. Moreover, if $f([0, \delta V(v^0)]) > 0$, $V^\varphi(\pi) > V^\varphi(v)$

ii) there exists $\gamma > 0$ such that for all unbiased $\nu \in N_\gamma(v^0)$ that follows a reservation wage policy, if φ is an equally stubborn underconfident prior, $V^\varphi(\nu) \geq V^\varphi(\varphi)$

⁵Bikhchandani and Sharma (1996) provide sufficient conditions on priors to ensure that searchers follow reservation wage rules.

Theorem 4 says that there exists an unbiased searcher who is almost certain about the truth and an overconfident searcher with equally stubborn priors who is better off. Second, it is overconfidence, and not an arbitrary bias, that makes the overconfident searchers better off. Underconfident searchers are still worse off.

4.2 The Costs of Overconfidence

Theorem 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 is the obvious one: 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 that searchers are stubborn (and keep making their original mistakes) needs to be added.

Consider an individual with priors v close to the degenerate v^0 . By continuity of V (see Corollary 10) one can make sure that, for almost any discount factor, the search rule of v resembles that of v^0 for a long period of time. Therefore, discounting ensures that V^φ , the true value of searching, is continuous at v^0 . Then, for v and π close to the degenerate v^0 and π^0 respectively, $V^\varphi(v^0) > V^\varphi(\pi^0)$ guarantees $V^\varphi(v) > V^\varphi(\pi)$. The result is summarized in the following theorem.

Theorem 5 *Fix any degenerate priors π^0 and v^0 . Assume $f(\delta V(v^0)) = 0$, $f(\delta V(\pi^0)) = 0$ and $V^\varphi(v^0) > V^\varphi(\pi^0)$. Then, there exists ϵ such that for all $\pi \in N_\epsilon(\pi^0)$ and $v \in N_\epsilon(v^0)$, $V^\varphi(v) > V^\varphi(\pi)$.*

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 $W = \{1, 2, 3, 4\}$, $\delta = .99$ and define $g, j \in P(W)$ by $g = (0, \frac{5}{12}, \frac{1}{2}, \frac{1}{12})$ and $j = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$. Also, for $1 \geq \epsilon \geq 0$, 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.⁶ ■

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 pessimistic. Then I show that, if there is an unbiased belief that is equally stubborn than the overconfident prior, the bias diminishes over time and the overconfident does not become pessimistic on average.

Suppose that v is unbiased. Then, by the law of iterated expectation, $E_f [m_{B(w,v)}] = f$. 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 $g = (\frac{3}{4}, \frac{1}{4}, 0)$, $f = (\frac{1}{2}, \frac{1}{2}, 0)$, $j = (0, 0, 1)$. Define priors π by $\pi(g) = \pi(j) = \frac{1}{2}$. Since only offers of 1 and 2 will occur, the posterior of π is always $g = (\frac{3}{4}, \frac{1}{4}, 0)$. Thus, although π is overconfident, he becomes underconfident with probability one. This implies that $f \triangleright E_f [m_{B(w,\pi)}]$. ■

In the example, there does not exist an unbiased belief that is equally stubborn than π . For any unbiased belief v , $E_{m_v} [m_{B(w,v)}] = m_v$ implies that $m_{B(2,v)} \succeq m_v = (\frac{1}{2}, \frac{1}{2}, 0) \triangleright (\frac{3}{4}, \frac{1}{4}, 0) = m_{B(2,\pi)}$, violating equal stubbornness. That is, while 2 is good news for v , it is “very” bad news for π , and that causes their order of optimism to be reversed. If there was an equally stubborn unbiased belief, π would remain optimistic 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 theorem is a generalization of the previous argument.

⁶In Burdett and Vishwanath (1988) this possibility is ruled out assuming that the cost of search is large, which ensures that only “very” low offers are rejected.

Proposition 6 *Fix any prior π such that there exist an equally stubborn unbiased belief. If π is overconfident $m_\pi \supseteq E_\varphi [m_{B(\omega^{t-1}, \pi)}] \supseteq E_\varphi [m_{B(\omega^t, \pi)}] \supseteq f$, for all t . If π is underconfident $f \supseteq E_\varphi [m_{B(\omega^t, \pi)}] \supseteq E_\varphi [m_{B(\omega^{t-1}, \pi)}] \supseteq m_\pi$, for all t .*

Proposition 6 states that, for overconfident priors for which there is an equally stubborn 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 below by the unbiased priors, which are a martingale.

A corollary of Proposition 6 is that beliefs are a martingale for unbiased priors. That is, the agent’s true average beliefs about the distribution that generates the offers does not change over time. This is not the usual “beliefs are a 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 a martingale” means that one can not 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. Assume, also, that the searcher starts each spell with equally stubborn priors. In addition, following Proposition 6, suppose that the overconfident’s prior at the beginning of today’s spell are dominated by his beliefs at the start of the last spell. Then, Theorem 2 ensures 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.

The results on welfare may also have implications for social policy. For example, consider the case of an unemployment insurance office. Since the welfare loss in Theorem 5 is derived exclusively from beliefs, policymaking only makes sense if the government knows more about the true distribution than searchers. The unemployment insurance office observes the same characteristics of the worker as the firm does and has been continuously receiving information

about job offers for a long time. Then, if it can be assumed to know the distribution that the unemployed worker will face, inducing workers, for example, to take jobs they do not want to accept, may be welfare enhancing. In traditional information economics models, where the individual knows more about himself than the office, inducing workers to take jobs is welfare decreasing.

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.

In closing, I note that all the results of the paper, except those on welfare, can be easily extended to the case of arbitrary $W \subset \mathbf{R}$.

7 APPENDIX

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

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

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

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

$$\int_{P(W)} h_n(g) \pi_n(dg) \equiv \int_{P(W)} \frac{y(g) g(w)}{m_{\pi_n}(w)} \pi_n(dg) \rightarrow \int_{P(W)} \frac{y(g) g(w)}{m_{\pi}(w)} \pi(dg) \equiv \int_{P(W)} h(g) \pi(dg)$$

The range r_n of each h_n is bounded. Then, since $\int_{P(W)} h_n(g) \pi_n(dg) = \int_{r_n} y \pi_n h_n^{-1}(dy)$ 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 $g_n(w_i) \rightarrow g(w_i)$. This, and continuity of y guarantee that $|y(g_n) g_n(w) - y(g) g(w)| \rightarrow 0$. 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 $\{\pi_n\}_1^\infty, \pi \in P^2(W)$, $\int_W y(B(w, \pi_n), w) m_{\pi_n} [dw] \rightarrow \int_W y(B(w, \pi), w) m_\pi [dw]$ 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 1: Proofs of continuity when search has stopped and of boundedness are trivial and will be omitted. Assume $\pi_n \Rightarrow \pi$. Since $\int_S y(s) C_{\pi_n} [ds] = \int_W y(B(w, \pi_n), w) m_{\pi_n} [dw]$, Lemma 9 completes the proof. ■

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

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

Lemma 11 Assume that $\pi \succeq v$ and that either π or v are monotonic. Then, for all t and $\omega \in \Omega$, $V(B(\omega^t, \pi)) \geq V(B(\omega^t, v))$

Proof: I will say that $y \in C(S)$ is non decreasing if $y(\pi, w) \geq y(v, w)$ whenever $\pi \succeq v$ and either π or v are 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 v is monotonic, $V(\pi) \geq \int v[B(w, v), w] m_\pi(dw) \geq V_T(v)$. The first inequality follows from equal stubbornness and non decreasingness of v . The second, because $v[B(w, v), w]$ is non decreasing in w for monotonic priors. The result follows because monotonicity and equal stubbornness are preserved by updating. For monotonic π the proof is symmetric. ■

Proof of 2: Given the optimal policy, $V(B(\omega^t, \pi)) \geq V(B(\omega^t, v))$ for all t and $\omega \in \Omega$ will complete the proof. The result follows from Lemma 11. ■

Proof of 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 $m_\pi = E_{m_\pi}[m_{B(w, \pi)}]$, $\pi \succeq B(w_1, \pi)$. Since π and $B(w_1, \pi)$ are equally stubborn, Lemma 11 ensures that $\delta V(\pi) \geq \delta V(B(w_1, \pi))$. Then, $w_2 > \delta V(\pi) \geq \delta V(B(w_1, \pi))$. Hence, in period 2, search continues only if w_1 occurs. Again, $B(w_1, \pi) \succeq B(w_1, w_1, \pi)$ and they are equally stubborn, so $\delta V(B(w_1, \pi)) \geq \delta V(B(w_1, w_1, \pi))$. Continuing in this manner, the result follows. ■

Lemma 12 For any Dirichlet $\pi = (\pi_1, \pi_2, \dots, \pi_n)$, $w_1 = \delta V(\pi)$ implies $V(\pi) = \int w m_\pi(dw)$

Proof: Since $w_1 > \delta V(B(w_1, \pi))$ implies that $V(\pi) = \int w m_\pi(dw)$, it will suffice to show that $V(\pi) > V(B(w_1, \pi))$. Then, $V(\pi) \geq \int \max\{w, \delta V(B((w_1, w_1), \pi))\} m_\pi(dw) > V(B(w_1, \pi))$. The first inequality follows from $B(w, \pi) \succeq B((w_1, w_1), \pi)$, equal stubbornness and Lemma 11. The second, because $\max\{w, \delta V(B((w_1, w_1), \pi))\}$ is strictly increasing and $\pi \succ B(w_1, \pi)$. ■

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

To prove *i*) I will find Dirichlet priors for unbiased and overconfident searchers. For $s > 0$, let v^s be a Dirichlet prior with parameter $(\frac{f_1}{s}, \frac{f_2}{s}, \dots, \frac{f_n}{s})$. For small $\gamma > 0$, let $f_\gamma \in P(W)$ be defined by $f_\gamma = (f_1 - \gamma, f_2 + \gamma, \dots, f_n)$ and let α^γ be degenerate in f_γ . For all $\gamma > 0$, $V(\alpha^\gamma) > V(v^0)$, so continuity of V guarantees that for small $\bar{\gamma}$, $\delta V(\alpha^{\bar{\gamma}}) \in (\delta V(v^0), w_2)$. Define $\pi^0 \equiv \alpha^{\bar{\gamma}}$ and, for $s > 0$, let π^s be a Dirichlet prior with parameter $(\frac{f_1 - \bar{\gamma}}{s}, \frac{f_2 + \bar{\gamma}}{s}, \dots, \frac{f_n}{s})$. Then, continuity of B and of V guarantee that there exists an S such that for all $s < S$, $\delta V(B(w_2, \pi^s)) \in (\delta V(v^0), w_2)$. 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 $s < S$, $\delta V(B(\omega^t, \pi^s)) \geq \delta V(B(\omega^t, v^s))$. So, to show that some overconfident is strictly better off than some unbiased, it suffices to prove that for some history with positive probability, the searcher with priors v^s accepts w_1 and π^s rejects it.

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

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

For each $r \in [0, 1]$, let r_x denote the true value of following the policy “in time t , if in the dyadic expansion of r the t th element is a 1, accept w iff $w \geq x$. If the t th element is a 0, accept 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’s 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 $t \leq T$ 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 $t \leq T$. 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 $|V^\varphi(\pi^0) - V^\varphi(\pi)| = |r_{\delta V(\pi^0)} - V^\varphi(\pi)| \leq \delta^T w_n$. 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 $(q_1, q_2 \dots q_n) \in P(W)$. Since $q_i = \delta V(\pi^0)$ for some $i < n$, let π^s be degenerate in $(q_1, \dots, q_i - \epsilon_s, \dots, q_n + \epsilon_s)$ for $\epsilon_s \downarrow 0$. Then, for all s , $V(\pi^s) > V(\pi^0)$ and for large s , $|V^\varphi(\pi^0) - V^\varphi(\pi^s)| = |1_{\delta V(\pi^0)} - 0_{\delta V(\pi^0)}| \neq 0$ ■

Proof of 5: $\{V(v^0), V(\pi^0)\} \cap \{w : f(w) > 0\} = \emptyset$ ensures that the condition for Lemma 13 is met, so V^φ is continuous both at π^0 and v^0 . ■

Proof of 6 : The part of overconfidence will be proved by induction. The other is analogous and will be omitted. Monotonicity and $m_\pi \succeq f$ guarantee that $\int_{W-\infty}^x m_{B(w,\pi)}(dt) m_\pi(dw) \leq \int_{W-\infty}^x m_{B(w,\pi)}(dt) f(dw)$ and thus, $m_\pi \succeq E_f[m_{B(w,\pi)}]$. By assumption, there exists v , such that $\int m_{B(w,\pi)} f(dw) \succeq \int m_{B(w,v)} f(dw) = \int m_{B(w,v)} m_v(dw) = m_v = f$.

Assuming $E_\varphi[m_{B(\omega^{t-1},\pi)}] \succeq f$, $E_{E_\varphi[m_{B(\omega^{t-1},\pi)}]}[m_{B(\omega^t,\pi)}] = E_\varphi[m_{B(\omega^{t-1},\pi)}]$ and Monotonicity guarantee that $E_\varphi[m_{B(\omega^{t-1},\pi)}] \succeq E_\varphi[m_{B(\omega^t,\pi)}]$. Finally, for unbiased v equally stubborn than π , $E_\varphi[m_{B(\omega^t,\pi)}] \succeq E_\varphi[m_{B(\omega^t,v)}] = f$. ■

References

- [1] Alpert, M. and H. Raiffa (1982): "A Progress Report on the Training of Probability Assessors," in *Judgement Under Uncertainty: Heuristics and Biases*, D. Kahneman, P. Slovic and A. Tversky eds. Cambridge and New York: Cambridge University Press.
- [2] Beaudry, P. and F. Portier (1998): "An Exploration into Pigou's Theory of Cycles," mimeo.
- [3] Benos, A. (1998): "Aggressiveness and Survival of Overconfident Traders," *Journal of Financial Markets*, forthcoming.
- [4] Bikhchandani, S. and S. Sharma (1996): "Optimal Search with Learning," *Journal of Economic Dynamics and Control*, **20**, 333-59.
- [5] Billingsley, P. (1968): *Convergence of Probability Measures*. New York: Wiley.

- [6] Burdett, K. and T. Vishwanath (1988): “Declining Reservation Wages,” *Review of Economic Studies*, **55**, 655-65
- [7] Camerer, C. (1997): “Progress in Behavioral Game Theory,” *Journal of Economic Perspectives*, **11**, No. 4.
- [8] De Groot, M.H. (1970): *Optimal Statistical Decisions*. New York: McGraw-Hill.
- [9] Denardo, E.V. (1967): “Contraction Mappings in the Theory Underlying Dynamic Programming,” *SIAM Review*, **9**, No. 2
- [10] Dubins, L. and L. Savage (1965): *How to Gamble if you Must*. New York: McGraw-Hill
- [11] Gervais, S. and T. Odean (1997): “Learning to be Overconfident,” WP, U.C. Davis.
- [12] Kalai, E. and E. Lehrer (1993): “Rational Learning leads to Nash Equilibrium,” *Econometrica*, **61**, No.5
- [13] Kiyotaki, N. (1988): “Multiple Expectational Equilibria Under Monopolistic Competition,” *Quarterly Journal of Economics*, **103**, 695-714
- [14] Kyle, A. and F.A. Wang (1997): “Speculation Duopoly with Agreement to Disagree: can Overconfidence survive the market test?,” *Journal of Finance*, **52**, 2073-90
- [15] Kohn, M.G. and S. Shavell (1974): “The Theory of Search,” *Journal of Economic Theory*, **9**, No. 2.
- [16] Manove, M. (1995): “Entrepreneurs, Optimism and the Competitive Edge,” Universitat Autònoma de Barcelona, WP 296.95
- [17] Manove, M. and A.J. Padilla (1999): “Banking (Conservatively) with Optimists,” mimeo.
- [18] Milgrom, P. (1981): “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics*, **12**, 380-91
- [19] Odean, T. (1996): “Volume, Volatility, Price, and Profit when all Traders are Above Average,” *Journal of Finance*, **53**, 6
- [20] Rothschild, M. (1974): “Searching for the Lowest Price When the Distribution of Prices is Unknown,” *Journal of Political Economy*, **82**, No. 4.
- [21] Shiryaev, A.N. (1984): *Probability*. Berlin: Springer-Verlag.