Pragmatic Beliefs and Overconfidence*

Hans K. Hvide†

Norwegian School of Economics and Business

June 26, 2000

Abstract

Several studies indicate that humans are overconfident about their own (relative) abilities. We propose a notion of pragmatic beliefs, and show through an example that this concept can shed light on why overconfidence emerges. Through the example, we also shed light on the idea that that 'bounded rationality' may arise endogenously in a game - without assuming complexity costs.

1 Introduction

Several studies from psychology and experimental economics indicate that agents are overconfident about their own (relative) abilities: For a variety of skills, agents seem to believe they are more skillful than they really are.[1] Although 'rational' theories of

*Thanks to Ken Binmore, Eddie Dekel, Ariel Rubinstein, Gaute Torsvik, and Bertil Tungodden for conversations on the topic of the paper. And thanks to the editor, a referee, and seminar participants for very useful comments and suggestions.

†Hans K. Hvide, Norwegian School of Economics and Business, Helleveien 30, N-5045 Bergen, Norway.

[1]The classic example is car-driving, as in Svensson (1981). Most of the literature on overconfidence is based on interview data. An exception is Camerer and Lovallo (1999), who reports on overconfidence in entry decisions in games of skill (with monetary incentives). An interesting example of overconfidence in market data appears in Asubel (1991), whose findings indicate that the credit card market is characterized by users of credit cards being overconfident about their future ability to avoid overdrawn accounts. For references to the psychologically minded literature on overconfidence, see Camerer and Lovallo.

belief formation, like the Bayesian paradigm, can easily be made consistent with some individuals in a population being overconfident, it is hard to make such theories consistent with most, or a majority of, a population being overconfident, and across different settings. Hence rational theories of belief formation seem unfit to explain overconfidence.

The purpose of the present paper is to propose a notion of pragmatic beliefs, and to suggest that this concept can shed light on why overconfidence can occur. In short, we attempt to show with a simple game-theoretic example that if agents form beliefs pragmatically, overconfidence can be the equilibrium outcome.

An agent that forms beliefs pragmatically forms beliefs that are (the most) useful to him. The pragmatic explanation of beliefs should be contrasted with 'rational' approaches to beliefs, which focus on agents that try to reflect the 'world as it is' in their beliefs (with tools like Bayes' law, the law of double expectation and so forth). As will become clear, pragmatic beliefs need only have a weak relation to 'the world as it is'. We also contrast the pragmatic explanation of beliefs to approaches assuming that beliefs are of intrinsic value to agents, as with wishful thinking and with cognitive dissonance.

To show how the idea of pragmatic beliefs can shed light on overconfidence, we consider a simple model of a job market. An agent is matched with a firm, and the firm makes a take-it or leave it offer. If the agent accepts the offer, the game ends, while if the agent rejects the offer, he gets an outside opportunity. The take-it or leave-it offer made by the firm will in equilibrium depend partly on the worker's productivity in the firm, and partly on the agent's beliefs about his outside opportunity, which is commonly known between the firm and the worker. The key feature of this example is that instead of taking the agent's beliefs as a primitive, we let the agent's beliefs about his outside opportunity be an endogenous variable. We show that if agents form beliefs pragmatically, then in equilibrium these beliefs will be inflated compared to the true distribution of the outside opportunity, and interpret this as overconfidence.

A way to interpret the result is that overconfidence is a way for the agent to obtain a first mover advantage towards the firm; overconfidence acts as if the first mover privilege is shifted from the firm to the agent. Although this analogy is useful, we also consider the case when both firms and workers form beliefs pragmatically, and show that the same type of overconfidence results carry through to this case.

An immediate criticism against the idea of pragmatic belief formation is that human agents are incapable of constructing beliefs that they do not really believe in. We should therefore stress from the outset that the main justification for pragmatic beliefs is dynamic; pragmatic beliefs in a population are formed either by this selection pressure towards agents that are born pragmatists, or by agents gradually learning that a certain way of forming beliefs is more rewarding than other ways. Hence we make no assumptions on the agents' awareness about their own overconfidence. This, and other criticisms against the notion of pragmatic beliefs will be discussed in a separate section.

The paper is structured as follows. In Section 2, we discuss the definition of 'overconfidence' and relate it to rationality. In section 3, we analyze the job market example, and in Section 4 we provide a discussion of the approach to beliefs outlined in this paper. Finally, in Section 5 we conclude.

1.1 Related Literature

The paper is related to several strands of the literature. In this section, we consider this relation in some detail. The section can be skipped on first reading.

First, the paper is related to a new school of thinking on bounded rationality, exemplified by Rubinstein (1997), and surveyed by Lipman (1995). A central feature of this school is that humans are information processors that are imperfect when compared to the standard of the rationality paradigm. For example, Fershtman and Kalai (1994) consider a setting where a manager has limited attention span, and the firm therefore concentrates in only a few markets, and Rubinstein (1993) considers a setting where consumers have heterogenous ability in understanding complex pricing schemes, and where the seller exploits this fact to profitably discriminate between consumers. Importantly, these models derive the behavioral implications of certain cognitive imperfections, taking these imperfections as exogenous. In contrast, our approach illustrates how bounded rationality - in the sense of imperfect information processing - can emerge endogenously from features of the physical description of the game.

Second, there is a link to a philosophical school known as Pragmatism. According to William James (1908), one of the founders of this school, 'beliefs are true to the extent

that they breed true in their consequences'. One interpretation, as emphasized e.g., by Russell (1945), is that agents should (or do) hold beliefs that have good consequences. While 'good' is a vague concept in most disciplines, economics seems particularly well-fit to include a notion of Pragmatism, since we can understand 'good' just as 'high utility' or 'high fitness'.[2] Therefore, the present paper may be seen as a modest attempt to link Pragmatism and Economics.

Third, since pragmatic beliefs are determined by what is useful, or beneficial, for the agent (rather than what is 'true'), the pragmatic approach to belief formation relates to models of cognitive dissonance. While an universally accepted definition of 'cognitive dissonance' does not exist, the notion of cognitive dissonance discussed by economists concerns agents whose utility depends directly on their beliefs (see e.g., Akerlof and Dickens, 1982, Rabin, 1994). Therefore, the agent may form beliefs that are at striking odds with 'reality' because he feels better that way. For example, an agent that cannot accept his limitations and convinces himself that he is other than he is to comfort himself, suffers from cognitive dissonance. In contrast, the link between beliefs and utility we explore is indirect; beliefs of an agent affect his utility only through affecting the actions of other agents. Specifically, we assume that agents are indifferent per se to which beliefs they hold.

Fourth, many authors acknowledge that material payoffs in games can be higher for agents that are less than fully rational than for fully rational agents (e.g., Gilboa and Samet, 1989), or might be less than fully rational (Kreps et al., 1986).[3] More closely related to our approach, Frank (1988), building on Schelling (1978), argues that emotions

[2] A standard line of argument against the Pragmatist position is that 'good' can be understood to mean 'true predictions', so that the Pragmatist position is impeccable, but vacuous (see e.g., Scruton, 1994).

[3] Let us briefly discuss these two approaches. Gilboa and Samet argue that agents able to handicap themselves through committing to use a simple strategy may 'blackmail' more rational players, and obtain a higher payoff. This approach is different to ours since the most plausible interpretation of their self-handicapping is (arguably) that of physically 'burning bridges', which has little to do with cognition or belief formation. The reputation model of Kreps et al. (1986), on the other hand, does not provide a story for how the 'irrational' agents come into the picture in the first place. In contrast, we argue that a simple learning mechanism can produce pragmatic agents.

may serve as a commitment device to e.g., enforce agreements that otherwise would suffer from asset-specificity problems. However, the role of cognitive capacity as a commitment device, and its relation to overconfidence, is not considered by Frank; in particular Frank assumes that strategic motives merely affects emotions; the rationality assumptions he makes are standard.

Fifth, there is a growing related literature that attempts to endogenize preferences through the theory of natural selection (e.g., Rogers, 1994, Robson, 1996). However, in contrast to the present approach the rationality assumptions made in this literature are standard; evolution is only allowed to work on preferences. In conclusion, while there are several attempts in the literature to endogenize emotional capacity and preferences through incentives (fitness), the literature has not considered the possibility that beliefs, or cognition more generally, could be explained by incentives.[4]

2 Overconfidence

In this section, we define and discuss some basic terms. Suppose for simplicity that ability, a, is a one-dimensional magnitude, that in principle can take any value between 0 and 1.[5] Furthermore, suppose that the agent has a probability distribution over this interval, interpreted as the agent's subjective beliefs about his abilities. For simplicity, assume that this distribution can be represented by a (differentiable) cumulative distribution, $F(..)$, with $F(0) = 0$, $F(1) = 1$, and density $f(..)$. Hence, from the viewpoint of the agent, his ability is a random variable, denoted by x. We define the agent's confidence level, c, as

[4]The only exception we are aware of is Robson (1999). However, while our starting point is to explain why human cognition falls short of the rationality ideal, Robson has the opposite starting point, namely to explain why humans are endowed with such a highly developed intelligence (compared to other species). Specifically, Robson models the 'Red Queen effect', that human intelligence can be explained by intra-species competition (i.e., competition between human genes), rather than inter-species competition. Since Robson's type of evolution best can be understood as working in the extreme long run, while our 'evolution' to pragmatic beliefs primarily should be understood as a learning phenomenon, the two approaches talk about different things rather than being inconsistent.

[5]It is assumed that the ability a has no direct relation to how the agents form beliefs about a. Hvide (1999b) considers the case when there is a circularity problem.

the mean of the distribution,

$$c = E(x) = \int_0^1 f(x)dx \tag{1}$$

We say that an agent is overconfident if,

$$c > a \tag{2}$$

and that the agent is underconfident if,

$$c < a \tag{3}$$

Hence the agent is overconfident if his confidence level exceeds his true ability, and underconfident if his confidence level is lower than his true ability.[6]

Notice first that overconfidence at individual level is perfectly consistent with unbounded rationality; an overconfident person could be someone that coincidentally has received a string of positive signals about his abilities, and hence has obtained a too high opinion about herself at 'misfortune'.[7]

Second, in the tradition of Kahneman and Tversky, experimental psychologists have for long been interested in 'overconfidence'. Somewhat confusingly, this literature applies two distinct meanings of the term 'overconfidence', and more or less interchangeably. Let us elaborate on that point since it is important to the originality of the project of the present paper. First, overconfidence$_1$, defined above, relates to 'hubris', or agents that conceive themselves as more able than they really are. Second, overconfidence$_2$, in

[6]Two comments to the definition. First, one alternative would be to define overconfidence as a confidence level that exceeds the level obtained with a rational interpretation of available data. In that case, overconfidence would imply bounded rationality, in contrast to under the proposed definition. The alternative definition would work equally well to the proposed definition in the example in the next section. However, the proposed definition seems to be less ambiguous than the alternative definition since the meaning of 'available data' is not always clear. Second, if one were to introduce a measure of overconfidence, higher moments of the belief distribution of the agent could also matter. For example, an agent that is certain that his ability is b (where $b > a$) may be viewed as more overconfident than an agent that also has b as his mean ability level, but with a non-degenerate belief distribution.

[7]Therefore, one may sometimes wish to define overconfidence as something that can occur on population level rather than at individual level, where the idea is that coincidences wash out at aggregate level (see Hvide 1999c, Lemma 1, for a result along these lines).

contrast, relates to agents that are 'too certain' about some event, through having a 'too tight' distribution of posteriors. For example, an agent is overconfident$_2$ if he believes that the stock price of share y tomorrow will be between \$15 and \$17 with 90% probability, while in reality the 90% confidence interval should be set between \$14 and \$18.

Since several papers in the finance literature explore the notion of overconfidence$_2$ (e.g., De Long et al. 1990, Kyle and Wang, 1996, and Daniel et al., 1998), it is natural to ask whether there is any logical relation between overconfidence$_1$ and overconfidence$_2$. Is there any? We think clearly not. For example, a stock analyst that has a too high opinion about his predictive abilities may believe that the volatility of a stock is higher than it really is, and set his confidence interval broader than a more experienced analyst. Hence the analyst is overconfident$_1$ and underconfident$_2$. Of course, it is also possible that the overconfident$_1$ analyst sets the confidence interval narrower than it should be, in which case he would be both overconfident$_1$ and overconfident$_2$. The point is that there is no clear relation between overconfidence$_1$ and overconfidence$_2$; they reflect different underlying phenomena.[8] In the extension of this point, let us make two comments. First, the finance literature assumes a clear relation between overconfidence$_1$ and overconfidence$_2$. For example, Daniel et al. (1998) states that, 'Our theory assumes that investors view themselves as more able to value securities than they actually are [overconfidence$_1$], so that they underestimate their forecast error variance [overconfidence$_2$] (p. 6). This is not a valid inference. Furthermore, from this type of statement, one could easily be tempted to conclude - invalidly - that the project of the present paper to a large part already has been undertaken in the finance literature.[9] Second, notice that since overconfidence$_2$ may simply be due to a too tight prior distribution, and hence consistent with the Bayesian

[8] In the stock market example, overconfidence$_1$ relates to a skewed first moment of a (subjective) probability distribution over one's own abilities, while overconfidence$_2$ relates to a skewed second moment of a (subjective) probability distribution over one's own forecast for some external event. Instead of predicting some external event, we could imagine that the agent were making overconfident$_2$ predictions about his own abilities. Neither in that case would there be any logical relation between overconfidence$_1$ and overconfidence$_2$.

[9] See De Long et al. (1990) and Kyle and Wang (1996) for papers showing that overconfident$_2$ agents can obtain a higher payoff than agents who are not.

paradigm,[10] while overconfidence$_1$ is hard to explain without appealing to some form of bounded rationality.

3 Pragmatic Beliefs

We define the pragmatic beliefs of an agent in a given situation as those beliefs that maximizes utility, or material payoff for that agent. In the present section we construct a job market example that shows how the notion of pragmatic beliefs can be useful in understanding overconfidence. The example will surpress how pragmatic beliefs are actually formed, and how agents that form beliefs pragmatically can be distinguished from agents that form beliefs rationally. These and other interpretational issues will be discussed in the next section.

There is one risk-neutral worker and a continuum of risk-neutral firms. At time 1, the worker is matched with a firm. The worker's productivity true productivity in that firm, denoted by z, depends on the match. Ex-ante, the distribution of productivities is given by $G(z)$, with density $g(z)$, where $z \in [0, 1]$. Then the time 1 firm makes a take-it or leave-it wage offer w_1 to the worker. If the worker accepts the offer, the game ends. If the worker rejects the offer, he moves on to time 2. At time 2, the worker is matched with a new firm, and receives an offer w_2. We assume that wages at time 2 are set competitively, such that the wage offers at time 2 follows the distribution $G(..)$.[11] Then the game ends (there is no recall of the time 1 offer). The decision of the time 1 firm which wage to offer the agent, and the decision of the worker is whether to accept the offer or not. Hence the agent faces a two-stage search problem, where the offer at time 1 is endogenous. For simplicity we assume that there is no discounting or search costs; hence the utility U of a worker simply equals his wage. Nothing essential in the example hinges on these assumptions.

The informational assumption we make is that the time 1 firm knows the worker's productivity in that firm, and moreover it knows $G(w_2)$. The worker, on the other hand,

[10] See Morris (1995).

[11] For example, the agent is matched with two randomly drawn identical firms at time 2. As will become clear, nothing hinges on this assumption.

has merely formed beliefs about $G(w_2)$, where these beliefs are denoted by $H(w_2)$. $H(w_2)$ has the same support as $G(w_2)$. Hence $H(w_2)$ reflects both the agent's beliefs about his productivity in the time 1 firm, and his beliefs about the distribution of wage offers at time 2. Notice that since the randomness inherit in $G(w_2)$ can be attributed to randomness in the matching technology, we can interpret the mean of $G(w_2)$, labeled $\bar{w}_2$, as the worker's true (average) ability. Let $\dot{w}_2$ be the worker's expectation of w_2, given his subjective beliefs, i.e., his confidence level. Then true ability equals $\bar{w}_2 = \int_0^1 g(w_2)dw_2$ and the agent's confidence level equals $\dot{w}_2 = \int_0^1 h(w_2)dw_2$.

Remark 1 If the agent is overconfident then $\dot{w}_2 > \bar{w}_2$, and if the agent is underconfident then $\dot{w}_2 < \bar{w}_2$.

Hence if the worker is overconfident, his confidence level exceeds his true (average) ability, and if the worker is underconfident, his confidence level is lower than his true (average) ability. If $\dot{w}_2 = \bar{w}_2$, say that the agent has a **neutral confidence level.**

Start out with the case with neutral confidence, i.e., when $\dot{w}_2 = \bar{w}_2$. Then, clearly, the worker's optimal strategy is to accept any offer above $\bar{w}_2$ at time 1.[12] The optimal strategy for the time 1 firm is to set,

$$\begin{aligned} w_1 &= \bar{w}_2 \text{ if } z > \bar{w}_2 \\ w_1 &< \bar{w}_2 \text{ if } z < \bar{w}_2 \end{aligned} \tag{4}$$

If the worker's productivity is sufficiently high, the firm offers the agent a wage equal to the agent's reservation wage. If not, the firm offers the worker a low wage (which is equivalent to the firm making no offer at all). Hence when the worker and the firm share the true beliefs about the probability distribution over wage offers at time 2, the expected utility of the worker is $\bar{w}_2$. Clearly the time 1 firm runs an expected surplus, since it will only hire the worker if the worker's productivity exceeds his wage.

Now consider the case when the worker's confidence level may differ from his true ability, i.e., when $\dot{w}_2 \neq \bar{w}_2$. How does the firm's optimal strategy change compared to the

[12] The agent may update his beliefs over time 2 offers when receiving the time 1 offer, but for simplicity we ignore this possibility.

benchmark (4)? Since the worker will not accept any offer below his confidence level $\dot{w}_2$, we get that the optimal strategy for the firm is in this case to set,

$$\begin{aligned} w_1 &= \dot{w}_2 \text{ if } z \geq \dot{w}_2 \\ w_1 &< \dot{w}_2 \text{ if } z < \dot{w}_2 \end{aligned} \tag{5}$$

We ask whether an agent with pragmatic beliefs will be under- or overconfident about his abilities, and hence about what he can get at time 2. First a formal definition of pragmatic beliefs.

Definition 3.1 *The pragmatic belief, $\dot{w}_2{}^*$, is the $\dot{w}_2$ that maximizes the agent's utility.*

Notice that from (5) we know that every $\dot{w}_2$ induces a unique equilibrium offer from the time 1 firm (and a unique acceptance strategy for the worker). Hence we can write the worker's utility on reduced form as a function of $\dot{w}_2$ and of $\bar{w}_2$ only, i.e., $U = U(\dot{w}_2, \bar{w}_2)$. Since $\bar{w}_2$ is a constant, $\dot{w}_2{}^*$ is well-defined.[13] We get the following result.

Proposition 1 *The pragmatic beliefs are overconfident.*

Proof. Denote the probability of the worker's productivity exceeding z for $P(z)$. We have that,

$$\dot{w}_2{}^* := \arg\max_{\dot{w}_2} U(\dot{w}_2, \bar{w}_2) = \arg\max_{\dot{w}_2} \{P(\dot{w}_2)\dot{w}_2 + (1 - P(\dot{w}_2))\bar{w}_2) \tag{6}$$

Hence the pragmatic belief finds the optimal trade-off between getting a high wage offer at time 1, and the probability of getting an offer at time 1 at all. Given the firm's optimal strategy from (2) we have that,

$$P(\dot{w}_2) = \int_{\dot{w}_2}^{1} g(z)dz \tag{7}$$

[13] Two comments to the definition of pragmatic beliefs. First, alternatively we could define pragmatic beliefs through $H(..)$ rather than through $\dot{w}_2$. In that case, all we would require is that the mean of $H(..)$ equals $\dot{w}_2$. Second, in games with multiple equilibria for at least some $\dot{w}_2$, we need to select between these equilibria in order to determine a unique pragmatic belief.

Since $\frac{\partial P(\dot{w}_2)}{\partial \dot{w}_2} = -g(\dot{w}_2)$, the first order condition becomes,

$$g(\dot{w}_2)[\bar{w}_2 - \dot{w}_2] + P(\dot{w}_2) = 0 \tag{8}$$

Since $P(..) > 0$, we have that $\dot{w}_2^* > \bar{w}_2$, and hence the pragmatic belief is overconfident. ■

Proposition 1 shows that a high belief about one's market value can to some extent be self-fulfilling, in that high expectations about future job offers forces the time 1 firm to give a higher offer to attract the worker than when the agent has lower expectations about his future.

Notice three features of the solution. First, the overconfident solution induces an inefficient allocation of workers: Since it is socially optimal that the worker starts working in any time 1 firm where his marginal productivity is at least $\bar{w}_2$, it is socially optimal that the worker has a neutral confidence level. Second, since $\dot{w}_2^*$ is below 1, there is a bound to the optimal degree of overconfidence. A worker that excessively overconfident will practically always reject the offer at time 1, and hence in expected terms get a payoff close to $\bar{w}_2$. Still, a small degree of overconfidence is better than no overconfidence at all (or even worse, underconfidence, in which case he would accept offers below $\bar{w}_2$), since his expected payoff is larger than $\bar{w}_2$ even for $\dot{w}_2$ very close to 1. We illustrate this point in a figure, calculated for $g(z) = 1$ (which gives $\bar{w}_2 = \frac{1}{2}$).

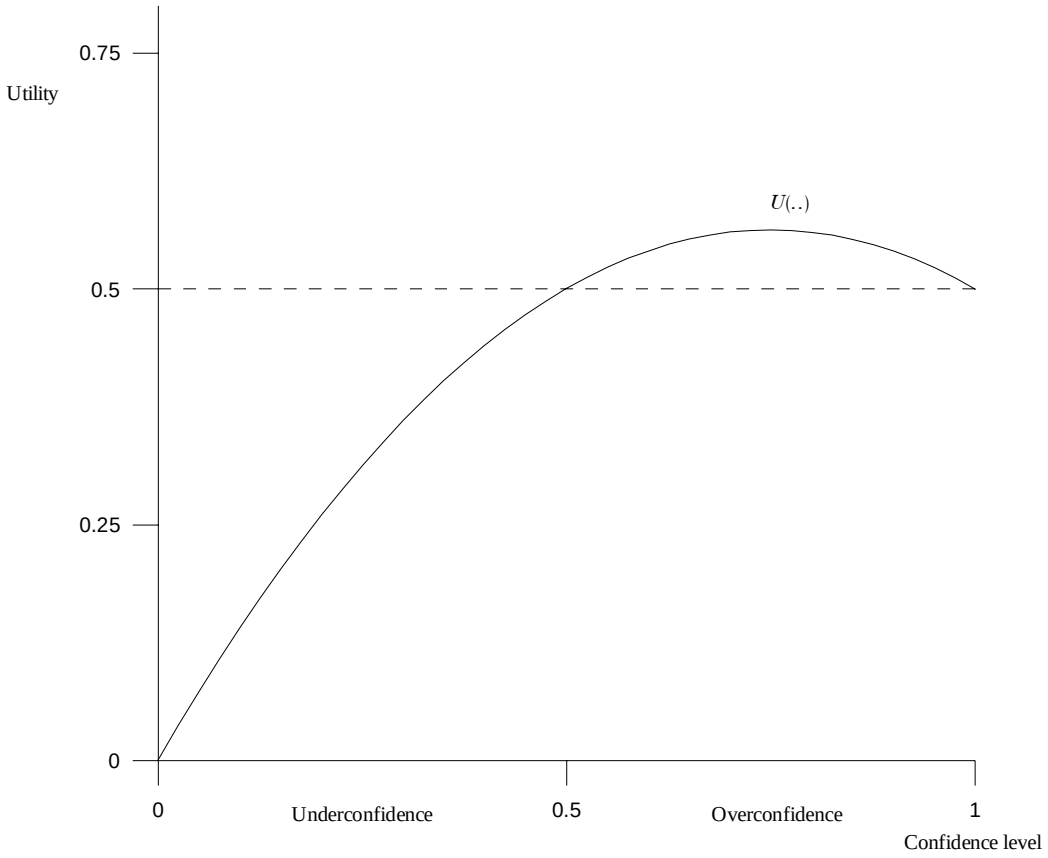

Figure 1

The figure depicts the expected utility of the worker as a function of his degree of confidence.[14] While underconfidence is severely punished in that the worker accepts a too low offer at time 1, a moderate degree of overconfidence ($\dot{w}_2 = \frac{3}{4}$) gives the agent the highest expected utility. When the agent becomes excessively overconfident ($\dot{w}_2 \to 1$), the expected utility converges from above to the expected utility under neutral confidence ($\dot{w}_2 = \frac{1}{2}$).[15]

The third thing to notice is that there is an alternative way to derive $\dot{w}_2^*$. Suppose that the worker (instead of the firm) makes the take-it or leave-it offer at time 1, and denote this offer by $\tilde{w}$. Conditional on the offer, the optimal strategy for the firm is then

[14] Surpressing $\bar{w}_2$, the reduced form utility is given by $U(\dot{w}) = -\dot{w}^2 + \frac{3}{2}\dot{w}$.

[15] If there are several stages in the search process, instead of only two, excessive overconfidence ($\dot{w}$ close to 1) would lead to excessively long search, but it would still hold that excessive underconfidence ($\dot{w}$ close to 0) is worse than excessive overconfidence.

to,

$$\begin{aligned} \text{accept the offer if } z &\geq \tilde{w} \\ \text{reject the offer if } z &< \tilde{w} \end{aligned} \tag{9}$$

In setting the optimal take-it or leave-it offer, the agent then trades-off the probability of getting the offer accepted with the size of the offer. Denote the probability of having the offer accepted by $Q(..)$. Then the optimal take-it or leave-it offer of the worker is the solution to,

$$\tilde{w}^* := \arg\max_{\tilde{w}} \{Q(\tilde{w})\tilde{w} + (1 - Q(\tilde{w}))\bar{w}_2) \tag{10}$$

The first order condition becomes,

$$\frac{\partial Q(\tilde{w})}{\partial \tilde{w}}[\bar{w}_2 - \tilde{w}] + Q(\tilde{w}) = 0 \tag{11}$$

but given the firm's optimal strategy, $Q(\tilde{w}) = \int_{\tilde{w}}^{1} g(z)dz = P(\tilde{w})$, and hence $\tilde{w}^* = w_2^*$. Consequently, the solution to the optimal belief problem is equivalent to the solution to the optimal take-it or leave-it offer, in terms of utility to the worker and profits to the firm. Hence we can interpret the possibility of overconfidence as a way of taking the first mover advantage from the firm to the worker.

Given the equivalence, a natural question is then whether the whole advantage of overconfidence boils down to a question about first mover advantage. To illustrate that more than that is at issue, in the Appendix we consider a twist to the story, where both parties at time 1 can form pragmatic beliefs. The conclusion here is that the overconfidence result survives also when both parties form beliefs pragmatically.

4 Discussion

In the example constructed in Section 2, overconfident beliefs (that are known to the firm) forces the firm to offer a higher wage than it would to an otherwise identical agent, except with a more neutral self-confidence. Hence overconfidence can be beneficial for an agent. It follows that agents with pragmatic beliefs obtain a higher payoff than those without

this faculty. In the present section we discuss whether we can expect pragmatic beliefs to be formed, and how beliefs possibly can be observable to firms.

There are two dynamic ways to justify a notion of pragmatic beliefs, either through long term evolution of cognition, in the spirit of modern evolutionary theory, or through learning. First, is perfectly possible to set up an evolutionary model where there is a selection pressure towards agents that are born pragmatists, in the sense that they are borne with a too high opinion about themselves, and are hardwired to retain this property of their belief structure as they grow up. However, since evolution by nature is backward-looking, it is not clear how adaption to social condition in the hunterer-gatherer society can make humans optimally equipped in a modern labor market. Therefore, although the evolutionary defense of pragmatic beliefs could be worth pursuing further, let us turn to the more easily argued for learning defense.

The learning defense of pragmatic beliefs takes as starting point that humans rather than being genetically equipped with overconfidence, gradually converges to pragmatic belief formation through a process of metalearning. The basic force behind this convergence is the law of reinforcement: agents converge to pragmatic information processing, i.e., any updating rule leading to pragmatic beliefs, because pragmatic information processing pays better off than other ways of processing information.

To make such metalearning plausible, we need the agent to repeatedly face the same type of problem, and for him to have a reasonable possibility to experiment with different ways of processing information.[16] The job market example is useful in pointing to such a context; every time a new contract is offered (or potentially is offered), the worker must weigh up the offer against how much he believes he is worth in the outside market. Moreover, the job market environment is sufficiently diffuse to make straightforward Bayesian learning improbable.

The crucial hypothesis in the argument chain leading to pragmatic beliefs is that in

[16]Notice that Bayesian learning presumes that the learning environment is specified to an extent that makes other types of learning appear illogical, or even inconsistent. In reality, information processing depends on subjective interpretation (e.g., how informative is the signal s about a?- a question that the agent is precluded from asking in the Bayesian model of belief updating), and where different ways of interpreting data can be subject to experimentation.

assessing the relative merit of different ways of processing information, the agent takes into account what pays rather than what is true.[17] For example, instead of minutiously balancing his job failures against the successes when assessing his ability, the agent (unconsciously) forgets the failures and remembers the successes. Or, more generally, agents learn that placing more emphasis on certain type of information about their abilities ('good news') than on a different type of information ('bad news') make them better off at the end of the day. The agent does not have to see the causal link himself, and - unless it should be profitable - the agents pay little attention to whether their confidence level is realistic or not. To conclude, we envisage a learning process underlying pragmatic beliefs similar in spirit to standard reinforcement learning; just as such learning implies that agents tend to drift towards better actions, pragmatic agents drift towards ways of processing information that is better for themselves, and hence towards pragmatic beliefs.[18]

Independently of how the pragmatic beliefs are reached, one may argue that beliefs cannot have a commitment effect on the behavior of other agents, or any other effect, because an agent's beliefs are not observable to other agents. In particular, how can overconfidence be distinguished from strategic bluffing - when a person knows his limitations well but have a talent for presenting themselves in a more favorable way than his abilities would warrant? We have two replies to this. First, beliefs, particularly self-beliefs, may indeed be observable to some extent. As described in detail by Frank (1988), there are many cues to detecting lies. One is to focus on the expressions that are least subject to conscious control, for example blushing or an elevated pitch of the voice. Of course, no cues are perfect, but the existence of them can induce a substantial cost of lying.[19] Notice

[17]The agent may define truth as what pays, in which case there would be no difference between a true position and a position that pays. Notice, however, that it is not clear why pragmatic agents should be interested in the concept of 'truth' at all. Unless it pays.

[18]This could imply that job market confidence level could be different from, say, confidence level at home. This seemingly implies that the belief system of the agent is inconsistent. However, although inconsistent through time, the belief system could be consistent at every point in time (which is probably what should be required), with the agent 'forgetting' his previous beliefs when entering a new situation, and then remember them again when entering the old situation again (for a paper that discusses logical properties of dynamic beliefs systems with forgetting, see Hvide, 1999a).

[19]As long as such cues are informative at all, with risk neutral agents it is simple to construct suitable

also that the cost of lying will be aggravated by a different incentive principle; since firms will make a higher profit on workers that have a correct level of confidence than on overconfident workers, firms will have an incentive to attempt to distinguish overconfidence from strategic bluffing.

Second, a more traditional explanation for how firms can infer the true beliefs of the workers is that workers undertake certain actions to credibly signal their confidence level. For example, schooling may have the property that an agent's expected cost of schooling is negatively correlated with his confidence level. In that case, there can exist a separating equilibrium where those with high confidence level undertake schooling, and those with low confidence level skip school.[20] In such a separating equilibrium, those who are truly overconfident can be distinguished from those that merely pretend to be overconfident.[21] To conclude, it is not necessarily true that true and fake overconfidence cannot be distinguished.

5 Conclusion

Which factors shape the manner in which agents create and update their beliefs about economic variables? The traditional answer to this question, imported from statistical theory, is that economic agents pursue some notion of 'truth' in their belief formation. This idea is clearly formulated in e.g., textbook game-theoretic approaches. The basic idea in the present paper is that the main force behind belief formation is incentives, and not any abstract notion of 'truth'. From this idea, we proposed the notion of pragmatic beliefs, where such beliefs are simply those that maximize the payoff in a given situation.

forcing schemes, such that only those who are truly overconfident will claim to have a (too) high confidence level.

[20] This is true e.g., in the signaling model of Weiss (1983).

[21] In such a separating equilibrium, overconfident agents may or may not receive a higher net material payoff than agents with neutral confidence (when the true cost of schooling is subtracted). Under risk neutrality, overconfident agents will be worse off than agents with neutral confidence. The intuition is that agents with neutral confidence can always imitate the overconfident agents, if the latter should get a higher payoff. Under risk aversion, however, this result does not necessarily hold true; under risk aversion overconfident agents may get a higher average material payoff than agents with neutral confidence, and hence be more fit. For a related point, see Waldman (1994).

To indicate that the notion of pragmatic beliefs can be useful, we considered the overconfidence puzzle established in experimental studies. This puzzle, that real life agents seems to be overconfident about their own abilities, can hardly be explained by any traditional, rational, notion of belief formation. Equipped with the notion of pragmatic beliefs, however, we showed that overconfidence can emerge as an equilibrium phenomenon, in that the pragmatic beliefs in a simple job market example were shown to be overconfident.

The example showed that overconfidence can be beneficial to an agent, due to strategic effects. However, it was silent on how agents emerge at overconfidence, and how firms can distinguish true from fake overconfidence. In Section 4, these questions were discussed. To model in a more general setting how pragmatic beliefs are formed and observed suggests an extension of the present work. Moreover, it could be interesting to see whether laboratory experiments could verify the hypothesis that agents in repeated settings tend to converge to pragmatic beliefs rather than to true beliefs about e.g., their own abilities.

6 References

Akerlof, G. and W. T. Dickens (1982). The Economic Consequences of Cognitive Dissonance. American Economic Review, 72, 307-19.

Asubel, L. M. (1991). The Failure of Competition in the Credit Card Market. American Economic Review, 81, 50-81.

Blume, L. and D. Easley (1992). Evolution and Market Behavior. Journal of Economic Theory, 58, 9-40.

Camerer, C. and D. Lovallo (1999). Overconfidence and Excess Entry: An Experimental Approach. American Economic Review, 89, 306-318.

Daniel, K., D. Hirshleifer and A. Subrahmanyan (1998). Investor Psychology and Security Market Under- and Overconfidence. Journal of Finance, 53, 1839-85.

De Bondt, W.F.M. and R. H. Thaler (1995). Financial Decision Making in Markets and Firms: A Behavioral Perspective. In: R. Jarrow, V. Maksimovic and W. T. Ziemba (eds): Finance Handbooks in Operations Research and Management Science 9, 385-410. North-Holland, Amsterdam.

Dekel, E. and S. Scotchmer (1999). On the Evolution of Attitudes toward Risk in

Winner-Take-All Games. Journal of Economic Theory, 87, 125-43.

De Long, B. J., A. Shleifer, L. H. Summers and R. Waldmann (1990). Noise Trader Risk in Financial Markets. Journal of Political Economy, 98, 703-38.

Elster, J. (1983). Sour Grapes. Studies in the Subversion of Rationality. Cambridge University Press.

Fershtman, C. and K. Judd (1987). Equilibrium Incentives in Oligopoly. American Economic Review, 77, 927-40.

Fershtman, C. and E. Kalai (1993). Complexity Considerations and Market Behavior. Rand Journal of Economics, 24, 224-35.

Frank, R. (1988). Passions within Reason. Norton Publishers, New York.

Gilboa, I. and D. Samet (1989). Bounded versus Unbounded Rationality: The Tyranny of the Weak. Games and Economic Behavior, 1, 213-21.

Golec, J. and M. Tamarkin (1995). Do Bettors Prefer Long Shots Because They are Risk-Lovers or are They just Overconfident? Journal of Risk and Uncertainty, 11, 51-64.

Hvide, H. K. (1999a). Bounds to Memory Loss. Theory and Decision, 46, 1-21.

Hvide, H. K. (1999b). Self-Awareness, Uncertainty, and Markets With Overconfidence. In: Machina, M. and B. Munier, 'Beliefs, Interactions and Preferences in Decision Making', 159-176. Kluwer Academic Publishers.

Hvide, H. K. (1999c). On the Informational Role of Education in the Allocation of Talent. Working Paper, Berglas School of Economics, Tel-Aviv University. Available at http:\\www.ssrn.com.

James, W. (1907/1975). Pragmatism: A New Name for Some Old Ways of Thinking. Harvard University Press.

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

Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1986). Reputation and Imperfect Information. Journal of Economic Theory, 27, 253-79.

Kyle, A. S. and F. A. Wang (1997). Speculation Duopoly with Agreement to Disagree: Can Overconfidence Survive the Market Test? The Journal of Finance, 52, 2073-90.

Lipman, B. (1995). Information Processing and Bounded Rationality: A Survey. Canadian Journal of Economics, 28, 42-67.

Morris, S. (1995). The Common Prior Assumption in Economic Theory. Economics & Philosophy, 11, 227-53.

Rabin, M. (1994). Cognitive Dissonance and Social Change. Journal of Economic Behavior and Organization, 23, 177-94.

Robson, A. J. (1996). A Biological Basis for Expected and Non-expected Utility. Journal of Economic Theory, 68, 397-424.

Robson, A. J. (1999). The Evolution of Intelligence and the Red Queen. Draft, University of Western Ontario. Available at http://publish.uwo.ca/~arobson/homepage.html.

Rogers, A. R. (1994). Evolution of Time Preference by Natural Selection. American Economic Review, 84, 460-81.

Russel, B. (1945). A History of Western Philosophy. Simon and Schuster, New York.

Rubinstein, A. (1995). On Price Recognition and Computational Complexity in a Monopolistic Model. Journal of Political Economy, 101, 473-84.

Rubinstein, A. (1997). Modeling Bounded Rationality. MIT Press.

Scruton, R. (1994). Modern Philosophy. An Introduction and Survey. Penguin Books.

Schelling, T. (1978). Micromotives and Macrobehavior. Norton, New York.

Svensson, O. (1981). Are We All Less Risky and More Skillful Than Our Fellow Drivers? Acta Psychologica, 47, 143-48.

Waldman, M. (1994). Systematic Errors and the Theory of Natural Selection. American Economic Review, 84, 482-497.

Weiss, A. (1983). A Sorting-cum-Learning Model of Education. Journal of Political Economy, 91, 420-442.

7 Appendix

Here, we consider the case when the time 1 firm also can construct a pragmatic belief. The purpose with this section is to show that the overconfidence result of Proposition 1 survives letting the firm as well as the agent form pragmatic beliefs.

We consider a situation where the worker forms a pragmatic belief about this distribution of wage offers at time 2, while the firm constructs a pragmatic belief about the agent's productivity in its own firm. We retain the assumption that both the firm and

the worker are risk-neutral. For simplicity it is assumed that also the agent knows his true productivity at time 1.[22] The timing of the game is now that the beliefs are first set simultaneously, then the model proceeds as before.

Notice that when both sides construct pragmatic beliefs, the pragmatic belief of the agent depends on the pragmatic belief formed by the firm, and vice versa. Denote the firm's belief about the worker by $\mathcal{MP}$, and the agent's belief by $\dot{w}_2$, as before. Say that the firm is overconfident (underconfident) about the ability of the worker if $\mathcal{MP} > z$ ($< z$).[23] Let us start out with the optimal bid by the firm. For a given $\mathcal{MP}$ and $\dot{w}_2$ we get the firm's optimal offer is,

$$\begin{aligned} w_1 &= \dot{w}_2 \text{ if } \mathcal{MP} > \dot{w}_2 \\ w_1 &< \dot{w}_2 \text{ if } \mathcal{MP} < \dot{w}_2 \end{aligned} \tag{12}$$

As before, the optimal strategy of the worker is to accept the offer if $w_1 = \dot{w}_2$ and to reject the offer if $w_1 < \dot{w}_2$. Given the equilibrium of the offer subgame, we solve for the equilibrium in pragmatic beliefs. We then have,

Proposition 2 If $z < \bar{w}_2$, both the firm and the worker have neutral confidence in equilibrium. However, if $z > \bar{w}_2$, the worker is overconfident in all Nash equilibria, the firm is underconfident, and the worker's utility is higher than under neutral confidence.

Proof. For a given $(\bar{w}_2, z)$ there are many Nash equilibria. We proceed to show that all these Nash equilibria involves the firm being underconfident about the ability of the worker, and the worker being overconfident. First we derive the agent's best response in terms of $\dot{w}_2$, as a function of $\mathcal{MP}$. First suppose that $\mathcal{MP} < \bar{w}_2$. In that case any

[22] This assumption makes things simpler, but qualitatively the same conclusion holds for the game where the agent forms his pragmatic belief based on a probability distribution of productivities at time 1.

[23] Notice that we allow for the firm to have a different belief about the agent's productivity than z, even if it is informed of z. This is a clearly somewhat contrived assumption. To make the example more realistic, we could model the firm as having slightly noisy information about z, and form (pragmatic) beliefs based on this information. However, since this more realistic version of the example is also more notationally demanding, without really adding anything, we keep the assumption that the firm may believe something different from what it is informed of.

$\dot{w}_2 > \mathcal{MP}$ is strictly better than $\dot{w}_2 < \mathcal{MP}$, since $\dot{w}_2 < \mathcal{MP}$ implies that the firm will set $w_1 = \dot{w}_2$, and the worker will accept the offer in spite of $\bar{w}_2 > \dot{w}_2$. Hence any $\dot{w}_2 > \mathcal{MP}$ is a best reply when $\mathcal{MP} < \bar{w}_2$. Now suppose that $\mathcal{MP} > \bar{w}_2$. In that case, the unique best reply will be to set $\dot{w}_2 = \mathcal{MP}$, since that will ensure the highest possible offer at time 1, and furthermore this offer will exceed $\bar{w}_2$. Now to the best response of the firm in terms of $\mathcal{MP}$, as a function of $\dot{w}_2$. For any $\dot{w}_2 > z$, it is clearly suboptimal for the firm to set $\mathcal{MP} > \dot{w}_2$, since the firm will then run a deficit when in the next turn offering $w_1 = \dot{w}_2$. It follows that any $\mathcal{MP} < \dot{w}_2$ is a best response of the firm in the case where $\dot{w}_2 > z$. Now to the case $\dot{w}_2 < z$. In that case it is a unique best response for the firm to set $\mathcal{MP} = \dot{w}_2$. It follows that for $z > \bar{w}_2$, the set of Nash equilibria are the points on the 45 degree line between ($\bar{w}_2$,$\bar{w}_2$) and (z,z), which are all characterized by the agent being overconfident about the distribution of wage offers at time 2, and the firm being underconfident about the worker's productivity. On the other hand, for $z < \bar{w}_2$, the unique Nash equilibrium is the point $(\bar{w}_2, z)$. Hence with probability 1/2 the equilibrium is one where the agent is overconfident and gets a higher utility than the case of correct confidence, and with probability 1/2 the equilibrium is one with correct degree of confidence. ■