Overconfidence and diversification

Heller, Yuval

Nuffield College Department of Economics, University of Oxford

18 February 2013

Overconfidence and Diversification

Yuval Heller

Department of Economics and Nuffield College, University of Oxford.
New Road, Oxford OX1 1NF, UK. Phone: +44-1865-278993, +44-7582-540427.

Abstract

Experimental evidence suggests that people tend to be overconfident in the sense that they overestimate the accuracy of their private information. In this paper we show that risk-averse principals might prefer overconfident agents in various strategic interactions because these agents help diversify the aggregate risk. This may help understanding why successful analysts and entrepreneurs tend to be overconfident. In addition, a different interpretation of the model presents a novel evolutionary foundation for overconfidence, and explains various stylized facts about this bias.

Key words: overconfidence, diversification, evolution. JEL Classification: C73, D82.

1 Introduction

In many experimental studies participants are asked to answer trivia questions, and to report the level of confidence (subjective probability) that they answered each of these questions correctly. The typical result in such experiments is that people are overconfident: their confidence systematically exceeds the true accuracy (as discussed in Section 5.2). That is, people overestimate the accuracy of their private information, personal judgment and intuition. Various evidence suggest that overconfidence substantially influences economic behavior of analysts (Friesen and Weller, 2006), investors (Berber and Odean, 2001), entrepreneurs (Cooper, Woo and Dunkelberg, 1988), managers (Rabin and Schrag, 1999; Gervais, Heaton and Odean, 2010), and consumers (Grubb, 2009).

Email address: `yuval.heller@economics.ox.ac.uk or yuval26@gmail.com` (Yuval Heller).

[1] This work is in partial fulfillment of the requirements for Ph.D. degree at Tel-Aviv University, and it was supported by the Israel Science Foundation (grant number 212/09). I would like to thank Eilon Solan for his careful supervision, and for the continuous help he has offered. I would also like to express my deep gratitude to Eddie Dekel, Tzachi Gilboa, Moti Michaeli, Zvika Neeman, Ariel Rubinstein, Assaf Romm, Ran Spiegler, Roee Teper, an anonymous referee, and the seminar participants at Tel-Aviv University, Hebrew University of Jerusalem, Israel Institute of Technology, Haifa University, Ben-Gurion University, Ecole Polytechnique, University of Southampton, Universitat Pompeu Fabra, and University of Oxford, for many useful comments, discussions and ideas.

In this paper we study the strategic interaction between a risk-averse principal and privately-informed agents, and show why the principal may prefer overconfident agents. For example, consider the following interaction between a CEO and analysts.

Example 1 *A risk-averse CEO of a venture capital fund is in charge of several analysts. Each analyst manages the investments of the fund in his area: he investigates several startup companies, and selects one of them. The fund invests money in the startup companies that were selected by the different analysts. Each analyst may either make: (1) a conventional choice - follow accepted guidelines and choose what a typical analyst would in such a situation, or (2) an unconventional choice based on his personal judgment and intuition. Each investment may either succeed or fail. The outcomes of conventional choices are positively correlated, as they all depend on the effectiveness of the accepted guidelines. The outcomes of unconventional choices are independent, as they are induced by diverse intuitions. Each analyst wishes to maximize the success probability of his investment. The CEO wishes to maximize the total number of successful investments, and each additional success has a smaller marginal payoff. Finally, the CEO can boost (or decrease) agents' confidence in making unconventional choices.*

Example 1 is formalized as follows. There is a strategic interaction between a risk-averse principal and many agents that includes two stages. At stage 1 the principal uses his managerial skills to influence agents' confidence in relying on their personal judgment. [2] For tractability, we assume that the principal may induce any bias function g, which determines how agents evaluate the accuracy of personal judgment: if the real accuracy is p, each agent believes it to be $g(p)$. Similar qualitative results would hold if the model is amended by adding a moderate cost to the principal's influence or by somewhat limiting its accuracy.

At stage 2, all agents receive signal q - the success probability of a conventional choice. In addition, each agent i privately receives signal p_i (evaluated as $g(p_i)$) - the success probability of an unconventional choice (the p_i-s are independent and identically distributed). Then each agent chooses whether to make a conventional choice or an unconventional one. Each agent who makes a conventional choice succeeds with probability q; outcomes of different agents who make conventional choices are positively correlated as they depend on a common unknown factor (the effectiveness of the accepted guidelines in Example 1). Each agent i who makes an unconventional choice succeeds with probability p_i independent of other agents.

Conventional choices bear a larger aggregate risk due to the positive correlation between the outcomes of different agents. This creates a conflict of interests between calibrated agents ($g(p_i) = p_i$) who maximize their probability of success, and the risk-averse principal. The principal has a tradeoff between two objectives: (1) maximizing the expected number of successes, and (2) reducing the variance in the number of successes. The first goal is fully consistent with the interests of calibrated agents. However, due to the second goal, the principal would like agents with p_i a little bit smaller than q to follow their less accurate personal judgment, in order to reduce the variance and achieve a better diversification of

[2] The principal can also influence agents' confidence by screening candidates: obtaining a signal about the confidence of each candidate, and adjusting the hiring priorities accordingly.

risk.

Our first result shows that if the number of agents is sufficiently large, then this conflict is optimally resolved by having overconfident agents. That is, there is a continuous and increasing bias function g^*, which always overestimates the perceived accuracy of agent's personal judgment ($g^*(p) > p$ for every $0 < p < 1$), such that if all agents have this bias function, it approximately induces the "first-best" outcome for the principal - the outcome he would achieve if he could receive all the private signals and directly control the actions of all agents. We further show that g^* is unique: all other bias profiles, including heterogeneous profiles in which agents have different bias functions, induce strictly worse outcomes. Our third result shows that the principal prefers more overconfident agents if: (1) he becomes more risk-averse, or (2) the correlation coefficient between the outcomes of conventional choices becomes larger. The intuition of both results is that both changes deepen the conflict of interests between the principal and calibrated agents, and more overconfidence is required to compensate for it.

The conflict of interests between the principal and the agents can also be resolved by having calibrated agents and affecting their payoffs with monetary incentives. For example, paying each agent: (1) some percentage of the firm's overall profit, or (2) a bonus when he makes an unconventional choice. Our model offers the same optimal level of conflict resolution using the non-standard mechanism of boosting agents' confidence. There are situations of economic interest in which the principal is limited in his ability to use monetary incentives, and our mechanism may be easier to implement. In particular, risk-averse CEOs are typically supervised by the firm's shareholders, which tend to risk-neutrality due to diversified portfolios. Such shareholders will not approve monetary mechanisms that reduce the aggregate risk by lowering the expected profit. Thus, in such cases, the CEO must rely only on informal mechanisms, such as boosting the agents' confidence.[3]

Another example for such situations is the case in which the principal is a risk-averse *angel investor* and the agents are entrepreneurs who work in related areas. The investor would like to encourage the entrepreneurs to make unconventional choices in order to reduce the aggregate risk. Using monetary incentives for this may be too expensive: if each entrepreneur holds a large share of his company, then large bonuses are required to encourage unconventional choices with smaller success probabilities. Contrary to this, it might be relatively easy to boost the confidence of inexperienced entrepreneurs. This presents a novel explanation for the high levels of overconfidence among entrepreneurs (see Cooper, Woo, and Dunkelberg, 1988; and Busenitz and Barney, 1997).[4]

Next we note that a different interpretation of the model presents a novel evolutionary foundation for overconfidence (as detailed in section 4). Consider a large population with several types, where each type induces a bias function for its members. In each generation, each

[3] In addition, monetary incentives are also constrained by the limited liability of the agents.

[4] This explanation has a unique falsifiable prediction: entrepreneurs in areas in which typical investors are individuals and area-specific funds would be more overconfident, than entrepreneurs in areas in which the typical investors are large multi-area funds or a government.

individual has to take either a conventional action or an unconventional one, and the outcome influences his fitness. The size of the each type in the next generation is determined by its average fitness (*replicator dynamics)*. A simple adaptation of an existing result (Robson, 1996) shows that in the long run a unique type prevails over the entire population: the type that maximizes the logarithm of the average fitness. That is, all individuals present the optimal confidence bias g^* of the principal-agent model, where the principal has a logarithmic utility.

Our last result is motivated by the evolutionary interpretation. It assumes that the principal's utility has constant relative risk aversion, and it shows that the optimal level of overconfidence is higher if: (1) there is a larger difference between the payoffs of success and failure; this result is in accordance with the experimental finding of Sieber (1974) which suggests that people are more overconfident with respect to more important decisions; and (2) the success probabilities of the unconventional choices tend to be lower (first-order stochastic dominance); this result is in accordance with the experimentally observed *hard-easy effect* (Lichtenstein, Fischhoff, and Phillips, 1982): the more difficult the task, the greater the observed overconfidence. In addition, we characterize conditions that induce the *false certainty effect* (Fischhoff, Slovic, and Lichtenstein, 1977): people are often wrong when they are certain about their private information.

This paper is structured as follows. Section 2 presents the model and Section 3 presents the results. The evolutionary interpretation is described in Section 4. In Section 5 we briefly describe several variants and extensions (which are detailed in the online appendix) and discuss the related literature. Formal proofs appear in the appendix.

2 Model

Our model includes seven parameters: $(I, \rho, f_q, f_p, L, H, h)$. $I = \{1, ..., n\}$ is a set of agents, and we assume that $n >> 1$ (see online appendix for discussing the cases in which the number of agents is small or endogenous). A typical agent is denoted by i or j. Each agent faces a choice between two actions: a_c and a_u. The former (resp., latter) is interpreted as making a conventional (resp., unconventional) choice by following accepted guidelines (resp., personal judgment). Parameter $0 < \rho < 1$ is the correlation coefficient between the outcomes of two agents who make conventional choices. If at least one of the agents makes an unconventional choice, then their successes are independent.

Parameters $L, H \in \mathbb{R}$ $(H > L)$ are the payoffs an agent may obtain: success yields H and failure yields L. In the principal-agent interactions (like Example 1) one can assume that $H = 1$ and $L = 0$. In the evolutionary interpretation (see Section 4) H (L) is the fitness of a successful (unsuccessful) agent. Function $h : [L, H] \to \mathbb{R}$ is the utility of the risk-averse principal. It is strictly concave and increasing and it satisfies decreasing absolute risk aversion (DARA). That is: (1) $h' > 0$, (2) $h'' < 0$, and (3) Arrow-Pratt coefficient $r_A(x) = -\frac{h''(x)}{h'(x)}$ is decreasing in x. (DARA assumption is required for part 2 of Theorem 1 and for Theorem 3.)

The unknown state of nature determines the value of the tuple of random variables

$$\left(\mathbf{q},(\mathbf{p}_i)_{i\in I},\xi_q,(\xi_{i,p},\xi_{i,q})_{i\in I}\right)\in\left([0,1]\times[0,1]^I\times\{0,1\}\times(\{0,1\},\{0,1\})^I\right).$$

The random variable $\mathbf{q}\sim f_q$ is the (unconditional) success probability of each agent who makes a conventional choice. For each $i\in I$, the variable $\mathbf{p}_i\sim f_p$ is the success probability of agent i if he makes an unconventional choice. The variables $\left(\mathbf{q},(\mathbf{p}_i)_{i\in N}\right)$ are independent. Distributions f_q and f_p are continuous probability density functions with full support: $\forall 0<p,q<1,\ f_p(p)>0,\ f_q(q)>0$. The full support assumption is not essential and it is used to simplify the presentation of the results.

Variable ξ_q is equal to 1 with probability $\mathbf{q}$ (and 0 otherwise). When $\xi_q=1$ (resp., $\xi_q=0$) the accepted guidelines are effective (resp., non-effective) and conventional choices have higher (resp., lower) success probability: each $\xi_{i,q}$ is equal to 1 with probability: $\sqrt{\rho}+\left(1-\sqrt{\rho}\right)\cdot\mathbf{q}$ (resp., $\left(1-\sqrt{\rho}\right)\cdot\mathbf{q}$). A conventional choice yields agent i high (low) payoff when $\xi_{i,q}=1$ ($\xi_{i,q}=0$). For each $i\in I$, $\xi_{i,p}$ is equal to 1 with probability $\mathbf{p}_i$ (and 0 otherwise). When $\xi_{i,p}=1$ (resp., $\xi_{i,p}=0$) the personal judgment of agent i is correct (resp., incorrect), and an unconventional choice would yield agent i high (resp., low) payoff. Variables $(\xi_{i,q},\xi_{i,p})_{i\in N}$ are independent conditionally on $(\mathbf{q},\xi_q)$, and variables $\left(\xi_q,(\xi_{i,p})_{i\in N}\right)$ are independent.

Note that without conditioning on ξ_q the success probability of a conventional choice is $\mathbf{q}$:

$$P(\xi_{i,q}=1|\mathbf{q}=q)=(1-q)\cdot(1-\sqrt{\rho})\cdot q+q\cdot(\sqrt{\rho}+(1-\sqrt{\rho})\cdot q)=q.$$

Also observe that conditionally on $\mathbf{q}$, the correlation coefficient between the outcomes of each two agents i,j who follow the public signal is ρ. That is, for each $q\in[0,1]$:

$$\rho(\xi_{i,q},\xi_{j,q}|\mathbf{q}=q)=\frac{\mathbf{E}(\xi_{i,q}\cdot\xi_{j,q}|\mathbf{q}=q)-\mathbf{E}(\xi_{i,q}|\mathbf{q}=q)\cdot\mathbf{E}(\xi_{j,q}|\mathbf{q}=q)}{\sqrt{var(\xi_{q,i}|\mathbf{q}=q)\cdot var(\xi_{q,i}|\mathbf{q}=q)}}$$

$$=\frac{q\left(\sqrt{\rho}+\left(1-\sqrt{\rho}\right)\cdot q\right)^2+(1-q)\left(\left(1-\sqrt{\rho}\right)\cdot q\right)^2-q^2}{\sqrt{q(1-q)q(1-q)}}=\frac{-\rho q^2+q\rho}{q(1-q)}=\rho.$$

The strategic interaction between the principal and the agents includes two stages. At stage 1 the principal (who has no information on the state of nature) chooses a profile of bias functions $(g_i)_{i\in I}$. Each function $g_i:[0,1]\to[0,1]$ determines the bias of agent i when estimating the accuracy of his own judgment. That is, if the true success probability of an unconventional choice based on personal judgment is p_i, then agent i mistakenly believes it to be $g_i(p_i)$ (see online appendix for allowing also biases with respect to the conventional choice.). As discussed in the introduction, the choice of the bias profile $(g_i)_{i\in I}$ is interpreted to be the result of using managerial skills to influence the confidence of the different agents. The principal knows all aspects of the model, while each agent is not aware that he has a confidence bias. (See the related experimental findings of Friesen and Weller (2006) which suggest that: (1) analysts are overconfident; (2) an analyst is not aware of his own bias; and (3) an analyst is aware that other analysts tend to be overconfident.)

After stage 1, all agents are publicly informed about the value of $\mathbf{q}$, and each agent i with bias function g_i, is privately misinformed that the value of $\mathbf{p}_i$ is $g_i(\mathbf{p}_i)$. At stage 2 each agent i chooses an action $a_i \in \{a_c, a_u\}$, where a_c (a_u) is interpreted as a conventional (unconventional) choice. The payoff of agent i is (see online appendix for discussing a safer conventional choice):

$$\mathbf{u}_i(a_c) = \begin{cases} H & \text{if } \xi_{i,q} = 1 \\ L & \text{if } \xi_{i,q} = 0 \end{cases}, \text{ and } \mathbf{u}_i(a_u) = \begin{cases} H & \text{if } \xi_{i,p} = 1 \\ L & \text{if } \xi_{i,p} = 0 \end{cases}.$$

Our assumption that f_p and f_q are continuous guarantee that the inequality $\mathbf{q} \neq g(\mathbf{p}_i)$ holds with probability 1. Thus, each bias profile $(g_i)_{i\in I}$ induces a strictly dominant strategy profile for each agent i: taking a conventional action if $\mathbf{q} > g_i(\mathbf{p}_i)$, and an unconventional one if $\mathbf{q} < g_i(\mathbf{p}_i)$ (and playing arbitrary if $\mathbf{q} = g(\mathbf{p}_i)$ - a 0-probability event). Let $\mathbf{u}_i(g_i) = u_i(g_i, \mathbf{p}_i, \mathbf{q}, \xi_q, \xi_{i,p}, \xi_{i,q})$ be the random payoff of agent i who uses this dominant strategy.

The principal is an expected utility maximizer with a payoff function that is strictly concave and increasing in the average agent payoff (or, equivalently, the number of successful agents):

$$u\left((g_i)_{i\in I}\right) = \mathbf{E}_{(\mathbf{p}_i)_{i\in I}, \mathbf{q}, (\xi_{i,p}, \xi_{i,q})_{i\in I}}\left(h\left(\frac{1}{n}\sum_{i\in I}\mathbf{u}_i(g_i)\right)\right).$$

3 Results

A few auxiliary definitions are required before presenting our results. Let the first-best payoff of the game, be the payoff that can be achieved by the principal when he obtains all the private signals $\mathbf{p}_i$-s and has full control over the actions of the agents. A bias profile *ϵ-induces the first-best payoff* if its payoff is as good as the first-best payoff up to ϵ. Bias profile $(g_i)_{i\in I}$ is *homogeneous* if all agents have the same bias function: $\forall i, j \in I,\ g_i = g_j$. With some abuse of notations, we identify function $g : [0,1] \to [0,1]$ with the homogeneous profile $(g)_{i\in I}$. We say that g *asymptotically induces the first-best payoff* if for every $\epsilon > 0$, there is large enough n_0 such that, for any game with at least n_0 agents, g ϵ-induces the first-best payoff.

The following theorem characterizes the unique optimal bias function.

Theorem 1 *There exists a unique bias function g^* that asymptotically induces the first-best payoff. Moreover, g^* satisfies the following properties:*

(1) Overconfidence: $g^*(p) > p$ for every $0 < p < 1$.
(2) Continuous and strictly increasing: $\frac{dg^*(p)}{dp} > 0\ \forall\, 0 < p < 1$, $g^*(0) = 0$, and $g^*(1) = 1$.
(3) g^* does not depend on the distribution f_q.

The intuition for Theorem 1 is as follows. There is a conflict of interest between calibrated agents $(g_i(p_i) = p_i)$ who maximize their probability of success, and the principal who wishes some agents with $\mathbf{p}_i < \mathbf{q}$ to make unconventional choices in order to achieve better risk diversification and to reduce the variance of the fraction of successes. The optimal action of agent i in the principal's first-best profile (if the principal could observe agents' private signals) generally depends on the entire realized profile of signals: $(\mathbf{p}_1, .., \mathbf{p}_n, \mathbf{q})$. However, when there are many agents, the realized empirical distribution of $(\mathbf{p}_1, .., \mathbf{p}_n)$ is very close to its prior distribution f_p. Thus, approximately, the first-best choice of agent i only depends on the realizations of $\mathbf{p}_i$ and $\mathbf{q}$. Specifically, for every q, there is some threshold level $g^{-1}(q) < q$, which is strictly increasing and continuous in q, such that it is approximately optimal for the principal if each agent i makes an unconventional choice if and only if $\mathbf{p}_i > g^{-1}(q)$. These thresholds construct the optimal bias function $g(p)$. This optimal level of overconfidence aligns the preferences of the principal and the agents. That is, the agents behave as if they have the payoff function of the principal.

Part 3 of Theorem 1 holds due to our simplifying assumption that all agents have the same success probability when making conventional choices. If agents were facing different values of q when making conventional choices, then our results would remain qualitatively similar but the optimal level of overconfidence would also depend on f_q.

Theorem 1 shows uniqueness in the set of homogeneous bias profiles. That is, it shows that any other homogeneous bias profile induces a worse outcome than g^*, given that the number of agents is sufficiently large. Theorem 2 extends the uniqueness also to the set of heterogeneous profiles. Bias profile $(g_i)_{i\in I}$ is *heterogeneous* if there is a set $Q \subseteq [0,1]$ with a positive Lebesgue measure such that, for each $q \in Q$, $\min_i (g_i)^{-1}(q) < \max_i (g_i)^{-1}(q)$. With some abuse of notation, we identify the bias profile $(g_i)_{i\in I}$ with the following bias profile in a game with $k \cdot |I|$ agents: agents $\{1, ..., k\}$ have bias function g_1, agents $\{k+1, ..., 2k\}$ have bias function g_2, ..., and agents $\{k \cdot (|I|-1)+1, ..., k|I|\}$ have bias function $g_{|I|}$. Theorem 2 shows that every heterogeneous profile can be replaced with an homogeneous profile that induces a strictly better outcome, given that the number of agents is sufficiently large.

Theorem 2 *Let $(g_i)_{i\in I}$ be an heterogeneous profile. Then there is $k_0 \in \mathbb{N}$ such that there is an homogeneous profile that induces a strictly better payoff than $(g_i)_{i\in I}$ in the game with $k \cdot |I|$ agents for every $k \geq k_0$.*

The intuition for Theorem 2 is as follows. Let g be a bias function (an homogeneous profile) that induces the same expected number of agents who follow the public signal as the profile $(g_i)_{i\in I}$ (for each q). Divide the set of agents into two disjoint subsets I_1 and I_2, where I_1 includes agents with bias functions that induce higher probabilities to follow the unconventional choice (relative to g), and I_2 is the complementary set. If the number of agents is sufficiently large, then the law of large numbers implies that replacing the heterogeneous bias profile $(g_i)_{i\in I}$ with the homogenous bias function g, would result in having several agents from I_1 moving from unconventional choices to conventional ones, and several agents from I_2 doing the opposite move. Observe that the "moving-out" I_1 agents have lower $\mathbf{p}_i$-s than the "moving-in" I_2 agents, as everyone uses the same threshold $(g^{-1}(q))$ when deciding whether or not to take the unconventional action. This implies that g induces a strictly better payoff (by using again the law or large numbers).

Our third result shows that the principal prefers more overconfident agents if: (1) he becomes more risk-averse, or (2) the correlation coefficient ρ becomes larger.

Theorem 3 *Let g_1^* (Resp., g_2^*) be the unique bias function that asymptotically induces the first-best payoff given the parameters $(\rho_1, f_{q,1}, f_{p,1}, L_1, H_1, h_1)$ (resp., $(\rho_2, f_{q,2}, f_{p,2}, L_2, H_2, h_2)$). Then $g_1^*(p) > g_2^*(p)$ for every $0 < p < 1$ in each of the following cases:*

(1) $h_1 = \psi \circ h_2$ where $\psi' > 0$ and $\psi'' < 0$ (and all other parameters are the same).
(2) $\rho_1 > \rho_2$ (and all other parameters are the same).

The intuition of Theorem 3 is as follows. If the principal becomes more risk-averse, then he gives more weight to reducing the variance of the number of successes. This deepens the conflict of interest with calibrated agents, and more overconfidence is required to align the preferences of the agents and the principal. Similarly, if the correlation coefficient becomes larger, this enlarges the aggregate risk that is induced by conventional choices, and it deepens the conflict of interest between the principal and calibrated agents.

Finally, we assume that the principal has a constant relative risk aversion (CRRA) utility:

$$h(x) = \begin{cases} \frac{x^{1-\phi}}{1-\phi} & \text{if } \phi > 0, \phi \neq 1, \\ ln(x) & \text{if } \phi = 1, \end{cases}$$

Let $D = \frac{H-L}{L}$ be the (normalized) potential gain: the ratio between the extra payoff that can be gained when succeeding $(H-L)$ and the minimal guaranteed payoff (L). Theorem 4 shows that assuming CRRA utility yields three additional results. The first result (parts 1-2) shows that the principal hires more overconfident agents if there is a larger potential gain. Larger potential gains are associated with more important decisions, and given the evolutionary interpretation presented in Section 4, this result fits the experimental findings that people tend to be more overconfident when making important decisions (e.g., Sieber, 1974). The second result (part 3) shows that the principal hires more overconfident agents if the success probabilities of unconventional choices become smaller, and this fits the experimentally observed *hard-easy effect* (see Section 5.2). The last result (part 4) shows that the induced overconfidence fits the experimentally observed *false certainty effect* (see Section 5.2). [5]

Theorem 4 *Let g_1^* (Resp., g_2^*) be the unique bias function that asymptotically induces the first-best payoff given the parameters $(\rho_1, f_{q,1}, f_{p,1}, L_1, H_1, h_1)$ (resp., $(\rho_2, f_{q,2}, f_{p,2}, L_2, H_2, h_2)$). Assume that both h_1 and h_2 satisfy CRRA. Then:*

(1) Each g_k^* ($k \in \{1,2\}$) depends on the payoffs only through its dependency on $D_i = \frac{H_k - L_k}{L_k}$.
(2) If $D_1 > D_2$ (and all other parameters are the same) then $g_1^*(p) > g_2^*(p)\ \forall\, 0 < p < 1$.
(3) If distribution $f_{p,2}$ has first order stochastic dominance over $f_{p,1}$ (and all other parameters are the same) then $g_1^*(p) > g_2^*(p)\ \forall\, 0 < p < 1$.

[5] Part (4) implies this effect only under the strong assumptions that D is high and ρ is close to 1. However, one can extend this result into a setup where $\mathbf{D}$ and ρ are random variables having a joint distribution with a positive weight on high values, and the principal is limited to inducing a single bias function $g^*(p)$ for all values of $\mathbf{D}$ and ρ.

(4) The ratio between the perceived and the true error probability of personal judgment $\left(\frac{1-g_k^*(p)}{1-p}\right)$ converges to $(1/(D_k+1))^{\phi} = \left(\frac{L_k}{H_k}\right)^{\phi}$ when both p and ρ converge to 1.

The intuition of the part (1) is that evaluations based on a CRRA utility are unaffected by scale. The intuition of part (2) is that larger potential gain, enlarges the aggregate risk of the conventional actions. This deepens the conflict of interest between the principal and calibrated agents, and more overconfidence is required to compensate for it. The intuition of part (3) is that the principal wishes that agents with the highest success probabilities would make unconventional choices. When lower success probabilities are more frequent, each accuracy level p_i is more likely to be one of the higher levels.

4 Evolutionary Interpretation

In this section we present a different interpretation of our model to show why overconfidence is a unique evolutionary stable behavior. Consider the following example about choices of individuals in prehistoric societies.

Example 2 *Human hunting in prehistoric hunter-gatherer societies is a skill-intensive activity. Human hunters use a wealth of information to make context-specific decisions during both the search and encounter phases of hunting. Acquiring good hunting skills requires many years of training and specialization (Kaplan et al., 2000). Consider an individual that has to choose in which hunting technique to specialize. The hunter may either make the conventional choice - specialize in the hunting technique that was most successful in the previous generation, or use his personal judgment to develop and specialize in a new unconventional hunting technique, which might better fit the changing environment or his own personal qualifications. Each agent may either succeed or fail, and this has large influence on his reproductive ability. The outcomes of individuals who specialize in the same hunting technique are positively correlated, while the outcomes of different hunting techniques are independent.*

Example 2 is formulated as follows. A large population of individuals includes several genetic types: $(T_1, ..., T_K)$. Each type k induces a (possibly random) bias function g_k for its members. In each generation, each agent faces an important decision that influences his fitness (reproduction ability), like the hunting method in Example 2. When making a decision the individual may either take a conventional action that follows accepted guidelines in his society, or follow his own judgment and take an unconventional action. At the beginning of each generation each individual i receives two signals $0 < p_i < 1$ and $0 < q < 1$. These signals have the same interpretation as in the principal-agent model: p_i (resp., q) is the independent (resp., positively correlated) success probability of an unconventional (resp., conventional) choice.

The chosen action and the state of nature leads each agent either to success or to failure: success (resp., failure) yields fitness H (resp., L). The size of each type (the number of its members) in the next generation is determined by *replicator dynamics* with a small positive mutation rate. That is, basically the new size of each type in the next generation is their size in the previous generation multiplied by their average fitness. In addition, each individual in the next generation has a small chance to be randomly assigned into a new type.

A well known argument in the evolutionary literature (see, Lewontin and Cohen, 1969; Samuelson, 1971; Robson, 1996) shows that with high probability in the long run a unique type prevails over the entire population: the type that maximizes the expectation of the logarithm of the average fitness in each generation. Formally, let $\xi_t = 1$ (resp., $\xi_t = 0$) be the event that the accepted guidelines are correct (resp., incorrect) in generation t (denoted by ξ_q in Section 2). Recall that $P(\xi_t = 1) = q$, and assume that ξ_t-s in different generations are independent. Let $\mathbf{u}_{i,k,t}$ be the number of offspring of agent i of type k in generation t, and let $\mathbf{m}_k(\xi_t) = \mathbf{E}(\mathbf{u}_{i,k,t}|\xi_t)$ be the expected number of offspring produced by an agent of type k conditional on ξ_t. Robson (1996, Theorem 2-iii) shows that if mutation rates are small, a long time elapsed, and the population is found to have avoided extinction, then the entire population is dominated by the type that maximizes:

$$\mathbf{E}(ln(\mathbf{m}_k)) = q \cdot ln(m_k(\xi_t = 1)) + (1-q) \cdot ln(m_k(\xi_t = 0)).$$

Observe that due to the law of large numbers, if the population of type k is large enough, then conditional on the value of ξ_t, the realized average fitness of the members of type k in generation t is very close to $\mathbf{m}_k(\xi_t)$: $\mathbf{m}_k(\xi_t) \approx \frac{1}{|T_k|}\sum_{i \in T_k} \mathbf{u}_{i,k,t}$. Thus, Robson (1996)'s result implies that in the long run nature selects the type that maximizes the expectation of the logarithm of the average fitness in each generation. Because of this, the long run limiting behavior that is the result of the evolutionary dynamics can be described as the bias profile that is directly chosen by a risk-averse principal with a logarithmic utility function. Thus, in the long run the homogeneous bias profile g^* is a unique evolutionary stable behavior, and all of the results of Section 3 hold in this setup as well.

In what follows we sketch out the intuition behind this result. Let n_k be the initial number of members of type k, and let $\mathbf{X}_{t,k} = \frac{1}{|T_k|}\sum_{i \in T_k} \mathbf{u}_{i,k,t}$ be the average fitness of type k in generation t. The size of each type T_k after M generations is equal to:

$$n_k \cdot \mathbf{X}_{1,k} \cdot ... \cdot \mathbf{X}_{M,k} = n_k \cdot e^{log\left(\mathbf{X}_{1,k} \cdot ... \cdot \mathbf{X}_{M,k}\right)} = n_k \cdot e^{log\left(\mathbf{X}_{1,k}\right) + ... + log\left(\mathbf{X}_{M,k}\right)}.$$

Assume that n_k is large enough and that $\mathbf{E}(log(\mathbf{X}_{t,k})) > 0$, then each $X_{t,k}$ is approximately identically distributed. Assuming that M is large enough then the size of type T_k after M generations is approximately (using the law of large numbers): $n_k \cdot e^{M \cdot \mathbf{E}\left(log\left(\mathbf{X}_{t,k}\right)\right)}$. This depends only on the expectation of the logarithm of the average fitness in each generation, and the type that maximizes this expression will expand exponentially faster than any other type.

At first glance, it might be puzzling that our dynamics is entirely based on individual selection, and yet natural selection does not choose agents who maximize the expected number of children. This is because natural selection "cares" for the number of offspring in the long run. This is not the same as maximizing the "short-run" expected number of children. A calibrated agent has an higher expected number of children than an agent with bias g^* but he also has a higher variance. Generations in which the realized average number of children of calibrated agents is small, substantially reduce their number of offspring in the long run.

Our results imply that all evolutionary histories induce overconfidence, but that there is a large variety in the level of induced overconfidence given different histories. In particular, it implies that a higher level of overconfidence is induced if the outcomes of the conventional choices have been more correlated through the history. This implication is in accordance with the experimental findings surveyed in Yates et al. (2002).

5 Discussion

5.1 Variants and Extensions

The online appendix presents several variants and extensions: (1) we allow agents to have bias also with respect to conventional choices, and we show that the results are similar; (2) we extend our results to a setup where conventional choices are safer (have a smaller variance); (3) we reformulate the model to capture overconfidence as underestimating the variance of a continuous noisy signal; (4) we extend our results to a setup where private information is costly, and each agent has to invest effort in improving the accuracy of his personal judgment; (5) we demonstrate that our results do not hold for a small number of agents, and we show that principals would prefer having many agents; (6) we show that our results also hold when agents are informed experts who recommend the principal which action to choose, and when the agents are risk-averse; (7) we discuss how incorporating unbiased random judgment errors would imply the experimentally observed underconfidence for easy tasks; and (8) we demonstrate how overconfidence can improve social welfare.

5.2 Related Literature

We conclude by discussing a small portion of the literature on overconfidence. The interested reader is referred to the surveys of Lichtenstein, Fischhoff and Phillips (1982), Griffin and Brenner (2004), and Skala (2008).

The term overconfidence has three different but related definitions in the literature. The most popular definition, also called *overestimation*, describes overconfidence as a systematic calibration bias, for which the assigned probability that the answers given are correct exceeds the true accuracy (see e.g., Oskamp, 1965; Lichtenstein, Fischhoff and Phillips, 1982). The second definition, also called *overprecision*, deals with excessive certainty regarding the accuracy of one's beliefs about an uncertain continuous quantity (see, e.g., Lichtenstein, Fischhoff and Phillips, 1982; Russo and Schoemaker, 1992). The third definition, also called *overplacement*, describes the phenomenon in which people believe themselves to be *better than average* (Svenson, 1981; Taylor and Brown, 1988). Empirical data suggests that overconfidence is not limited to the lab - it is also presented by experienced agents in real-life situations: geologists (Russo and Schoemaker, 1992), traders (Barber and Odean, 2001; Ben-David, Graham and Harvey, 2010), and analysts (Friesen and Weller, 2006).

Overconfidence has been studied in several financial and economic models. A few papers investigate motivational reasons for overconfidence: deceiving future self (Bénabou and Tirole, 2002), positive emotions improve performance (Compte and Postlewaite, 2004), "ego" utility from positive self-image (Köszegi, 2006; Weinberg, 2009), and confirmatory bias (Rabin and Schrag, 1999). Van den Steen (2004) and Santos-Pinto and Sobel (2005) demonstrate how it is possible to observe overconfidence among rational agents who choose actions based on unbiased random errors. Blume and Easley (1992), Wang (2001), Gervais and Odean (2001), and Khachatryan and Weibull (2011) study the conditions in which overconfidence can survive or even dominate in financial and economic markets. Finally, Odean (1998) investigates the influence of overconfident agents in financial markets, and Sandroni and Squintani (2007) study their influence in insurance markets.

In what follows we focus on overestimation, and describe some of its main experimental stylized facts. A main experimental finding is that the degree of overconfidence depends on the difficulty of the task - the *hard-easy effect*. The more difficult the task, the greater the observed overconfidence (Lichtenstein, Fischhoff, and Phillips, 1982; Moore and Healy, 2008): people present substantial overconfidence for hard tasks, moderate overconfidence for medium tasks, and moderate underconfidence for easy tasks. A few papers suggest that the hard-easy effect, and apparent overconfidence in general, may be the result of choosing unrepresentative hard questions in experiments and regression toward the mean (see, e.g., Erev, Wallsten and Budescu, 1994). Recent experiments demonstrate that people still present overconfidence, though to a less extent, when representative questions are used (which are randomly sampled from a natural set) and when unbiased judgmental random errors are taken into account in the analysis (see Budescu, Wallsten, and Au, 1997; Glaser, Langer and Weber, 2012). An additional finding is the *false certainty effect*: people are often wrong when they are certain in their private information. In the experiment of Fischhoff, Slovic, and Lichtenstein (1977) participants severely underestimated the probability they erred in seemingly easy questions. The participants had sufficient faith in their confidence judgments to be willing to stake money on their validity.

The evolutionary interpretation of our model is based on an adaptation of Robson (1996)'s result, which shows that Nature should design people to be risk averse with respect to aggregate risk and risk neutral with respect to idiosyncratic risk. Our main contribution beyond Robson's result is threefold: (1) we relate this risk attitude with overconfidence; (2) we demonstrate that the induced overconfidence satisfies empirically plausible comparative statics and fits various stylized facts; and (3) we show that the same mechanism helps to understand strategic (non-evolutionary) principal-agent interactions.

Recently and independently of our paper, Louge (2010) adapted Robson (1996)'s result to show that Nature should design people to have overconfidence-related biases: (1) a bias towards actions that defy "common wisdom", and (2) more extreme public information is required before disregarding private information. Louge applied this rule to *cascade* interactions and demonstrated that herds eventually arise, but the probability of herding on the wrong action is lower than with a rational rule. Our model differs from Louge (2010) in a few aspects: (1) we also deal with strategic principal-agent interactions; (2) our model explains various stylized facts about overconfidence; and (3) we allow for partial correlation between conventional choices, and for costly private signals (see online appendix).

We conclude by quickly surveying other evolutionary foundations to overconfidence. Waldman (1994) showed that "second-best" adaptations can be evolutionarily stable with sexual inheritance. In particular, he demonstrated that the combination of overconfidence with excess disutility from effort is a "second-best" adaptation. Similarly, Blume and Easley (1992) and Wang (2001) demonstrated that the combination of overconfidence and excess risk aversion are "second-best" adaptations. Contrary to that, in our model overconfidence induces the first-best outcome, and does not compensate for other errors. Finally, Bernardo and Welch (2001) show that an heterogeneous population that include both overconfident and calibrated agents may be stable if evolution combines both group and individual selection: being overconfident reduces the fitness of the individual, but it substantially improves the fitness of his group, by inducing positive information externality in a *cascade* interaction. Contrary to that, our model only relies on individual selection, and it does not include information externalities.

A Proofs

A.1 Preliminaries

The following lemma presents an equivalent formulation for DARA property.

Lemma 1 Let $h(y)$ be a strictly concave increasing function. Then $h(y)$ satisfies (strictly) decreasing absolute risk aversion (DARA) if and only if $f_a(y) = \frac{h'(y)}{h'(y+a)}$ is a strictly decreasing function of y for each $a > 0$.

PROOF. $h(y)$ satisfies DARA $\Leftrightarrow$ for every $y, a > 0$, $r_A(y) > r_A(y+a)$

$$\Leftrightarrow -\frac{h''(y)}{h'(y)} > -\frac{h''(y+a)}{h'(y+a)} \Leftrightarrow h''(y) \cdot h'(y+a) - h''(y+a) \cdot h'(y) < 0$$

$$\Leftrightarrow f'_a(y) = \frac{h''(y) \cdot h'(y+a) - h''(y+a) \cdot h'(y)}{(h'(y+a))^2} < 0. \square$$

A.2 Proof of Theorem 1

The proof includes two parts: (1) showing that the first-best outcome can be approximated by a bias function, and (2) characterizing this unique bias function.

Approximating the first-best payoff by a bias function

We begin by dealing with the "first-best" case in which the principal receives all private signals $(\mathbf{p}_i)_{i\in I}$ and public signal $\mathbf{q}$ and chooses the actions of all the agents. Without loss of generality the first-best strategy is a function ϕ that chooses a threshold $p = \phi(q, p_1, ..., p_n)$, such that each agent i with higher (resp., lower) accuracy level $p_i \geq p$ (resp., $p_i < p$) makes an unconventional (resp., conventional) choice. The expected payoff (u) of this threshold is:

$$\mathbf{E}\left(h\left(L + (H-L)\cdot\left(\frac{1}{n}\left(|\{i|p_i < p, \xi_{i,q} = 1\}| + |\{i|p_i \geq p,\ \xi_{i,p} = 1\}|\right)\right)\right)|q, (p_i)_{i\in I}\right).$$

Variables $(\xi_{i,q}, \xi_{i,p})_{i\in I}$ are conditionally independent given ξ_q. Assuming that the number of agents is large enough, the expected payoff is well approximated by

$$u = P(\xi_q = 1)\cdot h\left(L + \frac{H-L}{n}\cdot\left(\sum_{p_i<p} P(\xi_{i,q} = 1|\xi_q = 1) + \sum_{p_i\geq p} P(\xi_{i,p} = 1)\right)\right) +$$
$$P(\xi_q = 0)\cdot h\left(L + \frac{H-L}{n}\cdot\left(\sum_{p_i<p} P(\xi_{i,q} = 1|\xi_q = 0) + \sum_{p_i\geq p} P(\xi_{i,p} = 1)\right)\right) + o(\epsilon).$$

$$\Leftrightarrow\ u = q\cdot h\left(L + \frac{H-L}{n}\cdot\left((\sqrt{\rho} + (1-\sqrt{\rho})\cdot q)\cdot \#\{i|p_i < p\} + \sum_{p_i\geq p} p_i\right)\right) +$$
$$(1-q)\cdot h\left(L + \frac{H-L}{n}\cdot\left(((1-\sqrt{\rho})\cdot q)\cdot \#\{i|p_i < p\} + \sum_{p_i\geq p} p_i\right)\right) + o(\epsilon).$$

To simplify notation let $f = f_p$ and $F = F_p$. Assuming again that the number of agents is large enough, one can approximate the empirical distribution of the private signals $(p_1, ..., p_n)$ by their prior distribution f. This gives the following approximation:

$$u = q\cdot h\left(L + (H-L)\cdot\left((\sqrt{\rho} + (1-\sqrt{\rho})\cdot q)\cdot F(p) + \int_p^1 x\cdot f(x)\,dx\right)\right) + \tag{A.1}$$
$$(1-q)\cdot h\left(L + (H-L)\cdot\left(((1-\sqrt{\rho})\cdot q)\cdot F(p) + \int_p^1 x\cdot f(x)\,dx\right)\right) + o(\epsilon).$$

Consider the bias function $g^*(p)$ that is defined as follows: $p = (g^*)^{-1}(q)$ is the threshold that maximizes Eq. A.1 (neglecting the error term $o(\epsilon)$). By the above arguments, such a bias function ϵ-induces the first-best payoff.

Characterizing the unique optimal bias function $g^*(p)$

We now calculate the value of $p = (g^*)^{-1}(q)$ that maximizes Eq. A.1 (neglecting the error term $o(\epsilon)$). Observe first that p must be in the interval $\left[\left(1-\sqrt{\rho}\right)\cdot q, q\right]$ because for smaller (resp., larger) p-s following a_c (resp., a_u) strictly dominates following a_u (resp., a_c) as it yields a 1st-order (resp., 2nd-order) stochastic dominant payoff. To simplify notation let:

$$A_{p,q,\rho} = L + (H - L) \cdot \left(((1 - \sqrt{\rho}) \cdot q) \cdot F(p) + \int_p^1 x \cdot f(x)\, dx \right),$$

$$B_{p,q,\rho} = A_{p,q,\rho} + \sqrt{\rho} \cdot F(p) \cdot (H - L).$$

Substituting these into (A.1) yields: $u(p,q) = q \cdot h(B_{p,q,\rho}) + (1-q) \cdot h(A_{p,q,\rho})$. For every $0 < q < 1$ we find $p = (g^*)^{-1}(q)$ by derivation (one can verify that $g^*(0) = 0$ and $g^*(1) = 1$):

$$\frac{du}{dp} = q \cdot h'(B_{p,q,\rho}) \cdot \frac{\partial B_{p,q,\rho}}{\partial p} + (1-q) \cdot h'(A_{p,q,\rho}) \cdot \frac{\partial A_{p,q,\rho}}{\partial p}.$$

Assuming an internal solution ($\frac{du}{dp} = 0$) yields:

$$\frac{h'(A_{p,q,\rho})}{h'(B_{p,q,\rho})} = \frac{q \cdot \frac{\partial B_{p,q,\rho}}{\partial p}}{(1-q) \cdot \frac{\partial A_{p,q,\rho}}{\partial p}} = \frac{q \cdot \left(\sqrt{\rho} + \left(1 - \sqrt{\rho}\right) \cdot q - p\right)}{(1-q) \cdot \left(p - \left(1 - \sqrt{\rho}\right) \cdot q\right)}. \tag{A.2}$$

Observe that $A_{p,q,\rho}$ (resp., $B_{p,q,\rho}$) is strictly decreasing (resp., increasing) in p in the interval $\left(\left(1 - \sqrt{\rho}\right) \cdot q, q\right)$. Using the strict concavity of h, this implies that the left-hand side (l.h.s.) of Eq. A.2 is strictly increasing in p. Note that the right-hand side (r.h.s.) is strictly decreasing in p, and that for p close enough to $\left(1 - \sqrt{\rho}\right) \cdot q$ the r.h.s. is strictly larger than the l.h.s, while the opposite holds for p close enough to q. This implies that for each $0 < q < 1$ there is a unique solution $p = (g^*)^{-1}(q)$ to Eq. A.2 in $\left(\left(1 - \sqrt{\rho}\right) \cdot q, q\right)$, which is continuous in q.

Observe that $A_{p,q,\rho}$ is weakly increasing in q. By Lemma 1, it implies that the l.h.s. of Eq. A.2 is weakly decreasing in q. One can verify that the r.h.s. is strictly increasing in q. Using the observations of the previous paragraph, it implies that $p = (g^*)^{-1}(q)$ is strictly increasing in q. One can verify that $\frac{du}{dp} > 0$ for every $p < (g^*)^{-1}(q)$, and $\frac{du}{dp} < 0$ for every $p > (g^*)^{-1}(q)$. Thus, any other bias threshold $p \neq (g^*)^{-1}(q)$ would yield a strictly lower expected payoff. Finally, observe (Eq. A.2), that $g^*(p)$ does not depend on the distribution f_q.$\square$

A.3 Proof of Theorem 2

Let $\tilde{g}$ be the following bias function (homogeneous bias profile): for each $q \in [0,1]$, $(\tilde{g})(q)^{-1}$ is the unique solution to the following equation:

$$F\left((\tilde{g})^{-1}(q)\right) = \sum_{i \in I} \frac{1}{|I|} \left(F\left(g_i^{-1}(q)\right)\right).$$

That is, $\tilde{g}$ is a bias function that averages the heterogeneous profile $(g_i)_{i \in I}$. Fix any $0 < q < 1$ satisfying $\min_i (g_i)^{-1}(q) < \max_i (g_i)^{-1}(q)$. To simplify notation let $p_i = g_i^{-1}(q)$ and $\tilde{p} = \tilde{g}^{-1}(q)$. By the arguments given in the previous subsection, for each q, the expected payoff of the heterogeneous profile $(g_i)_{i \in N}$ in the game with $k \cdot |I|$ agents (for large enough k) is approximately given by

$$q\cdot h\left(L+(H-L)\cdot\left(\frac{1}{|I|}\sum_{i\in I}\left(\left(\sqrt{\rho}+\left(1-\sqrt{\rho}\right)\cdot q\right)\cdot F\left(p_i\right)+\int_{p_i}^{1}x\cdot f\left(x\right)dx\right)\right)\right)+$$
$$(1-q)\cdot h\left(L+(H-L)\cdot\left(\frac{1}{|I|}\sum_{i\in I}\left(\left(\left(1-\sqrt{\rho}\right)\cdot q\right)\cdot F\left(p_i\right)+\int_{p_i}^{1}x\cdot f\left(x\right)dx\right)\right)\right),$$

and the expected payoff of the homogeneous profile $\tilde{g}$ is approximately given by

$$q\cdot h\left(L+(H-L)\cdot\left(\left(\sqrt{\rho}+\left(1-\sqrt{\rho}\right)\cdot q\right)\cdot F\left(\tilde{p}\right)+\int_{\tilde{p}}^{1}x\cdot f\left(x\right)dx\right)\right)+$$
$$(1-q)\cdot h\left(L+(H-L)\cdot\left(\left(\left(1-\sqrt{\rho}\right)\cdot q\right)\cdot F\left(\tilde{p}\right)+\int_{\tilde{p}}^{1}x\cdot f\left(x\right)dx\right)\right).$$

As $F\left(\tilde{p}\right)=\sum_{i\in I}\frac{1}{|I|}\left(F\left(\tilde{p}_i\right)\right)$, the homogeneous profile yields a higher payoff if and only if

$$\frac{1}{|I|}\sum_{i\in I}\int_{p_i}^{1}xf\left(x\right)dx<\int_{\tilde{p}}^{1}xf\left(x\right)dx$$
$$\Leftrightarrow\ \frac{1}{|I|}\sum_{i\in I}\left(\int_{p_i}^{1}xf\left(x\right)dx-\int_{\tilde{p}}^{1}xf\left(x\right)dx\right)<0\ \Leftrightarrow\ \frac{1}{|I|}\sum_{i\in I}\int_{p_i}^{\tilde{p}}xf\left(x\right)dx<0$$
$$\Leftrightarrow\ \frac{1}{n}\sum_{i\in I}\left(F\left(\tilde{p}\right)-F\left(p_i\right)\right)\cdot\mathbf{E}\left(\mathbf{p}|\min\left(p_i,\tilde{p}\right)\leq\mathbf{p}\leq\max\left(p_i,\tilde{p}\right)\right)<0.\tag{A.3}$$

(Using the notation that $\int_a^b f\left(x\right)dx=-\int_b^a f\left(x\right)dx$ when $b<a$.) Observe that

$$\frac{1}{n}\sum_{i\in I}\left(F\left(\tilde{p}\right)-F\left(p_i\right)\right)=0,$$

and that $\mathbf{E}\left(\mathbf{p}|\min\left(p_i,\tilde{p}\right)\leq\mathbf{p}\leq\max\left(p_i,\tilde{p}\right)\right)$ is strictly increasing in p_i and strictly decreasing in $\left(F\left(\tilde{p}\right)-F\left(p_i\right)\right)$. This implies that (A.3) holds.

The above arguments show that for each q such that $\min_i\left(g_i\right)^{-1}(q)<\max_i\left(g_i\right)^{-1}(q)$, $\tilde{g}$ has higher expected value than $\left(g_i\right)_{i\in I}$, conditional on $\mathbf{q}=q$. The fact that $\left(g_i\right)_{i\in I}$ is a heterogeneous bias profile (i.e., that $\min_i\left(g_i\right)^{-1}(q)<\max_i\left(g_i\right)^{-1}(q)$ in a set with positive Lebesgue measure), implies that $\tilde{g}$ has higher expected value than $\left(g_i\right)_{i\in I}$ (without conditioning on the value of $\mathbf{q}$). By the law of large numbers, if the number of agents is sufficiently large then it implies that with high probability $\tilde{g}$ induces a strictly larger payoff than $\left(g_i\right)_{i\in I}$. □

A.4 Proof of Theorem 3

(1) Replacing h with $\psi\circ h$ in equation (A.2) yields:

$$\frac{\psi'\left(h\left(A_{p,q,\rho}\right)\right)\cdot h'\left(A_{p,q,\rho}\right)}{\psi'\left(h\left(B_{p,q,\rho}\right)\right)\cdot h'\left(B_{p,q,\rho}\right)}=\frac{q\cdot\left(\sqrt{\rho}+\left(1-\sqrt{\rho}\right)\cdot q-p\right)}{(1-q)\cdot\left(p-\left(1-\sqrt{\rho}\right)\cdot q\right)}.$$

Observe that the l.h.s. has become strictly larger (due to the concavity of ψ), while the r.h.s. is unchanged. By using the observations made earlier that the l.h.s (resp., r.h.s.) is weakly decreasing (resp., strictly increasing) in q, it implies that $g_1^*(p) > g_2^*(p)$.

(2) Observe that the r.h.s. of Eq. (A.2) is strictly decreasing in ρ, while the the l.h.s. of Eq. (A.2) is strictly increasing in ρ (as $\frac{\partial A_{p,q,\rho}}{\partial \rho} < 0$ and $\frac{\partial B_{p,q,\rho}}{\partial \rho} > 0$). By using the above argument, it implies that $g_1^*(p) > g_2^*(p)$. □

A.5 *Proof of Theorem 4*

(1) Let $D = \frac{H-L}{L}$. Placing $h'(x) = x^{-\phi}$ in the l.h.s. of (A.2) yields:

$$\frac{(A_{p,q,\rho})^{-\phi}}{(B_{p,q,\rho})^{-\phi}} = \left(\frac{A_{p,q,\rho} + \sqrt{\rho}\cdot F(p)\cdot (H-L)}{A_{p,q,\rho}}\right)^{\phi} = \left(1 + \frac{\sqrt{\rho}\cdot F(p)\cdot (H-L)/L}{A_{p,q,\rho}/L}\right) =^{\phi}$$

$$\left(1 + \frac{D\cdot\sqrt{\rho}\cdot F(p)}{1 + D\left(\left(\left(1-\sqrt{\rho}\right)\cdot q\right)\cdot F(p) + \int_p^1 x\cdot f(x)\,dx\right)}\right)^{\phi} \tag{A.4}$$

(2) Observe that this last expression strictly increases in D. By using the same argument as in the previous proof, it implies that a larger D induces more overconfidence.

(3) Assume that is, $F_2(p) < F_1(p)$ for every $0 < p < 1$. We have to show that $g_1^* \geq g_2^*$. Observe that (A.4) is larger when $f_{\mathbf{p},1}$ replaces $f_{\mathbf{p},2}$. This is because

$$\left(1 + \frac{D\cdot\sqrt{\rho}\cdot F_2(p)}{1 + D\left(\left(\left(1-\sqrt{\rho}\right)\cdot q\right)\cdot F_2(p) + \int_p^1 x\cdot f(x)\,dx\right)}\right)^{\phi} =$$

$$\left(1 + \left(\frac{1 + D\left(\left(\left(1-\sqrt{\rho}\right)\cdot q\right)\cdot F_2(p) + \int_p^1 (1-F_2(x))\,dx\right)}{D\cdot\sqrt{\rho}\cdot F_2(p)}\right)^{-1}\right)^{\phi} =$$

$$\left(1 + \left(\frac{D\left(\left(1-\sqrt{\rho}\right)\cdot q\right)\cdot F_2(p)}{D\cdot\sqrt{\rho}\cdot F_2(p)} + \frac{1 + D\int_p^1 (1-F_2(x))\,dx}{D\cdot\sqrt{\rho}\cdot F_2(p)}\right)^{-1}\right)^{\phi} \leq$$

$$\left(1 + \left(\frac{\left(\left(1-\sqrt{\rho}\right)\cdot q\right)}{\sqrt{\rho}} + \frac{1 + D\int_p^1 (1-F_2(x))\,dx}{D\cdot\sqrt{\rho}\cdot F_1(p)}\right)^{-1}\right)^{\phi} \leq$$

$$\left(1 + \left(\frac{\left(\left(1-\sqrt{\rho}\right)\cdot q\right)}{\sqrt{\rho}} + \frac{1 + D\int_p^1 (1-F_1(x))\,dx}{D\cdot\sqrt{\rho}\cdot F_1(p)}\right)^{-1}\right)^{\phi}.$$

Similar to the arguments in the proof of Theorem 3, $g_1^*(p) \geq g_2^*(p)$ for every p.

(4) Let both p and ρ converge to 1 (which implies that q also converges to 1). Substituting it in Eq. (A.4) completes the proof: $\left(\frac{H}{L}\right)^{\phi} = (1+D)^{\phi} \approx \frac{(1-p)}{(1-q)}$. □

References

[1] **Barber, Brad M., and Terrance Odean.** 2001. "Boys will be boys: gender, overconfidence, and common stock investment". *The Quarterly Journal of Economics*, 116(1): 261-292.

[2] **Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey.** 2012. "Managerial Miscalibration" http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1640552.

[3] **Bénabou, Roland, and Jean Tirole**. 2002. "Self-confidence and personal motivation" *Quarterly Journal of Economics,* 117(3): 871-915.

[4] **Blume, Lawrence, and David Easley**. 1992. "Evolution and Market Behavior" *Journal of Economic Theory*, 58, 9-40.

[5] **Bernardo, Antonio E., and Ivo Welch.** 2001. "On the evolution of overconfidence and entrepreneurs" *Journal of Economics & Management Strategy*, 10(3), 301-330.

[6] **Budescu, David V., Ning Du**. 2007. "Coherence and consistency of investors' probability judgments" *Management Science*, 53(11): 1731-1744.

[7] **Budescu, David V., Thomas Wallsten, and Wing Tung Au**. 1997. "On the importance of random error in the study of probability judgment. Part II: applying the stochastic judgment model to detect systematic trends" *Journal of Behavioral Decision Making,* 10: 173-188.

[8] **Busenitz, Lowel W., Jay B. Barney**. 1997. "Differences between entrepreneurs and managers in large organizations: biases and heuristics in strategic decision-making" *Journal of Business Venturing*, 12: 9-30.

[9] **Compte, Oliver, and Andrew Postlewaite**. 2004. "Confidence-Enhanced Performance" *The American Economic Review*, 94(5), 1536-1557.

[10] **Cooper, Arnold C., Carolyn Y. Woo, and William C. Dunkelberg**. 1988. "Entrepreneurs' perceived chances for success" *Journal of Business Venturing* 3, 97-108.

[11] **Erev, Ido, Thomas S. Wallsten, David V. Budescu**. 1994. "Simultaneous over and underconfidence: the role of error in judgment processes" *Psychological Review*, 101(3): 519-527.

[12] **Fischhoff, Baruch, Paul Slovic, and Sarah Lichtenstein.** 1977. "Knowing with certainty: the appropriateness of extreme confidence" *Journal of Experimental Psychology: Human Perception and Performance,* 3(4): 552-564.

[13] **Friesen Geoffrey, Paul A Weller**. 2006. "Quantifying cognitive biases in analyst earnings forecasts" *Journal of Financial Markets*, 9: 333-365.

[14] **Gervais, Simon, J.B. Heaton, Terrance Odean**. 2011. "Overconfidence, Compensation Contracts, and Capital Budgeting" *Journal of Finance*, 66(5), 1735-1777.

[15] **Gervais, Simon, Terrance Odean**. 2001. "Learning to be overconfident" *The Review of Financial Studies*, 14(1): 1-27.

[16] **Glaser, Markus, Thomas Langer, and Martin Weber**. 2012. "True overconfidence in interval estimates: evidence based on a new measure of miscalibration" *Journal of Behavioral Decision Making*, Forthcoming, DOI: 10.1002/bdm.1773.

[17] **Griffin, Dale, and Lyle Brenner.** 2004. "Perspectives on probability judgment calibration" In: *Blackwell Handbook of Judgment and Decision Making*, ed. Derek J. Koehler and Nigel Harvey, 177–99, Boston: Blackwell Publishing.

[18] **Howell, William C., and Sarah A. Burnett.** 1978. "Uncertainty measurement: a cognitive taxonomy" *Organizational Behavior and Human Performance* 22, 45-68.

[19] **Kaplan, Hillard, Kim. Hill, Jane Lancaster, A. Magdalena Hurtado.** 2000. "A theory of human life history evolution: diet, intelligence, and longevity" *Evolutionary Anthropology*, 9: 156-185.

[20] **Khachatryan Karen, and Jorgen Weibull.** 2011, "Entrepreneurial Overconfidence and Market Selection", mimeo, http://pzp.hhs.se/media/5064/khachatryan_duo.pdf.

[21] **Köszegi, Botond**. 2006. "Ego utility, overconfidence and task choice." *Journal of European Economic Association*, 4(4): 673-707.

[22] **Lewontin Richard C., and D. Cohen**. 1969. "On population growth in a randomly varying environment." *Proceedings of the National Academy of Sciences of USA*, 62: 1056-60.

[23] **Lichtenstein, Sarah, Baruch Fischhoff, and Lawrence D. Phillips**, 1982. "Calibration of probabilities: the state of the art to 1980." In *Judgment Under Uncertainty*: *Heuristics and Biases*, ed. D. Kahneman, P. Slovic, and A. Tversky, Cambridge: Cambridge University Press.

[24] **Louge, Fernando.** 2010. "Contrarianism and Caution in herding models: an evolutionary approach", https://sites.google.com/site/fernandolouge/research/herding_old.pdf.

[25] **Moore, Don A., and Paul J. Healy.** 2008. "The trouble with overconfidence" *Psychological Review* 2008, 115 (2): 502-517.

[26] **Odean, Terrance**. 1998. "Volume, volatility, price, and profit when all traders are above average" *The Journal of Finance*, LIII(7): 1887-1934.

[27] **Oskamp, Stuart.** 1965. "Overconfidence in case-study judgments" *Journal of Cognitive Psychology*, 29(3), 261-265.

[28] **Santos-Pinto, Luis, and Joel Sobel**. 2005. "A model of positive self-image in subjective assessments" *The American Economic Review,* 95(5): 1386-1402.

[29] **Rabin, Matthew, and Joel L. Schrag.** 1999. "First impressions matter: a model of confirmatory bias" *The Quarterly Journal of Economics*, 114: 37-81.

[30] **Robson, Arthur J**. 1996. "A biological basis for expected and non-expected utility" *Journal of economic Theory*, 68: 397-424.

[31] **Russo, Edward J., and Paul J.H. Schoemaker**. 1992. "Managing overconfidence" *Sloan Management Rev.* 33: 7-17.

[32] **Samuelson, Paul A.** 1971. "The 'fallacy' of maximizing the geometric mean in long sequences of investing or gambling" *Proceedings of the National Academy of Sciences* 68 , 2493-2496.

[33] **Sandroni, Alvaro, and Francesco Squintani**. 2007. "Overconfidence, insurance and paternalism" *The American Economic Review*, 97(5), 1994-2004.

[34] **Sieber, Joan E**. 1974. "Effects of decision importance on ability to generate warranted subjective uncertainty" *Journal of Personality and Social Psychology*, 30(5): 688-694.

[35] **Skala, Dorota.** 2008. "Overconfidence in psychology and finance - an interdisciplinary literature review" *Bank i Kredy*, 4: 33-50.

[36] **Svenson, Ola**. 1981. "Are we all less risky and more skillful than our fellow drivers?" *Acta Psychologica* 47, 143-148.

[37] **Taylor, Shelley E., Jonathan D. Brown**. 1988. "Illusion and well-being: a social psychological perspective on mental health" *Psychological Bulletin,* 103(2), 193-210.

[38] **Van Den Steen, Eric**. 2004. "Rational overoptimism (and other biases)" *The American Economic Review*, 94(4): 1141-1151.

[39] **Waldman, Michael.** 1994. "Systematic errors and the theory of natural selection" *The American Economic Review*, 84(3): 482-497.

[40] **Wang, F. Albert**. 2001. "Overconfidence, investor sentiment, and evolution" *Journal of Financial Intermediation* 10, 138-170.

[41] **Weinberg, Bruce A.** 2009. "A model of overconfidence" *Pacific Economic Review*, 14(4): 502-515.

[42] **Yates, J. Frank, Ju-Whei Lee, Winston R. Sieck, Incheol Choi, and Paul C. Price.** 2002. "Probability judgment across cultures." In: *Heuristics and Biases: The Psychology of Intuitive Judgment, ed.* T. Gilovich, D. Griffin and D. Kahneman, 271-291, Cambridge University Press.