Online Appendix for Overconfidence and Diversification

Yuval Heller

Department of Economics and Nuffield College, University of Oxford.

This appendix includes 9 sections. Section A illustrates how a typical optimal bias function g^* looks like. Sections B-I present several extensions and variants of our model. In Section B we allow agents to have bias also with respect to conventional choices, and we show that our results essentially remain the same. Section C extends the model to a setup where a conventional choice is safer (has a smaller variance). In Section D we reformulate the model to capture overconfidence as underestimating the variance of a continuous noisy signal. Section E extends 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. In Section F we demonstrate that our results do not hold for a small number of agents, and we show that principals would prefer to have many agents. Section G shows 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. In Section H we discuss how incorporating unbiased random judgment errors would imply the experimentally observed underconfidence for easy tasks. Finally, Section I demonstrates the relation between overconfidence ans social welfare.

A Illustration of an Optimal Bias Function

Figure A.1 demonstrates what a typical asymptotically-optimal bias function g^* looks like. The values of the parameters are as follows: $\rho = 1$ (perfect correlation between different agents who follow the accepted guidelines), uniform distribution for the accuracy of the private signal, payoffs are $H = 3$ and $L = 1$, and the principal's utility is logarithmic ($h(x) = ln(x)$).

B Bias With Respect to Conventional Choices

In the basic model we assume that agents can only have confidence bias with respect to unconventional choices, but not with respect to conventional choices. In this section, we observe that this assumption is without loss of generality.

Consider a more general model where the bias of each agent i is described by two functions $(g_{i,1}, g_{i,2})$ from $[0,1]$ to $[0,1]$, where $g_{i,1}$ is the bias with respect to unconventional choices

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

Figure A.1. An Example for an Asymptotically-Optimal Bias Function
Parameters: $\rho = 1,\ H = 3,\ L = 1,\ h(x) = ln(x),\ f_p \sim uniform((0,1))$

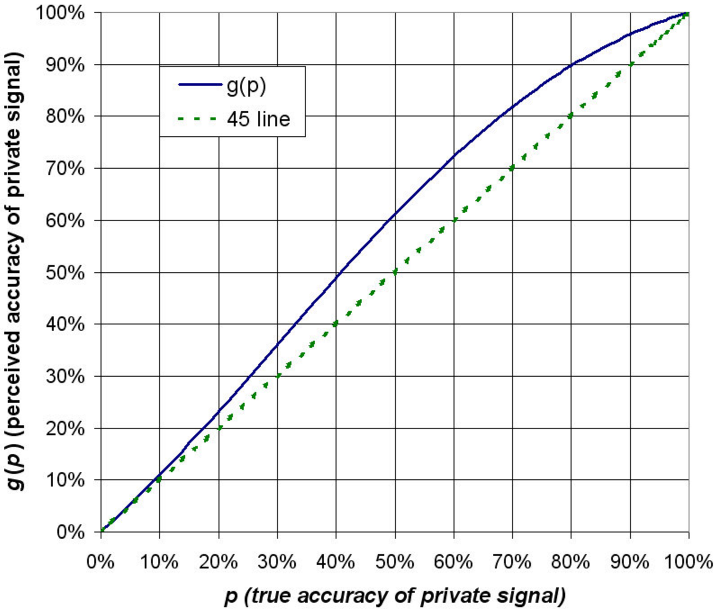

(accuracy p_i is perceived as $g_{i,1}(p_i)$) and $g_{i,2}$ is the bias with respect to conventional choices (accuracy q is perceived by agent i as $g_{i,2}(q)$). Observe that the choice of agent i only depends on the composite function $(g_{i,2})^{-1} \circ g_{i,1}$. This is because agent i chooses the conventional action if $g_{i,1}(\mathbf{p}_i) < g_{i,2}(\mathbf{q}) \Leftrightarrow (g_{i,2})^{-1} \circ g_{i,1}(\mathbf{p}_i) < \mathbf{q}$. This implies that our results remain the same in this extension. In particular, the asymptotically-optimal profile is such that each agent i has bias functions $(g_{i,1}, g_{i,2})$ that satisfy $(g_{i,2})^{-1} \circ g_{i,1} = g^*$, where g^* satisfies all the properties that were characterized in Theorems 1-4.

Thus, g^* is the excess bias in estimating the success probability of personal judgment relative to the bias in evaluating the success probability of accepted guidelines. This can be the result of either (or both) of the following reasons: (1) people are overconfident with respect to their personal judgment (and approximately calibrated with respect to the accepted guidelines); or (2) people are underconfident with respect to the accepted guidelines. We believe that the first interpretation is more plausible in most setups, and fits better the experimental stylized facts on overconfidence. Uncertainty regarding the success probability of the common guidelines is mainly external to the agent (Howell and Burnett, 1978). In many relevant applications, history of past successes and failures of following the common guidelines in the past is available to the agents and can be used to calibrate their beliefs and base them on more accurate set-based (frequency) evaluations. Uncertainty regarding the agent's own judgment is mostly internal. In most applications, an agent would have very limited information about the diverse intuitions that guided the unconventional choices of other agents in the past, and it may be hard for an agent to use this limited information to calibrate his personal judgment evaluation. Experimental evidence suggest that only internal uncertainty induces overconfidence (see, e.g., Budescu, and Du, 2007).

C Safer Accepted Guidelines

In the basic model we assume that conventional choices have the same variance as unconventional choices. However, in many applications, conventional choices might be safer - have a smaller variance. In this section, we demonstrate that this assumption can be relaxed without qualitatively affect our results.

Consider an extension of our model where a conventional choice has the following payoff: $\alpha \cdot (q \cdot H + (1-q) \cdot L) + (1-\alpha) \cdot H$ for success and $\alpha \cdot (q \cdot H + (1-q) \cdot L) + (1-\alpha) \cdot L$ for failure, where q is the success probability of the conventional choice (as in the basic model), and $0 \leq \alpha < 1$ denotes how safer is the conventional choice (relative to the unconventional one). In particular, $\alpha = 0$ is equivalent to the basic model where both choices have the same variance (assuming $q = p_i$). Observe, that α does not affect the expectation of the conventional choice $(q \cdot H + (1-q) \cdot L)$. The payoff of the unconventional choice remains the same as in the basic model (H for success and L for failure).

One can show that all of our results are qualitatively the same in this more general setup. As in the basic model, there is a conflict of interest between a calibrated agent who cares only for the expectation, and the principal who also cares for reducing the aggregate risk. For a single agent, an unconventional choice yields higher variance. However, when considering the aggregate choice of many agents, unconventional actions do not bear substantial aggregate risk for the principal due to the independence assumption. The principal still wants an agent with p_i a little bit smaller than q to make an unconventional choice, and thus it is still optimal for the principal that all agents would be overconfident (though, the induced optimal level of overconfidence $g^*(p) - p$ is decreasing in α).

D Modeling Overconfidence as Underestimating Variance

We modeled overconfidence as overestimating the accuracy of discrete private signals. Another common way to model overconfidence, especially in finance models (e.g., Odean, 1998), is underestimating the variance of continuous private signals. In this section, we demonstrate how our model can be reformulated to represent overconfidence in this way. For brevity, we only sketch a simple case that is analogous to the perfect correlation case ($\rho = 1$).

Let f_q and f_p be continuous distributions on $[0, M]$ ($M > 0$). The true unknown state of the world is $\left(\mathbf{o}_q, (\mathbf{o}_i)_{i\in I}, \sigma_q^2, (\boldsymbol{\sigma^2}_i)_{i\in I}\right)$, where: (1) $\mathbf{o}_q \in \mathbb{R}$ is interpreted as the optimal way to implement the accepted guidelines (explained below), and it is normally distributed with known expectation 0 and variance σ_q^2: $\left(\mathbf{o}_q | \sigma_q^2\right) \sim Normal\left(0, \sigma_q^2\right)$, and $\sigma_q^2 \sim f_q$; and (2) each $\mathbf{o}_i \in \mathbb{R}$ is the optimal way to follow agent i's intuition: $(r_{\mathrm{i}} | \sigma_i^2) \sim Normal\,(0, \sigma_i^2)$, and $\sigma_i^2 \sim f_p$. At the first stage of the interaction the principal chooses a bias profile for the agents: $(g_i)_{i\in I}$. Each function $g_i : \mathbb{R}^+ \to \mathbb{R}^+$ describes the bias function of agent i when he evaluates the accuracy of his intuition by estimating its variance σ_i^2. Then all agents publicly receive σ_q^2, and each agent i is privately misinformed that the variance of his intuition is $g_i\left(\sigma_i^2\right)$.

Finally, each agent i has to choose: (1) which direction to follow: his own intuition (a_o) or the accepted guidelines (a_c), and (2) how to implement the chosen direction, which is denoted by number $r_i \in \mathbb{R}$. The payoff of agent i who followed the accepted guidelines (his own intuition) and chose number r_i is decreasing in the distance with the unknown optimal value: $|r_i - \mathbf{o}_q|$ ($|r_i - \mathbf{o}_i|$). The payoff of the principal is a concave increasing function of the average payoff of the agents. Similar to Theorem 1, one can show that there is a unique optimal bias function which represents overconfidence: $g^* (\sigma_i^2) < \sigma_i^2$ for each $\sigma_i^2 \in \mathbb{R}^+$.

E Costly Private Signals

The basic model assumes that private signals are costless. In this section we relax this assumption and extend our results to a more general framework that allows private signals to have costs. In the extended model, an independent random variable $0 \leq \mathbf{t}_i \sim f_{\mathbf{t}} \leq 1$ is assigned to each agent $i \in N$. Variable $\mathbf{t}_i$ is interpreted as the effectiveness of agent i in improving the accuracy of his personal judgment.

After agents are publicly informed about the value of $\mathbf{q}$ (the success probability of a conventional choice), each agent is privately informed of $\mathbf{t}_i$. Then, each agent privately chooses an effort level $0 \leq e_i \leq 1$, and receives private signal $p_i = p(e_i, \mathbf{t}_i)$ - the success probability of an unconventional choice, where p is strictly increasing in both parameters and strictly concave in e_i. The payoff of each agent is either H (success) or L (failure) minus a cost of $(H - L) \cdot e_i$ for investing effort e_i. The rest of the model is the same as the basic model.

Let $p_{\mathbf{t}_i} \in [0, 1]$ be the unique number that maximizes $p(e_i, \mathbf{t}_i) - e_i$ (uniqueness holds due to concavity). The distribution of effectiveness levels $f_{\mathbf{t}}$ induces a unique distribution of maximizing accuracy levels $f_{p_{\mathbf{t}}}$. The following proposition asserts that our results also hold in this extended model, where $f_{p_{\mathbf{t}}}$ replaces $f_{\mathbf{p}}$.

Proposition 1 *The extended model with costly signals admits a unique asymptotically-optimal bias function g^*, which is the same as in the basic model with $f_{\mathbf{p}} = f_{p_{\mathbf{t}}}$.*

PROOF. We begin by calculating the first-best profile in a game with many agents $n >> 1$. For each $q \in [0, 1]$, there is some effectiveness value $t_0 = \alpha(q)$ such that the first-best payoff can be induced by all agents using the same threshold strategy: (1) agents with low effectiveness ($\mathbf{t}_i < t_0$) do not invest any effort and make a conventional choice, and (2) agents with high effectiveness ($\mathbf{t}_i \geq t_0$) invest some effort and make unconventional choices.

Consider an agent with high effectiveness: $\mathbf{t} \geq t_0$. His expected payoff from investing effort e is $L + (H - L) \cdot (p(e, t) - e)$. This is maximized in e_t^* that satisfies $\frac{d(p(e,t))}{de} = 1$ (a unique maximizer exists due to the strict concavity of $p(e, t)$). Let $p_t^* = p(e_t^*, t)$. For large enough n, if all agents with high effectiveness invest effort e_t^*, it ϵ-maximizes the principal's payoff (by the law of large numbers). Finally, let $p_0 = p_{t_0}^*$ be the success probability of an agent with threshold effectiveness value t_0. The choice of an optimal threshold t_0 is equivalent to the problem of finding the optimal threshold p_0 in Theorem 1. Thus the bias function g^* of the basic model is also optimal and unique in the extended model (with $f_{\mathbf{p}} = f_{p_{\mathbf{t}}}$). $\square$

F Small and Endogenous Number of Agents

In our model we assumed that the number of agents is large. Example 1 shows that Theorems 1-2 are not valid when the number of agents is small. It demonstrates: (1) an asymmetric bias profile that induces higher payoff than the best bias function; and (2) a first-best outcome which is strictly better than what can be induced by bias profiles.

Example 1 *There are two agents. The low payoff is zero ($L = 0$), the high payoff is one ($H = 1$). There is perfect correlation between agents who follow the accepted guidelines ($\rho = 1$). The distribution of each p_i is uniform in $(0, 0.5)$. The principal's utility $h(x)$ is $2x$ if $x < 0.5$ and 1 if $x \geq 0.5$.*[1] *That is, the principal wishes that at least one agent succeeds, and he does not care whether the other agent also succeeds. Consider the case in which $\mathbf{q} = 0.7$. One can see that the best bias function is one such that (approximately) $g^*(0.34) = 0.7$, and that it induces payoff 0.75. The principal can achieve a higher payoff of 0.775 by using the following optimal heterogeneous bias profile: one agent always make a conventional choice, while the other agent always takes an unconventional action. The principal's first best payoff is even higher - 0.8, and it is achieved by observing both p_1 and p_2, and choosing that the agent with the higher (lower) p_i makes an unconventional (conventional) choice.*

Next, we consider the case in which the principal can choose the number of agents he employs, and we show that it is optimal for the principal to hire many agents. Proposition 2 shows that the principal strictly prefers to hire $k \cdot n$ agents than n agents. The intuition is that having more agents enables the principal to achieve better diversification.

Proposition 2 *For each $n \geq 1$ and $k \geq 2$ the principal can induce a strictly better outcome when the number of agents is $k \cdot n$ than when it is n.*

PROOF. Let $(g_i)_{i \in I}$ be a bias profile in the game with $n = |I|$ agents. Recall that for each agent $i \in I$, $\mathbf{u}_i$ is the random payoff of agent i with bias function g_i, and that the principal's payoff is $h\left(\frac{1}{n}\sum_{i \in I}\mathbf{u}_i\right)$. Consider $(g_i)_{i \in I}$ as a profile in the game with $k \cdot n$ agents (where each k agents share one of the bias functions g_i). This profile induces the following payoff: $h\left(\frac{1}{n}\sum_{i \in I}\frac{1}{k}\sum_{j=1}^{k}\mathbf{u}_{(i-1)\cdot k+j}\right)$, where for each i, the variables $\left\{\left(\mathbf{u}_{(i-1)\cdot k+j}\right)_{j=1,\ldots,k}, u_i\right\}$ are identically distributed. Observe that $\frac{1}{n}\sum_{i \in I}\frac{1}{k}\sum_{j=1}^{k}\mathbf{u}_{(i-1)\cdot k+j}$ second-order stochastically strictly dominates $\frac{1}{n}\sum_{i \in I}\mathbf{u}_i$. By the concavity of h, it implies that the principal strictly prefers

[1] To simplify the example, we use a weakly concave and increasing function h and a distribution $f_{\mathbf{p}}$ without full support. The example can be adapted such that h would be strictly concave and increasing and $f_{\mathbf{p}}$ would have full support.

the outcome in the game with $k \cdot n$ agents. Thus, any outcome in the game with n agents is strictly dominated by an outcome in the game with $k \cdot n$ agents. $\square$

The following example shows that increasing the number of agents (but not multiplying it) may be bad for the principal.

Example 2 *(Example 1 revisited) Let $L = 0$, $H = 1$, $\rho = 1$, $f_{\mathbf{p}} \sim uniform\,(0, 0.5)$, $\mathbf{q} = 0.7$ and let the principal's utility $h\,(x)$ be $2x$ if $x < 0.5$ and 1 if $x \geq 0.5$. Recall that when there are two agents the principal can achieve payoff 0.775 by using an asymmetric bias profile: one agent always makes a conventional choice, while the other agent always makes an unconventional action. When there are three agents, the principal's best payoff is only 0.75, and it is achieved by having two agents always making a conventional choice, and the third agent always making an unconventional choice. The intuition why 3 agents are worse than 2 agents is that, the definition of utility h implies that the principal mainly cares that at least half of his agents succeed. It is easier to achieve this objective when there are only 2 agents rather then when there are 3 agents.*

Remark 1 *We can use Prop. 2 to demonstrate that our results do not depend on the assumption that there is a single principal. Consider a setup where there are several risk-averse principals and many agents, and that there is a small marginal cost for each additional hired agent. Due to the risk aversion of the principals and Proposition 2, each principal would choose to hire many agents, and all principals would prefer to hire overconfident agents.*

G Risk-Averse Agents and Experts

In the basic model the principal's utlity is a concave function of the average utility of the agents, and in this sense, the principal is assumed to be more risk-averse than the agents. This may seem implausible in some applications. However, this assumption can be relaxed without changing the results as follows (using the fact that each agent faces only two possible outcomes). We reinterpret $\mathbf{u}_i$ as a monetary payoff, and we allow the utility of agent i to be any monotone function of this monetary payoff: $h_i\,(\mathbf{u}_i)$, while the principal's utility depends on the average monetary payoff $h\left(\frac{1}{n}\sum_{i \in I}\mathbf{u}_i\right)$. Specifically, all our results (Theorem 1-4) hold if each agent is more risk averse than the principal (each h_i is more concave than $h\,(x)$).

Next, consider a variant of the basic model in which agents are only experts that give advices (but do not take actions). That is, at stage 2 each agent recommends an action (conventional or unconventional), and the principal chooses the profile of actions $(a_i)_{i \in N}$ based on these recommendations. That is, each agent i is an informed expert, who advises the principal what to do in his area of expertise. Each expert's payoff remains the same: a high payoff if the recommended action is successful, and a low payoff otherwise.

If all agents are calibrated ($g\,(p) = p$), then too many of them would advise the principal to take the conventional action (all experts i with $\mathbf{p}_i < \mathbf{q}$). The principal can gain higher payoff relative to the basic model, by violating some of these recommendations. However, his inability to separate agents with inaccurate private signals ($\mathbf{p}_i$ is substantially smaller than $\mathbf{q}$) from agents with relative accurate private signals limits his payoff.

Observe that this variant yields the same asymptotically-optimal bias function g^* as the basic model. This is because agents that follow g^* induce the principal's first-best payoff. Such agents behave as if they have the same utility as the principal. Thus, the principal will always choose to follow the recommendations of such g^*-biased experts.

H Unbiased Random Errors and Underconfidence for Easy Tasks

A main experimental finding about overconfidence (*hard-easy effect*) is that the degree of overconfidence is increasing in the difficulty of the task, and in particular, there is moderate underconfidence for easy tasks. At first glance it seems that our model gives a contradicting prediction: people present overconfidence for all tasks, including easy ones in which p_i is close to 1. In this section, we sketch how incorporating unbiased random errors into our model would imply an induced behavior which is consistent with this experimental finding (following similar analysis in the existing literature - see e.g. Erev, Wallsten and Budescu, 1994; Budescu, Wallsten, and Au, 1997).

Let π be the true difficulty level of a task. That is, when similar agents use their judgment to (independently) solve the task, they succeed with frequency π. It is reasonable to assume that an agent can only observe a noisy signal about the difficulty level: $s = \pi + \mathbf{e}$ where $\mathbf{e}$ is an unbiased random error ($E(\mathbf{e}) = 0$). In such a setup one should interpret p (i.e., p_i in our model) as the best (unbiased) estimation of the unknown difficulty level given the noisy signal: $p = E(\pi|s)$. The statistical phenomenon of *regression to the mean* implies that if the task is easy (π is close to 1) then the best estimation (p) is smaller than the true difficulty (π) due to: (1) boundary effects (π is bounded above by 1), and (2) assuming that the ex-ante distribution of π is single-peaked, a high signal s is more likely to include positive noise ($e > 0$). In a typical experiment the researcher can only observe π and $g(p(\pi))$ (agent's evaluation of his success probability) but not $p(\pi)$. Thus even when the true bias is always overconfidence ($g(p) > p$ for every p), then it is plausible that for easy tasks subjects in experiments would seem to be underconfident ($g(p(\pi)) < \pi$).

I Overconfidence and Social Welfare

Bénabou and Tirole (2002, Part I) survey various evidence that "in most societies, self-confidence is widely regarded as a a valuable individual asset", and raise the question: "why is a positive view of oneself, as opposed to a fully accurate one, see as such a good thing to have?". In what follows we briefly describe a variant of our model that implies a novel answer to their question: overconfidence is socially desirable because it increases social welfare.

Consider a society, where each agent i may either make a conventional choice or an unconventional one when choosing his productive activities. This decision influences agent i's productivity $\mathbf{x}_i$, which may be either high or low. The payoff of each agent is a function of his output $\mathbf{x}_i$ and the total output $\sum_j \mathbf{x}_j$: $u_i = h\left(\mathbf{x}_i, \sum_j \mathbf{x}_j\right)$. Function h is assumed to be strictly increasing and concave in both parameters. For example, this is the case if a fixed amount of each agent's output is taxed and is being used to produce a public good. Alternatively, it might be that the output of each agent has a direct positive externality on other agents.

Calibrated agents (without confidence-bias) would be conventional too often, and obtain an inefficient outcome, in which the variance of the total productivity $\sum_j \mathbf{x}_j$ is too high. A utilitarian social planner would act as if it were a risk-averse principal in our model. Such a planner would like to induce social norms in favor of moderate overconfidence, and as a result, moderate overconfidence in one's ability, would be socially regarded as a valuable asset.