

Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism

Philippe Jehiel

► **To cite this version:**

Philippe Jehiel. Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism. American Economic Review, American Economic Association, 2018, 108 (6), pp.1582-1597. halshs-01884380

HAL Id: halshs-01884380

<https://halshs.archives-ouvertes.fr/halshs-01884380>

Submitted on 15 Jul 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Investment strategy and selection bias: An equilibrium perspective on overoptimism*

Philippe Jehiel[†]

27th November 2017

Abstract

Investors implement projects based on idiosyncratic signal observations, without knowing how signals and returns are jointly distributed. The following heuristic is studied: Investors collect information on previously implemented projects with the same signal realization, and invest if the associated mean return exceeds the cost. The corresponding steady states result in suboptimal investments, due to selection bias and the heterogeneity of signals across investors. When higher signals are associated with higher returns, investors are overoptimistic, resulting in overinvestment. Rational investors increase the overoptimism of sampling investors, thereby illustrating a negative externality imposed by rational investors.

*Thanks to Roland Bénabou, Milo Bianchi, Jerker Denrell, Ignacio Esponda, David Hirshleifer, Charles Manski, Stephen Morris, Rani Spiegler, four anonymous reviewers, and seminar participants at the UCL behavioral game theory workshop, PSE, Princeton student workshop, Game theory French symposium, ESSET 2015, Bristol Game Theory workshop 2016, D-tea 2016 where preliminary versions of the ideas expressed here were presented.

[†]PSE, 48 boulevard Jourdan, 75014 Paris, France and University College London ; jehiel@enpc.fr

1 Introduction

A key aspect of entrepreneurial activity consists in deciding whether to make investments based on the observation of signals that can be thought of as investors' (initial) perceptions about the projects. In a Bayesian framework, the investor would know how signals, returns and costs are jointly distributed, and he would make the optimal investment decision using the standard Bayesian updating machinery. But, many investors (in particular those who are less experienced) would not know the joint distribution. It seems then natural that such investors would make their decisions using the data set they have access to, which I assume consists of return and cost data from previously implemented projects, as well as the perceptions (signals) that they get from these projects.¹ Specifically, I will be assuming that such investors use the following heuristic: They collect information on previously implemented projects delivering the same perception (signal) as in their own project, and they invest if the associated empirical mean return exceeds the cost. I study the steady states -referred to as equilibria-, generated by such a heuristic, in a simple model in which projects have homogeneous costs but heterogeneous returns, and I first consider the case in which all investors use this heuristic while allowing later for a mix with rational investors. A key observation is that assuming that signals are idiosyncratic across investors and that higher signal realizations are associated with higher returns, sampling investors are, in equilibrium, overly optimistic about how the mean return depends on the signal, thereby leading to systematic overinvestment as compared with the rational benchmark.

The overinvestment bias derived in the equilibrium with sampling investors is related to the selection bias implicit in the proposed heuristic, given that the samples considered by investors consist only of those projects that were previously implemented and not all projects. Had the investors also been able to collect data for non-implemented projects, the heuristic would have led to the correct assessments, and their investment decisions would have been optimal. The heuristic followed by the sampling investors can be viewed as reflecting a form of selection neglect given that it assumes investors do not correct for selection bias (as econometricians would do in the tradition of Heckman (1979)).²

¹Non-implemented projects are hard to access in part because such projects are generally not even recorded. In the analysis, it will be assumed that the signals generated about past projects or for current projects are governed by the same statistical distributions.

²Selection neglect as considered here has received experimental support when the sampling bias is not salient (Koehler and Mercer (2009) or when it is hard to adjust for it (Feiler et al., 2013) as I would argue is the case in many applications (see Manski (2004) on this). Camerer and Lovallo (1999) in their classic experiment on overconfidence in contests discuss reference group neglect according to which subjects fail to adjust their entry decision to the information that their competitors self-selected to skill-based

While several previous studies have noted the potential link between selection neglect and managerial decision biases (see in particular the survey by Denrell (2017)), a distinctive feature of the present approach is the equilibrium perspective on how the biased samples used by investors are formed: For any tentative distribution of implemented projects, the proposed heuristic pins down the investment strategy of sampling investors, which in turn gives rise to a distribution of implemented projects. In a steady state, the former and the latter distributions of implemented projects should coincide, thereby allowing for the study of the long run effects of the sampling heuristic.

The analysis is fairly simple under the monotone likelihood property (MLRP) assumption that requires a higher signal to be more representative of a higher return. In equilibrium investors use a cutoff rule, investing only when the signal they receive is above some threshold a^S . A sampling investor with signal a^S looks for past projects for which he gets the same signal a^S . Any such project he finds has the property that the investor whose project it was received a signal at least a^S (fixed point). In equilibrium, the average return of such projects (with one signal above a^S and another at a^S) must coincide with the cost. It exceeds the expected return of a project conditional upon having one signal at a^S (under MLRP) so that the investor overestimates the value of his project, leading to overinvestment.

While the derivations of the overoptimism and overinvestment biases seem intuitive, several comments are in order. First, there would be no bias if the signals received by all investors were perfectly correlated (conditional on the returns). Second, the overoptimism and overinvestment biases need not hold if MLRP were violated. These two observations illustrate that the overoptimism bias obtained here is not solely driven by the assumption that only previously implemented projects are accessible, as it relies also on the dispersion of signals across investors (a fairly natural and standard assumption in applied work) and the MLRP assumption, as well as on selection neglect. Third, the overinvestment found in equilibrium is less severe than the one that would arise if investors were sampling from projects decided by rational investors, the reason being that (under MLRP) the bias is all the more severe that the criterion used by others is more conservative. In other words, the equilibrium force here dampens the overinvestment bias without eliminating it. Fourth, under natural extra assumptions beyond MLRP, it turns out that the overoptimism bias is more pronounced for intermediate realizations of the signals, and that the welfare loss induced by the excessive investment is biggest for intermediate levels of informativeness of the signals. Finally, allowing for a mix of rational and sampling investors, I note that

contests. See also Esponda and Vespa (2017) and Enke (2017) for other recent experimental accounts of selection neglect.

the overoptimism and overinvestment biases of the sampling investors are more severe when the share of rational investors is greater, thereby illustrating a negative externality that rational investors impose on sampling investors.

An illustrative example: To illustrate the main findings, think of investors as having to decide whether or not to open a business. Businesses can be of two types. They are either lucrative, leading to profit \bar{x} , or poor, leading to profit \underline{x} , and each type of business is equally likely. The initial fixed cost c of opening a business is assumed to lie in between the two profit levels \underline{x} and \bar{x} . Before making their decisions, investors observe some signal about the type of their business, say about some characteristics of it. They also observe similar signals for all previously started businesses, and, conditional on the type of business, signals are independently distributed across investors (different investors focus on different characteristics). The signal realization can either be *Good*, *Medium* or *Bad* with a probability that depends on whether the business is lucrative or poor. When it is lucrative (resp. poor), the investor gets a signal that is either *Good* (resp. *Bad*) or *Medium* each with probability half. Thus, when the signal is *Good*, it is optimal to open the business, since a *Good* signal can only come from a lucrative business. Similarly, when the signal is *Bad*, it is optimal to not open the business, since a *Bad* signal can only come from a poor business. Assuming that $c > \frac{\underline{x} + \bar{x}}{2}$, it is optimal to not open the business when the signal is *Medium*, since given the symmetry of the problem, Bayesian updating would then tell the investor that the two profit levels \underline{x} and \bar{x} are equally likely.

Consider a sampling investor who would observe in his pool only businesses handled by rational investors. Since rational investors open their businesses only when their signal is *Good*, the pool of implemented businesses would all be lucrative. A sampling investor looking at such businesses would receive the signal *Medium* for half of them (remember signals are idiosyncratic). Accordingly, he would open the business upon receiving signal *Medium* for his own business, given that all the implemented businesses for which he gets the same signal *Medium* are lucrative. In the equilibrium with sampling investors only, an investor opens his business more often than in the rational case, but potentially less often than a sampling investor would do when surrounded with rational investors only. The reason why the investment decisions of sampling investors may be altered is that the presence of sampling investors results in the presence of poor businesses in the pool of implemented projects, and such a compositional effect reduces the pro-investment bias, even if it does not eliminate it, as implied by the main result of the paper. More precisely, within the proposed example, in the equilibrium with sampling investors only, when $c < \frac{2\bar{x} + \underline{x}}{3}$ investors open their businesses when they get signals *Good* or *Medium*, but

when $c > \frac{2\bar{x}+x}{3}$, only a fraction of businesses associated to signal *Medium* is implemented by sampling investors, and the perceived expected profit associated to that signal coincides exactly with the cost c in equilibrium.

Related literature: This paper is connected to several literatures. As already mentioned, previous studies have related the idea of sampling biases to decision biases. This includes the survivorship bias studied in the context of risk assessments by Denrell (2003), according to which failed projects are under-sampled. It also includes the mirror image upward censored sampling bias studied by Streufert (2000) in the context of assessing returns to schooling, according to which successful children are under-sampled in poor neighborhoods, as they tend to move to better locations. A key difference is that these papers unlike this paper do not consider an equilibrium approach to the bias (the sampled pool is not endogenously determined by the decisions of economic agents). Moreover, these studies do not allow for the possibility that the decision maker would receive private information before making his decision, and therefore they cannot compare the heuristic assessment with the assessment resulting from Bayesian updating. One can, of course, combine these various sampling biases depending on the application one has in mind (schooling vs managerial decisions), and analyze them in an equilibrium fashion as in this paper. While the survivorship bias would be expected to increase the overoptimism bias, the upward censorship bias could lead to a pessimism bias.

The literature on overconfidence has documented that entrepreneurs tend to be overly optimistic about their projects (see for example Cooper et al. (1988) or Malmendier and Tate (2005)), which has generally been used to justify that investors rely on subjective priors or attach excessive precision to the signals they receive (see for example Xiong (2013) or Daniel and Hirshleifer (2015) for such a use in finance models). This paper offers a different perspective suggesting that the overoptimism may be related to how informative the objective signals are, and also how experienced the surrounding investors are (where a more experienced investor is viewed as being rational in the present model).³

Finally, the equilibrium perspective of this paper can be viewed as belonging to a growing literature in behavioral game theory that has developed various solution concepts with mistaken expectations. These include the analogy-based expectation equilibrium

³Theoretical approaches to overconfidence that complement the one discussed in this paper include: 1) Rabin and Schrag (1999) who derive overconfidence from another psychological bias, the confirmation bias that leads agents to sometimes behave as if they had not made observations that go against their current beliefs, 2) Van den Steen (2004) who defines overconfidence as the subjective belief that one performs better than others, which Van den Steen derives from a revealed preference argument in a subjective prior world, and 3) several studies that derive overconfidence from motivated cognition purposes which include Bénabou and Tirole (2002), Köszegi (2006) or Bénabou (2015).

(Jehiel, 2005) to which the present study can be connected (see the WP version Jehiel, 2017),⁴ the cursed equilibrium (Eyster and Rabin, 2005), and the behavioral equilibrium (Esponda, 2008). As an approach related to selection neglect, this paper is closest to Esponda (2008), Esponda and Pouzo (2017), and Esponda and Vespa (2017) in that the biases arising in these papers, as well as in this paper, are due to the missing feedback on non-implemented projects or transactions. The environment of this paper and thus the mechanism leading to the resulting overoptimism bias are however different from these three models in that in the present study, the investor has to make an inference about what his observed signal implies for profitability rather than an inference from what others' actions imply about their signal (as in the adverse selection models considered by Esponda (2008) or also as in social learning environments) or an inference about the implications of his actions conditional on being pivotal (as in voting environments of the type considered in Esponda and Vespa, 2017).

The rest of the paper is organized as follows. Section 2 presents the investment problem. Section 3 analyzes the overoptimism and overinvestment biases arising in the equilibrium with sampling investors. It also discusses the effect of having a mix of rational and sampling investors. Section 4 concludes.

2 The Investment Problem

A large number of investors idealized as a continuum is considered. Each investor, assumed to be risk neutral, has to decide whether or not to invest in one project that is different for each investor. The cost of every project is c . The return of a project is random and can take values x in a set $X \subset \mathbb{R}$ (assumed to consist of finitely many values to avoid technical complications). Before making his decision, an investor knows the cost c but does not know the return realization x of his project. However, he observes a signal realization a for his project which can be thought of as representing his overall perception about the project. It takes values in (\underline{a}, \bar{a}) with $\underline{a} < \bar{a}$ (where I allow that $\underline{a} = -\infty$ and $\bar{a} = +\infty$).⁵ Based on a , the investor has to decide whether or not to invest. If the investor decides to invest, the project is implemented, it is observable by everyone, and after the implementation of the project takes place, the resulting return x is assumed to be publicly observable. Non-implemented projects are not observed. When an investor observes a previously implemented project, he can freely generate a signal that stands again for his

⁴See Spiegler (2017) for suggesting another link of this paper to Bayesian networks.

⁵In the sequel, for any continuous function $h(\cdot)$, I will refer to $\lim_{a \rightarrow \bar{a}} h(a)$ (resp. $\lim_{a \rightarrow \underline{a}} h(a)$) as $h(\bar{a})$ (resp. $h(\underline{a})$).

perception of that previously implemented project. That is, for every past implemented project, investor i observes both the return realization x and a signal (perception) a_i related to that project.

Returns and signals are generated similarly for all projects and for all investors. Importantly, I assume that for any project whose return realization happens to be x , the signals received by two different investors are independent draws from the same distribution (that typically depends on x). This is a simple and standard way of modelling the heterogeneity of observations among investors while allowing the signals to be informative about the return. For concreteness, one may think of the signal as being the sum of the return and an investor-specific realization of a noise term (with mean 0). Specifically, for each project, the probability that the return realization turns out to be x is $l(x) \geq 0$ with $\sum_{x \in X} l(x) = 1$. Conditional on the return realization x of a project, the signal realization a_i observed by any investor i about this project is assumed to be distributed according to the density $f(\cdot | x)$, assumed to be smooth with full support on (\underline{a}, \bar{a}) , and two different investors i and i' get two independent draws $a_i, a_{i'}$ from this distribution. Assuming that the distribution of a takes the form of a density will simplify the exposition of the analysis, but it is not required (the example in introduction assumes a can take finitely many values).

Importantly, I have in mind that investors do not know how the signal realization a and the return realization x are jointly distributed, i.e., they do not know $l(\cdot)$ nor $f(\cdot | \cdot)$. If they did, investors could find out the optimal investment strategy which consists in investing upon observing a when the expected mean return conditional on a , $E(x | a)$, denoted by $v^R(a)$, is no smaller than c , and not investing when $v^R(a) < c$, where $E(x | a)$ is derived from $l(\cdot)$ and $f(\cdot | x)$ by Bayes' law, i.e.,

$$E(x | a) = \frac{\sum_{x \in X} l(x) f(a | x) \cdot x}{\sum_{x \in X} l(x) f(a | x)}. \quad (1)$$

In the following, I will assume that the rational strategy requires that, for some signal realizations a , it is best to invest. That is, $\sup_{a \in (\underline{a}, \bar{a})} v^R(a) > c$.

Without the knowledge of how signals and returns are distributed, I assume that investors use the following heuristic, based on the data set consisting of all past implemented projects available to them. When getting a signal realization a for his current project, the investor collects information on all implemented projects in the past for which he gets the same signal realization a . Then he computes the empirical mean return in those projects (this only requires averaging the x observed in those projects for which this in-

vestor gets the same signal realization a), and he invests whenever the obtained empirical mean return is above the cost c , and he does not invest otherwise. I will consider the steady states of such a dynamic system, assuming that all investors follow the sampling heuristic while allowing later for the study of a mix with rational investors. I will refer to the resulting investment strategies as equilibria with sampling investors. In order to rule out trivial situations in which there would be no investment at all, I will also assume that whatever the observed signal there is a tiny probability (assumed to be the same for all signal realizations) that the decision maker invests, and I will let this probability tend to 0 in the analysis.

Formally, let $q(a)$ denote the (steady state) probability with which an investor observing a invests, and assume q is bounded away from 0 for some positive measure of signals. The probability of observing an implemented project with return x conditional on the signal a being in $A \subseteq (\underline{a}, \bar{a})$ is

$$\Pr(x \mid a \in A, \text{implemented}; q) = \frac{l(x) \Pr(a \in A, \text{implemented} \mid x; q)}{\sum_{x' \in X} l(x') \Pr(a \in A, \text{implemented} \mid x'; q)}$$

where

$$\Pr(a \in A, x' \mid \text{implemented}; q) = \Pr(a \in A \mid x') \int_{\underline{a}}^{\bar{a}} q(b) f(b \mid x') db,$$

given that a randomly drawn project with return realization x generates a signal $a \in A$ with probability $\Pr(a \in A \mid x)$ for the agent looking at such a project, and it is implemented with probability $q(b)$ by the agent in charge of this project whenever the latter agent receives signal b (which occurs according to the density $f(b \mid x)$).

Thus, the investor's perceived expected return conditional on a is

$$\hat{v}(a; q) = \frac{\sum_{x \in X} l(x) f(a \mid x) \int_{\underline{a}}^{\bar{a}} q(b) f(b \mid x) db \cdot x}{\sum_{x \in X} l(x) f(a \mid x) \int_{\underline{a}}^{\bar{a}} q(b) f(b \mid x) db}, \quad (2)$$

which results from the induced proportion of projects with return x in the pool of implemented projects associated to signal a . The following defines an equilibrium, making precise that the function $\hat{v}(a; q)$ is pinned down by the uniform trembling assumption in case there would be no investment.

Definition 1 *An investment strategy $q(\cdot)$ over (\underline{a}, \bar{a}) is an equilibrium with sampling investors if $\hat{v}(a; q) > c$ implies $q(a) = 1$, and $\hat{v}(a; q) < c$ implies $q(a) = 0$, where $\hat{v}(a; q)$ is*

defined to be $\lim_{n \rightarrow \infty} \widehat{v}(a; (1 - \frac{1}{n})q + \frac{1}{n})$.

The strategy just defined is the result of a fixed point. The probabilities $q(b)$ with which other investors choose to invest when getting signal b affect the subjective assessments $\widehat{v}(a; q)$, which in turn determine the probability with which an investor getting signal a is willing to invest. In equilibrium, these two probability mappings should be the same.

Comment. Given a project with return x , different investors were assumed to observe independent draws from $f(\cdot | x)$. If instead these draws were the same, $q(a)$ should replace $\int_{\underline{a}}^{\bar{a}} q(b)f(b | x)db$ in expression (2), and the assessment arising from the heuristic would boil down to the rational assessment (1) in the steady state when $q(a) > 0$.⁶ The decision bias derived below thus crucially relies on the idiosyncratic nature of the signals received by investors.

3 Overoptimism as a result of selection neglect

I analyze the above investment environment assuming that a higher signal realization is more representative of a higher return.

Assumption (MLRP): For any $a' > a$ and $x' > x$, it holds that: $\frac{f(a'|x')}{f(a|x')} > \frac{f(a'|x)}{f(a|x)}$.⁷

The rational assessment of the mean return as a function of a , $v^R(a)$, is given by (1), and under MLRP it is readily verified that it is an increasing function of a . Accordingly, let $a^R \in (\underline{a}, \bar{a})$ be uniquely defined by⁸

$$v^R(a^R) = c \text{ if } a^R > \underline{a} \text{ (and } v^R(a^R) \geq c \text{ if } a^R = \underline{a}).$$

A rational investor invests when $a > a^R$, and he does not when $a < a^R$.

To present the analysis of the equilibrium with sampling investors, it is convenient to introduce the function

$$H(a, z) = \frac{\sum_{x \in X} l(x)f(a | x)[1 - F(z | x)] \cdot x}{\sum_{x \in X} l(x)f(a | x)[1 - F(z | x)]} \quad (3)$$

⁶The trembling hand refinement implicit in Definition 1 would then guarantee that there is no bias, even if $q(a) = 0$.

⁷Assuming $f(\cdot | x)$ is smooth, this can be formulated as requiring that $\frac{\partial f(a|x)/\partial a}{f(a|x)}$ is increasing in x .

⁸The fact that $a^R < \bar{a}$ comes from the assumption that there is investment with positive probability in the optimal solution.

where $F(\cdot | x)$ denotes the cumulative of $f(\cdot | x)$. $H(a, z)$ is the expected value of the return x conditional on drawing one signal equal to a and one signal above z . Note that $v^R(a) = H(a, \underline{a})$ given that $F(\underline{a} | x) = 0$ for all x . Moreover, as shown in the Appendix:

Lemma 1 *Under MLRP, $H(\cdot, \cdot)$ is increasing in a and z .*

3.1 Equilibrium with sampling investors

The following Proposition shows that under MLRP there is a unique equilibrium with sampling investors. Letting q^S denote this equilibrium and letting $v^S(a)$ denote the equilibrium subjective assessment $\hat{v}(a; q^S)$ of a sampling investor getting signal a , one can state:

Proposition 1 *Under MLRP, there exists a unique equilibrium q^S with sampling investors. The equilibrium is such that for some threshold a^S , an investor chooses to invest if his observed signal realization a is above a^S and to not invest otherwise, where a^S is uniquely defined by*

$$H(a^S, a^S) = c \text{ if } a^S \in (\underline{a}, \bar{a}) \text{ (and } H(a^S, a^S) \geq c \text{ if } a^S = \underline{a}). \quad (4)$$

In the equilibrium with sampling investors, there is more investment than in the rational case, i.e., $a^S \leq a^R$. Moreover, sampling investors are overoptimistic, in the sense that $v^S(a) \geq v^R(a)$ for all a .

Proof of Proposition 1: Suppose that, in equilibrium, investment occurs with probability $q(a)$ when $a \in (\underline{a}, \bar{a})$ is observed. The perceived expected return $\hat{v}(a; q)$ of a project with signal realization a would then be given by (2). Since $a \mapsto \hat{v}(a; q)$ is increasing (for any function $q(\cdot)$) by MLRP, one can infer that investors must follow a threshold strategy, i.e. for some z , invest if $a > z$ and do not invest if $a < z$ where z (if interior) is defined by $\hat{v}(z; q) = c$. This also ensures that the subjective assessment $\hat{v}(a; q)$ takes the form of $H(a, z)$, for some z where z if interior must satisfy $H(z, z) = c$.

More precisely, given the assumption that there is some investment in the rational case, it follows that $H(\bar{a}, \underline{a}) > c$. An equilibrium with sampling investors must employ a threshold strategy z where the threshold z must satisfy

$$H(z, z) = c \text{ if } z \in (\underline{a}, \bar{a}) \text{ (} H(z, z) \geq c \text{ if } z = \underline{a} \text{ and } H(z, z) \leq c \text{ if } z = \bar{a}, \text{ respectively)}. \quad (5)$$

Given that $H(\bar{a}, \underline{a}) > c$, the monotonicity of $H(\cdot, \cdot)$ in the second argument implies that $H(\bar{a}, \bar{a}) > c$, and thus the latter case can be ignored. Suppose then that $H(\underline{a}, \underline{a}) < c$. The continuity of H ensures that there exists $z \in (\underline{a}, \bar{a})$ satisfying $H(z, z) = c$. Hence, there must exist $z > \bar{a}$ satisfying (5). Consider now $\bar{a} \geq z_1 > z_2 \geq \underline{a}$. Clearly, $H(z_1, z_1) > H(z_2, z_2)$ and (5) cannot be simultaneously satisfied for $z = z_1$ and z_2 . One concludes that there is only one equilibrium with positive investment, and that this equilibrium is a threshold equilibrium a^S where a^S is uniquely defined to satisfy (4). Observe also that there cannot be an equilibrium with $q(\cdot) \equiv 0$ because then $\lim_{n \rightarrow \infty} \hat{v}(a; (1 - \frac{1}{n})q + \frac{1}{n}) = H(a, \underline{a})$ and we have assumed that $\sup_a H(a, \underline{a}) = H(\bar{a}, \underline{a}) > c$.

Regarding the overinvestment bias, observe that when $a^R < \bar{a}$, one has $H(a^R, \underline{a}) \geq c$ and thus $H(a^R, a^R) \geq c$ by the monotonicity of H in its second argument. This, in turn, establishes the overinvestment bias $a^S \leq a^R$ using the monotonicity of $a \mapsto H(a, a)$.

Regarding overoptimism, the above analysis establishes that $v^S(a) = H(a, a^S)$. Given that $v^R(a) = H(a, \underline{a})$, the monotonicity of H in z (see Lemma 1) establishes that $v^S(a) \geq v^R(a)$ for all a , as required. **Q.E.D.**

Comment. While the MLRP assumption is natural and common (it is without loss of generality up to reordering of signals a if x can take only two values, and it holds whenever a is a noisy signal about x for many specifications of the noise distribution), one can show that the overoptimism and overinvestment biases need not hold without the MLRP assumption. This is so because ordering signals in an increase fashion according to the induced expected return (as derived from the Bayesian formula) and truncating the overall distribution of projects by censoring projects for which one drawn signal is below a threshold (in this ranking) no longer induces a distribution of returns that first-order stochastically dominates the overall distribution, thereby invalidating the argument used to show Proposition 1 (see the Appendix for an illustration of equilibrium underinvestment without MLRP).

Building on Proposition 1, it may be interesting to explore (under MLRP) how the overoptimism bias changes with the signal realization, and how the welfare loss derived in the equilibrium with sampling investors is affected by the informativeness of the signals. To this end, consider the following simplified setup. The cost c is normalized to 0. Return x can take two values, $\underline{x} = -1$ or $\bar{x} = 1$, with the same probability ($l(x) = 1/2$ for $x = -1$ and 1). Conditional on x , the signal a takes the form $a = x + \varepsilon$, where ε is the realization of a random variable that is symmetrically distributed around 0 (hence with mean 0) and such that the MLRP condition is satisfied. The density and cumulative of ε are denoted by $g_\sigma(\cdot)$ and $G_\sigma(\cdot)$ respectively, where σ denotes the variance that will be varied later on.

I also assume that $\frac{g_\sigma(a-\bar{x})}{g_\sigma(a-x)}$ approaches 0 (resp. $+\infty$) as a approaches \underline{a} (resp. \bar{a}). All these assumptions are satisfied when ε is drawn from a normal distribution with mean 0 and variance σ .

In such a setting, the rational strategy requires, given the symmetry, to invest if $a > 0$ and to not invest otherwise. That is, $a^R = 0$. The threshold signal a^S that arises in the equilibrium with sampling investors (see (4)) is characterized by

$$g_\sigma(a^S + 1)(1 - G_\sigma(a^S + 1)) = g_\sigma(a^S - 1)(1 - G_\sigma(a^S - 1)),$$

given that at signal realization a^S , the frequency of observed successful outcomes (that is proportional to $g_\sigma(a^S - 1)(1 - G_\sigma(a^S - 1))$) is equal to the frequency of observed failed outcomes (that is proportional to $g_\sigma(a^S + 1)(1 - G_\sigma(a^S + 1))$). As Proposition 1 implies, we know that $a^S < 0$. Define the overoptimism bias as the difference $v^S(a) - v^R(a)$ between the subjective assessment of the mean return in the sampling equilibrium and the rational assessment as a function of a . This bias is known to be positive by Proposition 1. As it turns out, the bias gets close to 0 for a close to \underline{a} or \bar{a} , and this follows because signals a close to \underline{a} (resp. \bar{a}) are very informative of x being \underline{x} (resp. \bar{x}) whenever $\frac{g_\sigma(a-\bar{x})}{g_\sigma(a-x)}$ approaches 0 (resp. $+\infty$), thereby making negligible for such signals the sampling bias that gives rise to overoptimism. This, in turn, implies that the overoptimism bias is biggest for intermediate realizations of a .

Noting that $\frac{1}{2}(g_\sigma(a-1) - g_\sigma(a+1))$ represents the expected value of the return conditional on a times the density of a , and letting $a^S(\sigma)$ represent the equilibrium threshold as a function of σ , welfare loss as a function of σ is given by:

$$WL(\sigma) = \int_{a^S(\sigma)}^0 \frac{1}{2}(g_\sigma(a+1) - g_\sigma(a-1))da,$$

which corresponds to the aggregate loss due to the suboptimal implementation of projects with signal realizations a falling in between the sampling equilibrium threshold $a^S(\sigma)$ and the rational threshold 0. Interestingly, it can be shown that $WL(\sigma)$ approaches 0 either when the signal is very informative (i.e., σ approaches 0) or when it is very uninformative (σ approaches $+\infty$). When signals are very informative, it is very likely for investor i to observe a signal a_i close to -1 whenever the return realization is $\underline{x} = -1$ and a signal a_i close to 1 whenever the return realization is $\bar{x} = 1$. This implies that the signal observations made by two different investors about the same project are very likely to be close to each other when signals are very informative, thereby explaining why the decision bias becomes small in this case. When signals are almost uninformative, investing or

not in the simplified setup is almost equally good, and thus the welfare loss can only be small.⁹ It follows that the welfare loss is biggest for intermediate levels of informativeness of the signals (under the normal distribution case, simulations show that $WL(\sigma)$ is single-peaked, see Jehiel, 2017).

Comment. As just illustrated, the degree of overoptimism arising in the equilibrium with sampling investors and the welfare consequences of it depend on the informativeness of signals. Such a dependence would not necessarily arise in the subjective prior approach to overoptimism, which typically puts no structure on how investors form their subjective prior and thus on how overoptimism varies with the primitives of the model.¹⁰

3.2 When rational investors exert negative externalities

Suppose the population of investors is mixed. A share $1 - \lambda$ of investors (referred to the sampling investors) proceeds as described in the main model: They observe a signal realization a for their project, sample all implemented projects in which they get the same signal realization a , and invest if the observed empirical mean return exceeds the cost c . A share λ of investors (referred to as rational investors) makes the optimal investment decision based on the observation of the signal realization. Signals and returns are distributed as in the main model.

Following the same logic as above, it is readily verified that sampling investors follow in equilibrium a threshold strategy that consists in investing in a project with signal realization a only if a exceeds a^* where a^* (when interior) is defined by¹¹

$$\frac{\sum_{x \in X} f(a^* | x)[(1 - \lambda)(1 - F(a^* | x)) + \lambda(1 - F(a^R | x))]l(x) \cdot x}{\sum_{x \in X} f(a^* | x)[(1 - \lambda)(1 - F(a^* | x)) + \lambda(1 - F(a^R | x))]l(x)} = c \text{ if } a^* \in (\underline{a}, \bar{a}) \quad (6)$$

To understand this expression, note that when a sampling investor makes an observation of another project, with probability λ she is facing a rational investor who invests

⁹The insight that with uninformative signals, the sampling equilibrium is efficient is more general though, since then there is no heterogeneity in the signals observed by investors.

¹⁰It may be mentioned here that within the subjective prior paradigm, Van den Steen (2004) makes the interesting and simple observation that as a consequence of a revealed preference argument, no matter how subjective priors are modelled, others' decisions always look (weakly) suboptimal from the subjective viewpoint of any agent, thereby leading to the systematic subjective belief that one performs better than others. Such a relative overoptimism bias would also arise in my setting. If one were to ask any given investor at an ex ante stage (i.e., before he receives a signal realization for his project) whether he thinks he would perform better than other investors, he would be affirmative. Relative overoptimism arises here for the very same reason highlighted by Van den Steen that the investor believes (based on his subjective assessment of the mapping between his signal and the investment decision) that he can screen projects better than others. See footnote 18 of Jehiel (2017) for elaborations on this.

¹¹When the left-hand side is no smaller than c at $a^* = \underline{a}$, then $a^* = \underline{a}$.

only if the signal realization observed by this investor is larger than the rational threshold a^R , and with probability $1 - \lambda$ she is facing another sampling investor who invests if the signal realization he observes is larger than a^* .¹² The left hand-side of (6) represents how a sampling investor subjectively assesses the mean return of a project with signal realization a^* , and it requires in equilibrium that if a^* is interior, this perceived mean return should be equal to the cost c .

Denote the above threshold a^* (that can be shown to be unique) by $a^S(\lambda)$. One has previously seen that when there are no rational investors ($\lambda = 0$), it holds that $a^S(0) \leq a^R$. The effect of λ on $a^S(\lambda)$ is unambiguously given by:

Proposition 2 *Under MLRP, the higher the share λ of rational investors, the more severe the pro-investment bias of sampling investors. That is, $a^S(\lambda)$ is weakly decreasing in λ , and for all λ , $a^S(\lambda) \leq a^S(0) \leq a^R$.*

The intuition behind Proposition 2, whose detailed proof appears in the Appendix, is simple. If an investor is surrounded with more rational decision makers, the decisions made by others are better, and thus when sampling from these to form an assessment regarding the profitability of the project it appears to the investor that the project is even more profitable. The selection bias is more severe, which leads the sampling investor to make a poorer decision. In other words, rational investors exert a negative externality on those investors who follow the sampling heuristic.

It is natural to consider the effect of an increase in λ on welfare. Given Proposition 2, an increase in λ deteriorates the welfare of sampling investors, but at the same time it increases the share of rational investors whose welfare is larger. Aggregating these two effects leads to ambiguous comparative statics in general. When the share of rational investors is sufficiently large, an increase in λ always enhances expected welfare (essentially because there are too few sampling investors who suffer from the negative externality imposed by rational investors). When the share of rational investors is sufficiently far from 1, the negative effect on sampling investors of increasing λ may dominate for some distributional assumptions. In this case, an increase in the share of rational investors results in an overall negative impact on expected welfare.¹³

Comment. Lerner and Malmendier (2013) observed that a higher share of peers with pre-MBA entrepreneurial background leads to lower rates of entrepreneurship post-MBA. Viewing the pre-MBA entrepreneurial activity as being random in nature, if a sampling post-MBA student is exposed to more pre-MBA cases, he would be subject to a less severe

¹²It should be stressed here that the sampling heuristic does not require any knowledge about λ .

¹³To illustrate, consider a two return \underline{x} , \bar{x} scenario with $\underline{x} < c < \bar{x}$ and $l(\underline{x}) = l(\bar{x}) = 1/2$. Simple

pro-entrepreneurial bias according to a logic similar to that developed in Proposition 2, thereby suggesting a selection neglect interpretation to Lerner and Malmendier's finding.¹⁴

4 Conclusion

The stylized model presented establishes that the combination of selection neglect and idiosyncratic signals can explain the overoptimism and overconfidence biases observed in entrepreneurial decisions even in a steady state. It can serve as a starting point to address further questions. For example, while I have assumed the economy has reached a steady state, it may be worth exploring more explicitly some dynamics. In the online appendix, I sketch a dynamic model in which investors of generation t sample from projects implemented by generation $t - 1$, formalizing in an extreme way a recency bias in the sampling procedure, and I observe that the overinvestment and overoptimism biases carry over in this dynamic system whether or not there is convergence to a steady state. In the online appendix, I also briefly consider a mixed population model consisting of rational and sampling agents who have to decide whether or not to become entrepreneur (with different outside options for different agents), where sampling agents of generation t form their view about how good it is to be an entrepreneur by sampling projects (startups) implemented by generation $t - 1$. I observe that such a dynamic system may lead to cycling between low entrepreneurship periods when sampling from (endogenously) less rational cohorts of entrepreneurs and high entrepreneurship periods when sampling from (endogenously) more rational cohorts of entrepreneurs, in agreement with the intuition of Proposition 2. Such an approach may pave the way to a richer study of entry and exit of entrepreneurs,

calculations yield that $\frac{d(WL)}{d\lambda}$ can be written as $-\frac{A+B}{2}$ where

$$\begin{aligned} A &= \int_{a^S(\lambda)}^{a^R} (f(a | \underline{x})(c - \underline{x}) - f(a | \bar{x})(\bar{x} - c)) da, \\ B &= (1 - \lambda) \frac{da^S(\lambda)}{d\lambda} (f(a^S(\lambda) | \underline{x})(c - \underline{x}) - f(a^S(\lambda) | \bar{x})(\bar{x} - c)). \end{aligned}$$

A (resp B) is shown to be positive (resp. negative) using the MLRP property, $f(a^R | \underline{x})(c - \underline{x}) = f(a^R | \bar{x})(\bar{x} - c)$, $a^S(\lambda) < a^R$ and $\frac{da^S(\lambda)}{d\lambda} < 0$. When λ is close to 1, B becomes negligible and thus $\frac{d(WL)}{d\lambda} < 0$. When λ is away from 1, A can be made small relative to B by having a sufficiently small probability that signal realizations a fall in $(a^S(\lambda), a^R)$ (this is consistent with MLRP which only requires that $\frac{f(a|\bar{x})}{f(a|\underline{x})}$

is increasing in a , but puts no restriction on how likely the various a are, realizing that $\frac{da^S(\lambda)}{d\lambda}$ is not sensitive to the overall probability that a falls in $(a^S(\lambda), a^R)$ but to the density of a around $a^S(\lambda)$). In such cases, $\frac{d(WL)}{d\lambda} > 0$ holds.

¹⁴Investors choosing randomly can equivalently be viewed as using a threshold rule $z = \underline{a}$ or $z = \bar{a}$ with some exogenously given probability. Accordingly, when there are more of these investors, the resulting bias is less severe for sampling investors.

and whether cycles can be sustained in more complex environments, in particular allowing economic agents to be longer-lived.

Appendix

Proof of Lemma 1. The monotonicity in a is a direct consequence of the MLRP condition given that $[1 - F(z | x)]$ does not depend on a . The monotonicity in z follows from the observation that under MLRP, the hazard rate $\frac{f(z|x)}{1-F(z|x)}$ decreases with x , and thus $x \rightarrow \frac{\frac{\partial}{\partial z}[1-F(z|x)]}{1-F(z|x)} = \frac{-f(z|x)}{1-F(z|x)}$ increases with x .

To show that under MLRP, $x_1 > x_2 \implies \frac{f(z|x_1)}{1-F(z|x_1)} < \frac{f(z|x_2)}{1-F(z|x_2)}$, observe that for all $z' > z$, one has:

$$f(z' | x_1)f(z | x_2) > f(z | x_2)f(z' | x_2).$$

Integrating both sides in z' from $z' = z$ to $z' = \bar{a}$ yields $(1 - F(z | x_1))f(z | x_2) > f(z | x_2)(1 - F(z | x_1))$ or $\frac{f(z|x_1)}{1-F(z|x_1)} < \frac{f(z|x_2)}{1-F(z|x_2)}$, as required. **Q.E.D.**

Example 1 (*Underinvestment bias without MLRP*) Four equally likely returns $x = -2, -1, 1, 2$. Three signal realizations $a = a_1, a_2, a_3$. Zero cost c . The distribution of a given x is summarized in the following table in which the number at the intersection of the a_i row and the x column -referred to as $p_i(x)$ - is the probability that signal a_i is drawn conditional on the return realization being x .

	2	1	-1	-2
a_1	0.1	0.4	0.1	0.24
a_2	0.1	0.31	0.5	0
a_3	0.8	0.29	0.4	0.76

The rational investment strategy requires that there is investment when $a = a_1$ or a_2 but not when $a = a_3$. The equilibrium with sampling investors takes the form: Invest when observing a_1 , invest with probability $\mu < 1$ such that $\sum_x p_2(x)(p_1(x) + \mu p_2(x)) \cdot x = 0$ when observing a_2 , and not invest when observing a_3 , resulting in less investment than in the rational case.

Proof of Proposition 2. Define $H(a, z, \lambda) = \frac{\sum_{x \in X} f(a|x)[(1-\lambda)(1-F(z|x)) + \lambda(1-F(a^R|x))]l(x) \cdot x}{\sum_{x \in X} f(a|x)[(1-\lambda)(1-F(z|x)) + \lambda(1-F(a^R|x))]l(x)}$.

Step 1. Under MLRP, H is increasing in a and z . It is decreasing in λ for $z \leq a^R$.

Proof of step 1. H increasing in a follows directly from MLRP. H increasing in z follows from the observation that $\frac{f(z|x)}{(1-\lambda)(1-F(z|x)) + \lambda(1-F(a^R|x))}$ is decreasing in x , which is proven in the same way as the decreasing hazard rate property.

To see this, integrate $f(a_1 | x_1)f(a_0 | x_0) \geq f((a_0 | x_1)f(a_1 | x_1)$ (which holds for all $a_1 \geq a_0, x_1 \geq x_0$) in a_1 from a_0 to \bar{a} and multiply by $1 - \lambda$ and integrate in a_1 from a^B to \bar{a} and multiply by λ to obtain that

$$\frac{f(a | x_0)}{(1 - \lambda)(1 - F(a^B | x_0)) + \lambda(1 - F(a | x_0))} \geq \frac{f(a | x_1)}{(1 - \lambda)(1 - F(a^B | x_1)) + \lambda(1 - F(a | x_1))}$$

as required.

H decreasing in λ for $z \leq a^R$ follows because $\frac{F(a^R|x) - F(z|x)}{\lambda(F(a^R|x) - F(z|x)) + 1 - F(z|x)}$ is increasing in x for $z \leq a^R$, which follows because $\frac{1 - F(a^R|x)}{1 - F(z|x)}$ is decreasing in x (which follows from the fact MLRP implies the first order stochastic dominance property noting that $\frac{F(a|x)}{1 - F(z|x)}$ is the cumulative of F conditional on x and a being no smaller than z and that MLRP still holds when we truncate the support of a).

Step 2. Proving that $a^S(\lambda)$ is smaller than a^R follows by noting that $H(a^R, a^R, \lambda) \geq H(a^R, 0, 0)$. Proving that $a^S(\lambda)$ is decreasing follows by noting that for an interior solution

$$H(a^S(\lambda), a^S(\lambda), \lambda) = c,$$

and thus if $\lambda' > \lambda$, $H(a^S(\lambda), a^S(\lambda), \lambda') \leq c$ (by the monotonicity of H in λ), which implies that $a^S(\lambda') \leq a^S(\lambda)$ (by the monotonicity of H in a and z). **Q.E.D.**

References

- [1] Bénabou R. (2015): 'The economics of motivated beliefs,' *Revue d'économie politique* 125, 665-685.
- [2] Bénabou R. and J. Tirole (2002): 'Self-confidence and personal motivation,' *Quarterly Journal of Economics* 117, 871-915.
- [3] Camerer C. and D. Lovallo (1999): 'Overconfidence and excess entry: An experimental approach,' *American Economic Review* 89, 306-318.
- [4] Cooper A., C. Woo, and W. Dunkelberg (1988): 'Entrepreneurs' perceived chances of success,' *Journal of Business Venturing* 3, 97-108.
- [5] Daniel, K. and D. Hirshleifer (2015): 'Overconfident investors, predictable returns, and excessive trading,' *Journal of Economic Perspectives* 29, 61-88.
- [6] Denrell, J. (2003): 'Vicarious learning, undersampling of failure, and the myths of management,' *Organization Science* 14, 227-243.
- [7] Denrell, J. (2017): 'Sampling biases explain decision biases,' forthcoming in the Oxford Handbook of group and organizational learning.
- [8] Enke, B. (2017): 'What you see is all there is,' mimeo Harvard University.
- [9] Esponda, I. (2008): 'Behavioral equilibrium in economies with adverse selection,' *American Economic Review*, 98, 1269-1291.
- [10] Esponda, I and D. Pouzo (2017): 'Conditional retrospective voting in large elections,' *American Economic Journal: Microeconomics*, 9(2), 54-75.
- [11] Esponda, I and E. Vespa (2017): 'Endogenous sample selection: A laboratory study,' *Quantitative Economics* forthcoming.
- [12] Eyster E. and M. Rabin (2005): 'Cursed equilibrium,' *Econometrica*, 73, 1623-1672.
- [13] Feiler, D. C., Tong, J. D., and Larrick, R. P. (2013): 'Biased judgment in censored environments,' *Management Science*, 59(3): 573 – 591.
- [14] Heckman J. (1979): 'Sample selection bias as a specification error,' *Econometrica* 47, 153-161.

- [15] Jehiel P (2005): 'Analogy-based expectation equilibrium,' *Journal of Economic Theory* 123, 81-104.
- [16] Jehiel P (2017): 'Investment strategy and selection bias: An equilibrium perspective on overoptimism,' discussion paper PSE number 2017-29.
- [17] Koehler, J. J., and Mercer, M. (2009): 'Selection neglect in mutual fund advertisements,' *Management Science*, 55(7), 1107–1121.
- [18] Köszegi, B. (2006): 'Ego utility, overconfidence, and task choice,' *Journal of the European Economic Association* 4, 673-707.
- [19] Lerner, J. and U. Malmendier (2013): 'With a little help from my (random) friends: Success and failure in post-business school entrepreneurship,' *Review of Financial Studies* 26, 2411-52.
- [20] Malmendier U. and G. Tate (2005): 'CEO overconfidence and corporate investment,' *Journal of Finance*, 60, 2661-2700.
- [21] Manski C. (2004): 'Measuring expectations,' *Econometrica* 72, 1329-1376.
- [22] Rabin, M. and J. L. Schrag (1999): 'First impressions matter: A model of confirmatory bias,' *Quarterly Journal of Economics* 114, 37-82.
- [23] Spiegel R (2017): 'Data monkeys: A procedural model of extrapolation from partial statistics,' *Review of Economic Studies* forthcoming.
- [24] Streufert P. (2000): 'The effect of underclass social isolation on schooling choice,' *Journal of Public Economic Theory* 2, 461-482.
- [25] Van den Steen E. (2004): 'Rational overoptimism (and other biases),' *American Economic Review* 94, 1141-1151.
- [26] Xiong, X. (2013): 'Bubbles, crises, and heterogeneous beliefs,' *Handbook for Systemic Risk*, edited by J-P Fouque and J. Langsman, Cambridge University Press, 663-713.

ONLINE APPENDIX (NOT FOR PUBLICATION)

Convergence to steady state

The main model assumes steady state. Embedding the above framework into a dynamic setting in which new cohorts of investors sample from previous cohorts of investors naturally leads to asking when we should expect to see convergence to steady state as considered in the main analysis. Another legitimate concern is whether the overoptimism bias identified in the main analysis would still arise in case there would be no convergence. To model the dynamics most simply, consider within the MLRP scenario discussed in the main model a sequence of time periods $t = 1, 2, \dots$. Assume that in every period $t > 1$ there is a new cohort of investors of the same mass who sample from the implemented projects handled by the cohort of investors living in period $t - 1$, and assume to fix ideas that in the first period investors choose to invest whatever signal they observe.

In such a dynamic setting, investors in period t would adopt a threshold strategy z_t specifying to invest if the observed signal realization a is above z_t and to not invest otherwise where the sequence of z_t would be characterized inductively by $z_1 = \underline{a}$ (since the first generation of investors was assumed to invest always) and for all $t > 1$, the threshold z_{t+1} would be uniquely defined by $H(z_{t+1}, z_t) = c$ (assuming $H(\underline{a}, z) < c < H(\bar{a}, z)$ for all z) where $H(\cdot, \cdot)$ is the function defined in Section 3 of the paper. It appears that z_2 coincides with a^R , and using the monotonicity of H , it can be shown by induction that the sequence $(z_{2k+1})_{k \geq 1}$ is weakly decreasing and satisfies $z_{2k+1} \geq a^S$ for all k while the sequence $(z_{2k})_{k \geq 1}$ is weakly increasing and satisfies $z_{2k} \leq a^S$ for all k where a^S is the equilibrium threshold defined in Proposition 1. Thus, $(z_{2k+1})_{k \geq 1}$ converges to z^* and $(z_{2k})_{k \geq 1}$ converges to z_* with $z_* \leq a^S \leq z^*$. If $z_* = z^* = a^S$ the system converges to the steady state described in Proposition 1. If $z_* < a^S < z^*$, the system converges to a limit two-period cycle in which in odd periods there is less activity as dictated by the threshold strategy z^* and in even periods there is more activity as dictated by the threshold strategy z_* . Whether the system converges or cycles depends on how the slope $\frac{\partial H / \partial z}{\partial H / \partial a}$ compares to 1. When it is uniformly lower than 1, (as is the case for the leading example with variance $\sigma = 1$), there is convergence. When it is larger than 1 in the neighborhood of $a = z = a^S$, the two-period limit cycle prevails.¹⁵

¹⁵If investors were sampling from all previous cohorts rather than just the most recent one, I suspect the convergence scenario would be made more likely (because such a sampling device would smoothen

It should be noted that in the above dynamics whether or not there is convergence, the overoptimism and overinvestment biases hold in every period (this follows from the monotonicity of H and the observation that $H(a^R, \underline{a}) = c$). Moreover, since $z_t \leq a^R$ for all t and $z_2 = a^R$, the monotonicity of H implies that the smallest z_t which corresponds to the most biased investment strategy is obtained in period 3 when the samples considered by the current cohort consist of projects handled by rational investors. In all subsequent periods, because sampled investors adopt suboptimal strategies, the sampling heuristic leads to less severe biases.

Cycling with heterogeneous investors

It is natural to combine dynamics as just considered with the possibility that investors could vary in their degree of sophistication, some of them being rational and others being subject to selection neglect as proposed in the main model. A full-fledged dynamic model along these lines would aim at endogenizing entry and exit of entrepreneurs, assuming for example entrepreneurs' sophistication vary with their experience. Analyzing such a model is clearly beyond the scope of this online appendix. Yet, in order to illustrate that some rich dynamics can be expected, consider the following stylized setting. In each period $t = 1, 2, \dots$ a new cohort of agents decides whether or not to become entrepreneur. Every entrepreneur faces the same distribution of projects as described above but agents may have different outside options assumed to be drawn independently across agents from a distribution with cumulative G . In every period, the share of rational agents is λ and the share of sampling agents is $1 - \lambda$. Let w^R denote the expected payoff a rational investor gets by becoming an entrepreneur (i.e., $w^R = E(\max v^R(a) - c, 0)$), and let $w^S(\lambda)$ denote the expected payoff a sampling investor subjectively expects to get when facing a share λ (resp. $1 - \lambda$) of rational (resp. sampling) investors.¹⁶ Rational agents become entrepreneur whenever their outside option falls below w^R , i.e. with probability $G(w^R)$. Sampling agents who would sample from a mix λ of rational investors and $1 - \lambda$ of sampling investors would become entrepreneur with probability $G(w^S(\lambda))$. Thus assuming the cohort of (sampling) agents in period t samples from the implemented projects in period $t - 1$, the share λ_t of rational investors in period t would follow the dynamic:

$$\lambda_t = \frac{\mu G(w^R)}{\mu G(w^R) + (1 - \mu) G(w^S(\lambda_{t-1}))}.$$

the reaction to previous behaviors), but more work is needed to establish this formally.

¹⁶With the notation previously introduced, $w^S(\lambda) = E[\max(H(a, a^S(\lambda)) - c, 0)]$ where the density of a is $f_\lambda(a) = \frac{\sum_{x \in X} f(a|x)[(1-\lambda)(1-F(a^S(\lambda)|x)) + \lambda(1-F(a^R|x))]l(x)}{\sum_{x \in X} ((1-\lambda)(1-F(a^S(\lambda)|x)) + \lambda(1-F(a^R|x)))l(x)}$.

As can be inferred from the above analysis, $w^S(\cdot)$ is increasing in λ . Thus, a higher share of rational investors in period t would lead more sampling agents to become entrepreneurs in period $t + 1$, which would result in a lower share of rational investors in period $t + 1$. Depending on the shape of G , such a dynamic system may either converge to a limit share λ^* of rational investors or lead to long term cycling between high and low shares (away and respectively above and below λ^*) of rational investors, corresponding respectively to low and high levels of entrepreneurial activity.¹⁷

¹⁷ λ^* is a solution to $\lambda^* = \frac{\mu G(w^R)}{\mu G(w^R) + (1-\mu)G(w^S(\lambda^*))}$ and if G has sufficient mass around $w^S(\lambda^*)$ one should expect cycling to emerge.