DISCUSSION PAPER SERIES

No. 3603

ARE WE ALL BETTER DRIVERS THAN AVERAGE? SELF-PERCEPTION AND BIASED BEHAVIOUR

Isabelle Brocas and Juan D Carrillo

INDUSTRIAL ORGANIZATION and PUBLIC POLICY

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP3603.asp**

ISSN 0265-8003

ARE WE ALL BETTER DRIVERS THAN AVERAGE? SELF-PERCEPTION AND BIASED BEHAVIOUR

Isabelle Brocas, Columbia University, ECARES and CEPR
Juan D Carrillo, Columbia Business School and CEPR

Discussion Paper No. 3603
October 2002

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 3603

October 2002

ABSTRACT

Are We All Better Drivers than Average? Self-Perception and Biased Behaviour*

This Paper studies a model where individuals have imperfect self-knowledge and learning is costly. It shows that the endogenous decision to collect information before taking an action creates a systematic and testable bias in the aggregate behaviour of agents in the economy. More precisely, individuals distort the information acquisition procedure so as to favour the possibility of undertaking the action that generates the highest benefits in some states, even if it also generates the biggest losses in some others. The Paper thus explains within a rational framework why 80% of individuals may perceive themselves as being brighter, better drivers and more able entrepreneurs than their average peer. Applications to biases in career choices and judicial decisions are discussed.

JEL Classification: A12 and D83
Keywords: bias in judgement, decision-making and self-perception

Isabelle Brocas
Department of Economics
Columbia University
1009B IAB
420 West 118th Street
New York
NY 10027
USA
Tel: (1 212) 854 4067
Email: ib2014@columbia.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=138720

Juan D Carrillo
Columbia Business School
Uris Hall 812
3022 Broadway
New York
NY 10027
USA
Tel: (1 212) 854 2455
Fax: (1 212) 662 8474
Email: jc2095@columbia.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=131469

Submitted 13 September 2002

1 Motivation

It has been long established in the Psychology literature that most individuals consider themselves better drivers than average, most college professors brighter than their colleagues, most entrepreneurs more capable to run a business than their competitors, etc. (see e.g. Gilovich, 1991). This bias in self-perception is reflected not only in opinions but, more importantly, in behaviors with their adverse welfare consequences: accidents, jealousies, business failures, etc.

Is this bias simply the result of an irrational cognitive process that protects the individual against his own fears? A recent strand of the behavioral economics literature has challenged that view and provided microeconomic foundations for some systematic biases in choices. The arguments are based on moral constraints (Rabin, 1995), hyperbolic discounting (Carrillo and Mariotti, 2000), anticipatory feelings (Caplin and Leahy, 2001a and 2001b) and preferences over beliefs (Koszegi, 2000).[1] However, since all these models depart in one way or another from the standard neoclassical paradigm, mainstream economists have been only moderately convinced by the conclusions.

In this paper, we argue that if each individual has imperfect self-knowledge and learning is costly, then the endogenous decision to collect information generates a *systematic* and *testable* bias in the aggregate behavior of the population.[2] More precisely, consider individuals who choose between a "safe" action that generates always an average utility and a "risky" action that generates either a high or a low utility depending on the (unknown) state of nature. When preliminary information reveals that the risky action is likely to yield a high utility, individuals are prone to stop the costly collection of news and take the risky action. By contrast, if the preliminary information suggests that the risky action is likely to yield a low utility, then individuals are inclined to keep accumulating evidence rather than stop and take the safe action. The asymmetry in the value of news

[1] See also Bénabou and Tirole (2001a and 2001b) for applications to self-confidence and personal rules. Other papers follow a different approach to the question: they assume irrational (biased) beliefs and study the economic consequences (see e.g. Manove (1999) or Manove and Padilla (1999) for the case of entrepreneurs who are intrinsically optimistic).

[2] Obviously, we do not claim that all biased behaviors are due to this cost. In this respect, the paper just adds one new element to the discussion, with the appeal that the premises are simple and standard and the implications are testable.

implies that, in equilibrium, if the collection of information is a choice variable then agents are more likely to take (both ex-post correctly and ex-post wrongly) risky actions than if information is fixed or flows exogenously. So, agents reach more often marked successes and dramatic failures than average outcomes. As a stylized example, suppose that being an employee is a relatively safe option whereas starting a private business generates great earnings to good entrepreneurs and big losses to bad entrepreneurs. Also, individuals have imperfect self-knowledge about their skills and can learn them at a cost before choosing their labor activity. Our model predicts that rational individuals with meager but positive evidence on their managerial skills will become entrepreneurs. By contrast, they will need a substantial amount of negative information about their managerial skills before deciding to become employees. As a result, in equilibrium, we will observe in the population more bad entrepreneurs who start their own company than good entrepreneurs who work for someone else.

Interestingly, the same analysis and the same conclusions can be applied to the following conceptually different but formally equivalent problem. Two individuals (e.g., two judges) decide between two actions (e.g., convicting or releasing an offender). They have imperfect information about the state of nature (e.g., the culpability of the suspect) and can acquire information at a cost (e.g., delaying the sentence). Suppose also that, for any given belief, the *difference in utility* between the two actions is the same for both individuals. This, in particular, means that they both make the same choices in the face of the same evidence (e.g., they both prefer to convict the suspect if his probability of being guilty is above a certain threshold and release him otherwise). If the individuals derive a different *total utility* for each action, then their endogenous choice to acquire information implies that, in equilibrium, we will observe systematically different biases in their decisions. More precisely, suppose that the action with highest and lowest payoff for one judge is to release and convict an innocent prisoner respectively, and for the other judge it is to convict and release a guilty prisoner respectively. Both formalizations of preferences seem equally reasonable. Yet, according to our model, the first judge will release guilty suspects with higher probability and convict innocent suspects with lower probability than the second judge. That is, in equilibrium, they will make different types of mistakes.

To sum up, the paper shows that individuals will bias their information acquisition procedure so as to increase the likelihood of taking the action that may potentially generate the highest benefit: start a (successful) business or release a (innocent) suspect. Stated this way, it might seem trivial. However, note that with this strategy, individuals are also increasing the likelihood of getting the biggest losses: start a (unsuccessful) business or release a (guilty) suspect. The overall consequences for aggregate behavior are immediate.

The plan of the paper is the following. We first present a model with two agents who have imperfect information about the state of nature. For any given belief, the difference in expected utility between the two actions is the same for both agents (section 2). We show that their different incentives to acquire information affects their behavior and expected payoffs (section 3). We then argue that the model and the results immediately extend to the case of one agent with imperfect self-knowledge who learns about his own preferences and manipulates his own choices (section 4). Last, we provide some concluding comments (section 5).

2 A model of biased behavior

2.1 States, actions and utilities

The model has the following characteristics. Agents can be of two types ($i \in \{1, 2\}$). A type-i agent chooses an action $\gamma_i \in \{a, b\}$. His utility $u_i(\cdot)$ depends on his action γ_i and the state of the economy $s \in \{A, B\}$. Agents have imperfect knowledge about the state s. We denote by p the prior belief shared by all agents that the true state is A, that is $\Pr(A) = p$ and $\Pr(B) = 1 - p$.

Type-1 and type-2 agents have different utility functions (i.e. different representation of preferences). However, we will assume that for any given belief $p \in [0, 1]$, they both have the same difference in expected utility between the two possible actions. This means not only that they have the same preferred action when confronted to the same evidence, but also that they have the same willingness to pay in order to have the freedom of choosing which action they take. We will say that these two types of agents are *identical in behavior and utility difference* (IBUD). The property is summarized as follows.

Definition *Agents 1 and 2 are* IBUD *if and only if* $E\Big[u_1(a) - u_1(b)\Big] = E\Big[u_2(a) - u_2(b)\Big]$ *for all* p *(which, in particular, implies* $\arg\max_{\gamma_1} E[u_1(\gamma_1)] = \arg\max_{\gamma_2} E[u_2(\gamma_2)]$ *for all* p*).*

For expositional purposes, we will focus on the following specific utility functions for types 1 and 2:

$$u_1(a) = \begin{cases} h & \text{if} \quad s = A \\ -h & \text{if} \quad s = B \end{cases} \quad \text{and} \quad u_1(b) = \begin{cases} -l & \text{if} \quad s = A \\ l & \text{if} \quad s = B \end{cases}, \tag{1}$$

$$u_2(a) = \begin{cases} l & \text{if} \quad s = A \\ -l & \text{if} \quad s = B \end{cases} \quad \text{and} \quad u_2(b) = \begin{cases} -h & \text{if} \quad s = A \\ h & \text{if} \quad s = B \end{cases}, \tag{2}$$

with $h > l > 0$, so that action a has the greatest variance in payoffs for agent 1 and action b the greatest variance for agent 2. Given (1) and (2), the IBUD property translates into:

$$E\Big[u_i(a) - u_i(b)\Big] = (h+l)(2p-1) \ \ \forall\, i \ \ \Rightarrow \ \ \gamma_i = a \ \text{ if } \ p > 1/2 \ \text{ and } \ \gamma_i = b \ \text{ if } \ p < 1/2 \ \ \forall\, i.$$

Figure 1 provides a graphical representation of these utilities.

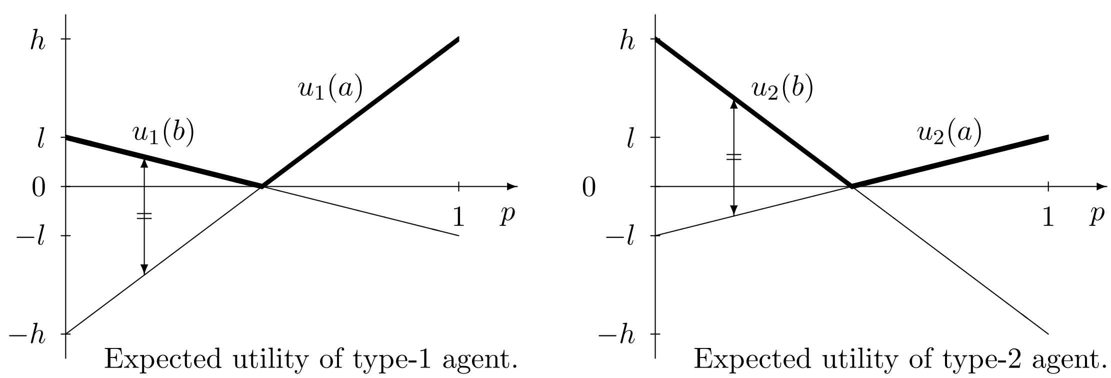

Figure 1. Utility representations for type-1 and type-2 agents.

The theory on reveal preferences says that utility functions are only modelling devices used to represent orders of preferences; actions and (maybe) payments incurred in order to choose actions are often the observable variables from which we deduce the preferences of individuals. Therefore, one might think that two agents with different utility representations but who satisfy the IBUD property (as types 1 and 2) should be indistinguishable, as long as the set of choices is restricted to a and b. This intuition is correct either when information is freely available or exogenously given, but what happens when we allow individuals to decide how much costly information they collect?

2.2 Information

In order to answer this question, we need to introduce the information acquisition technology. Learning is formalized in the simplest way. We denote by $\tau_{i,t}$ the decision of agent i at a given date $t \in \{0, 1, ..., T-1\}$, where T is finite but arbitrarily large. At each date, his options are either to take the optimal action conditional on his current information ($\tau_{i,t} = \gamma_i \in \{a, b\}$) or to wait until the next period ($\tau_{i,t} = w$). The action is irreversible, so if the agent undertakes it, then payoffs are realized and the game ends. Waiting has costs and benefits. On the one hand, the delay implied by the decision to wait one more period before acting is costly. We denote by δ (< 1) the discount factor. On the other hand, the agent obtains between dates t and $t+1$ one signal $\sigma \in \{\alpha, \beta\}$ imperfectly correlated with the true state. Information allows the agent to update his belief and therefore to improve the quality of his future decision. As long as the agent waits, he keeps the option of undertaking action a or b in a future period, except at date T in which waiting is not possible anymore so the agent's options are reduced to actions a and b.[3] The relation between signal and state is the following:

$$\Pr[\,\alpha \mid A\,] = \Pr[\,\beta \mid B\,] = \theta \quad \text{and} \quad \Pr[\,\alpha \mid B\,] = \Pr[\,\beta \mid A\,] = 1 - \theta,$$

where $\theta \in (1/2, 1)$ captures the accuracy of information: as θ increases, the informational content of a signal σ also increases (when $\theta \to 1/2$ signals are completely uninformative, and when $\theta \to 1$ one signal perfectly informs the agent about the true state).[4]

Suppose that a number n_α of signals α and a number n_β of signals β are released during the $n_\alpha + n_\beta$ periods in which the agent waits. Using standard statistical methods,

[3] A finite horizon game ensures the existence of a unique stopping rule at each period that can be computed by backward induction. By setting T arbitrarily large we can determine the limiting properties of this stopping rule.

[4] It is formally equivalent to increase θ or to assume that the agent obtains more than one signal per-period.

it is possible to compute his posterior belief about the state of the economy:

$$\begin{aligned}\Pr(A \mid n_\alpha, n_\beta) &= \frac{\Pr(n_\alpha, n_\beta|A)\Pr(A)}{\Pr(n_\alpha, n_\beta|A)\Pr(A) + \Pr(n_\alpha, n_\beta|B)\Pr(B)} \\ &= \frac{\theta^{n_\alpha - n_\beta} \cdot p}{\theta^{n_\alpha - n_\beta} \cdot p + (1-\theta)^{n_\alpha - n_\beta} \cdot (1-p)}\end{aligned}$$

The relevant variable which will be used from now on is $n \equiv n_\alpha - n_\beta \in \mathbb{Z}$, that is the difference between the number of signals α and the number of signals β. Also, we define the posterior probability $\mu(n) \equiv \Pr(A \mid n_a, n_b)$.[5] Before solving the game, we want to provide some stylized examples that illustrate the meaning of the IBUD property, the utility representations (1) and (2), and the costly decision to acquire information.

2.3 Examples

1. Court judgment (Civil law). A type-i judge has to decide whether to release (a) or convict (b) an offender. The offender is either innocent (A) or guilty (B). The prior probability of his being innocent is $p = \Pr(A)$. The judge can acquire information about the culpability of the accused (signals σ) at the cost of delaying his sentence. According to (1) and (2), for any belief p, the differential in utility between convicting and releasing the offender is the same for both types of judges (IBUD property). In particular, both prefer to release the prisoner if his probability of being innocent is greater than 1/2 and convict him otherwise. The main difference is that for a type-1 judge releasing the suspect is more risky than convicting him ($u_1(a) \in \{-h, h\}$ vs. $u_1(b) \in \{-l, l\}$). The opposite is true for a type-2 judge.[6] Under these circumstances, should the offender be concerned about which type of judge handles his case?

2. Career choice. An individual has to choose between accepting an academic position (a) and starting his own consulting company (b). His career achievements depend on the skills he has (A or B), where A represents an agent with a comparative advantage for research and B an agent with a comparative advantage for management. The individual has imperfect information about his own skills and can improve his self-knowledge at the

[5] Some properties of $\mu(n)$ are: (i) $\lim_{n \to -\infty} \mu(n) = 0$, (ii) $\lim_{n \to +\infty} \mu(n) = 1$, and (iii) $\mu(n+1) > \mu(n) \ \ \forall n$.

[6] In other words, judge 1 is happiest when he releases an innocent prisoner and judge 2 is happiest when he convicts a guilty prisoner.

expense of delaying the career choice. He is concerned both with prestige (an issue with highest variance in academia) and financial remuneration (an issue with highest variance in consulting). Equations (1) and (2) represent the preferences of two individuals who are mostly sensitive to prestige and to money, respectively. However, since both types are IBUD, will they –on average– take the same career decisions?

3 Information acquisition and optimal decision-making

3.1 Option value of waiting and optimal stopping rule

In this section we show that under costly acquisition of information two IBUD agents may behave in a systematically different way. Given the information revelation structure presented in section 2.2, agents face a trade-off between delay and information. This type of models has been extensively analyzed in the literature on investment under uncertainty (for a comprehensive overview see e.g. Dixit and Pindyck, 1994). Following the standard techniques developed in this literature, we first determine the value function V_i^t that a type-i agent maximizes at date t. It can be written as:

$$V_1^t(n) = \begin{cases} \max\Big\{h(2\mu(n)-1), \delta\Big[\nu(n)V_1^{t+1}(n+1) + (1-\nu(n))V_1^{t+1}(n-1)\Big]\Big\} & \text{if } \mu(n) \geqslant \frac{1}{2} \\ \max\Big\{l(1-2\mu(n)), \delta\Big[\nu(n)V_1^{t+1}(n+1) + (1-\nu(n))V_1^{t+1}(n-1)\Big]\Big\} & \text{if } \mu(n) < \frac{1}{2} \end{cases} \tag{3}$$

$$V_2^t(n) = \begin{cases} \max\Big\{l(2\mu(n)-1), \delta\Big[\nu(n)V_2^{t+1}(n+1) + (1-\nu(n))V_2^{t+1}(n-1)\Big]\Big\} & \text{if } \mu(n) \geqslant \frac{1}{2} \\ \max\Big\{h(1-2\mu(n)), \delta\Big[\nu(n)V_2^{t+1}(n+1) + (1-\nu(n))V_2^{t+1}(n-1)\Big]\Big\} & \text{if } \mu(n) < \frac{1}{2} \end{cases} \tag{4}$$

where $\nu(n) = \mu(n)\theta + (1-\mu(n))(1-\theta)$. In words, at a given date t and for a given difference of signals n that implies a posterior $\mu(n) > 1/2$, type-1 agent chooses between taking action a with expected payoff $h\mu - h(1-\mu)$ or waiting. In the latter case, signal α (resp. β) is received with probability ν (resp. $1-\nu$) and the value function in the following period $(t+1)$ becomes $V_1^{t+1}(n+1)$ (resp. $V_1^{t+1}(n-1)$), discounted at the rate δ. For $\mu(n) < 1/2$, the argument is the same except that the optimal action if the agent does not wait is b with payoff $-l\mu + l(1-\mu)$. A similar reasoning applies to a type-2 agent. Given (3) and (4), we can determine the optimal strategy for each type.

Lemma 1 *For all $\delta < 1$ and $h > l > 0$, there exist $(n^*_{1,t}, n^{**}_{1,t}, n^*_{2,t}, n^{**}_{2,t})$ at each date t s.t.:*

(i) $\tau_{1,t} = b$ *if* $n \leqslant n^*_{1,t}$, $\tau_{1,t} = a$ *if* $n \geqslant n^{**}_{1,t}$ *and* $\tau_{1,t} = w$ *if* $n \in (n^*_{1,t}, n^{**}_{1,t})$.

(ii) $\tau_{2,t} = b$ *if* $n \leqslant n^*_{2,t}$, $\tau_{2,t} = a$ *if* $n \geqslant n^{**}_{2,t}$ *and* $\tau_{2,t} = w$ *if* $n \in (n^*_{2,t}, n^{**}_{2,t})$.

Given the symmetry of types 1 and 2, we have $\mu(n^{**}_{1,t}) = 1-\mu(n^*_{2,t})$ *and* $\mu(n^*_{1,t}) = 1-\mu(n^{**}_{2,t})$. *Also,* $1/2 < \mu(n^{**}_{1,t}) < 1 - \mu(n^*_{1,t})$ *(so* $\mu(n^*_{1,t}) < \mu(n^*_{2,t}) < 1/2 < \mu(n^{**}_{1,t}) < \mu(n^{**}_{2,t})$*).*

Proof. See Appendix. □

The idea is simple. As in the standard literature on information and the option value of waiting, agents trade-off the costs of delaying their choice between actions a and b with the benefits of acquiring a more accurate information. When $\mu(n) > 1/2$, waiting becomes more costly as n increases, because delaying the action one more period reduces the expected payoff by an amount proportional to $2\mu(n) - 1$. Similarly, when $\mu(n) < 1/2$, waiting becomes more costly as n decreases, because delaying the action reduces the expected payoff by an amount proportional to $1 - 2\mu(n)$. In other words, at each date t, there are two cutoffs $\mu(n^{**}_{1,t}) > 1/2$ and $\mu(n^*_{1,t}) < 1/2$ for a type-1 agent. When $\mu \geqslant \mu(n^{**}_{1,t})$, the individual is "reasonably confident" that the true state is A, and when $\mu \leqslant \mu(n^*_{1,t})$, he is "reasonably confident" that the true state is B. In either case, the marginal gain of improving the information about the true state is offset by the marginal cost of a reduction in the expected payoff due to the delay. As a result, he strictly prefers to stop learning and undertake his optimal action. For intermediate beliefs, that is when $\mu(n) \in (\mu(n^*_{1,t}), \mu(n^{**}_{1,t}))$, the individual prefers to keep accumulating evidence (the same argument holds for a type-2 agent).

The most interesting property of these cutoffs is that $1/2 < \mu(n^{**}_{1,t}) < 1 - \mu(n^*_{1,t})$ or, equivalently, $\mu(n^{**}_{1,t}) - 1/2 < 1/2 - \mu(n^*_{1,t})$. It means that the confidence of a type-1 agent on the true state being A when he takes action a is smaller than his confidence on the true state being B when he takes action b. The opposite is true for a type-2 agent. Comparing the two agents, we conclude that a type-1 agent will need fewer evidence in favor of a in order to decide to stop collecting news and take action A and more evidence in favor of b in order to stop collecting news and take action B than a type-2 agent.

The intuition for this result is simply that, given the delay associated to the accumulation of evidence, the marginal cost of learning is a function of the agent's expected payoff. Therefore, for a type-1 individual, it is proportional to $h(1-\delta)$ when $\mu > 1/2$ and to $l(1-\delta)$ when $\mu < 1/2$. Thus, other things being equal, it is relatively less interesting to keep experimenting when the currently optimal action is a rather than b. The argument for a type-2 agent is symmetric. The shape of these cutoffs is graphically represented in Figure 2.

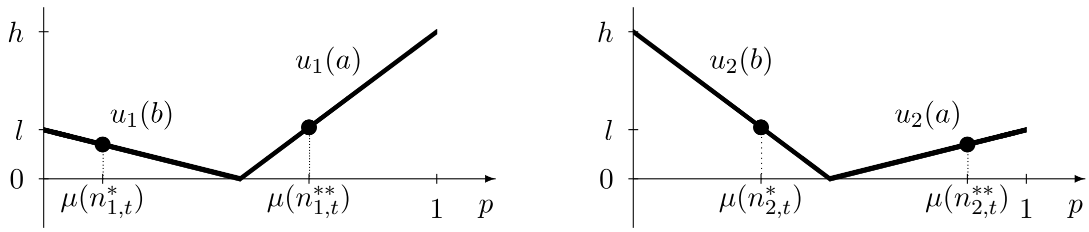

Stopping rule for type-1 agent.

Stopping rule for type-2 agent.

Figure 2. Stopping rules for type-1 and type-2 agents.

Suppose now that $T \to \infty$, so that $n^*_{i,t} \to n^*_i$ and $n^{**}_{i,t} \to n^{**}_i$ for all t. Denote by $\Pr(\tau_i = \gamma_i \mid s)$ the probability that a type-i individual eventually undertakes action γ_i ($\in \{a,b\}$) when the true state is s ($\in \{A,B\}$). Also, let $\mu^{**} \equiv \mu(n^{**}_1)$ and $\mu^* \equiv \mu(n^*_1)$ (which means that $\mu(n^{**}_2) = 1-\mu^*$ and $\mu(n^*_2) = 1-\mu^{**}$). Suppose that one type-1 and one type-2 agent start with the same prior belief $p \in (1-\mu^{**}, \mu^{**})$. Each agent chooses the amount of information collected before undertaking an action and the signals obtained by each of them are independent. Their optimal stopping rule is given by Lemma 1. In the main proposition of this paper, we compare the relative probabilities that each agent undertakes action a and action b.

Proposition 1 *For all $p \in (1-\mu^{**}, \mu^{**})$, $\delta < 1$, $h > l > 0$ and when $T \to \infty$ we have:*

(i) $\Pr(\tau_1 = a \mid s) > \Pr(\tau_2 = a \mid s)$ *for all s, which means that* $\Pr(\tau_1 = a \mid B) > \Pr(\tau_2 = a \mid B)$ *and* $\Pr(\tau_1 = b \mid A) < \Pr(\tau_2 = b \mid A)$.

(ii) $\frac{\partial \Pr(\tau_1 = a \mid s)}{\partial h} > 0 > \frac{\partial \Pr(\tau_2 = a \mid s)}{\partial h}$ *and* $\frac{\partial \Pr(\tau_1 = a \mid s)}{\partial l} < 0 < \frac{\partial \Pr(\gamma_2 = a \mid s)}{\partial l}$ *for all s.*

Proof. Part (i) is a direct consequence of $\mu(n_2^{**}) > \mu(n_1^{**})$ and $\mu(n_2^*) > \mu(n_1^*)$. Part (ii) results from the fact that, also by Lemma 1, $\frac{\partial n_1^*}{\partial h} < 0$, $\frac{\partial n_1^{**}}{\partial h} < 0$, $\frac{\partial n_1^*}{\partial l} > 0$, $\frac{\partial n_1^{**}}{\partial l} > 0$ and by symmetry $\frac{\partial n_2^*}{\partial h} > 0$, $\frac{\partial n_2^{**}}{\partial h} > 0$, $\frac{\partial n_2^*}{\partial l} < 0$, $\frac{\partial n_2^{**}}{\partial l} < 0$.

These comparative statics are sufficient for the purpose of our analysis. However, for the reader interested, the analytical expressions of the probabilities $\Pr(\tau_i \mid s)$ are derived in Brocas and Carrillo (2002, Lemma 1).[7] They are given by: $\Pr(\tau_1 = a \mid A) = \frac{p-\mu^*}{\mu^{**}-\mu^*}\frac{\mu^{**}}{p}$, $\Pr(\tau_1 = a \mid B) = \frac{p-\mu^*}{\mu^{**}-\mu^*}\frac{1-\mu^{**}}{1-p}$, $\Pr(\tau_2 = a \mid A) = \frac{p-(1-\mu^{**})}{\mu^{**}-\mu^*}\frac{1-\mu^*}{p}$, $\Pr(\tau_2 = a \mid B) = \frac{p-(1-\mu^{**})}{\mu^{**}-\mu^*}\frac{\mu^*}{1-p}$. Naturally, the results of Proposition 1 can be easily confirmed using these expressions. □

Proposition 1 shows that, even if type-1 and type-2 agents are IBUD, they will make systematically different choices, at least in a stochastic sense. As shown in Lemma 1, a type-1 agent is relatively more likely to stop collecting news when the preliminary evidence points towards the optimality of action a than when it points towards the optimality of action b (i.e. when the first few signals are mainly α and not β). Stated differently, the evidence in favor of A needed to induce a type-1 agent to stop accumulating pieces of news and take action a is smaller than the evidence in favor of B needed to induce him to stop accumulating news and take action b. The opposite is true for a type-2 agent. As a result, in equilibrium, a type-1 agent is more likely to take action a by mistake (i.e. when the true state is B) and less likely to take action b by mistake (i.e. when the true state is A) than a type-2 agent (part (i) of the proposition). Note that the endogenous choice to acquire information is crucial for this result: by definition of IBUD, the two types of agents would take action a with the same probability if the number of signals they receive were exogenously given. Last, as the difference between the two maximum payoffs $(h-l)$ increases, the likelihood that the two agents behave differently also increases: type-1 takes more often action a by mistake and less often action b by mistake whereas type-2 takes less often action a by mistake and more often action b by mistake (part (ii) of the proposition).

[7] The paper uses similar techniques to analyze a different issue. It studies a principal/agent model with incomplete contracts where the former controls the flow of public information and the latter undertakes an action that affects the payoffs of both. The paper characterizes the rents of the principal due to his ability to control information.

How should we interpret this result? In terms of our examples, it means that a type-1 judge releases guilty suspects more often and convicts innocent suspects less often than a type-2 judge. Hence, an offender will strictly prefer to face a type-1 rather than a type-2 judge, independently of his culpability. As for the second example, individuals mostly concerned with prestige will (mistakenly) follow academic careers more often and (mistakenly) start their own consulting company less often than individuals mostly interested in the financial remuneration of the job.[8] The difference in their career choices becomes more pronounced as the difference in the utility between being a successful researcher and being a successful entrepreneur increases. The paper thus provides one reason for the propensity of individuals to make different types of errors.[9]

Note that agents select a stopping rule that increases the probability of taking the action with highest payoff ($\tau_1 = a$ when $s = A$ for a type-1 and $\tau_2 = b$ when $s = B$ for a type-2 agent, which both imply a utility h). This might seem a trivial conclusion. However, it is not that straightforward. In fact, the other side of the coin is that, with this strategy, agents are de facto also increasing the probability of making the mistakes that are most costly ($\tau_1 = a$ when $s = B$ for a type-1 and $\tau_2 = b$ when $s = A$ for a type-2 agent, which both generate utility $-h$). In our examples, judge-1, who is most willing to release innocent individuals but most averse to release guilty ones, engages in both practices more often than judge-2. Also, financially motivated individuals start their own company and become rich (highest utility) or ruined (lowest utility) with higher probability than individuals more concerned with prestige.

Last, we would like to make clear that when we refer to a "biased" behavior, we do not mean that agents are fooled, deceived or misled. Obviously, this is not possible in our setting given that agents are rational both in their acquisition of information and in their choice of action. What happens instead is that few evidence in favor of A is enough to induce a type-1 agent to undertake action a whereas substantial evidence in favor of

[8]Mistakenly means in the first case "conditional on having a comparative advantage for management" (state B) and in the second case "conditional on having a comparative advantage for research" (state A).

[9]Note that mistakes are unlikely to persist indefinitely if, for example, individuals keep learning about their skills after the selection of their career and the occupational choice is never completely irreversible. However, our qualitative insights are not affected by this remark.

B is needed to convince him to undertake action b, and the opposite is true for a type-2 agent. So, suppose that there is a population of type-1 judges, each of them handling one different case. All suspects are guilty with ex-ante probability 1/2. The model predicts that a majority of judges will end up releasing their suspects under (weak) evidence of their being innocent and a minority of judges will end up convicting their suspects under (strong) evidence of their being guilty (again, the opposite holds if the population is composed of type-2 judges).

Technically, the point is simply that the endogenous decision to acquire information cannot affect the first-order moment of beliefs (i.e., the average belief in the population always coincides with the true average). However, it may influence the higher-order moments and, in particular, the *skewness* of the distribution of beliefs. Given a limited set of actions, two populations whose distribution of beliefs have the same average but different skewness will exhibit different aggregate behaviors.

3.2 Comparative statics

Whether the effects highlighted in Proposition 1 are of considerable importance is clearly an empirical issue, interesting but beyond the scope of this paper. We simply want to provide simple numerical examples that give an idea of the propensity of agents to make different types of mistakes. Consider the extreme situation in which $h > 0$ and $l \to 0$.[10] From the proof of Proposition 1, the probability that a type-i agent makes a mistake is:

$$\Pr(\tau_1 = a \mid B) = \frac{p}{1-p} \times \frac{1-\mu^{**}}{\mu^{**}} \quad \text{and} \quad \Pr(\tau_1 = b \mid A) \to 0$$

$$\Pr(\tau_2 = a \mid B) \to 0 \quad \text{and} \quad \Pr(\tau_2 = b \mid A) = \frac{1-p}{p} \times \frac{1-\mu^{**}}{\mu^{**}}$$

A type-1 agent will never take action b mistakenly, and a type-2 agent will never take action a mistakenly. Simple comparative statics about the likelihood of taking the "wrong" action given a prior probability p and a stopping posterior μ^{**} are illustrated in Figure 3.

[10]This means that $n_1^* \to -\infty$, $n_2^{**} \to +\infty$ and therefore $\mu^* \to 0$. This assumption is by no means necessary. However, it makes the comparative statics much more neat.

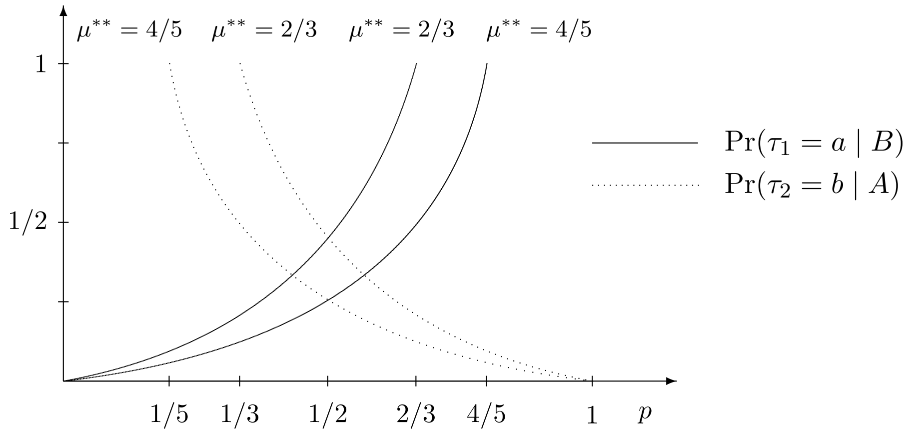

Figure 3. Frequency of mistakes by type-1 and type-2 agents.

As an illustrative case, suppose that half the convicts are guilty ($p = 1/2$) and the stopping rule is such that $\mu^{**} = 2/3$. A type-1 judge releases 3/4 of the suspects: all the innocents (1/2 of the population) and 1/2 of the guilty ones (1/4 of the population). Hence, he only convicts 1/2 of the guilty suspects (1/4 of the population). The conviction pattern for judge-2 is the opposite. Last, note that μ^{**} is increasing in δ, and $\lim_{\delta \to 1} \mu^{**} = 1$. This means simply that, if individuals are infinitely patient, the cost of waiting vanishes. It then becomes optimal for both types to be perfectly informed before choosing any action, and therefore no mistakes occur in equilibrium.

3.3 Welfare implications

Suppose that a welfare maximizing principal can ask several individuals their opinion about which action a or b should be taken. For simplicity, we assume that each individual is interested in maximizing the probability of providing the correct appraisal (a if A and b if B), independently of whether the suggestion is followed by the principal or not. Such behavior is rational if individuals have career-concerns and their payoff is only a function of the accuracy of their suggestion. In this setting, each agent's optimal rule for the acquisition of information coincides with that of Lemma 1.[11]

If individuals receive independent signals about the true state, then increasing the number of agents can only decrease the probability of an incorrect decision. However, recall that the two types of agents have different biases in the errors they commit. Suppose that the principal can choose the proportion of type-1 and type-2 agents. Is it optimal to select all agents of the same type or to have appraisals from agents of both types?

To answer this question, we consider the simplest version of our model (although it can be generalized in a number of dimensions). We denote by γ_i^j the recommendation made by the jth agent of type-i. We suppose that $l \to 0$, so that $\Pr(\gamma_1^j = b \mid A) = 0$ and $\Pr(\gamma_2^j = a \mid B) = 0$ for all j. The total number of agents is fixed and equal to n. The principal chooses x, the number of type-1 agents ($n - x$ being the number of type-2 agents). Last, the principal's sole concern is to minimize the probability of a mistake (i.e., it is equally costly to take action a when $s = B$ than action b when $s = A$). If we denote by $\gamma_p \in \{a, b\}$ the action taken eventually by the principal, we have the following result.

Proposition 2 *If $p < 1/2$, then $x = n$. Furthermore, $\gamma_p = a$ if and only if $\gamma_1^j = a$ $\forall j \in \{1, ..., n\}$. Last, $\Pr(\gamma_p = b \mid A) = 0$ and $\Pr(\gamma_p = a \mid B) = \left(\frac{p}{1-p} \times \frac{1-\mu^{**}}{\mu^{**}}\right)^n$.*

If $p > 1/2$, then $x = 0$. Furthermore, $\gamma_p = b$ if and only if $\gamma_2^j = b$ $\forall j \in \{1, ..., n\}$. Last, $\Pr(\gamma_p = b \mid A) = \left(\frac{1-p}{p} \times \frac{1-\mu^{**}}{\mu^{**}}\right)^n$ and $\Pr(\gamma_p = a \mid B) = 0$.

Proof. Given $l \to 0$, we have $\Pr(\gamma_1 = b \mid A) = 0$ and $\Pr(\gamma_2 = a \mid B) = 0$, so the only possible error arises when all type-1 agents announce $\gamma_1^j = a$ ($j \in \{1, ..., x\}$) and all type-2 agents announce $\gamma_2^k = b$ ($k \in \{1, ..., n - x\}$). The remaining question is whether, if this happens, the principal will take action a or action b.

- Suppose that the principal minimizes costs with $\gamma_p = a$. The expected loss is then:

$$L_A(x) = \Pr(B) \cdot \prod_{j=1}^{x} \Pr(\gamma_1^j = a \mid B) \cdot \prod_{k=1}^{n-x} \Pr(\gamma_2^k = b \mid B) = (1-p)\left(\frac{p}{1-p} \times \frac{1-\mu^{**}}{\mu^{**}}\right)^x$$

So, conditional on taking $\gamma_p = a$, the principal optimally sets $x = n$, and the loss is:

$$L_A(n) = (1-p)\left(\frac{p}{1-p} \times \frac{1-\mu^{**}}{\mu^{**}}\right)^n \quad (5)$$

[11] By contrast, if individuals were rewarded as a function of the quality of the final decision taken, then they should integrate the behavior of the other agents in their decision to acquire information (and, possibly, free-ride accordingly). The optimal stopping rule would then be modified.

- Suppose that the principal minimizes costs with $\gamma_p = b$. The expected loss is then:

$$L_B(x) = \Pr(A) \cdot \prod_{j=1}^{x} \Pr(\gamma_1^j = a \mid A) \cdot \prod_{k=1}^{n-x} \Pr(\gamma_2^k = b \mid A) = p\left(\frac{1-p}{p} \times \frac{1-\mu^{**}}{\mu^{**}}\right)^{n-x}$$

So, conditional on taking $\gamma_p = b$, the principal optimally sets $x = 0$, and the loss is:

$$L_B(0) = p\left(\frac{1-p}{p} \times \frac{1-\mu^{**}}{\mu^{**}}\right)^{n} \tag{6}$$

Last, from (5) and (6): $L_A(n) \lesseqgtr L_B(0) \Leftrightarrow (1-p)\left(\frac{p}{1-p}\right)^n \lesseqgtr p\left(\frac{1-p}{p}\right)^n \Leftrightarrow p \lesseqgtr 1/2$. □

Proposition 2 states that even if the principal can choose the source of information, a systematic bias in his choice is likely to persist. The idea is simple. Since the principal dislikes equally both types of errors, he will select agents so as to minimize their likelihood of committing a mistake, independently of the nature. Given $l \to 0$, type-1 (resp. type-2) agents can never mistakenly recommend action b (resp. a). We also know from Proposition 1 that the likelihood of providing an incorrect appraisal is inversely proportional to the distance between the prior belief and the posterior at which the agent decides to stop collecting evidence and recommends an action (formally, $\mu^{**} - p$ for a type-1 agent and $p - (1 - \mu^{**})$ for a type-2 agent). Hence, if $p < 1/2$, type-1 agents are relatively less likely to mislead the principal than type-2 agents, so it is optimal to select only this type of individuals. The opposite is true when $p > 1/2$. Overall, fewer mistakes will occur as we increase the number of agents who provide an appraisal. However, the systematic bias in the final decision is going to persist.

4 Bias in self-reported behavior

As mentioned in the introduction, there is extensive evidence that individuals report systematic errors in self-judgement (concerning their intellectual, driving or entrepreneurial capacity, for example). However, contrary to a common view, this does not necessarily imply irrationality of individuals. In fact, we claim that a simple reinterpretation of the costly information gathering model developed in section 2 can account for such systematic self-biased behavior, even in a perfectly rational world.

To illustrate this idea, consider the following stylized situation.[12] There is a continuum of individuals of mass 1 in the economy who do not know their entrepreneurial ability. Individuals are highly able (state A) with probability p and completely unable (state B) with probability $1-p$. They decide whether to start their own company (action a) or to become an employee (action b). The revenues of employees are normalized to 0. Being self-employed generates a revenue h to an able entrepreneur and $-h$ to an unable one. This model matches the utility formulation given by equation (1) when $l=0$. Suppose that individuals can, at a cost, collect information about their own skills (start first a small business, share their projects with experts, etc.). The optimal stopping rule of each individual is going to be such that $\mu^{**} \in (1/2, 1)$ and $\mu^* \to 0$. Using the results of Proposition 1, this means that all able entrepreneurs and a fraction of the unable ones will start their own business (formally, $\Pr(a \mid A) = 1$ and $\Pr(a \mid B) = p(1-\mu^{**})/(1-p)\mu^{**}$). Hence, the total fraction of self-employed agents in the economy will exceed that of capable ones ($\Pr(a) = p \times \Pr(a \mid A) + (1-p) \times \Pr(a \mid B) = p/\mu^{**} > p$). In the same vein, one might find reasonable to assume that, on average, driving yields higher utility than using public transportation if the individual has driving skills but lower utility if the individual does not have these skills (due e.g. to his increased probability of having an accident). If this is the case, our model predicts that we will observe more bad drivers using their car than good drivers taking public transportation.

The point of this exercise is simply to demonstrate that, when costly information is available, individuals will systematically bias their choices towards the alternative *with the potential* to generate highest payoff (running a business, driving to work).[13] Can we label this as "self-influence"? In a sense, the choice of this specific stopping rule in the acquisition of information can be seen as a (conscious or unconscious) attempt towards a self-bias: the individual is convincing himself that he has good reasons to take one particular action. However, we have also proved that such stopping rule is optimal for a rational person given his structure of payoffs and the costs of learning. Whether we think

[12] The example is only adopted for expositional purposes and therefore must not be taken literally. Also, the caveats mentioned in footnotes 2 and 9 apply.

[13] Note that the results also hold if entrepreneurial ability is valuable also as an employee, as long as it is relatively less important.

of it as self-influence or not, the main message is that in a rational world where information is scarce, the endogenous choice to collect evidence yields a rich set of predictions for the observed behavior of individuals.

5 Concluding remarks

This paper has shown how a majority of (rational) individuals can perceive themselves as being "better" (e.g. entrepreneurs) than average, and behave accordingly. The starting point is simply that the cost of acquiring information is sometimes a delay in the choice between the different actions or an increase in the probability of the project becoming obsolete. In both cases, the loss is proportional to the expected payoff that the individual would obtain if he rather decided to act immediately. Hence, it is endogenously more costly to keep accumulating evidence when the information points towards the state in which the individual can obtain the highest payoff than when it points towards other states. So, the individual will be satisfied with fewer information in favor of the alternative he likes most than with information in favor of the other ones. As a result, he is more likely to believe that he is capable of obtaining the highest rewards when in fact he is not than to believe that he is not capable when in fact he is.

Naturally, it would be absurd for a number of reasons to conclude that this explanation can account for all biases documented in the psychology literature. First, because the ingredients of our model are not relevant in many settings.[14] Second, because some aggregate beliefs are impossible to reconcile with statistical inference.[15] And third, because the behavioralist explanations reviewed in the introduction seem to do a good job in some situations. Yet, we feel that adding this extra element to the discussion can be very useful if we want to improve our understanding of the reasons, means and situations where individuals systematically distort their self-perception and actions.

[14] Among other things, the main cost of information has to be proportional to the expected payoff (a delay or a probability of losing the project). Also, stakes have to be sufficiently small, otherwise the incentives of individuals to become perfectly informed before choosing their optimal action will crowd-out other motivations. Last, incomplete information has to be a crucial element at play.

[15] Bayesian inference can be compatible with more than half of the population believing to be above average. However, as explained earlier, the average belief cannot exceed the true average.

References

1. Bénabou, R. and J. Tirole (2001a), "Self-confidence: Intrapersonal Strategies", forthcoming in *Quarterly Journal of Economics*.

2. Bénabou, R. and J. Tirole (2001b), "Personal Rules", *mimeo*, Princeton.

3. Brocas, I. and J.D. Carrillo (2002), "A Theory of Influence", *mimeo*, Columbia.

4. Caplin, A. and J. Leahy (2001a), "Psychological Expected Utility Theory and Anticipatory Feelings", *Quarterly Journal of Economics*, 116, 55-80.

5. Caplin, A. and J. Leahy (2001b), "The Supply of Information by a Concerned Expert", *mimeo*, NYU and Boston University.

6. Carrillo, J.D., and T. Mariotti (2000), "Strategic Ignorance as a Self-Disciplining Device", *Review of Economic Studies*, 67, 529-544.

7. Dixit, A.K., and R.S. Pindyck (1994), *Investment Under Uncertainty*, Princeton University Press.

8. Gilovich, T. (1991), *How We Know What Isn't So*, New York: Free Press

9. Koszegi, B. (2000), "Self-Image and Information Acquisition", *mimeo*, MIT.

10. Manove, M. (1999), "Entrepreneurs, Optimism, and the Competitive Edge", *mimeo*, CEMFI and Boston University.

11. Manove, M., and A.J. Padilla (1999), "Banking (conservatively) with Optimists", *RAND Journal of Economics*, 30, 324-350.

12. Rabin, M. (1995), "Moral Preferences, Moral Constraints, and Self-Serving Biases", *mimeo*, UC Berkeley.

Appendix: Proof of Lemma 1

Type-1 agent.

Date T. Denote $V_1^T(n) = \max\{h(2\mu(n)-1); l(1-2\mu(n))\}$ and let:

$$Y_1^t(n) = V_1^t(n) - h(2\mu(n)-1) \quad \text{and} \quad W_1^t(n) = V_1^t(n) - l(1-2\mu(n)).$$

For $t = T$, we have $Y_1^T(n) = \max\{0; (h+l)(1-2\mu(n))\}$ and $W_1^T(n) = \max\{0; (h+l)(2\mu(n)-1)\}$. Since $\mu(n)$ is increasing in n, $W_1^T(n)$ is non-decreasing and $Y_1^T(n)$ is non-increasing in n. Besides, $\lim_{n\to+\infty}\mu(n) = 1$ and $\lim_{n\to-\infty}\mu(n) = 0$, so there exists $\overline{n}$ defined by $\mu(\overline{n}) = 1/2$ such that for all $n > \overline{n}$ then $\tau_{1,T} = a$, and for all $n < \overline{n}$ then $\tau_{1,T} = b$.

Date $T-1$.
Case-1: $n \geqslant \overline{n}$. $V_1^{T-1}(n) = \max\{h(2\mu(n)-1); \delta\nu(n)V_1^T(n+1) + \delta(1-\nu(n))V_1^T(n-1)\}$ and

$$Y_1^{T-1}(n) = \max\{0, -(1-\delta)h(2\mu(n)-1) + \delta\nu(n)Y_1^T(n+1) + \delta(1-\nu(n))Y_1^T(n-1)\}$$

where $Y_1^{T-1}(n)$ is defined on $(\overline{n}, +\infty)$. Since $\nu(n)$ is increasing in n and $Y_1^T(n)$ is non-increasing in n, we can check that the right-hand side (r.h.s.) of $Y_1^{T-1}(n)$ is decreasing in n, and therefore there exists a cutoff $n_{1,T-1}^{**}$ such that for all $n > n_{1,T-1}^{**}$ then $\tau_{1,T-1} = a$, and for all $n \in [\overline{n}, n_{1,T-1}^{**})$ then $\tau_{1,T-1} = w$. To solve the previous equation, the cutoff has to be such that $n_{1,T-1}^{**} + 1 \geqslant \overline{n}$ and $n_{1,T-1}^{**} - 1 < \overline{n}$, and therefore it is the solution of:

$$0 = h \cdot f(n_{1,T-1}^{**}, \delta) - l \cdot g(n_{1,T-1}^{**}, \delta)$$

where $f(n_{1,T-1}^{**}, \delta) \equiv 2\mu(n_{1,T-1}^{**}) - 1 - \delta\nu(n_{1,T-1}^{**})(2\mu(n_{1,T-1}^{**}+1) - 1)$ and $g(n_{1,T-1}^{**}, \delta) = \delta(1-\nu(n_{1,T-1}^{**}))(1-2\mu(n_{1,T-1}^{**}-1))$. Differentiating with respect to h, l and δ we have:[16]

$$\begin{aligned}
\frac{\partial n_{1,T-1}^{**}}{\partial h}\left[l \cdot g_n(n_{1,T-1}^{**}, \delta) - h \cdot f_n(n_{1,T-1}^{**}, \delta)\right] &= f(n_{1,T-1}^{**}, \delta) \\
\frac{\partial n_{1,T-1}^{**}}{\partial l}\left[h \cdot f_n(n_{1,T-1}^{**}, \delta) - l \cdot g_n(n_{1,T-1}^{**}, \delta)\right] &= g(n_{1,T-1}^{**}, \delta) \\
\frac{\partial n_{1,T-1}^{**}}{\partial \delta}\left[l \cdot g_n(n_{1,T-1}^{**}, \delta) - h \cdot f_n(n_{1,T-1}^{**}, \delta)\right] &= h \cdot f_\delta(n_{1,T-1}^{**}, \delta) - l \cdot g_\delta(n_{1,T-1}^{**}, \delta)
\end{aligned}$$

Given $f(n_{1,T-1}^{**}, \delta) > 0$, $g(n_{1,T-1}^{**}, \delta) > 0$, $l \cdot g_n(n_{1,T-1}^{**}, \delta) - h \cdot f_n(n_{1,T-1}^{**}, \delta) < 0$, $h \cdot f_\delta(n_{1,T-1}^{**}, \delta) - l \cdot g_\delta(n_{1,T-1}^{**}, \delta) < 0$, we finally have:

$$\frac{\partial n_{1,T-1}^{**}}{\partial h} < 0, \quad \frac{\partial n_{1,T-1}^{**}}{\partial l} > 0, \quad \frac{\partial n_{1,T-1}^{**}}{\partial \delta} > 0.$$

[16] The subscripts n and δ denote a partial derivative with respect to that argument.

Case-2: $n \leqslant \overline{n}$. $V_1^{T-1}(n) = \max\{l(1-2\mu(n)); \delta\nu(n)V_1^T(n+1) + \delta(1-\nu(n))V_1^T(n-1)\}$ and

$$W_1^{T-1}(n) = \max\{0, -(1-\delta)l(1-2\mu(n)) + \delta\nu(n)W_1^T(n+1) + \delta(1-\nu(n))W_1^T(n-1)\}$$

where $W_1^{T-1}(n)$ is defined on $(-\infty, \overline{n})$. Since $\nu(n)$ is increasing in n and $W_1^T(n)$ is non-decreasing in n, we can check that the r.h.s. of $W_1^{T-1}(n)$ is increasing in n, and therefore there exists a cutoff $n^*_{1,T-1}$ such that for all $n \in (n^*_{1,T-1}, \overline{n}]$ then $\tau_{1,T-1} = w$, and for all $n < n^*_{1,T-1}$ then $\tau_{1,T-1} = b$. This cutoff has to be such that $n^*_{1,T-1} + 1 > \overline{n}$ and $n^*_{1,T-1} - 1 \leq \overline{n}$, so it is solution of:

$$0 = l \cdot x(n^*_{1,T-1}, \delta) - h \cdot y(n^*_{1,T-1}, \delta)$$

where $x(n^*_{1,T-1}, \delta) = 1 - 2\mu(n^*_{1,T-1}) - \delta(1-\nu(n^*_{1,T-1}))(1-2\mu(n^*_{1,T-1}-1))$ and $y(n^*_{1,T-1}, \delta) = \delta\nu(n^*_{1,T-1})(2\mu(n^*_{1,T-1}+1)-1)$. Again, differentiating with respect to h, l and δ we have:

$$\begin{aligned}
\frac{\partial n^*_{1,T-1}}{\partial h}\left[l \cdot x_n(n^*_{1,T-1}, \delta) - h \cdot y_n(n^*_{1,T-1}, \delta)\right] &= y(n^*_{1,T-1}, \delta) \\
\frac{\partial n^*_{1,T-1}}{\partial l}\left[h \cdot y_n(n^*_{1,T-1}, \delta) - l \cdot x_n(n^*_{1,T-1}, \delta)\right] &= x(n^*_{1,T-1}, \delta) \\
\frac{\partial n^*_{1,T-1}}{\partial \delta}\left[l \cdot x_n(n^*_{1,T-1}, \delta) - h \cdot y_n(n^*_{1,T-1}, \delta)\right] &= h \cdot y_\delta(n^*_{1,T-1}, \delta) - l \cdot x_\delta(n^*_{1,T-1}, \delta)
\end{aligned}$$

Given $y(n^*_{1,T-1}, \delta) > 0$, $x(n^*_{1,T-1}, \delta) > 0$, $l \cdot x_n(n^*_{1,T-1}, \delta) - h \cdot y_n(n^*_{1,T-1}, \delta) < 0$, $h \cdot y_\delta(n^*_{1,T-1}, \delta) - l \cdot x_\delta(n^*_{1,T-1}, \delta) > 0$, we finally have:

$$\frac{\partial n^*_{1,T-1}}{\partial h} < 0, \quad \frac{\partial n^*_{1,T-1}}{\partial l} > 0, \quad \frac{\partial n^*_{1,T-1}}{\partial \delta} < 0.$$

The proof is completed using a simple recursive method.[17]

Case-1: $n \geqslant \overline{n}$. $V_1^{t-1}(n) = \max\{h(2\mu(n)-1); \delta\nu(n)V_1^t(n+1) + \delta(1-\nu(n))V_1^t(n-1)\}$ and

$$\begin{aligned}
Y_1^t(n) &= \max\{0, -(1-\delta)h(2\mu(n)-1) + \delta\nu(n)Y_1^{t+1}(n+1) + \delta(1-\nu(n))Y_1^{t+1}(n-1)\} \\
Y_1^{t-1}(n) &= \max\{0, -(1-\delta)h(2\mu(n)-1) + \delta\nu(n)Y_1^t(n+1) + \delta(1-\nu(n))Y_1^t(n-1)\}
\end{aligned}$$

Suppose that the following assumptions **(A1)-(A5)** hold.

[17] For the reader unfamiliar with this method, the technique is very simple. Basically, we have already proved that some properties (that will be labelled below as **(A1)-(A5)** or **(A1')-(A5')** depending on whether $n \leqslant \overline{n}$ or $n \geqslant \overline{n}$) hold at dates T and $T-1$. The second step consists in assuming that these properties hold at a given date $t \in \{1, ..., T-1\}$, that we leave unspecified. If, starting form this assumption, we are able to prove that the properties also hold at $t-1$, then we have proved that the properties hold for all $t \in \{0, ..., T\}$.

(A1): $Y_1^t(n)$ is non-increasing in n and there exists $n_{1,t}^{**}$ such that $\tau_{1,t} = a$ if $n > n_{1,t}^{**}$ and $\tau_{1,t} = w$ if $n \in [\overline{n}, n_{1,t}^{**})$.
(A2): $Y_1^t(n) \geqslant Y_1^{t+1}(n)$ and therefore $n_{1,t}^{**} > n_{1,t+1}^{**}$.
(A3): $Y_1^t(n,h) \leqslant Y_1^t(n,h')$ if $h > h'$ (and therefore $\partial n_{1,t}^{**}/\partial h < 0$).
(A4): $Y_1^t(n,l) \geqslant Y_1^t(n,l')$ if $l > l'$ (and therefore $\partial n_{1,t}^{**}/\partial l > 0$).
(A5): $Y_1^t(n,\delta) \geqslant Y_1^t(n,\delta')$ if $\delta > \delta'$ (and therefore $\partial n_{1,t}^{**}/\partial \delta > 0$).

Given **(A1)**, the r.h.s. of $Y_1^{t-1}(n)$ is decreasing in n, so $Y_1^{t-1}(n)$ is non-increasing in n. Therefore, there exists a unique cutoff $n_{1,t-1}^{**}$ such that for all $n > n_{1,t-1}^{**}$ then $\tau_{1,t-1} = a$, and for all $n \in [\overline{n}, n_{1,t-1}^{**})$ then $\tau_{1,t-1} = w$. Also, given **(A2)**, the r.h.s. of $Y_1^{t-1}(n)$ is greater or equal than the r.h.s. of $Y_1^t(n)$ and therefore $Y_1^{t-1}(n) \geqslant Y_1^t(n)$. Overall, both **(A1)** and **(A2)** hold at date $t-1$. Furthermore, $n_{1,t-1}^{**} > n_{1,t}^{**}$. Now, denote:

$$Y_1^{t-1}(n,h) = \max\{0, -(1-\delta)h(2\mu(n)-1) + \delta\nu(n)Y_1^t(n+1,h) + \delta(1-\nu(n))Y_1^t(n-1,h)\}$$

$$Y_1^{t-1}(n,h') = \max\{0, -(1-\delta)h'(2\mu(n)-1) + \delta\nu(n)Y_1^t(n+1,h') + \delta(1-\nu(n))Y_1^t(n-1,h')\}$$

By **(A3)**, if $h > h'$ then $Y_1^t(n+1,h) \leqslant Y_1^t(n+1,h')$ and $Y_1^t(n-1,h) \leqslant Y_1^t(n-1,h')$. Therefore, $Y_1^{t-1}(n,h) \leqslant Y_1^{t-1}(n,h')$. This means that **(A3)** holds at date $t-1$ and, as a consequence, that $\partial n_{1,t-1}^{**}/\partial h < 0$. Using a similar reasoning, it is immediate that **(A4)** and **(A5)** also hold at $t-1$ and therefore that $\partial n_{1,t-1}^{**}/\partial l > 0$ and $\partial n_{1,t-1}^{**}/\partial \delta > 0$.

Case-2: $n \leqslant \overline{n}$. $V_1^{t-1}(n) = \max\{l(1-2\mu(n)); \delta\nu(n)V_1^t(n+1) + \delta(1-\nu(n))V_1^t(n-1)\}$ and

$$W_1^t(n) = \max\{0, -(1-\delta)l(1-2\mu(n)) + \delta\nu(n)W_1^{t+1}(n+1) + \delta(1-\nu(n))W_1^{t+1}(n-1)\}$$
$$W_1^{t-1}(n) = \max\{0, -(1-\delta)l(1-2\mu(n)) + \delta\nu(n)W_1^t(n+1) + \delta(1-\nu(n))W_1^t(n-1)\}$$

Suppose that the following assumptions **(A1')-(A5')** hold.

(A1'): $W_1^t(n)$ is non-decreasing in n and there exists $n_{1,t}^*$ such that $\tau_{1,t} = b$ if $n < n_{1,t}^*$ and $\tau_{1,t} = w$ if $n \in (n_{1,t}^*, \overline{n}]$.
(A2'): $W_1^t(n) \geqslant W_1^{t+1}(n)$ and therefore $n_{1,t}^* < n_{1,t+1}^*$.
(A3'): $W_1^t(n,h) \geqslant W_1^t(n,h')$ if $h > h'$ (and therefore $\partial n_{1,t}^*/\partial h < 0$).
(A4'): $W_1^t(n,l) \leqslant W_1^t(n,l')$ if $l > l'$ (and therefore $\partial n_{1,t}^*/\partial l > 0$).
(A5'): $W_1^t(n,\delta) \geqslant W_1^t(n,\delta')$ if $\delta > \delta'$ (and therefore $\partial n_{1,t}^*/\partial \delta < 0$).

Given **(A1')**, the r.h.s. of $W_1^{t-1}(n)$ is increasing in n, so $W_1^{t-1}(n)$ is non-decreasing in n. Therefore, there exists a unique cutoff $n_{1,t-1}^*$ such that for all $n < n_{1,t-1}^*$ then $\tau_{1,t-1} = b$, and for all $n \in (n^*_{1,t-1}, \overline{n}]$ then $\tau_{1,t-1} = w$. Also, given **(A2')**, the r.h.s. of $W_1^{t-1}(n)$ is greater or equal than the r.h.s. of $W_1^t(n)$ and therefore $W_1^{t-1}(n) \geqslant W_1^t(n)$. Overall, both **(A1')** and **(A2')** hold at date $t-1$. Furthermore, $n^*_{1,t-1} < n^*_{1,t}$. Now, denote:

$$W_1^{t-1}(n,h) = \max\{0, -(1-\delta)l(1-2\mu(n)) + \delta\nu(n)W_1^t(n+1,h) + \delta(1-\nu(n))W_1^t(n-1,h)\}$$

$$W_1^{t-1}(n,h') = \max\{0, -(1-\delta)l(1-2\mu(n)) + \delta\nu(n)W_1^t(n+1,h') + \delta(1-\nu(n))W_1^t(n-1,h')\}$$

By **(A3')**, if $h > h'$ then $W_1^t(n+1,h) \geqslant W_1^t(n+1,h')$ and $W_1^t(n-1,h) \geqslant W_1^t(n-1,h')$. Therefore, $W_1^{t-1}(n,h) \geqslant W_1^{t-1}(n,h')$. This means that **(A3')** holds at date $t-1$ and, as a consequence, that $\partial n^*_{1,t-1}/\partial h < 0$. Using a similar reasoning, it is immediate that **(A4')** and **(A5')** also hold at $t-1$ and therefore that $\partial n^*_{1,t-1}/\partial l > 0$ and $\partial n^*_{1,t-1}/\partial \delta < 0$.

Type-2 agent.

From equations (1) and (2), it is immediate to notice that type-1 and type-2 agents are fully symmetric. Therefore, if at date t there exists $n^{**}_{1,t}$ s.t. $\tau_{1,t} = a$ if $n > n^{**}_{1,t}$ and $\tau_{1,t} = w$ if $n \in [\overline{n}, n^{**}_{1,t})$, then there also exists $n^*_{2,t}$ s.t. $\tau_{2,t} = b$ if $n < n^*_{2,t}$ and $\tau_{2,t} = w$ if $n \in (n^*_{2,t}, \overline{n}]$. Furthermore, $n^*_{2,t}$ is such that $\overline{n} - n^*_{2,t} = n^{**}_{1,t} - \overline{n}$, that is $\mu(n^{**}_{1,t}) = 1 - \mu(n^*_{2,t})$. Similarly, if at date t there exists $n^*_{1,t}$ s.t. $\tau_{1,t} = b$ if $n < n^*_{1,t}$ and $\tau_{1,t} = w$ if $n \in (n^*_{1,t}, \overline{n}]$, then there also exists $n^{**}_{2,t}$ s.t. $\tau_{2,t} = a$ if $n > n^{**}_{2,t}$ and $\tau_{2,t} = w$ if $n \in [\overline{n}, n^*_{2,t})$. Furthermore, $n^{**}_{2,t}$ is such that $n^{**}_{2,t} - \overline{n} = \overline{n} - n^*_{1,t}$, that is $\mu(n^*_{1,t}) = 1 - \mu(n^{**}_{2,t})$.

Note that if $h = l$, then for all t we have $\mu(n^*_{1,t}) = 1 - \mu(n^{**}_{1,t})$ and $\mu(n^*_{2,t}) = 1 - \mu(n^{**}_{2,t})$. As a result, $n^*_{2,t} = n^*_{1,t} < \overline{n}$ and $n^{**}_{2,t} = n^{**}_{1,t} > \overline{n}$. Also, we know that $\frac{\partial n^{**}_{1,t}}{\partial h} < 0$ and $\frac{\partial n^*_{1,t}}{\partial h} < 0$ (which, again by symmetry, implies that $\frac{\partial n^*_{2,t}}{\partial h} > 0$ and $\frac{\partial n^{**}_{2,t}}{\partial h} > 0$). Therefore, for all $h > l$ we have $n^*_{1,t} < n^*_{2,t} < \overline{n} < n^{**}_{1,t} < n^{**}_{2,t}$.

Summing up, when $\delta < 1$, $h > l > 0$ and $T \to +\infty$, we have $n^*_1 < n^*_2 < \overline{n} < n^{**}_1 < n^{**}_2$ where $\mu(n^{**}_1) = 1 - \mu(n^*_2)$ and $\mu(n^*_1) = 1 - \mu(n^{**}_2)$. Moreover, $\frac{\partial n^{**}_1}{\partial h} < 0$, $\frac{\partial n^{**}_1}{\partial l} > 0$, $\frac{\partial n^{**}_1}{\partial \delta} > 0$, $\frac{\partial n^*_1}{\partial h} < 0$, $\frac{\partial n^*_1}{\partial l} > 0$, $\frac{\partial n^*_1}{\partial \delta} < 0$ and $\frac{\partial n^*_2}{\partial h} > 0$, $\frac{\partial n^*_2}{\partial l} < 0$, $\frac{\partial n^*_2}{\partial \delta} < 0$, $\frac{\partial n^{**}_2}{\partial h} > 0$, $\frac{\partial n^{**}_2}{\partial l} < 0$, $\frac{\partial n^{**}_2}{\partial \delta} > 0$.