

Motivating Public Information Acquisition

Andres Zambrano *
Universidad de los Andes

September 10, 2016

Abstract

Successful innovators have become billionaires by generating breakthrough technologies that are later widely adopted in the society. This paper proposes a model to explain why this is (constrained) efficient. I consider an environment where a population of agents must choose between pursuing a risky project or a safe project. Before taking the decision, each agent can acquire information about the risky project by exerting unobservable effort. Then an aggregate statistic of the acquired information in the economy becomes public and each agent picks a project. In equilibrium a free-riding problem arises and the amount of agents acquiring information decreases compared to the efficient allocation. I then study the optimal contract designed by a social planner who wants to maximize social welfare but does not observe individual efforts. The optimal contract divides the population between nonexperimenters and a few experimenters that exert high effort, thus substituting the extensive margin of information by its intensive margin. It also splits the total returns of the risky project among experimenters when the unknown project yields significantly greater returns than the safe project. Therefore, it explains why paying so much to successful innovators is optimal.

Keywords: Information Acquisition, Public Information, Contract, Technology Adoption

JEL codes: D82, D83, D86, O31

*E-mail: ja.zambrano@uniandes.edu.co. Please download the most recent version at <https://sites.google.com/site/andreszam/>. The author is indebted to Andrew Atkeson and Hugo Hopenhayn for their helpful advice during all these years. I also gratefully acknowledge the comments from Simon Board, Jernej Copic, Fernando Leibovici, Moritz Meyer-ter-Vehn, Ichiro Obara, John Riley, Venky Venkateswaran, Pierre-Olivier Weill, and Bill Zame, as well as participants of the Macro Lunch seminar, Macro Proseminar, IO Proseminar and Theory Proseminar at UCLA, Workshop in Theory at Universidad de Chile, and Universidad del Rosario. Funding from the Colombian Central Bank and NBER is also acknowledged. Remaining errors are mine.

1 Introduction

Recognized innovators such as Bill Gates and Steve Jobs have become billionaires by creating breakthrough technologies that are later widely used in the society. Could we achieve similar levels of innovation without having to pay them so much for such new technologies? Why is it that other innovators who create more marginal improvements are not paid as much? (Ebersberger et al., 2008; Marsili and Salter, 2005) This paper proposes a model of delegated expertise with multiple agents to explain these features.

The model features a population of risk neutral agents who must decide between pursuing a safe project with known returns and a risky project with unknown returns. Each agent decides simultaneously how much effort to exert to obtain a signal about the risky returns; such effort determines the precision of their signal. Then, an aggregate signal, which precision is the sum of the individual efforts, is revealed to all the agents and the agents decide simultaneously between the safe and the risky project.

I first study the equilibrium of this game and compare it with the efficient outcome, which can be implemented by a social planner that observes individual efforts. Whereas the first best is characterized by lots of agents investing in the smallest possible precision, in the equilibrium there are fewer agents investing also in the smallest possible precision. The separation of the agents occurs endogenously since there is a minimum scale effort in case an agent decides to experiment, that assures the existence of the equilibrium and prevents everybody to become an experimenter.¹ The equilibrium features free-riding and the overall amount of available information is smaller. Agents prefer to wait until others incur in the cost of experimentation since the next stage they will observe the aggregate signal and take a more informed decision. Hence, the probability of choosing a new and better technology decreases with respect the first best because of the underlying incentives.

¹Instead of imposing a minimum effort for experimenters, one can impose a fixed cost for experiments yielding the same result.

The second contribution of this paper is to characterize optimal contracts that increase the information in this economy. I take a mechanism approach where the social planner wants to maximize social welfare but does not observe individual efforts. Optimal contracts under budget balance² splits the population between nonexperimenters and a few experimenters that must exert high effort, and proposes a wage contingent on the type of agent.

The optimal wage must balance two countervailing effects. On one hand, there is a moral hazard effect represented by a likelihood ratio, as is usual in hidden action problems, comparing (marginally) the probability of undertaking the risky project when more effort is exerted. Although effort in this environment does not generate greater expected returns directly, it does increase the expected returns indirectly by improving decision making. Therefore, the likelihood ratio suggests that contracts should be monotone increasing in the return of the risky project for experimenters.

On the other hand, there is a free riding effect represented also by the same likelihood ratio, expressed this time in discrete terms, but with the opposite sign. The planner has now some pressure to reward experimenters for low returns since otherwise experimenters will prefer to become non-experimenters, decreasing the amount of information in the economy and increasing the probability of being rewarded for low outcomes.

However, as the spillovers of the innovation become large, i.e. the size of the population increases, the moral hazard effect dominates. The reason is that the previous deviation is no longer profitable since such returns have to be split among too many people. Therefore, under limited liability, the optimal contract should split the total output among nonexperimenters if the risky technology is not significantly better than the safe one to provide them incentives to not become experimenters. And output should be split among experimenters if the risky technology is chosen and its return is significantly better than the return of the safe technology, thus imposing more risk on them.

²If budget balance is not imposed, then the first best is attainable.

Since agents can learn more accurately about the risky technology by having more people experimenting or by increasing the precision per experimenter, the planner substitutes the extensive margin by a greater intensive margin to provide more powerful incentives. Moreover, it pays the experimenters the whole surplus of the society when the project becomes very successful. This constrained efficient outcome resembles the features exposed at the beginning where a few successful innovators will get paid a large amount of money, whereas not so successful entrepreneurs, and the rest of the society will not be paid as much. The result has implications for optimal taxation for innovators, payment schemes to encourage technology adoption and patent licensing.

1.1 Literature Review

Models of delegated expertise were first proposed by Lambert (1986) and Demski and Sappington (1987) in the context of a single principal and a single agent. I adapt this framework to study incentives to innovation in presence of free-riding. An extension to two agents was done by Gromb and Martimort (2007) who explore the problem in presence of vertical and horizontal collusion. However, they study the case where the gathered information is not observable and the agents must be encouraged to report it truthfully.

The acquisition of information is also related to bandit problems where an agent can learn about the return of a project by undertaking it as in Manso (2011). This approach was extended to two agents by Ederer (2008). This paper departs from this framework by enriching the information acquisition process and allowing agents to invest in the precision of their signals, thus exploiting the substitution between the intensive and the extensive margin. Bonatti and Horner (2011) study moral hazard in teams over time where the return of a project is unknown and effort determines the rate of arrival of the return. My setup is different in that individuals invest one time on a signal before deciding to undertake the risky project, and in that I solve for optimal contracts to improve the equilibrium.

Papers that study experimentation in teams using a bandit structure include

Bolton and Harris (1999), Keller et al. (2005), and Klein and Rady (2011). Closer to our paper is Malueg and Tsutsui (1997), on their model firms compete over time by investing in R&D on an uncertain project. All these mentioned papers characterize equilibria when actions are observable whereas our main question is how to improve the allocation when actions are private information. On the other hand, Rosenberg et al. (2007), Hopenhayn and Squintani (2011), and Murto and Valimaki (2011) study problems where actions chosen by agents are observed and outcomes are unobserved. This is the opposite case as the model proposed here where signals are perfectly observed and actions are private information.

The next section introduces the setup for one agent. The third section studies the equilibrium and compares it with the first best allocation. The fourth section characterizes the optimal contract when the precision chosen by each agent is unobservable. In the next section I discuss how to implement the contract in different real-world applications. In the last section I conclude.

2 One Agent Setup

Consider the case of a risk agent who has to decide between a safe project with known returns and a risky project with unknown returns.³ The safe project has a net return $y_s > 0$. There is also a risky one whose return $y_r \in [0, \bar{y}]$ is unknown, with $\bar{y} > y_s$. The agent has a nondegenerate prior belief $g(y_r)$ over the unknown return with finite mean μ_0 .

The agent can generate information about the risky project by acquiring a continuum of $e > e_{\min}$ independent signals at a cost $C(e)$; this cost can be associated with R&D expenditures or the cost of running trials. Assume the cost function satisfies $C(0) = 0$ and $C(e_{\min}) = c$, is increasing and is strictly convex in e . The minimum scale effort e_{\min} is associated to a fixed cost of starting the process of experimentation.

Each independent signal x_k is drawn from the distribution $f(x_k|y_r)$, for each

³The individual can be in fact risk averse or risk lover, just let the returns perceived by the agent be measured in utils and let the agent maximize a Von Neuman-Morgenstern utility function

$k \in [0, e]$. We will refer to e as the precision of the information and denote by x a sufficient statistic of the signals. Let the conditional pdf and cdf of x be denoted by $f(x|y_r, e)$ and $F(x|y_r, e)$, respectively, with support $[\underline{x}, \bar{x}]$.⁴ Assume both functions are twice differentiable in e and x . Note, however, that the prior distribution of y_r is independent of e . Let the sufficient statistics be ordered, following Milgrom (1981), a signal x is more favorable than signal x' if the posterior distribution $g(y_r|x, e)$ first order stochastically dominates the posterior distribution $g(y_r|x', e)$.

The agent faces the following problem:

$$\max_e \mathbb{E}_x \left[\max_{j_x \in \{s, r\}} \mathbb{E}_{y_{j_x}} [y_{j_x}|x, e] \right] - C(e)$$

where j_x is the project chosen by the principal when x is observed. Since there are two stages, we proceed to solve the individual's problem using backward induction. That is, I will first determine which project is going to be chosen given the information acquired. Second, I find the optimal precision e given the agent decides to experiment. Finally, I characterize when the agent decides to experiment as a function of y_s and c .

Without loss of generalization, let $x = \mathbb{E}_{y_r} [y_r|x, e]$ be the posterior mean of the risky project.⁵ The individual will choose the risky project if $x > y_s$, thus the payoff of the second period is given by $\max\{x; y_s\}$. Note this is a convex function of x . The value of experimentation is defined as the ex ante expectation of the utility in the second period, that is

$$U(e) = \mathbb{E}_x [\max\{x; y_s\}]$$

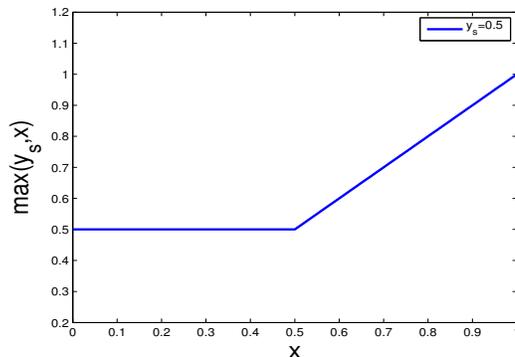
⁴Alternatively, the acquisition of information can be modelled as the purchase of a signal x with precision e defined in the Blackwell (1951) sense. That is experiment X is more precise than experiment X' if you can mimic signal X' by adding noise to signal X . Formally, experiment X is sufficient for (more precise than) experiment X' if for every $x' \in X'$ there is a probability distribution over X , $g(x; x')$ where $x \in X$, such that $\int g(x; x') f(x|y) dx = f(x'|y)$ for any y .

⁵Since signals are ordered, the posterior mean will be a monotone transformation of the signal. Thus the distribution of the posterior mean will be a transformation of the distribution of the sufficient statistic. Hence, we can let $\underline{x} = 0$ and $\bar{x} = \bar{y}$

From the previous properties we can prove the following lemma:

Proposition 1. *The value of experimentation $U(e)$ is greater than $\max\{\mu_0; y_s\}$, and is differentiable, strictly increasing, and bounded in e*

Figure 1: **Utility in second period**



The intuition of the proposition is as follows. The expected posterior mean $\mathbb{E}_x[x]$ is just the prior mean μ_0 , which means that the learning process is a Martingale and assures that with enough signals or with a sufficiently large precision we would eventually learn. Therefore, as e approaches infinity, the sufficient statistic perfectly reveals the unknown return, i.e. it is a known transformation of the unknown return. In other words, the limit conditional distribution $\lim_{e \rightarrow \infty} f(x|y_r, e)$ will be degenerate at y_r . Hence, the unconditional distribution of the posterior mean converges to the prior distribution. Moreover, the distribution $f(x|e)$, second order stochastically dominates the distribution $f(x|e')$, for all $e < e'$. Since the utility is convex in x , then the individual prefers a higher e .

If at time 0 the individual chooses to experiment, then she will choose precision e to maximize:

$$\max_e U(e) - C(e) \tag{1}$$

Unfortunately we cannot assure this is a concave problem since $U(e)$ may be convex for low values of precision (Moscarini and Smith, 2002). However, since the

option value is bounded and the cost is strictly convex, we can assure the existence of a solution to the problem. The next proposition characterizes the solution.

Proposition 2. *A solution to problem (1) exists. If the solution is interior, the optimal precision e^* is characterized by*

$$C_e(e^*) = U_e(e^*) \quad (2)$$

Moreover, if $\int_0^y F(x|e) dx$ is strictly concave in e , then condition (2) is also sufficient and the maximum is unique.

The strict concavity condition implies that even though increasing the precision generates a mean preserving spread on the distribution, such spread becomes smaller as the precision becomes larger. Such condition was also suggested by Szalay (2009) to motivate the use of the first order approach in a procurement problem. The next lemma shows how the optimal precision depends on the prior belief.

Proposition 3. *If $\int_0^y F(x|e) dx$ is concave in e then the optimal precision e^**

1. *Achieves a unique maximum when $y_s = \mu_0$*
2. *Is strictly increasing in y_s as long as $y_s < \mu_0$*
3. *Is strictly decreasing in y_s as long as $y_s > \mu_0$.*

At the beginning of the first period, the agent will decide to experiment if

$$U(e^*) - C(e^*) \geq \max\{y_s; \mu_0\}$$

Let the maximized objective function of the agent be denoted by $V(y_s, c) = \max\{U(e^*) - C(e^*); y_s; \mu_0\}$.

Proposition 4. *The function $V(y_s, c)$ is nondecreasing and convex in y_s and the principal decides to experiment, $e^* > 0$, when $c \leq \hat{c}$ and $y_s \in (a_c, b_c) \subseteq (0, \bar{y})$, where*

$\mu_0 \in (a_c, b_c)$. Moreover, such interval is decreasing in c , that is $(a_c, b_c) \subset (a_{c'}, b_{c'})$ for any $c < c' < \hat{c}$, with $(a_0, b_0) = (0, \bar{y})$ and $a_{\hat{c}} = \mu_0 = b_{\hat{c}}$.

Even if beliefs are relatively pessimistic the individual decides to acquire information because of the potential gain represented by the value of experimentation. Indeed the precision chosen is increasing in the beliefs when they are pessimistic and achieves a maximum when ex ante the two projects have the same expected mean. When beliefs start to be optimistic, precision decreases, until a point where the agent does not have any more incentives to acquire information. The lower is the fixed cost c , the greater is the interval over which the agent decides to experiment.

3 Team Problem with Multiple Agents

Suppose there is a population of N risk neutral agents who must choose between a risky project and a safe project. Let the agents appropriate the whole return of the project they choose. All agents share the same prior beliefs and can acquire information about the risky project by exerting costly effort. Assume the information gathered becomes public and other agents can use it to update beliefs. Since information is now a public good there will be free-riding in equilibrium and there will be less information than in the first best.

The utility functions for each agent and the returns for each project are given as before. Of particular importance will be the fixed cost incur by experimenters since it will generate a natural partition of the population between experimenters and non-experimenters. The fixed cost will also assure the existence of an equilibrium when the size of the population goes to infinity. Assume that the sufficient statistic obtained from the information acquisition of all agents is publicly observed by the agents. In this scenario agents decide whether to experiment and obtain costly extra signals or use the available signals from others.

I consider again the case of two stages. In the first stage the individuals decide simultaneously whether to acquire information or not. If an agent i decides to experiment, she must also choose how many signals to acquire, e_i at a cost $C(e_i)$. Assume

signals acquired by different individuals are independent and identically distributed with pdf $f(x_k|y_r)$. At the end of the period each individual i observes the overall sufficient statistic x with pdf given by $f\left(x|y_r, \sum_{j=1}^N e_j\right)$. At the beginning of the second stage each agent updates her beliefs and decides which project to pursue. Since all signals are public, in equilibrium everybody will take the same decision in the final period.

3.1 Equilibrium

Suppose first that each agent appropriates the return from the chosen project and thus the payoff for agent i is given by $y - C(e_i)$. In order to simplify the analysis I will focus on equilibria where experimenters choose the same level of precision. Let the number of experimenters be denoted by n . It is important to note that not every agent will be necessarily an experimenter in equilibrium, a nonexperimenter will set $e_i = 0$. Let e_{-i} be the sum of precision chosen by all individuals except i . Therefore each agent is willing to solve the following problem:

$$\max_{e_i} E_x \left[\max_{j_x \in \{s,r\}} \mathbb{E}_{y_{j_x}} [y_{j_x} | x, e_i + e_{-i}] \right] - C(e_i)$$

Following the analysis of the previous section, we will define an option value of experimentation for an agent i that this time will depend on the overall precision chosen by all the experimenters:

$$U(e_i + e_{-i}) = \mathbb{E}_x \left[\max_{j_x \in \{s,r\}} \mathbb{E}_{y_{j_x}} [y_{j_x} | x, e_i + e_{-i}] \right]$$

Therefore, the ex-ante utility for a non-experimenter is given by $U(e_{-i})$. On the other hand, the ex-ante utility for experimenters is given by $\max_{e_i} \{U(e_i + e_{-i}) - C(e_i)\}$. Assume first that agents also observe the overall precision, hence the appropriate concept for equilibrium is a Subgame Perfect Nash equilibrium (SPE).⁶ The number of

⁶Alternatively we can assume that the overall precision is unobservable, then the appropriate concept is a Perfect Bayesian Nash equilibrium (PBE). However, there exists a PBE with the same

equilibria depends on the properties of the cost function and the distribution of the posterior expected mean. Assume for the rest of the paper that $\int_0^y F(x|e_i + e_{-i}) dx$ is strictly concave to obtain a unique solution (recall Proposition (2)).

Definition 5. *A experimenter-symmetric SPE is defined by the number of people experimenting, n^* , and the symmetric precision of the signals e^* , such that nobody has incentives to deviate, that is*

- *Non-experimenters do not want to deviate:*

$$U(n^*e^*) \geq U(e_i + n^*e^*) - C(e_i) \text{ for any } e_i$$

- *Experimenters do not want to deviate:*

$$U(n^*e^*) - C(e^*) \geq U((n^* - 1)e^*)$$

and

$$e^* = \arg \max_{e_i} \{U(e_i + (n^* - 1)e^*) - C(e_i)\}$$

In this experimenter-symmetric equilibrium the overall precision will be n^*e^* . These two variables will play the same role in the learning process since an increase in either one will have the same effect (in terms of elasticities) on the value of experimentation. In fact, this function will have the same properties with respect the overall precision ne as the ones described in the previous section.

As before, we will use backward induction to solve for the agent's behavior. In a symmetric equilibrium where n people experiment in the first period, an experimenter will choose the optimal precision e^* such that

$$C_e(e^*) = U_{e_i}(e_i + (n^* - 1)e^*)$$

payoffs and actions on the equilibrium path as the ones in the SPE we are interested in. Thus we use this equilibrium concept for the sake of simplicity.

The optimal precision inherits the same properties as the ones found in the previous section. However, we now have an interesting relationship between precision, overall precision, and the number of experimenters as summarized in the following lemma.

Lemma 6. *The optimal precision e^* is decreasing in n , but the overall precision ne^* is increasing in n .*

The intuition behind this lemma is that although n and e have the same relative effect on the value of experimentation, the convexity of the cost induces an imperfect complementarity between the two components of the overall precision. If the cost of the precision were to be linear, the overall precision would remain unchanged since an increase in the number of signals will be exactly offset by the decrease in the precision. The convexity of the cost function implies that effort becomes less responsive to changes in n .

The monotonic response of the overall precision to n implies that $U(n^*e^*) > U((n^* - 1)e^*)$ since the value of experimentation is monotone increasing in the overall precision. However, these increments will become smaller after some threshold because of the concavity of such function. Therefore, the equilibrium number of experimenters n^* is given by the equation⁷

$$U(n^*e^*) - C(e^*) \geq U((n^* - 1)e^*)$$

Lemma 7. *In a experimenter-symmetric SPE, every agent will experiment if $N < \hat{N}$, and the number of experimenters will be independent of N as long as $N > \hat{N}$*

The concavity and boundedness of the option value of experimentation jointly with the fixed cost of experimentation implies that the number of experimenters will be finite in equilibrium even as the size of the population goes to infinity. On the other hand, the optimal precision is set such that its marginal cost is greater than the

⁷In differentiable terms this condition is equivalent to $C(e^*) \geq U_n(n^*e^*)$, with equality if $n > 0$.

average cost. However, as the total amount of signals acquired increases (the total number of experimenters increase), the optimal precision converges to the minimum efficient scale, which ensures its finiteness. Note the similarity of the problem solved by a firm facing perfect competition with fixed costs.

3.2 First Best

Suppose now that there exists a social planner who wants to maximize the ex-ante total welfare of this economy. Given that n agents are experimenting, ex-ante total welfare is then defined as:

$$W = NU \left(\sum_{i=1}^n e_i \right) - \sum_{i=1}^n C(e_i)$$

Note that in the aggregate, the number of experimenters are associated with a linear cost, whereas the precision of the signals have a convex one. Note also that the social planner must consider the externality generated by the signals by multiplying the value of experimentation by the number of agents in the economy. Given a number of experimenter n , the social planner chooses e^{FB} such that

$$C_e(e^{FB}) = NU_e(ne^{FB})$$

The social planner will also increase the number of experimenters n^{FB} as long as⁸

$$C(e^{FB}) \leq N [U(n^{FB}e^{FB}) - U((n^{FB} - 1)e^{FB})]$$

Lemma 8. *In the first best, the number of experimenters goes to infinity as N goes to infinity, but its proportion $\frac{n^{FB}}{N}$ goes to zero. The first best precision remains finite and converges to the minimum efficient scale.*

⁸In differentiable terms this condition is equivalent to $C(e^{FB}) \geq NU_n(n^{FB}e^{FB})$, with equality if $n^{FB} > 0$.

Note again how N increases the marginal value of experimentation, this implies that in the first best there is more experimentation n than in equilibrium. Moreover, the number of experimenters grows without bound as the size of the population goes to infinity. However, n does not increase as fast as N because of the concavity and boundedness of $U(\cdot)$. Note also that the central limit theorem implies that the speed of convergence of the learning process is \sqrt{ne} and thus is not optimal to increase n as fast as N .

On the other hand, while an increase in N has a direct positive effect on e pushing it to infinity, there is also an indirect effect of N coming through n which is growing bigger and decreases e . It turns out that these effects approximately offset and the first best precision converges to the point where average cost is minimized, which is close to the same precision chosen in equilibrium.⁹

Because of the greater marginal benefit and the linearity of the costs associated to the number of experimenters, a social planner decides to choose a greater n than the one obtained in equilibrium. On the other hand, because of the convexity of the cost associated to the precision, the optimal precision remains finite and close to the equilibrium one. In other words, the social planner decides to increase n and maintain fixed e since their relative effect on the option value of experimentation is the same but is more costly to increase e .

The first best can be implemented under budget balance if the individual precision is observed. For example, by distributing the surplus from adopting the new technology among the experimenters when they choose the first best precision, but not when they deviate, the first best is implemented. To see this note that if experimenters follow the suggested first best precision their payoff is given by

$$\frac{N}{n^{FB}} U(n^{FB} e^{FB}) - C(e^{FB})$$

whereas the payoff for nonexperimenters or experimenters that deviate from the

⁹The discreteness of n is what prevents them to be equal, but as N grows large the choice of n resembles the case of a continuum number of experimenters.

first best precision is 0. From the conditions obtained for the first best we know that the payoff for obedient experimenters will be greater than 0 and thus they will not want to deviate. Note also that because the suggested wages are nondecreasing the agents will be willing to adopt the best technology.

4 Constrained Efficiency

Suppose there exists a social planner that wants to maximize ex-ante total welfare as before, but this time it does not observe the precision chosen by experimenters. In addition assume both agents have limited liability as before and that there must be budget balance. Under this restriction, the wages suggested in the last subsection do not implement the first best since agents would have incentives to choose a lower precision because they do not internalize the social gains since they are split among all the experimenters. Assume that the social planner only observes the final output and designs a wage schedule for experimenters and nonexperimenters as a function of the observed returns. Let $w(y)$ and $v(y)$ be the wages for an experimenter and a nonexperimenter when the observed return is y .

The timing is as follows. First the social planner offers the menu of contracts. Then each agent chooses a contract. Agents who decided to experiment then choose the precision of their signals. At the end of the stage every agent observes the aggregate sufficient statistic. Then agents update beliefs given the reported information on how many people experimented, their precision, and the observed aggregate signal. Then agents take a definitive decision over which technology to use. After the final output is realized the social planner pays the promised contract.

Solving the problem using backward induction, agents will choose the risky project whenever the aggregate signal x is greater than the safe return y_s . Otherwise they choose the safe project. To assure the agents exert the effort suggested by the social planner, the latter must solve the following problem:

$$\max_{n, \{e_i\}_{i=1}^n, w(y_r), w(y_s), v(y_r), v(y_s)} NU \left(\sum_{i=1}^n e_i \right) - \sum_{i=1}^n C(e_i) \quad (3)$$

subject to

$$\mathbb{E}_x [\mathbb{E}_{y_{j_x}} [w(y_{j_x}) | e_i + e_{-i}]] - C(e_i) \geq \mathbb{E}_x [\mathbb{E}_{y_{j_x}} [v(y_{j_x}) | e_{-i}]] \quad (4)$$

$$\mathbb{E}_x [\mathbb{E}_{y_{j_x}} [v(y_{j_x}) | e_{-i}]] \geq \mathbb{E}_x [\mathbb{E}_{y_{j_x}} [w(y_{j_x}) | e + e_{-i}]] - C(e) \text{ for any } e \quad (5)$$

$$e \in \arg \max \mathbb{E}_x [\mathbb{E}_{y_{j_x}} [w(y_{j_x}) | e + e_{-i}]] - C(e) \quad (6)$$

$$w(y_j), v(y_j) \geq 0 \text{ for } j = r, s \quad (7)$$

$$nw(y_j) + (N - n)v(y_j) \leq Ny_j \text{ for } j = r, s \quad (8)$$

Constraint (4) is the individual rationality constraint for experimenters that prevents them from not exerting any effort and free ride the public information. Constraint (5) is the individual rationality constraint for the nonexperimenter and it states that the individual is better off by not investing in information rather than acquiring a costly extra signal. The incentive compatibility constraint (6) assures the individual is willing to invest in the prescribed precision. Finally, constraints (7) and (8) are the limited liability and budget balance constraints. The latter constraint will be binding since we are maximizing total welfare, thus we can substitute the wage for a nonexperimenter by $v(y_j) = \frac{Ny_j - nw(y_j)}{N - n}$.

Lemma 9. *If (4) and (6) are satisfied, then (5) is satisfied.*

The Lemma suggests that if experimenters are already optimizing then nonexperimenters are not willing to deviate since it would be more costly to acquire information than the benefit from becoming experimenters. Using (8) into the remaining constraints, taking into account that once the signal is realized the agents choose the risky project whenever $x > y_s$, and substituting (6) by the first order approach,¹⁰ the principal's problem is simplified to:

¹⁰The first order approach is valid if $F(x|y_r, e)$ is convex in e and wages are monotone nondecreasing.

$$\max_{n, \{e_i\}_{i=1}^n, w(y_r), w(y_s)} NU \left(\sum_{i=1}^n e_i \right) - \sum_{i=1}^n C(e_i) \quad (9)$$

subject to

$$\int_{\underline{y}}^{\bar{y}} w(y_r) (1 - F(y_s|y_r, e_i + e_{-i})) dy_r + w(y_s) F(y_s|e_i + e_{-i}) - C(e) \geq \int_{\underline{y}}^{\bar{y}} \frac{Ny_r - nw(y_r)}{N - n} (1 - F(y_s|y_r, e_{-i})) dy_r + \frac{Ny_s - nw(y_s)}{N - n} F(y_s|e_{-i}) \quad (10)$$

$$- \int_{\underline{y}}^{\bar{y}} w(y_r) F_e(x, y_r|e_i + e_{-i}) dy_r dx + w(y_s) F_e(y_s|e_i + e_{-i}) - C_e(e_i) \geq 0 \quad (11)$$

$$0 \leq w(y_j) \leq \frac{N}{n} y_j \text{ for } j = s, r \quad (12)$$

Let λ and δ be the Lagrange multipliers for the first two constraints, respectively. The problem is linear on wages and thus there is a bang-bang solution that is bounded by the limited liability constraints. Rearranging the derivative with respect to wages we obtain

$$1 + \left(\frac{n-1}{N-n+1} \right) \left(\frac{1 - F(y_s|y_r, e_{-i})}{1 - F(y_s|y_r, e_i + e_{-i})} \right) - \frac{\delta}{\lambda} \left(\frac{F_{e_i}(y_s|y_r, e_i + e_{-i})}{1 - F(y_s|y_r, e_i + e_{-i})} \right) \quad (13)$$

As it is common in moral hazard problems, the last term is a likelihood ratio that compares the probability of undertaking the risky project when more effort is exerted and the individual increases the precision of the signal. Recall that higher effort induces a second order stochastically dominated unconditional distribution, that is a mean-preserving spread. The numerator is zero in the interior only if

$y_s = y_r$ and positive (negative) if y_r is greater (smaller) than y_s , thus satisfying the single crossing property. Therefore, by rewarding the agent only for high returns, the principal is encouraging her to exert a greater effort.

The second term represents the free-riding trade-off. It is also represented by a likelihood ratio expressed in discrete terms, but this time is a decreasing function of y_r since the numerator has a lower precision. The rationale for having an element decreasing in the returns is that by rewarding the experimenter for extreme outcomes, the experimenter is having incentives to become a nonexperimenter by reducing the overall precision and increasing the probability of being rewarded.

Lemma 10. *There exists a \hat{N} such that, for $N \geq \hat{N}$, optimal wages $w(y_r)$ are characterized by a cutoff $z > y_s$ such that*

$$w(y_r) = \begin{cases} \frac{N}{n}y_r & \text{if } y_r \geq z, \\ 0 & \text{otherwise.} \end{cases} \quad (14)$$

As the population grows, the free-riding effect decreases since the ratio of experimenters converges to zero, thus their relative reward is bigger compared to the one for nonexperimenters. Therefore the incentives to become nonexperimenters disappear and the optimal wage for experimenters becomes monotone nondecreasing in the risky return. The optimal contract suggests that the experimenter is encouraged to increase the variance of the posterior mean by increasing her payoff in realizations that are much better than the returns from the safe project. On the contrary, the total output will be split among non-experimenters when the risky return is not significantly better than the safe return. Note that if the risky return is similar to the safe return, nonexperimenters will appropriate the whole surplus. If this were not to be the case, then nonexperimenters would be better off by choosing the safe project even if the aggregate signal suggests the opposite, and welfare would decrease.

5 Implementation

The analyzed problem can be understood from the perspective of optimal taxation when individuals generate information externalities. The contract suggests that the planner should subsidize experimenters when new projects are adopted and generate sufficiently high returns, otherwise they should be heavily taxed.

In the context of technology adoption, the existence of farmer cooperatives can be used to implement this contract. For example, National Cereal Boards play an important role on the coordination of farmers within a country. In Kenya this organism has played the role of creditors, promoters of research and market development, and regulators (Raikes, 1994). Although in recent decades their policies have been oriented towards a free market, the board used to lobby in the government for the determination of prices. An institution like this one could potentially reward farmers differently according to their willingness to try new technologies.

Another example of such organizations is the National Federation of Coffee Growers of Colombia.¹¹ The federation groups more than 500 thousand colombian families that produce coffee. Its mission is to represent the interests of coffee growers, create social programs to improve the quality of life of the producers, investment in research and knowledge transfer as well as in promotion and advertising, and the commercialization of coffee.

Within this last objective is what is considered by them their most significant service called the Purchase Guarantee Policy. This policy involves the setting of a minimum price at which the coffee should be sold. If no buyers are willing to acquire the product at this price, they commit to purchasing it. This price is public and constitutes a reference point for the market. Hence, the federation's role could be used to differentiate the price paid to experimenting growers who increase the speed of adoption of new technologies given this is in its interests.

The proposed contract can also be implemented within firms. The level of inno-

¹¹<http://www.federaciondecafeteros.org/particulares/en/>

vation or the adoption of better technologies or practices are reflected in the value of the firm. Therefore stock options and profit sharing strategies possess similar properties as our contract. Stock option programs give workers the right to buy company's shares at a fixed price for a given period of time. These will be only exercised if the market price is higher than the strike price originally agreed to. Usually stock options are used as a long-term motivator and the employee is constrained on exercising the option after certain time. Similarly, the firm could constrain the option exercise until the market price crosses a threshold, and thus implementing the proposed contract. Likewise, profit sharing is also used as a long-term motivator where individuals are entitled to a percentage of the profits of a firm after a given period. To implement the contract the firm could set a threshold on the profits such that the workers can only claim her share if profits are greater than such level.

The optimal contract might be also interpreted as a patent policy to encourage innovation. It suggests that patents should only be given if it is shown that the new technology is significantly better than the previous one, and not for marginal improvements. However, this result cannot be interpreted as a restriction on the use of new technology as often happens with patents. In other words, the optimal contract does not allocate the property rights of the new technology. On the contrary, it encourages the adoption of the new technology by all the population, while rewarding innovators with the surplus they generated, suggesting an optimal pricing policy.

6 Conclusions

The paper analyses the problem faced by innovators who can acquire costly information before choosing between a known project and an unknown one, but such information is also observed by others for free. In this context, information is a public good and thus a free-riding problem arises. In equilibrium there will be less experimenters than in the first best since agents do not internalize the social benefits of experimentation.

The first best level of experimentation can be implemented when the number of

signals acquired by the agents is observed; however, this is not necessarily the case when such investment is not observed. I derive the optimal contract when the amount or precision of the revealed information is unobserved and experimentation cannot be enforced. The optimal contract suggest that experimenters should be given the whole surplus if new technology is significantly better than the previous one. The intuition for this result is that experimenters must increase the number of signals acquired to increase the probability of being rewarded with the surplus.

The conclusions of the model are robust to the case where agents have heterogeneous beliefs. In this case, agents that ex-ante are more indifferent between the risky and the safe project should be chosen as experimenters since they have more intrinsic motives to experiment.

References

- S. Athey. Monotone comparative statics under uncertainty. *The Quarterly Journal of Economics*, 117(1):187, 2002.
- D. Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability. July 31-August 12, 1950. Statistical Laboratory of the University of California, Berkeley.*, volume 1, pages 93–102, Berkeley, CA, 1951. University of California Press.
- P. Bolton and C. Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999. ISSN 1468-0262.
- A. Bonatti and J. Horner. Collaborating. *The American Economic Review*, 101(2): 632–663, 2011.
- J.S. Demski and D.E.M. Sappington. Delegated expertise. *Journal of Accounting Research*, pages 68–89, 1987.
- J.L. Doob. *Stochastic processes*. Wiley publications in statistics. Wiley, 1953. ISBN 9780471218135. URL <http://books.google.com/books?id=KvJQAAAAMAAJ>.

- B. Ebersberger, O. Marsili, T. Reichstein, and A. Salter. Fortune favours the brave:: The distribution of innovative returns in finland, the netherlands and the uk. *Structural Change and Economic Dynamics*, 19(4):357–362, 2008.
- F. Ederer. *Launching a Thousand Ships: Incentives for Parallel Innovation*. Working Paper, 2008.
- D. Gromb and D. Martimort. Collusion and the organization of delegated expertise. *Journal of Economic Theory*, 137(1):271–299, 2007.
- H.A. Hopenhayn and F. Squintani. Preemption games with private information. *The Review of Economic Studies*, 78(2):667, 2011.
- G. Keller, S. Rady, and M. Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005.
- N. Klein and S. Rady. Negatively correlated bandits. *The Review of Economic Studies*, 78(2):693, 2011.
- R.A. Lambert. Executive effort and selection of risky projects. *The Rand Journal of Economics*, pages 77–88, 1986.
- D.A. Malueg and S.O. Tsutsui. Dynamic r&d competition with learning. *The RAND Journal of Economics*, 28(4):751–772, 1997.
- G. Manso. Motivating innovation. *The Journal of Finance*, 66(5):1823–1860, 2011.
- O. Marsili and A. Salter. ÔinequalityÕof innovation: skewed distributions and the returns to innovation in dutch manufacturing. *Economics of Innovation and New Technology*, 14(1-2):83–102, 2005.
- P.R. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.
- G. Moscarini and L. Smith. The law of large demand for information. *Econometrica*, pages 2351–2366, 2002. ISSN 0012-9682.

- P. Murto and J. Valimaki. Learning and information aggregation in an exit game. *Review of Economic Studies forthcoming*, 2011.
- P. Raikes. Business as usual: Some veal food markets in kenya. *Sociologia Ruralis*, 34(1):26–44, 1994.
- D. Rosenberg, E. Solan, and N. Vieille. Social learning in one-arm bandit problems. *Econometrica*, 75(6):1591–1611, 2007.
- M. Rothschild and J.E. Stiglitz. Increasing risk: I. a definition. *Journal of Economic theory*, 2(3):225–243, 1970.
- D. Szalay. Contracts with endogenous information. *Games and Economic Behavior*, 65(2):586–625, 2009.

A Appendix

Proof for Proposition 1. First, to show that $U(e) > \max\{\mu_0, y_s\}$ integrate by parts the value of experimentation to obtain

$$\begin{aligned}
 U(e) &= y_s + \int_{y_s}^{\bar{y}} (1 - F(x|e)) dx \\
 &= y_s + \int_0^{\bar{y}} (1 - F(x|e)) dx - \int_0^{y_s} (1 - F(x|e)) dx \\
 &= \mu_0 + \int_0^{y_s} F(x|e) dx
 \end{aligned} \tag{15}$$

Thus the value of experimentation is greater than $\max\{\mu_0, y_s\}$. The value of experimentation is also differentiable because the conditional distribution $f(x|y_r, e)$ is assumed differentiable. To prove that is strictly increasing in the precision, I will prove that the distribution $f(x|e)$ second order stochastically dominates the distribution $f(x|e')$ whenever $e < e'$. Let x , x' and x'' be the sufficient statistics for the first e signals, the e' signals, and the additional $e' - e$ signals, respectively.

Note that by independence $f(x'|e') = f(x|e) f(x''|e' - e)$. Define the corresponding conditional means as:

$$\mu = \int_0^{\bar{y}} y f(y|x, e) dy$$

$$\mu' = \int_0^{\bar{y}} y f(y|x', e') dy$$

First note that by the law of iterated expectations

$$\mathbb{E}_x[x] = \mu_0 = \mathbb{E}_{x'}[x']$$

Therefore the sequence of posterior means is a Martingale. Now, using independence we know that

$$\begin{aligned} \mathbb{E}_{x'}[x'|x] &= \int_0^{\bar{y}} \int_0^{\bar{y}} y f(y|x', e') f(x''|e' - e) dy dx'' \\ &= \int_0^{\bar{y}} \int_0^{\bar{y}} y f(y, x''|x, e') dy dx'' \\ &= \int_0^{\bar{y}} y f(y|x, e) dy \\ &= x \end{aligned}$$

Therefore x' is a mean preserving spread of μ . Rothschild and Stiglitz (1970) show that this is equivalent to having

$$\int_0^a F(x|e) dx \leq \int_0^a F(x|e') dx$$

for all $a \in [0, \bar{y}]$.

Using (15) we obtain $U(e) \leq U(e')$. On the other hand, the fact that the posteriors are Martingale imply that when the number of independent signals becomes large enough, the posterior mean will converge almost surely to the true y_r (Doob, 1953). Formally, we have that $\lim_{e \rightarrow \infty} f(x|y_r, e) = 1$ if $x = y_r$ and 0 otherwise. Thus the limit unconditional distribution of the posterior is given by

$$\begin{aligned} \lim_{e \rightarrow \infty} f(x|e) &= \lim_{e \rightarrow \infty} \int_0^{\bar{y}} f(x|y_r, e) g(y_r) dy_r \\ &= \int_0^{\bar{y}} \lim_{e \rightarrow \infty} f(x|y_r, e) g(y_r) dy_r \\ &= g(x) \end{aligned}$$

where the second line is obtained using uniform convergence, which in turn is obtained from the almost surely convergence and the differentiability of the distribution (Ascoli's theorem). Thus the value of experimentation is bounded by $\mu_0 + \int_0^{y_s} G(y) dy < \infty$ \square

Proof for Proposition 2. Define a compact domain for e where the upper and lower bound are given by the largest and the smallest e such that $C(e) = U(e)$, respectively. If there is no such e , the individual is better off by not experimenting ($e = 0$). Since the objective function is differentiable we know a maximum exists using Weierstrass Theorem. Moreover, using the Intermediate Value Theorem we know an interior optimum is characterized by

$$U_e(e) - C_e(e) = 0$$

The latter condition is also sufficient if the problem is concave, which is the case

when $U_{ee}(e) = \int_0^{y_s} F_{ee}(x|e) dx \leq 0$ since the cost is convex by assumption. The maximum is unique if the latter inequality is strict. \square

Proof for Proposition 3. Suppose the agent decides to experiment, and thus is in an interior solution. Using the implicit function theorem we know

$$\frac{\partial e^*}{\partial y_s} = -\frac{F_e(y_s|e)}{\int_0^{y_s} F_{ee}(x|e) dx - C_{ee}(e)}$$

The numerator is strictly negative, then the sign of the derivative will be determined by the denominator. The ordering of the signals implies that the conditional distribution $f(x|y_r, e)$ is logsupermodular, thus the unconditional cdf $F(x|e)$ is also logsupermodular (see Athey (2002)). Therefore for every e , all the unconditional cdfs will cross uniquely at μ_0 . To see this just consider the case when $e = 0$ and the distribution is degenerate at μ_0 . Hence, $F_e(\mu_0|e) = 0$. Also note that by the second order stochastic dominance $F_e(y_s|e) > 0$ whenever $y_s < \mu_0$, and $F_e(y_s|e) < 0$ whenever $y_s > \mu_0$. Hence the optimal effort achieves a maximum when $y_s = \mu_0$, and is increasing (decreasing) for smaller (greater) y_s .

\square

Proof for Proposition 4. Using the envelope theorem we obtain

$$\frac{\partial}{\partial y_s} (U(e^*) - C(e^*)) = F(y_s|e^*) \geq 0$$

which implies that the value is nondecreasing in y_s since the other alternatives are also nondecreasing in y_s . The second derivative is given by

$$\frac{\partial^2}{\partial^2 y_s} (U(e^*) - C(e^*)) = f(y_s|e^*) + F_e(y_s|e^*) \frac{\partial e}{\partial y_s}$$

The first element is always positive since it is a density function. The second term is also positive by the single crossing property and Proposition (3). Thus the

function is strictly convex in y_s . Since the other alternatives are also convex in y_s , then the maximum of convex functions is also convex.

In the presence of fixed costs, $U(e^*) - C(e^*)$ as a function of y_s will cross at most once each of the outside options. It could cross once the constant μ_0 from below since $U(e^*) - C(e)$ is increasing in y_s . It could cross once y_s from above since its first derivative with respect to y_s are between 0 and 1. When there is no fixed cost, $c = 0$, $U(e) = \mu_0$ and $e^* = 0$ at $y_s = 0$. On the other hand, when $y_s = \bar{y}$, then $U(e) = y_s$ and again $e^* = 0$. Therefore the principal always prefers to hire an agent for any interior y_s .

Since $U(e^*) - C(e^*)$ is linear in c , there exists a \hat{c} such that $U(e^*) - C(e^*) = \mu_0 = y_s$. Thus, for any $c < \hat{c}$, there exists $a_c, b_c \in (0, \bar{y})$ such that $U(e^*) - C(e) > \mu_0$ for any $y_s > a_c$, and $U(e^*) - C(e) > y_s$ for any $y_s < b_c$. Obviously it must be the case that $\mu_0 \in (a_c, b_c)$. Note that a_c and b_c are increasing and decreasing in c , respectively, precisely because the function crosses from below and above each of the corresponding outside options. Finally, for any $c > \hat{c}$, the interval is empty and the principal never experiments.

□