

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/139085>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

The Wisdom of the Crowd when Acquiring Information is Costly¹

Jacob Glazer, University of Warwick and Tel Aviv University

Ilan Kremer, University of Warwick and The Hebrew University

Motty Perry, University of Warwick

21/06/2020

ABSTRACT

We consider a sequential investment process that is characteristic of crowdfunding platforms, among other contexts. Investors wish to avoid the cost of information acquisition and thus prefer to rely on information acquired by previous investors. This may lead to phenomena similar to an information cascade. We characterize the optimal policy that balances between the incentive to acquire information and the optimal investment decision. The policy is based on time-varying transparency levels such that it may be worthwhile to conceal some information in some periods.

¹We would like to thank the editor of this journal for his guidance and encouragement. Special thanks are due to a referee for his highly constructive suggestions, which led us to rethink and rewrite parts of the paper. We believe that his/her thoughtful comments greatly improved the paper.

1 Introduction

Platforms that implement the "wisdom of the crowd" are becoming increasingly important in the new internet economy. They are used to collect information from the public, process it and redistribute it among the public to be acted upon. Examples include Google Maps, Yelp, and others. However, there are numerous cases in which the information first needs to be acquired by the crowd members, and an important task of an information disclosure mechanism is to induce the right number of agents to incur the cost of doing so. One example is crowdfunding platforms where potential investors can conduct their own (costly) investigation before deciding whether or not to invest. Another is product review platforms where consumers can share information about the products, in addition to the information provided by the platform.

When acquiring information is costly an important task of an optimal information disclosure mechanism is to induce the right number of agents to incur the cost. To the best of our knowledge, little attention has been devoted to the question of how to design such mechanisms, and it is our intention to take the first step in closing that gap.²

We study a stylized model in which an uninformed principal hosts a single project that is either a "lemon" or "good". Potential agents arrive sequentially, and each of them decides whether or not to invest in the project without knowing its quality. Agents cannot observe the preceding choices made by other agents, although they can learn about them from the principal, who observes each agent's decision and decides how much of that information to reveal and to whom. After receiving whatever information is provided by the principal and before deciding whether or not to invest, each agent can conduct his own costly exploration of the project. This takes the form of acquiring a noisy but informative binary signal indicating whether the project is a "lemon" or "good". Each agent's objective is to maximize his expected profit whereas the principal's objective is to maximize the present value of the sum of all agents' profits. Thus, the principal's benefit from each agent's exploration is larger than that of the agent himself. Consequently, and since exploration is costly, a conflict of interest may arise between the principal and the agent as to whether the agent should carry out an exploration. It is assumed that an agent's decision whether to explore and the outcome of the exploration are unobservable to the principal

² An exception in a different context is Gershkov and Szentes (2009); see below for a comparison to our model.

and any other agent. Thus, the only information available to the principal is each agent's final decision whether or not to invest. Finally, the principal can commit ex-ante to a disclosure mechanism.

The principal faces the following trade-off: if a policy of full disclosure is adopted, an inefficient cascade may emerge and an insufficient number of agents will conduct the costly investigation.³ If, on the other hand, the principal chooses not to reveal any information, an excessive number of agents may end up incurring the cost of exploration and in addition some may end up making the wrong investment decision. The optimal disclosure policy must walk a fine line between providing information to agents and creating incentives for them to acquire costly (and noisy) information.

After describing the first-best mechanism and showing that it is not implementable, we move on to characterize the second-best mechanism which we refer to as Recommend & Correct. It consists of at most two phases of communication between the mechanism and each agent. In Phase 1, the mechanism randomizes between the following recommendations: *invest*, *pass* or *explore* and the agent follows the recommendation (along the equilibrium path). If the mechanism's recommendation is *explore* then the agent does so and reports the outcome. In Phase 2, the mechanism randomizes between recommending *invest* or *pass*.

The fact that the mechanism randomizes in both phases involves a welfare loss since the mechanism does not necessarily recommend the best action, given the information it possesses. The randomization is necessary in order to make the mechanism incentive-compatible and in particular to induce the agent to explore whenever it is recommended that he do so. The exact randomization of the optimal mechanism, however, crucially depends on the signals' information structure, as will become evident from the two canonical cases we study in detail. In one case, the "bad" signal (and only that signal) is fully revealing, while in the other, signals are noisy and symmetric.

The characterization of the optimal mechanism enables us to derive a number of insights regarding disclosure policy when obtaining information is costly to the agents. In general, at every point in time the mechanism can be in one of three states, the duration of which depends on the particular information structure and can be random. The mechanism enters the third and final state in which there is no conflict of interest between the mechanism and the agent after the mechanism has obtained enough information. From this point onward,

³The inefficient cascade can go in both direction: either all investors invest in a lemon or none of them invest in a good project.

it recommends to all agents either to invest or to pass. The mechanism starts from a state in which there is no conflict of interest, since both the principal and the agent are interested in exploration. It is in the second state where the conflict of interest arises, since the agent, if he had the information that the mechanism possesses, would have chosen to either invest or pass whereas the mechanism would have rathered that he explore. As a result, the mechanism must hide some information from the agent, resulting in a welfare loss.

Even though the model is theoretical in nature, we believe that its insights are relevant to several real-life situations, such as crowdfunding. The internet has revolutionized the investment industry, enabling potential agents to rely on the wisdom of the crowd as a source of information about potential investment opportunities. There is, however, widespread agreement among regulators and researchers that insufficient and asymmetric information lead to major market failures in the crowdfunding industry⁴. In what follows, we attempt to shed some light on how information disclosure policies should be designed in the crowdfunding market. The insights gained may be of value to both platform owners who must decide how to design their information disclosure policies and regulators who must decide what information should be disclosed by the principals and when.⁵

2 The Model

A project is either a good project (G) or a lemon (L). There are infinitely many periods, denoted by $t = 1, 2, 3, \dots$. In each period, a different agent chooses whether to invest a predetermined fixed amount in the project or to pass. Hereafter, we will refer to the agent in period t as "agent t ".⁶ If the project is good, then it is optimal to invest and if it is a lemon, then it is optimal to pass. We normalize the agent's payoffs so that the optimal action yields a payoff of one and the other action yields a payoff of zero. An agent knows his place in line but cannot observe the actions (nor the payoff) of the agents who moved

⁴For a comprehensive survey of the crowdfunding markets see Agrawal, Catalini and Goldfarb (2013) and Belleamne, Omrani and Peitz (2015). Discussion of the market failures in this industry can also be found in Howe (2008), Burtch et al. (2016), Block and Koellinger (2009), and Geva et al. (2106).

⁵Our reading of the literature on the regulation of crowdfunding is that regulators have so far taken a very direct and perhaps overly simplistic approach to what the optimal disclosure policy should look like.

⁶If agents have no information about their position in line, then a first-best policy (at least in terms of information acquisition) can be implemented.

previously.⁷

The common prior belief that the project is good is denoted by p_0 . Before making the investment decision, the agent can conduct at most one costly exploration in an attempt to discover the true nature of the project. This is modeled formally by assuming that agent t can obtain an informative signal about the state of the world at a cost of $c > 0$. We denote the choice of whether or not to obtain such a signal by *explore*. We assume that the signal has two possible realizations, s^g and s^l . Conditional on the event that the project is a lemon, the signal s^l is realized with probability q_l and the signal s^g is realized with probability $1 - q_l$. Similarly, conditional on the project being good, s^g is realized with probability q_g and s^l is realized with probability $1 - q_g$. We further assume that $q_g + q_l > 1$ and therefore, upon observing the signal s^g the probability of the project being a lemon decreases, whereas upon observing the signal s^l it increases. We say that an agent *follows his signal* if he invests after observing the signal s^g or passes after observing s^l .

2.1 The Principal

There exists a principal who observes the investment decisions made by the agents, but does not observe whether an agent has acquired a signal (and accordingly nor the realization of the signal) nor the agent's payoff. The principal's objective is to maximize the sum of the discounted present value of all agents' payoffs, taking into account the costs of acquiring a signal. We let $\delta \in (0, 1)$ denote the principal's discount factor. The principal can commit to a mechanism that specifies the messages to be conveyed to each agent as a function of the history (to be defined formally below).⁸ The mechanism chosen by the principal is known to all agents.

Let p_t denote the principal's belief at the beginning of period t that the project is good. Prior to the arrival of agent 1, the principal holds the same prior as the agent, namely, $p_1 = p_0$. In subsequent stages, p_t will depend on the information possessed by the principal at the beginning of the period. If, at the

⁷An alternative and perhaps more natural assumption would be that the agent's payoff is 1 when he invests in a good project, -1 when he invests in a lemon and zero when he does not invest. Since we care only about the difference in payoffs between the two actions, our assumption is equivalent to such a setup. In both payoff structures, conditional on the project being good, the benefit from investing compared to not investing is one. Conditional on the project being a lemon, the benefit from not investing as compared to investing is also one.

⁸In particular, monetary transfers to and from the principal are not permitted. In future work, it might be interesting to allow for such transfers in order to improve the incentives.

beginning of period t and on the basis of his information, the principal believes that n_g agents have observed the signal s^g and n_l have observed the signal s^l , then his belief will be:

$$p_t(n_g, n_l) = \left\{ \frac{q_g^{n_g} (1 - q_g)^{n_l} p_1}{q_g^{n_g} (1 - q_g)^{n_l} p_1 + (1 - q_l)^{n_g} q_l^{n_l} (1 - p_1)} \right\}. \quad (1)$$

Recall, however, that the principal does not observe the signals but can only infer them from the decisions made by the agents.

2.2 The First-Best Solution

The first-best mechanism refers to the principal's optimal policy in the (hypothetical) case in which he can decide whether each agent should purchase a signal prior to the investment decision and in which he can observe the realization of all signals acquired by the agents.

As will be shown below, the first-best solution is characterized as a stopping rule with two cutoffs, p_g and p_l . The principal starts by requiring the first agent to explore and continues doing so with subsequent agents as long as the posterior p_t is in the interval $[p_l, p_g]$. Furthermore, for each agent who acquired a signal, the principal, after observing the realization of that signal, instructs the agent whether to invest or to pass.

The probability p_g has the property of being the lowest posterior for which the likelihood that the project is good is high enough such that acquiring a signal is no longer worthwhile and the principal is better off requiring the agent (and all subsequent agents) to invest. Similarly, the probability p_l has the property of being the highest posterior at which the likelihood that the project is good is low enough that acquiring a signal is no longer worthwhile and the principal is better off requiring the agent (and all subsequent agents) to pass. This is formally stated in the following proposition (the proof of which is provided in the appendix):

Proposition 1 *The first-best solution is given by two cutoffs, p_g and p_l , such that $p_g > p_l$ and:*

- (i) *if $p_t \in [p_l, p_g]$, then agent t acquires a signal and invests if and only if $p_{t+1} \geq 0.5$;*
- (ii) *if $p_t \leq p_l$, then agent t passes (without acquiring a signal); and*
- (iii) *if $p_t > p_g$, then agent t invests (without acquiring a signal).*

2.3 The Agent's Choice

In order to demonstrate the infeasibility of the first-best mechanism, we describe the agent's problem and the conflict of interest between him and the principal. Let μ_t denote agent t 's belief that the project is good when he is contemplating taking an action. Obviously, μ_t may depend on the information provided to the agent by the principal. For the first agent, we have $\mu_1 = p_0$.

Consider some agent t whose belief is μ_t . If the agent decides to invest without acquiring a signal, then his expected payoff is μ_t and if he decides to pass, then his expected payoff is $1 - \mu_t$.

As for the agent's decision whether or not to acquire a signal, observe first that if no further information is revealed to the agent after acquiring a signal, then the agent will find it optimal to do so only if he plans to follow it. In such a case, agent t 's expected payoff from acquiring a signal and following it is:

$$\mu_t q_g + (1 - \mu_t) q_l - c.$$

We can now make the following claim:

Proposition 2 *If no further information is revealed to agent t after acquiring a signal, his expected utility-maximizing decision is given by:*

$$d_t(\mu_t) = \left\{ \begin{array}{l} \text{invest if } \mu_t > \mu_g \\ \text{pass if } \mu_t < \mu_l \\ e \text{ if } \mu_l \leq \mu_t \leq \mu_g \end{array} \right\} \quad (2)$$

where μ_g solves

$$\mu_g q_g + (1 - \mu_g) q_l - c = \mu_g$$

and μ_l solves

$$\mu_l q_g + (1 - \mu_l) q_l - c = 1 - \mu_l.$$

Therefore,

$$\mu_g = \frac{q_l - c}{1 - q_g + q_l} \quad \text{and} \quad \mu_l = \frac{1 + c - q_l}{1 + q_g - q_l}.$$

Claim 1 *In order for the model to be non-trivial and in particular to ensure that in equilibrium some agents acquire signals (i.e., $\mu_g > \mu_l$), we make the following assumption:*

Assumption 1: $0 < c < \frac{q_g + q_l - 1}{2}$ and $p_0 \in [\mu_l, \mu_g]$.

Remark 1 (i) If the above assumption fails, the first agent as well as those that follow him, will refuse to acquire information and there is no mechanism that can convince them otherwise. (ii) If acquiring a signal is costless ($c = 0$), then the principal's problem becomes trivial. He can simply ask each agent for the realization of the signal and pass the information on to the other agents (or alternatively recommend the optimal action to them based on the information). This mechanism is incentive-compatible and implements the first-best solution. (iii) The assumption that $p_0 \in [\mu_l, \mu_g]$ guarantees that an agent with no information other than the prior belief p_0 (as in the case of the first agent) will acquire a signal.

Based on Proposition 1 and Proposition 2, we can now discuss the conflict of interest between the principal and the agents. Assume that the platform adopts a policy of full disclosure. That is, at the beginning of every period, the principal reveals all the information it has gathered so far to the agent. In such a case, it is clear that $\mu_t = p_t$ for all t , and the agent will find it optimal to acquire a signal if and only if $p_t \in [\mu_l, \mu_g]$. However, when $p_g > p_t > \mu_g$ it is optimal for the agent to invest without acquiring a signal, even though it is socially optimal to acquire a signal given the potential benefit to future agents and similarly for the case in which $p_l < p_t < \mu_l$.

Since a policy of full disclosure cannot implement the first best, the problem then becomes to find the (second) best mechanism that the principal can implement in order to maximize social welfare. In the following section, we present a general characterization of the second-best mechanism.

3 Recommend & Correct Mechanisms

In this section, we first describe a set of direct mechanisms and then argue that without loss of generality we can focus on the optimal mechanism within it. We refer to these mechanisms as *Recommend & Correct* mechanisms, (hereafter: *R&C* mechanisms). The rationale for restricting our attention to *R&C* mechanisms follows a similar logic to that of the proof of the revelation principle in Myerson (1986). In an *R&C* mechanism, each period consists of at most the following two phases:

Phase 1: The principal recommends (privately and possibly randomly) to agent t whether to *invest*, *pass* or *explore* (acquire a signal). If the principal recommends *invest* or *pass*, the mechanism proceeds to agent $t + 1$. Otherwise,

the principal continues together with agent t to Phase 2.

Phase 2: Agent t (privately) reports the signal's outcome (s^g or s^l) to the principal which then recommends to agent t (privately and possibly randomly) whether to *invest* or to *pass*.

We restrict our attention to mechanisms that are characterized by Incentive Compatibility (*IC*) namely that the agent finds it optimal to follow the principal's recommendation and if he acquires a signal, then he reports it truthfully. Note that, off equilibrium, if the agent has decided not to acquire a signal and since the principal cannot observe that decision, the agent in Phase 2 can lie and report a fake realization (s^g or s^l). In this event, the agent can wait in order to acquire a signal only after hearing the principal's recommendation. As we argue below, the role of mixing (in both phases) is to guarantee that the *R&C* mechanism satisfies *IC*.

Following Myerson (1986), in any mechanism, and after some communication back and forth between the principal and the agent, the agent decides whether to *invest*, *pass* or *explore*. This can be implemented in Phase 1 of our mechanism. After acquiring a signal, there is again some communication between the parties that leads to an investment decision by the agent as a function of the private information possessed by the principal and the agent. This outcome can be directly implemented in Phase 2. The following proposition (the proof of which is straightforward and therefore omitted) states that restricting attention to an *R&C* mechanism is without loss of generality.

Proposition 3 *The optimal R&C mechanism is optimal within the set of all mechanisms.*

3.1 The Optimal R&C Mechanism

We now provide a characterization of the optimal *R&C* mechanism. In particular, the following proposition describes the properties that every optimal *R&C* mechanism satisfies, regardless of the information structure. The exact randomization that the mechanism assigns to every action after every history depends on the specific information structure. Two canonical cases will be examined in the following sections.

Proposition 4 *Let p_t be the principal's posterior at the beginning of period t , and let \tilde{p}_t be the principal's posterior after hearing t 's report in Phase 2. The optimal *R&C* mechanism is given by:*

Phase 1:

(i) If $p_t \in [\mu_l, \mu_g]$, then the mechanism recommends *explore*.

(ii) If $p_t > \mu_g$, then the mechanism randomizes between recommending *explore* and *invest*, and if $p_t < \mu_l$, then the mechanism randomizes between recommending *explore* and *pass*.

Phase 2:

(iii) If agent t reports s^g and $\tilde{p}_t > 1/2$, then the mechanism recommends *invest*. If agent t reports s^l and $\tilde{p}_t < 1/2$, then the mechanism recommends *pass*.

(iv) If agent t reports s^g and $\tilde{p}_t < 1/2$, then the mechanism randomizes between recommending *pass* and recommending *invest*. If agent t reports s^l and $\tilde{p}_t > 1/2$, then the mechanism randomizes between recommending *pass* and recommending *invest*.

An informal argument will suffice to establish Proposition 4. We start with Phase 1. To establish (i), suppose by contradiction that the mechanism recommends to agent t to invest or to pass when $p_t \in [\mu_l, \mu_g]$. If instead it recommends *explore*, then this will increase not only agent t 's payoff (since when $p_t \in [\mu_l, \mu_g]$ agent t will wish to explore), but also that of all future agents since the mechanism will have more information about the state of the world.

To prove (ii), assume that $p_t \in [0, 1] \setminus [\mu_l, \mu_g]$. Recall that in this case if the platform applies a policy of full transparency, then the agent will not explore. It follows that to incentivize the agent to explore, the mechanism may need to randomize. We need to show that there is no history following which the mechanism assigns a strictly positive probability to all three recommendations, namely *invest*, *pass*, and *explore*. Assume that $p_t > \mu_g$ (a similar argument applies when $p_t < \mu_l$). We shall argue that the optimal mechanism does not recommend *pass* in this case.

Suppose, by contradiction, that the mechanism recommends *pass* with a strictly positive probability. Consider the following modification of the mechanism: instead of recommending *pass*, the mechanism recommends *invest*. To see why the modified mechanism is incentive-compatible, notice first that the recommendation *explore* is incentive-compatible (since the modified mechanism recommends *explore* for the same set of histories as the original one). The recommendation to invest is also incentive-compatible under the modified mechanism, since the agent knows that the mechanism recommends *invest* only when $p_t > \mu_g$. Finally, observe that the modified mechanism yields a higher payoff to

both the agent and the mechanism since $p_t > 1/2$.

We can now proceed to Phase 2. First, observe that there is no conflict of interest between the mechanism and agent t , in this phase, since it is optimal for both that the agent makes the right investment decision. This implies that when agent t reports s^g but $\tilde{p}_t < 1/2$ or when he reports s^l but $\tilde{p}_t > 1/2$, the mechanism would like to correct the agent. However, since the mechanism needs to ensure that the agent follows the recommendation to acquire a signal in Phase 1, it may not always correct the agent in these cases. If the mechanism did always correct the agent, then the agent might choose not to explore in Phase 1, to report a fake realization, and to make his decision only after hearing the mechanism's recommendation in Phase 2. Because not correcting the agent is costly, the principal will randomize just enough to make exploration incentive-compatible for the agent. The exact randomization that the mechanism assigns to an action after every history depends on the specific information structure.

To understand the construction in Phase 2, it is useful to consider a simple example in which if the principal does not randomize in Phase 2, then the agent has no incentive to acquire a signal after hearing the recommendation *explore*.

Example 1 ⁹ Suppose that $p_1 = \mu_1 = 0.55$, $q_l = q_g = 0.7$, and $c = 0.1$. Suppose also that agent 1 has already acquired a signal and truthfully reported it to the principal, thus allowing us to focus on agent 2. Notice that the principal now has more information about the project than agent 2 (because he possesses the information acquired by agent 1). Agent 2 will then choose one of two possible strategies:

1. Acquire a signal and report it truthfully.
2. Do not acquire a signal and randomize between reporting s^g and s^l (say, with equal probability).

Notice that signals affect the principal's belief. Thus:

- If both agents report s^g , then $\tilde{p}_2 > p_1 > 1/2$ and the principal recommends invest.

- If both agents report s^l , then $\tilde{p}_2 < 1/2$ and the principal recommends pass.

- If one agent reports s^g and the other reports s^l , then $\tilde{p}_2 = p_1$ and the principal recommends invest.

- It is straightforward to show that agent 2's expected payoff is 0.621 under the first strategy and 0.625 under the second. Thus, agent 2's payoff is increased by not acquiring a signal. In order to correct for this and make it

⁹We would like to thank the referee for suggesting this example.

incentive-compatible for the agent to explore whenever the principal makes that recommendation, the principal will randomize between correcting the agent (in the sense of Proposition 4, iv) and not correcting him. Since not correcting the agent is costly, the principal will randomize just enough to make exploration incentive-compatible for the agent. In this example, the principal optimally corrects the agent with probability 0.96 and does not with probability 0.04.

In what follows, we present two canonical information structures and characterize their optimal mechanism. To accomplish this, we first define the general mechanism more formally in the following subsection.

3.2 The R&C Mechanism: A Formal Representation

To define the mechanism more formally, let h_t denote a history of actions observed by the principal at the beginning of period $t + 1$. The history h_t consists of all the observable actions taken by the agents, as well as the principal's, in all periods $t' \leq t$. Thus, h_t is the relevant history prior to the principal recommending an action m_{t+1} to agent $t + 1$ in Phase 1. Let H_t denote the set of all possible histories of length t induced by the mechanism. The principal's recommendation policy in Phase 1 of period $t + 1$ is history-dependent, and we write $\gamma_{t+1} : H_t \rightarrow [0, 1]$ to denote a function mapping the possible histories onto the probability of a recommendation to invest. Similarly, $\lambda_{t+1} : H_t \rightarrow [0, 1]$ and $\varepsilon_{t+1} : H_t \rightarrow [0, 1]$ are defined for the recommendation to pass and to acquire a signal, respectively.

Let $\tilde{h}_{t+1} = (h_t, m_{t+1}, r_{t+1})$ denote the history at the beginning of period $t + 1$ amended by the principal's recommendation m_{t+1} and the agent's report r_{t+1} . Let \tilde{H}_{t+1} denote the set of all possible histories \tilde{h}_{t+1} . The principal's recommendation policy in Phase 2 of period $t + 1$ is a function $\nu_{t+1} : \tilde{H}_{t+1} \rightarrow [0, 1]$ where $\nu_{t+1}(\tilde{h}_{t+1})$ is the probability that the mechanism recommends that the agent invest.

We can now formally define a mechanism as a pair $(\{M_t\}_{t=1}^\infty, \{\nu_t\}_{t=1}^\infty)$ where

$$M_t = (\gamma_t, \lambda_t, \varepsilon_t) \mid \gamma_t, \lambda_t, \varepsilon_t : H_{t-1} \rightarrow [0, 1]^3, \quad (3)$$

and, for all $h_{t-1} \in H_{t-1}$,

$$\gamma_t(h_{t-1}) + \lambda_t(h_{t-1}) + \varepsilon_t(h_{t-1}) = 1,$$

and

$$\nu_t : \tilde{H}_t \rightarrow [0, 1]. \quad (4)$$

Knowing the mechanism and understanding the law of motion of the principal's beliefs, the agent forms beliefs about the principal's prior (i.e., a distribution over p_t) after hearing the principal's recommendation. Let $\mu_t(m_t)$ denote agent t 's beliefs that the project is good after hearing the recommendation $m_t \in \{invest, pass, explore\}$ in Phase 1, and let $\tilde{\mu}_t(\tilde{m}_t)$ denote agent t 's beliefs that the project is good after hearing the recommendation $\tilde{m}_t \in \{invest, pass\}$ in Phase 2.

As will become clear in the following sections, the optimal $R\&C$ mechanism results in a welfare loss relative to the first-best outcome for two reasons: agents explore in states where it is socially optimal for them to either invest or pass, and agents invest or pass when it is socially optimal for them to explore. Nonetheless, the optimal $R\&C$ mechanism can significantly improve welfare relative to a policy of full disclosure. The size of the welfare gain generally depends on the parameters of the problem. It can be seen, however, that the larger the social discount factor, the larger the gain will be.

In what follows, we fully characterize the optimal mechanism in two canonical cases.

4 Asymmetric Signals

We first consider the case of strong asymmetry between the two signals, in the sense that one of them is fully revealing. More precisely, assume that conditional on the project being good, the realization of the signal is s^g with probability 1, (i.e., $q_g = 1$), whereas conditional on the project being a lemon, the realization of the signal is s^l with probability q (i.e., $q_l = q$) and s^g with probability $1 - q$. Clearly, upon observing the signal s^g the posterior that the project is a lemon decreases, whereas upon observing the signal s^l the project is a lemon with probability 1.

It is easy to see that in this case:

$$\mu_g = \frac{q - c}{q} \quad \text{and} \quad \mu_l = \frac{1 + c - q}{2 - q}.$$

The following assumption is an application of Assumption 1 to this case:

Assumption 1': $0 < c < q/2$ and $p_0 \in [\mu_l, \mu_g]$.

The following proposition characterizes the first-best cutoffs in the asymmetric case. The proof of which (similar to Wald (1947)) is provided in the appendix.

Proposition 5 *The first-best solution is given by the cutoff:*

$$p_g = 1 - c \frac{1 - \delta}{q} = \mu_g + c\delta/q$$

such that:

- (i) if $p_t \in [p_1, p_g]$, then agent t acquires a signal and follows it;
- (ii) if $p_t = 0$, then agent t passes (without acquiring a signal).

The following definitions will be useful in what follows:

Consider some history of length t , along which all agents explore, obtain the signal s^g , and choose to invest. Let \hat{t} denote the first period along such a history for which the principal's posterior is strictly above μ_g . That is,

$$\hat{t} = \min \left[t \mid \frac{p_1}{p_1 + (1 - p_1)(1 - q)^{t-1}} > \mu_g \right]. \quad (5)$$

Thus, $\hat{t} - 1$ is the maximal number of (consecutive) s^g signals following which the principal's belief is weakly below μ_g . That is, if agent \hat{t} knows that all the previous agents received the signal s^g , then he will still want to explore, but if he also receives the signal s^g and the history is known to agent $\hat{t} + 1$, then agent $\hat{t} + 1$ will prefer to invest without further exploration.

Similarly, let

$$\bar{t} = \max \left[t \mid \frac{p_1}{p_1 + (1 - p_1)(1 - q)^{t-1}} \leq p_g \right]. \quad (6)$$

As above, $\bar{t} - 1$ is the maximal number of (consecutive) s^g signals following which the principal's posterior is weakly below p_g . Even if all the agents prior to agent \bar{t} received the signal s^g , the principal will still want agent \bar{t} to explore. However, if agent \bar{t} also receives the signal s^g , then the principal will want agent $\bar{t} + 1$ and all subsequent agents to invest without any further exploration. Since $\mu_g \leq p_g$, we know that $\hat{t} \leq \bar{t}$. Hereafter, we will assume that $\hat{t} < \bar{t}$.¹⁰

¹⁰Note that if $\hat{t} = \bar{t}$, then a policy of full disclosure is optimal and the first-best is trivially achieved. This might be the case if the signals are very precise.

4.1 The Optimal Mechanism

Some additional notation and discussion are required prior to the presentation of this section's main theorem. We define a "modified" cost of exploration function $c^*(p)$ as follows:

$$c^*(p) = c + \frac{\max\{0, (p - \mu_g)\}}{\mu_g} (1 + c - q). \quad (7)$$

Note that for $p \leq \mu_g$, and as shown above, there is no conflict of interest between the principal and the agents, i.e. $c^*(p) = c$; otherwise, $c^*(p)$ is monotonic in p . As will be discussed below and shown formally in the proof of the theorem, the solution to the second-best problem is, in fact, the solution to a modified first-best problem in so far as the cost of exploration is $c^*(p)$ rather than c . The modified cost of exploration function, $c^*(p)$, captures not only the real cost of exploration but also the "implied" cost resulting from the need to make the agent's choice to explore incentive-compatible when $\mu_g < p_t$.

Using this modified cost function, we can define a modified cutoff, p_g^* , in much the same way as we defined p_g for the original first-best problem. That is, p_g^* is the point at which the principal's posterior probability is sufficiently high so that exploration becomes too costly (under the modified first-best mechanism) when the cost is $c^*(p_g^*)$. Specifically,

$$p_g^* = 1 - c^*(p_g^*) \frac{1 - \delta}{q},$$

and note that $p_g^* \leq p_g$. Similarly we define t^* by substituting p_g^* for p_g in the definition of \bar{t} , such that:

$$t^* = \max \left[t \mid \frac{p_1}{p_1 + (1 - p_1)(1 - q)^{t-1}} \leq p_g^* \right]. \quad (8)$$

That is, $t^* - 1$ is the maximal number of (consecutive) s^g signals following which the posterior is weakly below p_g^* .

In what follows, we first prove that in the asymmetric case the principal can make do with a simple *R&C* mechanism, which will be referred to as a recommendation mechanism. We then analyze the optimal recommendation mechanism.

5 A Recommendation Mechanism Without Loss of Generality

In a recommendation mechanism, each agent receives a single incentive-compatible recommendation. Therefore, a recommendation mechanism is a special case of the *R&C* mechanism in which the second stage is degenerate, namely $\nu_t(\tilde{h}) \equiv 0$ for all \tilde{h} . In what follows, we explain why limiting the recommendation mechanism does not involve a loss of generality.

Suppose, by contradiction, that there exists a history \tilde{h}_{t-1} , following which the mechanism recommends (with some positive probability) that the agent acquire a signal and that after the agent reports s^g (s^l) the principal (in Phase 2) recommends the action *pass* (*invest*). Note, however, that all histories following which it is optimal for the agent to ignore his signal are those in which, had he known the history, his choice would have been not to acquire a signal.

It follows that if at the beginning of period t the agent assigns a positive probability to such a history, then he is better off avoiding the cost of exploration, reporting a fake signal, waiting to receive the principal's recommendation, and then choosing whether to acquire a signal and follow it or to act without acquiring a signal.

6 The Optimal Mechanism: an Informal Description

The optimal private mechanism consists of the following three stages:

Stage 1: During this stage, which starts in period 1 and lasts for \hat{t} periods, agents receive a recommendation to explore unless one of them chooses the action *pass*. If, during this stage, an agent t' chooses the action *pass*, the mechanism recommends that all agents t , $t' < t \leq \hat{t}$, choose the action *pass* without acquiring a signal. Therefore, in this stage, the recommendation policy coincides with that of the first-best mechanism (although the recommendations are now private). Furthermore, if an agent receives a recommendation to explore during this period, then he must conclude that all the previous agents acquired a signal and that it was s^g ; if he receives a recommendation to *pass*, he must conclude that one of the previous agents observed the signal s^l .

It is important to note that for all t in this stage (and, as will be shown later, all other stages as well), the principal's posterior at the end of period t is

either 0 or $p_1/(p_1 + (1 - p_1)(1 - q)^t) > p_1$. Thus, the distribution of p_t consists of only two points. This turns out to be one of the main differences between the asymmetric and symmetric cases, namely that the support of the distribution of \tilde{p}_t consists of many points in the later case as opposed to only two in the former case.

Stage 2: This stage starts at the beginning of period $\hat{t} + 1$ and concludes at the end of period t^* , where, in general, $\hat{t} < t^* < \bar{t}$ (see the definitions in 5, 6 and 8). Since $t > \hat{t}$, the principal, in order to induce the agents to explore, must hide some information from the agents by, for example, committing to the following (randomizing) policy: if $p_t > 0$ (i.e., if none of the agents have yet observed the signal s^l), then the principal will recommend that agent t explore with probability one, namely, $\varepsilon_t^*(p_t) = 1$. If $p_t = 0$, then the principal will randomize between recommending that agent t choose the action *pass* and recommending that he choose the action *explore*, namely, $\varepsilon_t^*(0) = 1 - \lambda_t^*(0) > 0$. The magnitude of $\varepsilon_t^*(0)$ is chosen such that agent t 's posterior, after receiving the recommendation to explore (i.e., $\mu_t(e)$), will be exactly μ_g , and he will follow the recommendation.

Stage 3: This stage starts in period $t^* + 1$ and goes on forever. During it, none of the agents acquire a signal. The principal recommends that all agents choose the action *invest* if none of the agents who moved in Stage 1 or Stage 2 chose the action *pass*; otherwise, it recommends that all agents choose the action *pass*.

In sum, the second-best mechanism differs from the first-best mechanism in two respects: (i) exploration is conducted in the second stage and the action *invest* may be chosen by some of the agents, even after the principal has already learned that the project is a lemon, and (ii) exploration may be terminated and the action *invest* recommended earlier than in the first-best case, in the case that $p_t = p_g^* < p_g$.

A key insight of this mechanism emerges in Stage 2, in which the principal randomizes after receiving the signal s^l . In this environment, if the principal adopts a policy of full transparency, then the agent will acquire too little information and will invest too early, whereas the principal would like to be more convinced about the state of the world before stopping information acquisition and recommending *invest* to all agents hereafter. The principal therefore commits to recommending that the agent acquire information with a positive probability following reports of the signal s^l (even though he knows that the state is L and therefore this information is of no value). The agent who receives a

recommendation to acquire information becomes more pessimistic and therefore is induced to acquire information. In the case of the optimal mechanism, the principal's recommendation must make the agent just indifferent between acquiring information and investing. After hearing the recommendation *explore*, agent t faces three possible scenarios:

- (i) The state of the world is G . All previous agents observed s^g and invested. Following this history, the principal recommends *explore* with probability 1.
- (ii) The state of the world is L . The principal has observed s^g in every previous period. Conditional on L , this history occurs with probability $(1-q)^{t-1}$ and following it, the principal recommends *explore* with probability 1.
- (iii) The state of the world is L . The principal has observed at least one agent taking the action *pass*. Conditional on L , this occurs with probability $1 - (1-q)^{t-1}$. In this case, the principal recommends *explore* with probability $\varepsilon^t(0)$ and *pass* with probability $(1 - \varepsilon^t(0))$. The probability $\varepsilon^t(0)$ solves:

$$\frac{p_1}{p_1 + (1-p_1)[(1-q)^{t-1} + (1 - (1-q)^{t-1})\varepsilon^t(0)]} = \mu_g$$

which is the agent's indifference condition.

6.1 The Optimal Mechanism: a Formal Statement and Proof

Recall that $\gamma_t(p_t)$ denotes the probability of the principal recommending that agent t choose the action *invest* conditional on $\tilde{p}_t = p_t$ and similarly for $\lambda_t(p_t)$ and $\varepsilon_t(p_t)$. Let $p_t^* = p_1/[p_1 + (1-p_1)(1-q)^{t-1}]$.

Proposition 6 *The optimal mechanism M^* consists of three stages:*

Theorem 1 Stage 1: *For all $t \leq \hat{t}$:*

$$\lambda_t^*(0) = 1, \text{ otherwise } \varepsilon_t^*(p_t) = 1.$$

Stage 2: *For all t such that $\hat{t} < t \leq t^*$,*

$$\lambda_t^*(0) = 1 - \varepsilon_t^*(0) = \frac{p_t^*}{p_t^* - p_1} \frac{(\mu_g - p_1)}{\mu_g}; \text{ otherwise } \varepsilon_t^*(p_t) = 1.$$

Stage 3: *For all t such that $t > t^*$,*

$$\lambda_t^*(0) = 1; \text{ otherwise } \gamma^*(p_t) = 1.$$

While the formal proof is relegated to the appendix, we will present its main idea and the intuition behind it.

The logic behind Stage 1 is based on the fact that during the first few periods of the mechanism, no conflict of interest exists between the principal and the

agents. They all prefer the action *pass* if they know that one of the agents received the signal s^l , and they also find it worthwhile to acquire a signal if they know that all previous agents received the signal s^g (i.e. their belief is in $[\mu_l, \mu_g]$). Thus, the optimal policy is based on full transparency as long as $t \leq \hat{t}$. It is important to note that because messages are private, there is no need to take into account the effect of this transparency on future agents. This is formally established in a series of simple claims presented in the appendix.

Stage 2 is more interesting because it highlights a potential conflict of interest between the agents and the principal. Following the principal's policy in Stage 1, the agents know that the principal's belief in Stage 2 is either $p_t = 0$ or $p_t > \mu_g$. When $p_t > \mu_g$ but is small enough, the principal would like the agents to acquire a signal whereas the agents would like to choose the action *invest*. To overcome this obstacle and make exploration incentive-compatible, the principal must recommend exploration after some histories, even when it already knows that the project is a lemon. This strategy of the principal leaves agents with sufficient uncertainty about the exact history so that their posterior remains in $[\mu_l, \mu_g]$. The strategy is, costly however since it means that agents acquire a signal when the state is already known, and moreover they may choose the action *invest* when the principal already knows that the project is a lemon. To minimize these costs, the mechanism randomizes in such a way that when the agent receives a recommendation to acquire a signal, his posterior is exactly μ_g . More precisely, because agent t knows that the principal's prior, p_t , is either 0 or $p_1/(p_1 + (1 - p_1)(1 - q)^{t-1}) \equiv p_t^*$, the principal must assign enough weight to exploration when $p_t = 0$ so that the agent's posterior, conditional on hearing the recommendation to explore, is equal to μ_g . This weight is exactly $\frac{p_1}{p_t^* - p_1} \frac{(p_t^* - \mu_g)}{\mu_g}$ (see the appendix for details). Once we establish that exploration beyond period \hat{t} implies setting the agent's posterior $\mu_t(e)$ to exactly μ_g , it remains to determine which agents, apart from agent \hat{t} , will receive a recommendation to explore and with what probability.

A key component of the proof is the relationship between the optimal mechanism and the first-best mechanism in a hypothetical environment, in which the cost of exploration is a function of the principal's belief, given by $c^*(p)$, as defined in equation 7. The cost $c^*(p)$ internalizes the additional cost of equating the agent's posterior, conditional on hearing the recommendation to explore, to μ_g . In this hypothetical environment, the more likely it is that the principal believes that the project is good, the greater will be the cost for him to acquire the signal. In the formal proof of Theorem 6, it is shown that the optimal

incentive-compatible mechanism in our problem can be obtained by solving the first-best problem in the hypothetical environment described above. The similarity between the two problems arises from the fact that the "indirect" cost that the principal has to incur, if it wishes agents to explore when $p_t > \mu_g$, appears as a direct cost in the hypothetical first-best problem. For every $p_t > \mu_g$, the cost of exploration in the hypothetical first-best environment, i.e. $c^*(p)$, is equal exactly to the direct cost of exploration plus the indirect cost of inducing agents to explore in our original problem.

The solution to the modified first-best problem determines the threshold p_g^* and which agent is the last to explore (i.e., t^*) in our original problem. Furthermore, the solution also shows that for $p_t^* > 0$ the optimal policy does not involve any randomization. The mechanism recommends *explore* for $t \leq t^*$ and *invest* for $t > t^*$.

By showing that the second-best mechanism can be presented as a modified first-best, we prove that no randomization takes place when $p_t > \mu_g$. However, unlike the original first-best, the modified mechanism does involve randomization when $p_t = 0$. The modified cost function captures the cost of this randomization exactly.

7 Symmetric Signals

The symmetric signals case is defined by $p_0 = 1/2$ and $q_g = q_l = q$. First, note that in this case (and given our assumption about the agents' payoffs) the first-best thresholds, p_g and p_l , are also symmetric, i.e. $p_l = 1 - p_g$. When the signals are symmetric, the beliefs conditional on signal realizations $\{n_l, n_g\}$ depend *only* on the difference $n_g - n_l$. Observe that if the mechanism is fully transparent, then a cascade will emerge after agent 1 since as soon as $|n_g - n_l| \geq 1$ no agent will find it beneficial to explore. However, and as will be shown below, there are histories in which the optimal mechanism calls for many agents to explore. Before arriving at the main result of this section, we prove that when signals are symmetric, there always exists a symmetric mechanism. A mechanism is symmetric if when it recommends *invest* for some history with some probability, then it recommends *pass* with the same probability for the mirror image of that history.

Proposition 7 *When the information structure is symmetric, there exists an optimal R&C mechanism that is symmetric.*

Proof. Assume by contradiction that a symmetric optimal mechanism does not exist. Let $(\{M_t\}_{t=1}^\infty, \{\nu_t\}_{t=1}^\infty)$ be an optimal asymmetric mechanism. Assume that this mechanism is symmetric up to and including agent t . We prove the proposition by constructing an optimal mechanism that is symmetric up to and including agent $t+1$.

Let $(\{M'_t\}_{t=1}^\infty, \{\nu'_t\}_{t=1}^\infty)$ be the "mirror image" of $(\{M_t\}_{t=1}^\infty, \{\nu_t\}_{t=1}^\infty)$ and observe that from the symmetry of the information structure it follows that $(\{M'_t\}_{t=1}^\infty, \{\nu'_t\}_{t=1}^\infty)$ is also optimal. Consider the following mechanism which is symmetric for the first $t+1$ agents: for all agents $t' \leq t$, the mechanism is as in the above two cases; for agent $t+1$, the mechanism, at the beginning of period $t+1$, randomizes equally between $(\{M_t\}_{t=1}^\infty, \{\nu_t\}_{t=1}^\infty)$ and $(\{M'_t\}_{t=1}^\infty, \{\nu'_t\}_{t=1}^\infty)$, and the realization is revealed to all future agents but not to agent $t+1$.

Clearly, the principal's payoff under this mechanism is identical to that of the two asymmetric mechanisms. Furthermore,, incentive compatibility is trivially satisfied for every agent $t' \neq t+1$. To see that it is also satisfied for agent $t+1$, consider the recommendation $a \in \{\text{invest}, \text{pass}, \text{explore}\}$ in Phase 1. It was incentive-compatible for agent $t+1$ to follow it in $(\{M_t\}_{t=1}^\infty, \{\nu_t\}_{t=1}^\infty)$ and in $(\{M'_t\}_{t=1}^\infty, \{\nu'_t\}_{t=1}^\infty)$, and therefore it is also incentive-compatible for him to do so when he does not know which of them is being played.

Since this result is true for every t , there is no "maximal" t beyond which the mechanism is asymmetric.

Proposition 8 When the information structure is symmetric, there exists an optimal R&C mechanism that never randomizes in Phase 1, i.e. either $\gamma_t(h_{t-1}) = 1$ or $\lambda_t(h_{t-1}) = 1$ or $\varepsilon_t(h_{t-1}) = 1$.

Proof. Assume that for some posterior, p_t , the mechanism randomizes between invest and explore with some strictly positive probability. Then, by symmetry, the mechanism also randomizes between pass and explore for the posterior $1-p_t$ with the same probability. Consider the modified mechanism according to which when the posterior is p_t or $1-p_t$ the mechanism recommends explore with probability 1. To see that the modified mechanism is also incentive-compatible, observe that conditional on the mechanism recommending explore the expected probability value of p_t is the same as in the original mechanism. Also observe that the recommendation to invest or pass is incentive-compatible since $p_t \notin [\mu_i, \mu_g]$ with probability 1 by Proposition 4. ■

■

8 Related Literature

Our model is designed to balance between the potential cost of an information cascade and the benefit of revealing information. The literature on information cascades originated with Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992) and was later developed by others (e.g., Lones Smith and Peter Sørensen (2003)). A few papers, however, have studied the design of mechanisms to reveal information in the contexts of information cascades. The two that we are aware of are Sgroi (2002) who evaluates - from the perspective of a social planner - the strategy of forcing a subset of agents to make decisions early and Smith, Sørensen, and Tian (2014) who conduct a welfare analysis of the herding model to show that the efficient outcome can be decentralized by rewarding individuals if their successor mimics their action.

Several papers have studied the role of costly information in social learning (Burguet and Vives (2000), Hendricks et al. (2012) and Mueller-Frank and Pai (2016)). Ali (2018) is the first one to study the impact of costly information in a herding model. None of these papers, however, have examined the design of optimal information disclosure when acquiring information is costly.

The current paper is also part of the literature on mechanism design without monetary transfers. Gershkov and Szentes (2009)'s voting model shares the sequential feature of our model. In the optimal mechanism, the social planner asks voters randomly and sequentially to invest in information gathering and to report the resulting signal. As in our model, the planner does not observe whether the agent invests in acquiring information, but in contrast to our model, the planner makes only one decision and all agents share the same payoff. Martimort and Aggey (2006) consider the problem of communication between a principal and a privately informed agent in the absence of monetary incentives.

Some other related papers are Ely (2016) who studied disclosure of information by a principal in a dynamic setup. Even though our mechanism and that of Ely (2016) share some common features, the situations we analyze are quite different.

Another related paper is Ely, Frankel, and Kamenica (2013) who analyze the optimal way in which to reveal information over time in order to maximize the expected suspense or surprise experienced by a Bayesian passive agent. Also relevant are Kamenica and Gentzkow (2011) and Rayo and Segal (2010) who consider optimal disclosure policies in a static environment, in which a principal wishes to influence an agent's choice by sending the correct message.

As in our model, Kremer, Mansour, and Perry (2014) and Che and Horner (2014) also study a mechanism design problem in which the planner wishes to aggregate information. In contrast, our paper relates to a different scenario, and the resulting optimal policies differ significantly.

References

- [1] S.N. Ali, S.N. (2018): "Herding with Costly Information," *Journal of Economic Theory*, 713–729.
- [2] Abhijit, V. and A. Banerjee (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797– 817.
- [3] Agrawal, A., C. Catalini and A. Goldfarb (2013): "Some Simple Economics of Crowdfunding," NBER working paper 19133.
- [4] Belleflamme P., N. Omrani and M. Peitz (2015): "The Economics of Crowdfunding Principals," working paper.
- [5] Block, J., and P. Koellinger (2009): "I Can't Get No Satisfaction—Necessity Entrepreneurship and Procedural Utility," *Kyklos* (62:2), pp. 191-209.
- [6] Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.
- [7] Burguet, R. and X. Vives (2000): "Social Learning and Costly Information Acquisition" *Economic Theory* 15(1):185-205.
- [8] Burtch, G., S. Carnahan and B. Greenwood (2016): "Can You Gig It? An Empirical Examination of the Gig-Economy and Entrepreneurial Activity," Ross School of Business Paper No. 1308.
- [9] Che, Y.K. and J. Horner (2014): "Optimal Design for Social Learning," working paper.
- [10] Ely, J., A. Frankel, and E. Kamenica (2013): "Suspense and Surprise," working paper.
- [11] Ely, J.C. (2015): "Beeps," *American Economic Review*, forthcoming.

- [12] Freedman, S., and G.Z. Jin (2008): "Do Social Networks Solve Information Problems for Peer-to-Peer Lending? Evidence from Prosper.com.," working paper.
- [13] Freedman, S. and G.Z. Jin (2011): "Learning by Doing with Asymmetric Information: Evidence from Prosper.com.," NBER Working Paper, No. 16855.
- [14] Geva H., O. Barzilay and G. Oestreicher-Singer (2016): "A Potato Salad with a Lemon Twist: Using Supply-Side Shocks to Study the Impact of Low-Quality Actors on Crowdfunding Principals," Collier School of Management, Tel Aviv University.
- [15] Gershkov, A. and B. Szentes (2009): "Optimal Voting Schemes with Costly Information Acquisition," *Journal of Economic Theory*, 144, 36–68.
- [16] Golder, P. N., and G. J. Telli. (2004): "Growing, Growing, Gone: Cascades, Diffusion, and Turning Points in the Product Life Cycle Marketing," *Science* 23(2) 207–218.
- [17] Hendricks, K, A. Sorensen, and T. (2012): "Observational Learning and Demand for Search Goods," *American Economic Journal: Microeconomics*, Vol. 4 No. 1, 1-31.
- [18] Howe, J. (2008): "Crowdsourcing: How the Power of the Crowd is Driving the Future of Business," New York, NY: Random House.
- [19] Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.
- [20] Kirmani, A., and Rao, A.R., (2000): "No Pain, No Gain: A Critical Review of the Literature on Signaling Unobservable Product Quality," *Journal of Marketing* (64:2), pp. 66–79.
- [21] Kremer, I., Y. Mansour, and M. Perry (2014): "Implementing the Wisdom of the Crowd," *Journal of Political Economy*, Vol. 122, No. 5, 988-1012.
- [22] Martimort, D. and S. Aggey (2006): "Continuity in Mechanism Design without Transfers," *Economics Letters*, 93, 182–189.
- [23] Mueller-Frank, M., and Mallesh M. Pai (2016): "Social Learning with Costly Search," *American Economic Journal: Microeconomics*, Vol. 8, No. 1, 83-109.

- [24] Myerson R. (1986): "Multistage Games with Communication," *Econometrica* 54, 323-358.
- [25] Rayo, L. and I. Segal (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949-987.
- [26] Sgroi, D. (2002): "Optimizing Information in the Herd: Guinea Pigs, Profits and Welfare," *Games and Economic Behavior*, 39, 137-166.
- [27] Skrzypacz, A. and J. Horner (2012): "Selling Information," working paper.
- [28] Smith, L. and N. Sorensen (2003): "Pathological Outcomes of Observational Learning," *Econometrica* Volume 68, No 2, 247-524.
- [29] Smith, L., N. Sorensen, and J. Tian (2014): "Informational Herding, Optimal Experimentation, and Contrarianism," working paper.
- [30] Wald, A. (1947): *Sequential Analysis*. New York: John Wiley and Sons. ISBN 0-471-91806-7.

9 Appendix

9.1 Proof of Proposition 1

Proof: When the principal does not face incentive constraints, the expected payoff in every period depends only on p_t , namely its prior belief in that period. Hence, the first-best solution can be viewed as a dynamic optimization where the state variable is p_t . In principle, the principal can randomize, at least in some periods, over the three alternatives, i.e. *invest*, *pass*, and *explore*, but such a randomization will never be necessary. If the principal randomizes between, say, two actions, then his expected payoff is simply the average of the payoffs for the two actions. Thus, we can conclude that the first-best mechanism is deterministic.

By assumption, it is optimal for the first few agents (or perhaps just the first one) to explore and hence it is also optimal for the principal to do so during the first few periods. Since $c > 0$, we know that there exists a posterior probability p_t large enough so that at this posterior it would not be optimal, from the principal's perspective, for an agent to acquire a signal, even if it is

fully informative. This is particularly the case if:

$$\frac{p_t}{(1-\delta)} = \frac{1}{(1-\delta)} - c \rightarrow p_t \geq 1 - c(1-\delta).$$

Similarly, for small enough p_t , it is socially optimal to pass.

Note also that if for some posterior p_t , the first-best mechanism chooses the action *invest* or *pass* without exploration, then it will do so in all subsequent periods. This follows immediately from the fact that if, at period t , the action *invest* or *pass* is chosen, then $p_{t+1} = p_t$, and therefore the optimal action in $t+1$ is the same.

The above argument implies that the payoff from choosing the action *invest* when the posterior is p_t is $p_t/(1-\delta)$ and similarly the payoff from choosing the action *pass* when the posterior is p_t is $(1-p_t)/(1-\delta)$. Thus, the payoffs from *invest* and *pass* are linear in p_t and are larger than those from choosing to explore in the neighborhood of 1 and 0, respectively. By assumption, the payoff from *explore* is higher than that from either *pass* or *invest* when p_t is close to 1/2. To complete the argument, we now show that the payoff from choosing to explore in period t is convex in p_t . Namely,

$$\alpha V(e_{1t} | p_{1t}) + (1-\alpha)V(e_{2t} | p_{2t}) \geq V(e' | \alpha p_{1t} + (1-\alpha)p_{2t})$$

where

$$p'_t = \alpha p_{1t} + (1-\alpha)p_{2t}, \text{ for } 0 \leq \alpha \leq 1,$$

and $V(e | p)$ is the principal's payoff in period t when the prior is p and he chooses to explore in period t and to continue optimally thereafter. Note that:

$$\begin{aligned} V(e' | p'_t) &= p'_t V(e' | p'_t, G) + (1-p'_t) V(e' | p'_t, L) = \\ &= \alpha p_{1t} V(e' | p'_t, G) + (1-\alpha) p_{2t} V(e' | p'_t, G) + \\ &+ \alpha(1-p_{1t}) V(e' | p'_t, L) + (1-\alpha)(1-p_{2t}) V(e' | p'_t, L) \\ &\leq \alpha V(e_{1t} | p_{1t}) + (1-\alpha) V(e_{2t} | p_{2t}). \end{aligned}$$

9.2 Proof of Proposition 5

By assumption, it is optimal for the first agent to explore and hence it is also optimal for the principal to do so in the first period. Also, by assumption and since the signal s^l reveals the project to be a lemon, we know that if, at some period t , the principal observes agent t choosing the action *pass* (after obtaining

a signal), it is optimal for all subsequent agents to also choose the action *pass*. The only remaining question is what the principal should do in a given period if, in all previous periods in which agents chose to explore, the realization of the signal was s^g (i.e., all agents chose the action *invest*). Since $c > 0$, we know that there exists a prior probability p_t large enough so that at this prior it would not be optimal from the principal's perspective for an agent to acquire a signal, even if the signal is fully informative (i.e., $q = 1$). Thus (and recalling the updating rule 1), there exists a period \bar{t} large enough so that if all agents observe the signal s^g in all periods $t \leq \bar{t}$, it is optimal for the principal to choose the action *invest* in all periods $t > \bar{t}$. It remains to be shown that if for some $t < \bar{t}$ the first-best mechanism chooses the action *invest* without exploration, then it will do so in all subsequent periods. This last claim follows immediately from the fact that if at period t the action *invest* is chosen, then $p_{t+1} = p_t$ and therefore the optimal action in $t + 1$ is also *invest*.

Notice that p_g is the solution to the following equation:

$$\frac{p}{(1-\delta)} = \frac{p}{(1-\delta)} + (1-p)\frac{q}{(1-\delta)} - c.$$

The LHS is the principal's payoff from choosing the action *invest* in the current period and in all periods thereafter, whereas the RHS is the principal's payoff from exploring one more time and following the signal in all periods thereafter when the prior is p . If $p < p_g$, the RHS is greater than the LHS and therefore the principal is better off exploring at least one more time. It is left to show that if $p > p_g$, the optimal action is *invest*. Assume, by contradiction, that there exists some $p' > p_g$ at which the optimal action is to acquire a signal. Since we know that there exists a posterior large enough such that it is optimal for the mechanism to choose *invest*, there must be $p'' \geq p'$ at which it is optimal for the principal to acquire only one more signal and to follow it in all periods thereafter. This, however, leads to a contradiction since the LHS of the equation above increases with p at a higher rate than the RHS and hence it must be the case that it is better to terminate exploration a period earlier.

9.3 Proof of Proposition 6

The proof is accomplished by proving several claims which, taken together, characterize the optimal recommendation mechanism M^* . Hereafter, we refer only to p_t for which, given the mechanism in place, there is a positive probability

that the principal holds that belief, i.e., $\pi_t^*(p_t) > 0$. Note that given Assumption 1, every incentive-compatible mechanism must recommend that agent 1 explore, simply because that is what he is going to do in any case. The following claim states that for small enough t , when it is common knowledge that there is no conflict of interest between agent t and the principal, a policy of full revelation is adopted and it is recommended that the agent choose the best action from his perspective.

Claim 2 *For all $t \leq \hat{t}$, if the principal's prior at t is strictly positive (i.e., $p_t > 0$), then $\varepsilon_t^*(p_t) = 1$; otherwise $\lambda_t^*(0) = 1$.*

Proof. By the definition of \hat{t} , for all $t \leq \hat{t}$, it is common knowledge that $\Pr(\tilde{p}_t \in [\mu_l, \mu_g] \cup \{0\}) = 1$ and hence the agent's optimal choice is e if he knows that $\tilde{p}_t \neq 0$ and *pass* otherwise. This is also the first-best choice. Consequently, if the agent is fully informed, then he will follow the first-best choice strategy. Assume, by contradiction, that the optimal IC mechanism is such that there exists some agent $t' \leq \hat{t}$ who is not fully informed about the moves of all preceding agents and as a result does not choose the action that is optimal for him. Consider a modified mechanism under which agent t' is informed about the moves of his predecessors and in all periods thereafter the principal ignores the additional information obtained in period t' and instead follows the original mechanism. Clearly, this modified policy yields a higher level of social welfare, a contradiction. ■

A consequence of Claim 2 is that the optimal mechanism essentially reveals p_t to agent t in the first \hat{t} periods, and unless s^l is observed by one of the agents, all agents $t \leq \hat{t}$ acquire a signal. The following corollary follows directly from Claim 2 together with the posterior's law of motion and the assumption that $p_1 \in [\mu_l, \mu_g]$.

Corollary 1 *For $t \leq \hat{t}$, $\Pr(\tilde{p}_t \in [\mu_l, \mu_g] \cup \{0\}) = 1$ and for $t > \hat{t}$, $\Pr(\tilde{p}_t = 0) > 0$ and $\Pr(\tilde{p}_t \in (0, \mu_g]) = 0$.*

Claim 2 and Corollary 1 allow us to restrict the search for the optimal recommendation mechanism to histories in which $p_t = 0$ or $p_t > \mu_g$.

Claim 3 *For all t , the principal never recommends the action invest if he knows that the project is a lemon, i.e., $\gamma_t^*(0) = 0$, and for $p_t > \mu_g$ the principal never recommends the action pass, i.e., $\lambda_t^*(p_t) = 0$.*

Proof. Suppose that at some t we have $p_t = 0$ and $\gamma_t^*(0) > 0$. Consider the modified mechanism such that $\gamma_t'(0) = 0$ and $\lambda_t'(0) = \lambda_t^*(0) + \gamma_t^*(0)$. Such a modification will certainly increase agent t 's utility and will not affect the other agents' payoffs since the posterior beliefs in all periods thereafter do not change. A similar proof shows that $\lambda_t^*(p_t) = 0$ for all $p_t > \mu_g$. ■

An immediate consequence of Corollary 1 is that for $t > \hat{t}$, agent t knows that some previous agents (and, in particular, agent 1) chose e and with some positive probability acquired the signal s^l . Thus, agent t knows that either the principal has learned that $\tilde{p}_t = 0$, i.e., the project is a lemon, or the principal's posterior is $\tilde{p}_t > \mu_g$. While we know from Claim 3 that $\gamma_t^*(0) = 0$, this does not imply that $\lambda_t^*(0) = 1$.

Claim 4 *Consider period $t > \hat{t}$. If $\varepsilon_t^*(p_t) > 0$ for $p_t > \mu_g$, then $\varepsilon_t^*(0) > 0$ and $\mu_t^*(e) = \mu_g$.*

Proof. If $\pi_t^*(p_t) > 0$, then from Claim 1 we know that $\Pr(\tilde{p}_t \in (0, \mu_g]) = 0$ for some $p_t > \mu_g$. The proof of the first part of the claim follows immediately from the fact that in order for the mechanism to be incentive-compatible, it must be the case that the agent's posterior, conditional on the recommendation to explore, i.e., $\mu_t^*(e)$, is such that $\mu_t^*(e) \in [\mu_l, \mu_g]$. Now assume, by contradiction, that $\mu_t^*(e) < \mu_g$. It is then possible to decrease $\varepsilon_t^*(0)$ by a small amount and increase $\lambda_t^*(0)$ by that same amount. This increases agent t 's utility without affecting the distribution of \tilde{p}_τ^* for all $\tau > t$, a contradiction. ■

Taken together, the above claims summarize the incentive-compatible constraints within which the optimization is carried out. Starting with the prior $p_1 \in [\mu_l, \mu_g]$, the random variable \tilde{p}_t is either zero or above p_1 . When $\tilde{p}_t = 0$, the principal recommends either *pass* or *explore*; similarly, when $\tilde{p}_t > \mu_g$, the principal recommends either *invest* or *explore*; and when $\tilde{p}_t \in [\mu_l, \mu_g]$, the principal recommends *explore*. Moreover, whenever the recommendation is *explore*, it must be that $\mu_t(e) \in [\mu_l, \mu_g]$ (where $\mu_t(\text{explore})$ is the agent's posterior following the recommendation explore). When $t > \hat{t}$, the agent knows that $\Pr((\tilde{p}_t > \mu_g) \cup (\tilde{p}_t = 0)) = 1$. The expected value $\mu_t(\cdot)$ is over all possible p_t 's

given the mechanism in place and, in order to provide agent $t > \hat{t}$ with the incentive to explore when $\tilde{p}_t > 0$, the mechanism must also recommend exploration when $\tilde{p}_t = 0$ so that $\mu_t(\text{explore})$ will be in the region where the agent is willing to explore, i.e., $[\mu_l, \mu_g]$. Recommending exploration when $\tilde{p}_t = 0$ is costly, due to the unnecessary cost of exploration when the principal already knows that the project is a lemon and because the agent may obtain the signal s^g and choose the wrong action. To minimize these "costs", exploration, when $\tilde{p}_t = 0$, is recommended with the smallest probability necessary, which implies that the agent's posterior becomes exactly μ_g (i.e., $\mu_t(\text{explore}) = \mu_g$).

With this in mind, we can now derive the optimal mechanism by solving a modified first-best problem in which the cost of acquiring a signal includes not only c but also the implied cost to maintain the agent's posterior, i.e. $\mu_t(\text{explore})$, at μ_g . As will be shown, this cost is monotonic in p , and hence we can employ the same technique used in the solution to the original first-best problem to establish that, for $t > \hat{t}$ and $p_t > \mu_g$, the optimal solution is deterministic, namely, either $\varepsilon_t^*(p_t) = 1$ or $\gamma^*(p_t) = 1$. Before proving this result formally, the following discussion will be helpful.

Let $\bar{\varepsilon}_t(p) = \pi_t(p)\varepsilon_t(p)$ denote the "ex ante" probability of exploration at p in period t . As discussed above, the efficient way to satisfy the IC constraints implies that:

$$\mu_t^*(\text{explore}) = \frac{\sum_{\tilde{p}_t} p \bar{\varepsilon}_t(p)}{\sum_{\tilde{p}} \bar{\varepsilon}_t(p)} = \frac{\sum_{\tilde{p}_t > 0} p \bar{\varepsilon}_t(p)}{\bar{\varepsilon}_t(0) + \sum_{\tilde{p}_t > 0} \bar{\varepsilon}_t(p)} = \mu_g,$$

which can be written as:

$$\bar{\varepsilon}_t(0) = \frac{\sum_{\tilde{p}_t > 0} p \bar{\varepsilon}_t(p) - \mu_g \sum_{\tilde{p}_t > 0} \bar{\varepsilon}_t(p)}{\mu_g}.$$

Thus, if for some $p > 0$, exploration at p (i.e., $\bar{\varepsilon}_t(p)$) is increased by one unit, this will have a direct cost of c and an indirect cost of:

$$H(p) = \frac{p - \mu_g}{\mu_g} [c + (1 - q)],$$

which is the cost of increasing exploration when the project is already known to be a lemon. Note that $H(p)$ is increasing in p .

We can now show that the optimal mechanism can be viewed as the solution to a modified first-best problem in which the cost of exploration is:

$$c^*(p) = c + H(p).$$

Using the same line of reasoning and as in the solution of the (original) first-best mechanism, we denote by p_g^* the solution to the following equation:

$$p = 1 - c^*(p) \frac{1 - \delta}{q}.$$

We note that: (i) since $c^*(p)$ is non-decreasing, p_g^* is uniquely defined by the above equation; and (ii) $\mu_g < p_g^* < p_g$. We can therefore conclude that the solution to the first-best problem with a modified cost function is given by:

$$\begin{aligned} \gamma_t(p) &= 1 \text{ for } p > p_g^*, \\ \varepsilon_t(p) &= 1 \text{ for } 0 < p < p_g^* \\ \lambda_t(0) &= 1 \text{ for } t \leq \hat{t} \text{ and for } t > t^*. \end{aligned} \tag{9}$$

In what follows, we first formally define the planner's (second-best) problem for the case of $p_t > \mu_g$. We then show that the IC constraints can be plugged into the planner's objective function, thus transforming the second-best problem into a first-best problem with the modified cost function.

Claim 5 *For every $\tau > \hat{t}$, the optimal mechanism must satisfy the following maximization problem (referred to as the SB problem hereafter):*

$$\underset{\{\gamma_t(p)\}_{t=\tau}^\infty, \{\lambda_t\}_{t=\tau}^\infty}{Max} V_\tau =$$

$$\sum_{p > \mu_g} \pi_\tau(p) (\gamma_\tau(p)p + (1 - \gamma_\tau(p)) (p + (1 - p)q - c)) + \pi_\tau(0)(\lambda_\tau + (1 - \lambda_\tau)(q - c)) + \delta V_{\tau+1}$$

subject to:

$$(I) V_t = \sum_{p > \mu_g} \pi_t(p) (\gamma_t(p)p + (1 - \gamma_t(p)) (p + (1 - p)q - c)) + \pi_t(0)(\lambda_t + (1 - \lambda_t)(q - c)) + \delta V_{t+1},$$

for $t = \tau + 1, \tau + 2, \tau + 3, \dots$

$$(II) \pi_t(p) = \pi_{t-1}(p)\gamma_{t-1}(p) + \pi_{t-1}(p_{t-1}^{-1}(p))(1 - \gamma_{t-1}(p_{t-1}^{-1}(p)))(p_{t-1}^{-1}(p) + (1 - p_{t-1}^{-1}(p))(1 - q)),$$

for $p > \mu$ and $t = \tau + 1, \tau + 2, \tau + 3, \dots$

$$(III) \sum_{p > \mu_g} [(\mu_g - p)\pi_t(p)(1 - \gamma_t(p)) + \pi_t(0)(1 - \lambda_t)\mu_g] = 0,$$

for $t = \tau, \tau + 1, \tau + 2, \tau + 3, \dots$ and

$$(IV) \pi_t(0) = 1 - \sum_{p > \mu_g} \pi_t(p), \gamma_t(p) \in [0, 1], \lambda_t \in [0, 1],$$

for $t = \tau, \tau + 1, \tau + 2, \tau + 3, \dots$

where:

(i) $\pi_\tau(p)$ is the distribution of possible beliefs at period τ , given the mechanism;

(ii) $p_t^{-1}(p)$ is the inverse of the function

$$p(p_t) = \frac{p_t}{p_t + (1 - p_t)(1 - q)},$$

which is the probability that the posterior is p given that the prior was p_t and the agent explored and acquired the signal s^g ;

(iii) the function V_t is the principal's expected present value, given a mechanism M , for all $t > \tau$;

(iv) the second constraint specifies the evolution of the distribution of the random variable \tilde{p}_t , given the mechanism M ; and

(v) the third constraint specifies the incentive-compatible constraint, guaranteeing that agent t 's posterior is exactly μ_g when the mechanism recommends explore.

Proof. From the incentive-compatible constraint (iii), we obtain that, for every $t \geq \hat{t}$,

$$\pi_t(0)(1 - \lambda_t) = \sum_{p > \mu_g} \frac{(p - \mu_g)}{\mu_g} \pi_t(p)(1 - \gamma_t(p)). \quad (10)$$

Plugging (10) into V_t , we obtain that for all $t \geq \hat{t}$,

$$V_t = \sum_{p > \mu_g} \pi_t(p) (\gamma_t(p)p + ((1 - \gamma_t(p)) (p + (1 - p)q - c^*(p))) + \pi_t(0) + \delta V_{t+1},$$

where

$$c^*(p) = c + \frac{(p - \mu_g)}{\mu_g}(1 + c - q).$$

Thus, the SB problem can be simplified as follows (and is hereafter referred to as the SB' problem):

$$\begin{aligned} \text{Max}_{\{\gamma_t(p)\}_{t=\tau}^{\infty}} V_{\tau} = \\ \sum_{p > \mu_g} \pi_{\tau}(p) (\gamma_{\tau}(p)p + ((1 - \gamma_{\tau}(p)) (p + (1 - p)q - c^*(p))) + \pi_{\tau}(0) + \delta V_{\tau+1}, \end{aligned}$$

subject to:

$$(I) V_t = \sum_{p > \mu_g} \pi_t(p) (\gamma_t(p)p + ((1 - \gamma_t(p)) (p + (1 - p)q - c^*(p))) + \pi_t(0) + \delta V_{t+1},$$

for $t = \tau + 1, \tau + 2, \tau + 3, \dots$,

$$\begin{aligned} (II) \pi_t(p) &= \pi_{t-1}(p)\gamma_{t-1}(p) + \\ &\pi_{t-1}(p_{t-1}^{-1}(p))(1 - \gamma_{t-1}(p_{t-1}^{-1}(p)))(p_{t-1}^{-1}(p) + (1 - p_{t-1}^{-1}(p))(1 - q)) \end{aligned}$$

for $p > \mu_g$ and $t = \tau + 1, \tau + 2, \tau + 3, \dots$, and

$$(III) \pi_t(0) = 1 - \sum_{p > \mu_g} \pi_t(p), \gamma_t(p) \in [0, 1], \lambda_t \in [0, 1],$$

for $t = \tau, \tau + 1, \tau + 2, \tau + 3, \dots$

To complete the proof, we can now show that the $\{\gamma_t(p)\}_{t=\tau}^{\infty}$ which solves the SB' problem above also solves a modified first-best problem in which the cost of exploration is $c^*(p)$. Notice that for every period $\tau > \hat{t}$, the solution to the modified first-best problem is given by the solution to our original SB problem with the following adjustments: c is replaced everywhere by $c^*(p)$, the IC constraint (III) is deleted, and $\lambda_t(0) = 1$ in all periods. It follows that the solution to the modified first-best problem is identical to that of the SB'

problem. Thus, the optimal mechanism (i.e., the solution to the original SB problem) does not randomize at $p_t > \mu_g$ and either $\varepsilon_t(p) = 1$ (for $0 < p < p_g^*$) or $\gamma_t(p) = 1$ (for $p > p_g^*$).

■