

When half the truth is better than the truth: Theory and experiments on aggregate information cascades*

Antonio Guarino, Heike Harmgart and Steffen Huck[†]

Abstract

We introduce a new model of *aggregate information cascades* where only one of two possible actions is observable to others. When called upon, agents (who decide in some random order that they do not know) are only informed about the total number of others who have chosen the observable action before them. This informational structure arises naturally in many applications. Our most important result is that only one type of cascade arises in equilibrium, the aggregate cascade on the observable action. A cascade on the unobservable action never arises. We tested our theory in the laboratory, to check whether it can be expected to be of empirical relevance. The answer is, luckily, affirmative. Our results may have important policy consequences. Central agencies, for example in the health sector, may optimally decide to *withhold* information from the public.

1 Introduction

A central agency in health policy must decide how to disclose information on the adoption of a new treatment. One possibility is to inform the doctors on how many others have already decided to adopt the new treatment. Another is to inform them on how many have considered doing it but have judged

*We thank Giuseppe Della Corte and Giulia Sestrieri for useful comments. We gratefully acknowledge financial support from the ESRC and the Leverhulme Trust.

[†]Guarino: Department of Economics and ELSE, University College London (e-mail: aguarino@ucl.ac.uk); Harmgart: European Bank for Reconstruction and Development (EBRD) (e-mail: harmgarh@ebrd.com); Huck: Department of Economics and ELSE, University College London (e-mail: s.huck@ucl.ac.uk).

that it is preferable to stick to the old practice. A third possibility is to reveal both, the number of doctors in favor of the new practice and the number of physicians in favor of the old one. Can the way the information is disclosed make a difference for the diffusion of the new treatment? Suppose the agency is uncertain about the effects of the new treatment and considers as the worst case scenario the situation in which the new treatment is widely adopted while ultimately resulting in worse health outcomes than the old treatment, for instance because of side effects. Which disclosure policy should the agency employ?

Intuitively, one would think that the disclosure of *all* available information should maximize social welfare. This only holds, however, if there are no externalities and a doctor's decision about a treatment for a patient is obviously a case with huge externalities. While the life of a patient may matter a lot to the physician it surely matters more to the patient, at least in most cases.

In their seminal paper on information cascades, Bikhchandani et al. (1992) have argued that the adoption of medical procedures is often based on fairly weak information and that in many cases doctors tend to imitate others. As an example they cite the widespread use of tonsillectomies in the sixties and seventies and argue that it was essentially an information cascade. In an informational cascade, agents rationally neglect their own private information, i.e., they choose the same action independently of the information they receive (for instance because they follow the decisions of the predecessors). In the sixties and seventies, according to Bikhchandani et al. (1992), the sheer fact that the majority of physicians employed the procedure overrode any private information individual doctors might have had against tonsillectomies. And this was a “wrong cascade”—a cascade that generated the worst outcome, since it eventually turned out that tonsillectomies did more harm than good.

One of the questions we raise in this paper is whether a central agency that can influence the transmission information may want to decide to *withhold* some information. And indeed it turns out that this might be the case. This is due to the following simple result: If agents have access to information about both how many others have adopted and how many have not, both types of cascades are possible—cascades where everybody adopts and cascades where everybody does not adopt (this is, of course, what we know from the literature

already). If, however, the agency decides only to inform about how many have adopted, then *there is only one type of cascade*—one where everybody adopts. And if the agency only reports about those who have decided *not* to adopt, there is also only one type of cascade, one where everybody decides *not* to adopt. Thus, if the really bad outcome is the one where everybody herds on the new treatment while the new treatment is, in fact, worse than the old, the agency can avoid that this happens—by withholding the information about how many others so far have decided to adopt.

While information about medical procedures appears to be a particularly appealing policy-relevant example for the theory we develop here, there are, in fact, many other applications. At the core of our paper is a simple model of social learning where agents have to make binary decisions, like adopting a new treatment or not, or making an investment or not. Our crucial assumption (where we deviate from the previous literature) is that when an agent makes his decision, he can only observe the total number of others who have already taken *one* of the two available actions, for example, the total number of others who have already decided to adopt the new treatment (while he cannot observe the number of those who have previously contemplated the choice but opted otherwise).

This structure appears very natural for many examples that have previously been discussed in the literature on social learning. A restaurant goer, who must decide whether or not he wants to dine at a particular restaurant he stands in front of, may be able to peer through the window to see how many others decided to have dinner there, but he can only speculate about how many others stood before the same door and decided to pass. A similar example is that of a manager of a venture capital firm who discusses a project with an inventor who needs capital to develop a new product. Say, the inventor has already secured funds from two other venture capital firms. Then we may expect that the inventor gladly mentions this to the manager with whom he negotiates. The manager who will have some private information about the viability of the project will, of course, also extract some information from knowing that there were also two other firms who thought the project was good. On the other hand, since we can safely assume that the inventor will not tell the manager about how often he was turned down, the manager can only guess how many other firms had information telling them that the

project was bad.¹

In all these examples, agents who have to decide between two options have only aggregate information about one of the two options while in all models present in the literature so far, a decision maker has access to information about the *individual* choices of others who decided before him. The standard model of social learning (Banerjee, 1992, and Bikhchandani et al., 1992), for instance, contemplates a sequence of binary decisions which are all observable. Agent n knows whether each predecessor in the sequence, from agent 1 to agent $n - 1$, decided in favor of one option or the other. We find that our set up captures quite naturally some situations frequently arising in social interactions. Like in the case of the restaurant goer or the venture capital firm, in many circumstances, a decision maker can gather some aggregate information (how many agents have already adopted, invested, chosen a restaurant), but he can rarely observe all the individual decisions. Clearly, if the decision is binary, knowing the number of agents who have made a certain decision also helps to update on the number of agents who have made the opposite decision. But this is not equivalent to knowing it. And we can show that it makes an important difference. In particular, we can show, as we have mentioned already above, that with aggregate information there can only be one type of information cascade, a cascade on the observable action.

In the standard sequential model of Banerjee (1992) and Bikhchandani et al. (1992) different types of cascades can arise. If the decision is binary, say, between investing and not, there can be cascades where, from a certain point onwards, *all* decision makers decide *to invest*, as well as cascades where, from a certain point onwards, all decisions makers decide *not to invest*. At a first glance, one could think that this is the case in our set up, too. If a restaurant goer sees many people in a restaurant, he could disregard his information and just join the crowd; and if he sees the restaurant empty, he could decide to pass on the restaurant independently of his signal, too. We can prove that, on the contrary, in our model only the first cascade is possible. In equilibrium it cannot be that there is a cascade on the unobservable action, e.g., that

¹This example could also be extended to the market for syndicated loans where several banks jointly offer funds to a borrowing firm. See for example Sufi (2007) for an empirical analysis of the effect of information provision between several lenders and the borrower on the syndicate structure of the contract.

the restaurant remains empty although some people have read good reviews about it.

In many cases, one of the available actions arises naturally as the observable action. In one of the examples above it is the number of investment decisions (as opposed to the number of declined investments). In another it is the number of decisions to dine in a particular restaurant (as opposed to the number of passes). There are, however, important cases where *third parties* may have the power to decide what kind of information is provided to agents. This is the case from which we started: the disclosure policy of a health agency. The worst case scenario for the agency is that in which agents start to herd on the wrong action. This induces the biggest welfare loss. The central agency may not know what the wrong action is (if it did, it could simply announce it) but often the welfare losses will be asymmetric: the case where the new treatment is indeed more effective but everybody sticks to the old one may be better or worse than the opposite case, where the old treatment is more effective but everybody switches to the new. Since the central agency can *choose* which action to make observable, it can rule out one of the two erroneous cascades—by withholding information. In summary, we like to argue that the aggregate information set up that we introduce here has not only several intriguing properties—some of which are in stark contrast to the predictions of the standard model—it also has potentially important policy implications.

We extend the literature on social learning not only theoretically but also experimentally. After establishing our theoretical results, we implement our set up in a laboratory experiment. The aim of this experiment is to have a first “reality check,” since we believe that a theory like ours is more appealing if it is not completely off the mark. Previous models of social learning, starting with the standard sequential model with all observable actions, have been extensively tested in the laboratory, with results that are sometimes favorable and sometimes less so. For our theory we find support in the laboratory, although some interesting anomalies emerge. In just two simple treatments we find that the main comparative statics go all in the right directions. In particular, while we observe cascades on the observable action, cascades on the unobservable actions either do not occur (in one treatment) or rarely occur (in the other).

Other papers in the social learning literature have studied what happens when we remove the strong assumption that agents can observe the entire history of individual decisions. Smith and Sørensen (1998) study a sequential decision model in which agents can only observe unordered random samples from predecessors' actions (e.g., because of word of mouth communication). With unbounded private signals, eventually, in their model complete learning obtains. Similarly, Banerjee and Fudenberg (2004) present a model in which, at every time, a continuum of agents choose a binary action after observing a sample of previous decisions (and, possibly, of signals on the outcomes). This can be interpreted as a model of word of mouth communication in large populations. The authors find sufficient conditions (on the sampling rule, etc.) for herding to arise, and conditions for all agents to settle on the correct choice. Çelen and Kariv (2004) extend the standard model of sequential social learning by allowing each agent to observe the decision of his immediate predecessor only. The prediction of these authors is that behavior does not settle on a single action. Long periods of herding can be observed, but switches to the other action occur. As time passes, the periods of herding become longer and longer, and the switches more and more rare. Finally, Larson (2006) is close to our paper in that it analyzes a situation in which agents observe the pooled average action of a population of their predecessors (before making a choice in a continuous action space). In contrast to our work, the focus of the study is not on whether a cascade occurs or not, but on the speed of learning (since the continuous action space guarantees that complete learning eventually occurs).

The remainder of the paper is organized as follows. In Section 2 we introduce the formal model. We present its equilibrium analysis in Section 3. Section 4 contains an example. Section 5 describes the experiment and its results, and Section 6 concludes.

2 The Model

In our economy there are n agents who have to decide in sequence whether or not to take up a certain option. For convenience, we shall refer to this choice as the decision about whether or not to *invest*. Time is discrete and indexed by $t = 1, 2, \dots, n$. Each agent makes his choice only once in the sequence. Agent i 's ($i = 1, 2, \dots, n$) action space is given by $\{0, 1\}$, where 1 is interpreted

as investment. Player i 's action is denoted by $I_i \in \{0, 1\}$. An agent's payoff π_i depends on his choice and on the true state of the world $\omega \in \{0, 1\}$. The prior probability of $\omega = 1$ is $r \in (0, 1)$. If $\omega = 1$ agent i receives a payoff of 1 if he chooses to invest, and a payoff of zero otherwise; vice versa if $\omega = 0$. That is,

$$\pi_i = \omega I_i + (1 - \omega)(1 - I_i).$$

The sequence in which agents make their choices is randomly determined before the first agent makes a decision, and agents are, w.l.o.g., (re-)numbered according to their positions: agent i chooses at time i only. All sequences are equally likely. The agents are, however, *not* informed about which sequence has been chosen. Furthermore, they do not know their own position in the sequence. When called upon, agent i is only informed about the total number of agents before him who have decided to invest. In other words, the decision to invest is assumed to be the only *observable* action. This means that, while the aggregate number of investments is observable, each individual decision to invest or not is not publicly known. We denote the total number of agents who have invested before agent i by T_i , i.e., agent i is informed about $T_i = \sum_{j=1}^{i-1} I_j$. In addition to observing T_i , each agent i receives a private signal $\sigma_i \in \{0, 1\}$ that is correlated with the true state ω . In particular, we assume that each agent receives a symmetric binary signal distributed as follows:

$$\Pr(\sigma_i = 1 \mid \omega = 1) = \Pr(\sigma_i = 0 \mid \omega = 0) \equiv q.$$

Note that, conditional on the state of the world, the signals are i.i.d.. We shall refer to $\omega = 1$ as the “good state” and to $\omega = 0$ as the “bad state.” A signal pointing in the direction of the good state ($\sigma_i = 1$) shall be called a “good signal” and a signal pointing in the opposite direction ($\sigma_i = 0$) a “bad signal.” We assume that $1 > q > r$ and that $r + q > 1$. These conditions ensure that, in the one-agent case, an agent would invest after a good signal but not after a bad signal, which renders the problem interesting. Note that these two conditions also imply that $q > \frac{1}{2}$, i.e., that the signal respects the monotone likelihood ratio property. Finally, the signal is not perfectly informative, which makes social learning possible and relevant.

Agent i 's information set is, therefore, represented by the couple (T_i, σ_i) . An agent's strategy \mathfrak{J}_i maps (T_i, σ_i) into an action, i.e.,

$$\mathfrak{J}_i : \{0, 1, 2, \dots, n - 1\} \times \{0, 1\} \rightarrow \{0, 1\}.$$

An agent's mixed strategy induces, for each (T_i, σ_i) , a probability with which the agent invests. We denote the probability with which agent i invests after observing (T_i, σ_i) by $\mathcal{I}_i(T_i, \sigma_i)$.

To conclude the description of our model, it is useful to introduce the notion of an *aggregate information cascade*. The definition is virtually identical to the standard definition of information cascade, with the characteristic that histories are summarized by the *aggregate statistic* T_i .

Definition 1 *An aggregate information cascade (AIC) occurs when, along the equilibrium path, there is a decision after which all agents choose an action independently of their signal. In particular:*

In an aggregate up cascade (AUC) there is a critical trigger value T^{UP} such that if $T_k = T^{UP}$ all agents from k onwards choose to invest regardless of their private signals. Consequently, there is some k such that $T_{k+j} = T_k + j$ for all $j = 1, \dots, n - k$.

In an aggregate down cascade (ADC) there is a critical trigger value T^{DOWN} such that if $T_k = T^{DOWN}$ all agents from k onwards choose not to invest regardless of their private signals. Consequently, in an ADC there is some k such that $T_{k+j} = T_k$ for all $j = 1, \dots, n - k$.

We are now ready to start analyzing the equilibrium decisions in our economy.

3 Equilibrium Analysis

The ultimate goal of our analysis is to understand the social learning process that occurs in our economy. Each agent can learn about the true state of the world from the aggregate information that he receives about other agents' choices. This can lead to better decisions. On the other hand, it may be that also in our economy, as in the canonical model of social learning of Banerjee (1992) and Bikhchandani et al. (1992), there is room for information cascades. In such a case, the process of information aggregation will not be efficient. We will show that, indeed, "up cascades" of investments are possible even in our set up, as they are in the canonical model. In contrast, "down cascades" of non-investments never occur in equilibrium.

We shall restrict the entire analysis to symmetric Perfect Bayesian Nash

equilibria (PBNEs). For convenience, we shall sometimes drop the qualification and simply speak of an “equilibrium.”²

To start our analysis, it is convenient to focus first on the case of $T_i = 0$, in which an agent observes that no one has invested before him. At a first glance, the decision problem in such a situation appears to be fairly complicated. If the agent knew that $T_i = 0$ simply because he is the first decision maker, then he should certainly follow his private signal, since that is the only information available. If, instead, he knew that he is not the first decision maker, then he could decide not to invest independently of the signal, as other agents have already chosen the non-investment option. Intuitively, one might think that $T_i = 0$ is pretty bad information if there are many players. Suppose that n is very large and you observe that nobody has invested before you. But at the same time your own private signal is good. Would you trust your own signal? Of course, the answer to this question would depend on the other agents’ strategies. While the problem is made hard due to the fact that the agent does not know his position in the sequence, it is made easier due to the fact that the only thing that matters about other agents’ strategies is what these specify for the very same case of $T_i = 0$.

To attack the problem, let us start with the following definition:

Definition 2 *An initially-pure equilibrium (IPE) is an equilibrium that prescribes pure actions for $T_i = 0$ and both possible signal realizations $\sigma_i = 0$ and $\sigma_i = 1$.*

Note that there can be mixing in an IPE after observing $T_i > 0$. The definition of an IPE just excludes the cases in which an agent mixes after observing $T_i = 0$. We are able to establish some results that focus on $T_i = 0$. First, we prove that in any IPE agents must follow their signal after observing $T_i = 0$: there cannot exist IPEs in which an agent plays independently of his signal or plays against it.

Lemma 1 *In any IPE, an agent follows his own signal if he observes that nobody has invested so far, i.e., $\mathcal{I}_i(0, \sigma_i) = \sigma_i$ for all i .*

²Our economy is represented by a symmetric game and there is nothing in the environment that could help agents to coordinate on an asymmetric outcome. Therefore, the restriction to symmetric equilibria is very natural.

Proof We prove this by contradiction. Suppose that for $T_i = 0$ agents choose either to invest always or never (independently of their private signals). Consider the latter possibility first, i.e., consider a pure-strategy equilibrium with $\mathcal{I}_i(0, 0) = \mathcal{I}_i(0, 1) = 0$. Then, along the equilibrium path, nobody ever invests and, for any agent $i = 1, \dots, n$, $T_i = 0$. Hence, $T_i = 0$ does not reveal any information on the true state of the world. Since the posterior probability that $\omega = 1$ is still r , agent i is better off by following his informative signal σ_i . Next, consider the case of investment after $T_i = 0$, i.e., an equilibrium with $\mathcal{I}_i(0, 0) = \mathcal{I}_i(0, 1) = 1$. In this case, along the equilibrium path, only the first agent in the sequence observes that nobody else has invested before. That is, $T_i = 0$ if and only if $i = 1$. Hence, after observing $T_i = 0$ agent i knows that he is the first agent in the sequence and, thus, should follow his signal. Finally, suppose that for $T_i = 0$ agents choose to play against their private information, i.e., consider a pure-strategy equilibrium with $\mathcal{I}_i(0, \sigma_i) = 1 - \sigma_i$. Then, along the equilibrium path, after observing $T_i = 0$, agent i knows that he is either the first in the sequence or other agents before him have received good signals. In both cases, he should follow his signal. ■

While we have shown that in any IPE an agent who observes zero investments should follow his signal, it remains unclear whether such equilibria exist. The next lemma identifies a necessary and sufficient condition under which an IPE does indeed exist.

Lemma 2 *An IPE exists if and only if $r \geq \frac{1-q^n}{2-(1-q)^n-q^n}$.*

Proof We first prove that it is indeed optimal for an agent i to follow his own good signal after $T_i = 0$ provided that everybody else follows his signal after $T_i = 0$, and that the condition stated in the lemma holds. (Notice that what another agent j does for $T_j > 0$ is irrelevant for agent i 's optimal choice of $\mathcal{I}_i(0, \sigma_i)$). Assuming such behavior of others, an agent i who observes $T_i = 0$ and $\sigma_i = 1$ attaches to the good state a posterior of

$$\Pr(\omega = 1 \mid T_i = 0, \sigma_i = 1) = \frac{rq \sum_{j=1}^n (1-q)^{j-1}}{rq \sum_{j=1}^n (1-q)^{j-1} + (1-r)(1-q) \sum_{j=1}^n q^{j-1}}.$$

He will follow his good signal if this posterior is at least $1/2$, i.e., if

$$rq \sum_{j=1}^n (1-q)^{j-1} \geq (1-r)(1-q) \sum_{j=1}^n q^{j-1}.$$

Solving for the sums and rearranging the terms, we get the condition in the lemma. To complete the proof we have to show that an agent i who assumes that the others play according to the rules stated in the lemma and who observes $T_i = 0$ and $\sigma_i = 0$ does *not* invest, i.e., we need that

$$\begin{aligned} \Pr(\omega = 1 \mid T_i = 0, \sigma_i = 0) = \\ \frac{r(1-q) \sum_{j=1}^n (1-q)^{j-1}}{r(1-q) \sum_{j=1}^n (1-q)^{j-1} + (1-r)q \sum_{j=1}^n q^{j-1}} < \frac{1}{2}, \end{aligned}$$

or

$$r(1-q) \sum_{j=1}^n (1-q)^{j-1} < (1-r)q \sum_{j=1}^n q^{j-1},$$

which can be written as

$$\frac{r}{(1-r)} < \frac{q^2}{(1-q)^2} \frac{1-q^n}{1-(1-q)^n}. \quad (1)$$

Since $r < q$ we also have $\frac{r}{1-r} < \frac{q}{1-q}$. Hence, inequality (1) holds if $\frac{q}{1-q} \frac{1-q^n}{1-(1-q)^n} > 1$. This can be rewritten as $2q > 1 + q^{n+1} - (1-q)^{n+1}$ which is obviously true for $q > 1/2$. ■

Notice that the condition imposed in the lemma is always fulfilled if $r \geq 1/2$, i.e., when the good state is initially at least as likely as the bad state, an IPE always exists.

We now turn our attention to Perfect Bayesian Nash equilibria that are not initially pure. The next lemma trivially follows from Bayesian updating. We state it formally because we shall need it later on. The lemma after that shows that, in an equilibrium that is not an IPE, agents who observe $T_i = 0$ never invest if their signal is bad, but will invest with some positive probability if their signal is good.

Lemma 3 (i) *In any equilibrium, $\mathcal{I}_i(T_i, 1) \geq \mathcal{I}_i(T_i, 0)$ for all T_i .*

(ii) *In any equilibrium, if $0 < \mathcal{I}_i(T_i, 0) < 1$ then $\mathcal{I}_i(T_i, 1) = 1$, and if $0 < \mathcal{I}_i(T_i, 1) < 1$ then $\mathcal{I}_i(T_i, 0) = 0$ for all T_i .*

Proof In equilibrium, each agent will infer the same information from observing a particular value of T_i . Whatever the posterior induced by just observing T_i , it follows immediately from Bayes' rule that an agent who has an additional good signal will be more optimistic about the good state than an agent with a bad signal. The first part of the proposition results from this consideration and from expected payoff maximization. The second part follows from the same argument and the additional observation that mixing requires indifference. ■

Lemma 4 *In any equilibrium that is not an IPE, $\mathcal{I}_i(0,0) = 0$ and $0 < \mathcal{I}_i(0,1) < 1$ for all i .*

Proof Given Lemma 3 we just need to rule out an equilibrium with $0 < \mathcal{I}_i(0,0) < 1$ and $\mathcal{I}_i(0,1) = 1$. For an agent to be indifferent between investing and not after observing $T_i = 0$ and $\sigma_i = 0$ we need $\Pr(\omega = 1 | T_i = 0, \sigma_i = 0) = 1/2$. Using Bayes' rule, this can be re-written as

$$r \Pr(T_i = 0, \sigma_i = 0 | \omega = 1) = (1 - r) \Pr(T_i = 0, \sigma_i = 0 | \omega = 0),$$

or

$$r \sum_{j=1}^n (1 - q)^j (1 - p)^{j-1} = (1 - r) \sum_{j=1}^n q^j (1 - p)^{j-1},$$

where p denotes the probability with which all other agents who see $T_i = 0$ and $\sigma_i = 0$ invest. Rewriting this as

$$\sum_{j=1}^n [(r(1 - q)^j - (1 - r)q^j) (1 - p)^{j-1}] = 0$$

makes it obvious that there is no $p > 0$ that solves the equation: since $q > \max\{\frac{1}{2}, r\}$ the left-hand side is strictly negative for any positive p . ■

Having characterized equilibria that are not initially pure, we must discuss whether they exist. The next lemma introduces a necessary and sufficient condition for such mixed-strategy equilibria to exist.

Lemma 5 *(i) Mixed-strategy equilibria with $\mathcal{I}_i(0,0) = 0$ and $0 < \mathcal{I}_i(0,1) < 1$ for all i exist if and only if there is a $p \in (0,1)$ that solves*

$$r[1 - (1 - pq)^n] = (1 - r)[1 - (1 - p(1 - q))^n].$$

(ii) If such a p exists, it is unique.

(iii) A mixed strategy equilibrium does not exist for $r \geq \frac{1}{2}$.

Proof The first part of the lemma follows from observing that agent i 's indifference between investing and not investing after observing $T_i = 0$ and $\sigma_i = 1$ requires $\Pr(\omega = 1 \mid T_i = 0, \sigma_i = 1) = 1/2$. If all other agents $j \neq i$ use $I_j(0, 0) = 0$ and $I_j(0, 1) = p$, after applying Bayes' rule and some algebraic manipulation, this equality becomes

$$rq \sum_{j=1}^n (1 - pq)^{j-1} = (1 - r)(1 - q) \sum_{j=1}^n (1 - p(1 - q))^{j-1},$$

which is equivalent to the equation in the lemma.

Both the RHS and the LHS of the equation in the lemma are equal to 0 for $p = 0$, and are strictly increasing and strictly concave in p . Furthermore, at $p = 0$, the derivative of the LHS with respect to p is strictly greater than the derivative of the RHS, since $r + q > 1$. This implies that the two functions cross at most once for $p > 0$ which establishes the second part of the lemma. Finally, note that for $r \geq \frac{1}{2}$ the LHS is strictly greater than the RHS for any value of p , which proves the last part of the lemma. ■

This lemma completes our characterization of equilibrium decisions after observing $T_i = 0$. In the following proposition we summarize what we have learned so far.

Proposition 1 (i) *If $r \geq 1/2$ agents who observe $T_i = 0$ follow their signal in all equilibria.*

(ii) *If $\frac{1-q^n}{2-(1-q)^n-q^n} < r < 1/2$ there is an equilibrium where agents who observe $T_i = 0$ follow their signal but there may also be other (mixed-strategy) equilibria where agents who observe $T_i = 0$ follow their signal if it is bad and mix if it is good.*

(iii) *If $r \leq \frac{1-q^n}{2-(1-q)^n-q^n}$ there can only be equilibria where agents who observe $T_i = 0$ follow their signal if it is bad and mix if it is good.³*

Proof The proposition follows immediately from the four previous lemmas and the observation that $\frac{1-q^n}{2-(1-q)^n-q^n} < 1/2$. ■

³Notice that the third part of the proposition touches on an existence problem. For obvious reasons we have restricted our analysis to symmetric equilibria—in case of bad priors these may fail to exist.

Our analysis essentially shows that, when facing a situation with no previous investments, an agent should either follow his signal or use a mixed strategy (only if the signal is good). An agent should never decide independently of his signal, neither should he decide against it. This clearly indicates that we should not observe a “down cascade” where all agents choose not to invest. In other words, to go back to one of our examples, a restaurant will not stay empty forever only because it is empty when it opens. While this puts already a lot of structure on the equilibrium solution of our game, we still need to investigate what happens for different values of the aggregate investment T_i .

To this purpose, we establish in the next step an intuitive monotonicity result, according to which a higher value of T_i is always good news: when an agent observes a higher number of investments made before him, he cannot be less willing to invest himself. Once this monotonicity lemma is established, we will be able to prove two fundamental results about aggregate cascades.

Lemma 6 *In any equilibrium, if $T_i^l < T_i^h$ then $\mathcal{I}_i(T_i^l, \sigma_i) \leq \mathcal{I}_i(T_i^h, \sigma_i)$ for both $\sigma_i = 0$ and $\sigma_i = 1$. In particular, if $0 < \mathcal{I}_i(T_i^l, \sigma_i) < 1$, then $\mathcal{I}_i(T_i^h, \sigma_i) = 1$.*

Proof *See Appendix. ■*

While this lemma seems very intuitive (how could a fuller restaurant be worse news than an emptier?) it is actually not trivial to prove it. At the very core of the proof there is, however, some very basic logic operating. Essentially, it is the earlier monotonicity result (in Lemma 3) which is driving this one. Agents with good signals are more likely to invest than agents with bad signals. Good signals are more likely to be generated in the good state than in the bad state. Hence, T_i grows, on average, “faster” in the good state than in the bad state. Hence, the higher T_i the more confident can we be about being in the good state.

Equipped with Lemma 6 we are now ready to state our two main propositions that characterize which forms of cascades will or will not arise. In particular, we will see that aggregate down cascades *never* arise, while aggregate up cascades are *always* part of an equilibrium.

Proposition 2 *(i) In any equilibrium, $\mathcal{I}_i(0, 1) > 0$, and $\mathcal{I}_i(T_i, 1) = 1$ for all $T_i > 0$, i.e., an agent with a good signal always invests with positive probability*

(and invests with probability one after observing at least another investment) and an ADC never occurs in equilibrium.

(ii) In any equilibrium, there can be at most one T^{MIX} for which $0 < \mathcal{I}_i(T^{MIX}, 0) < 1$. For all $T_i < T^{MIX}$ agents with bad signals follow their signal and do not invest. For all $T_i > T^{MIX}$, an AUC occurs in which agents invest independently of their signal.

Proof The first part of the proposition follows from Proposition 1 and Lemma 6. The second part follows again from Lemma 6. ■

The first part of the proposition clearly implies that there are no cascades on the unobservable action. In particular, after observing at least one investment, agents with a good signal always invest. Incidentally, we note that such a result just comes from an equilibrium argument. One could imagine that, when facing a “low” value of T_i , in order to make his decision, agent i should consider all possible sequences and attach a probability to the event that he is the first in the sequence, or the second, etc. After all, a low number of investments may merely come from the fact that only few agents had the opportunity to invest so far, in which case the low value of T_i should be considered good news. Or it could arise from many agents having the option of investing but few only using it, in which case the low T_i should be viewed as bad news. All this inference process could be quite complicated. Our analysis solves the problems by just invoking some equilibrium arguments.

The second part of the proposition hints at the possible role of aggregate up cascades. But from all we have established so far it could be that $T^{MIX} \geq n$, i.e., that agents always follow their signal (or mix) such that an AUC never arises. The next proposition, however, shows that AUCs do arise—and are, in fact, part of *any* equilibrium.

Proposition 3 *AUCs are part of any equilibrium. In particular, in any equilibrium $\mathcal{I}_i(T_i, \sigma_i) = 1$ for all $T_i > \frac{n}{2}$.*

Proof Consider an agent i who observes $T_i > \frac{n}{2}$ and suppose he *knew* that he were the last agent in the sequence. Further suppose there were no AUC. If $T_i = T^{MIX}$, then for $T_i = T^{MIX} + 1$, an AUC would occur by Proposition 2. If, instead, $T_i \neq T^{MIX}$, then, due to Lemma 6, this

	$\sigma_i = 0$	$\sigma_i = 1$
$T_i = 0$	0	<i>mixed</i> , 1
$0 < T_i < T^{MIX}$	0	1
$T_i = T^{MIX}$	<i>mixed</i>	1
$T_i > T^{MIX}$	1	1
$T_i \geq T^{UP}$	1	1

Table 1: Structure of all equilibria. Entries indicate whether the agent mixes or invests with probability 0 or 1.

agent knows that there were at least T_i good signals and no more than $n - T_i - 1$ bad signals. Hence, even if this agent's own signal is bad, he knows that there were altogether more good signals than bad signals and he will decide to invest. Of course, agent i can't be sure that he really is the last agent. But if he isn't, this means that there were *fewer* bad signals so far, while he can still be sure that there were T_i good signals. Hence, an agent who observes $T_i > n/2$ will always invest and, thus, trigger an AUC. ■

The value $\frac{n}{2}$ is just an upper bound for the critical mass of observable choices that triggers an AUC. Depending on the parameters' values, AUCs may well be triggered earlier. But AUCs are indeed part of *all* equilibria. Of course, this does not necessarily imply that AUCs will actually be triggered, since there is always the possibility of sufficiently many bad signals occurring such that the critical T_i that triggers an AUC may not be reached.

We summarize the structure of equilibria in Table 1. The rows in the table indicate possible values of T_i , while the columns indicate the two possible signal realizations. Note that not necessarily all the values of T_i in the table exist. In particular, T^{MIX} might not exist. For this reason, the last two rows of the table have the same entries. Notice also that if T^{MIX} exists, then $T^{UP} = T^{MIX} + 1$. In any case, the basic structure of all equilibria is captured in the table and is nicely monotonic.

4 An example

It is now the moment to illustrate our theory through a simple example. The example will highlight some properties of the equilibrium analysis and will be

the basis for the experiment that follows.

Consider the case in which $n = 3$ and $r \geq 1/2$. From Proposition 1 we know that $I_i(0, \sigma_i) = \sigma_i$ and from Proposition 2 we know that $I_i(2, \sigma_i) = 1$ and that $I_i(1, 1) = 1$. But what should agents do after observing $T_i = 1$ and $\sigma_i = 0$. As is clear from the results illustrated in Table 1, this depends on further conditions on r and q . Let us first check under which conditions agent i rationally follows his bad signal. Recall that we are analyzing a symmetric equilibrium, therefore suppose each other agent j chooses $I_j(1, 0) = 0$. Then it is optimal for agent i to do the same if his posterior for the good state is not bigger than $1/2$, i.e., if

$$\frac{r[q(1-q) + 2q(1-q)^2]}{r[q(1-q) + 2q(1-q)^2] + (1-r)[q(1-q) + 2q^2(1-q)]} \leq \frac{1}{2}$$

which is equivalent to

$$q \geq 2r - \frac{1}{2} \equiv \underline{q}.$$

Similarly, $I_i(1, 0) = 1$ is optimal if

$$r[q(1-q) + q(1-q)^2] \geq (1-r)[q(1-q) + q^2(1-q)].$$

which is equivalent to

$$q \leq 3r - 1 \equiv \bar{q}.$$

Note that $\underline{q} \leq \bar{q}$. Hence, in the case in which $\underline{q} > 0$ and $\bar{q} < 1$, we obtain three equilibrium regions. For $q < \underline{q}$ there is a unique pure-strategy equilibrium in which $I_i(1, 0) = 1$ and an AUC starts with $T_i = 1$. For $q > \bar{q}$ there is a unique pure-strategy equilibrium in which $I_i(1, 0) = 0$ and an AUC starts only with $T_i = 2$. Finally, for $\underline{q} \leq q \leq \bar{q}$ both the two pure-strategy equilibria exist and there is a mixed-strategy equilibrium as well—with $I_i(1, 0) = \frac{1+2q-4r}{q-r}$.

5 Experimental Evidence

5.1 Design and Procedures

In order to check whether our theory can reasonably be expected to be of any empirical relevance we conducted a laboratory experiment. We implemented the three-agent example that we analyzed above for two sets of parameters. In Treatment A, the prior for the state being good is $r = 1/2$, and the signal precision is $q = 0.7$. In Treatment B, the prior is $r = 3/4$, and the signal

	Treatment A		Treatment B	
	$\sigma_i = 0$	$\sigma_i = 1$	$\sigma_i = 0$	$\sigma_i = 1$
$T_i = 0$	0	1	0	1
$T_i = 1$	0	1	1	1
$T_i = 2$	1	1	1	1

Table 2: Equilibrium investments in both treatments.

precision $q = 0.8$. As one can verify from the previous section, for both treatments theory predicts unique equilibria, shown in Table 2. The only difference between the two treatments concerns the critical value of T_i that triggers an AUC. In Treatment A an AUC is only triggered for $T_i = 2$, while in Treatment B an AUC starts already with $T_i = 1$.

The experiment was fully computerized and run at the ELSE laboratory at UCL in the Spring 2005. We recruited subjects from the College’s undergraduate population across all disciplines. They had no previous experience with this experiment. At the beginning of the sessions, we gave written instructions (reported in Appendix B) to all subjects which were also read aloud. Then subjects had the opportunity to ask clarifying questions, which were answered privately.

Several subjects participated in each session of the experiment. They were matched in multiple groups of three. Altogether, we observed 57 subjects who were randomly assigned to the two treatments, with 11 groups in Treatment A, and 8 groups in Treatment B.

The experiment consisted of 15 rounds of game playing and groups remained constant over time. We employed the strategy method, i.e., subjects had to indicate whether or not they would invest for each possible combination of T_i and σ_i . They submitted their strategies by clicking on radio buttons to indicate their yes/no decisions in a table with a structure similar to one of the panels in Table 2. Once everybody had submitted a strategy, the computer drew the state of the world, the sequence in which subjects played and each subject’s private signal. Using subjects’ strategies, the computer then determined the outcome and each subjects received feedbacks on their values σ_i and T_i (i.e., on which of the six choices was relevant for their payoff) as well as on the state ω and the resulting payoff. They did not receive feedback

about other subjects' signals or outcomes, thus avoiding any possibilities for imitative learning from round to round (Apesteguia *et al.* 2007).

A subject would earn money by investing in the good state or not investing in the bad state. More precisely, we gave £3 for a correct decision and nothing for a wrong decision. At the end of the experiment, after the 15 rounds, the computer randomly chose three rounds, one from each third of the experiment. Each subject's payment depended on how he performed in the selected rounds. The final payment was equal to the sum of the payoffs in these three rounds plus an amount of £4 given as show up fee. Subjects were paid in private, immediately after the experiment.

5.2 Results

Table 3 summarizes the results using an identical format to Table 2. For each combination of σ_i and T_i the table shows the mean investment rate, with standard errors in parenthesis. For most cells, average investments are actually quite close to equilibrium predictions. Also, standard errors are rather small, indicating that behavior was rather homogenous across groups.

Let us now examine some specific cells of the table. We start by considering the case of $T_i = 0$ and $\sigma_i = 1$. For this case our theory predicts that, in both treatments, subjects should always invest. In Treatment *B*, the data comes very close to the prediction, with an investment rate of 86%. In Treatment *A*, the result is less strong: we observe a majority of investments, but we are far from the full investment equilibrium outcome. Nevertheless, it should be noted that with 56% investments, ADC rarely occur in this treatment too. The probability of no investments despite three good signals in this treatment is just 17%. Although this result is less impressive than that for Treatment *B*, overall the experiment seems to indicate that aggregate down cascades not only do not arise in equilibrium, they also do not (or rarely) occur in the laboratory.

Let us now move on to the case of $T_i = 1$ and $\sigma_i = 0$. In such a case, our model predicts two different types of behavior: in Treatment *A* subjects should abstain from investing, while full investment should occur in Treatment *B*. The theory captures the comparative statics well, in that we do observe a substantial and significant difference between the two treatments. The doubling of the investment rate is significant at the 5% level (two-sided MWU

	Treatment A		Treatment B	
	$\sigma_i = 0$	$\sigma_i = 1$	$\sigma_i = 0$	$\sigma_i = 1$
$T_i = 0$	0.15 (0.03)	0.56 (0.03)	0.21 (0.09)	0.86 (0.05)
$T_i = 1$	0.24 (0.04)	0.87 (0.03)	0.41 (0.06)	0.94 (0.03)
$T_i = 2$	0.68 (0.07)	0.85 (0.03)	0.76 (0.09)	0.94 (0.04)

Table 3: Average investment rates in both treatments. Standard errors in parentheses.

test using group averages). It should be noticed, however, that while theory predicts an AUC in Treatment *B*, the actual investment rate is only slightly above 40%. While this result is somehow disappointing for the theory, the subjects' reluctance to invest is perhaps better understood if we look at the case of $T_i = 2$.

According to our model, an AUC should have arisen in both treatments after two investments and subjects should have always invested, independently of their signals. It is worth pointing out that $T_i = 2$ is, in some sense, an easier case for the subjects than $T_i = 1$. When $T_i = 2$ subjects know for sure in which position they are in the sequence and they also know the entire history, namely that everybody before them invested. Therefore, it should be easier for them to infer that, in equilibrium, both other subjects received a good signal (in Treatment *A*) or at least one other subject did (in Treatment *B*). We observe that the majority of choices are in line with the theoretical prediction. Subjects invested despite their bad signal in about 70% of the cases in both treatments, with a slightly higher percentage in Treatment *B*. Nevertheless, also for the case of $T_i = 2$ subjects are more reluctant to invest than they should be in theory. This result is not a peculiarity of our experiment, but is reminiscent of previous findings in the experimental social learning literature. The data for $T_i = 2$ are indeed comparable to the results obtained in previous experiments on the Bikhchandani et al. (1992) sequential learning model. Since subjects know for sure in which position they are in the sequence and they also know the entire history of investment decisions, they are almost in

the position of a subject who plays third in the sequential model and observes two predecessors investing, or plays later in the sequence and observe two more investments than passes. One could arguably believe that the decision process is actually simpler in the sequential learning model, since in that model the equilibrium arguments are somehow more straightforward. Despite this, even in the experiments based on the sequential learning model (see, for example, the nice overview in Kübler and Weizsacker, 2005, or Weizsacker’s (2006) meta analysis), there is a lower investment rate than the equilibrium prediction. For instance, Dominitz and Hung (2004) observe that after two identical decisions the frequency with which the third player in the sequence conformed to the same action (despite an opposite signal) is below 60%. Similarly, Kübler and Weizsacker (2004) report an investment rate of 65%. On balance, we do find evidence for the occurrence of AUCs in at least the same extent as there is evidence for information cascades in the standard sequential setup.⁴

It is interesting to analyze the deviations from equilibrium play in a little more detail. For that purpose we compute the beliefs that result from actual behavior. These are the beliefs on the state of the world being good or bad that a Bayesian agent would form (for each possible observed investment level and private signal realization), knowing the frequency with which subjects actually invested in any contingency of the experiment. We compute these beliefs separately for each group that has taken part in the experiment. Hence, for each group we have a set of six empirical beliefs for the six possible combinations of σ_i and T_i . In Table 4 we show the averages of these empirical beliefs over groups. For instance, after observing $\sigma_i = 1$ and $T_i = 2$ a Bayesian agent should conclude that the probability of the state being good in Treatment B is 98%. It is noteworthy that, given these empirical beliefs, a subjects should always choose the equilibrium actions. In other words, the incentive structure that subjects endogenously create in the lab is identical with the equilibrium incentive structure. Whenever it pays to invest in equilibrium it

⁴As Kübler and Weizsacker (2005) point out, an empirical regularity in cascade experiments is that the frequency with which subjects follow the crowd is increasing in the crowd’s size. For this reason, in comparing our results with those of previous experiments we have only considered the case of the third player in the sequence. Even not restricting to this specific case and considering the average behavior, however, the conclusion would not change. For instance, in their seminal paper, Anderson and Holt (1997) report an average occurrence of cascades of 73%.

	Treatment A		Treatment B	
	$\sigma_i = 0$	$\sigma_i = 1$	$\sigma_i = 0$	$\sigma_i = 1$
$T_i = 0$	0.26	0.66	0.32	0.88
$T_i = 1$	0.37	0.76	0.52	0.95
$T_i = 2$	0.52	0.86	0.72	0.98

Table 4: Beliefs about good state given actual behavior.

also pays to invest in the experiment, given other subjects' behavior.

We also compute the (marginal) incentives to behave optimally given the empirical beliefs (see Table 5).⁵ The main purpose of computing these incentives is to see how they correlate with subjects' actual behavior. For this exercise, we compute such beliefs and incentives also separately, for each group that has taken part in the experiment. Hence, for each group we have a set of six empirical beliefs and incentives for the six possible combinations of σ_i and T_i . In Figure 1 we plot observed error rates (i.e., the fraction of non-optimal choices given the empirical beliefs) as a function of the imputed incentives. Each dot in that scatter plot represents an observation from one group for a (σ_i, T_i) -combination. In all, the scatter plot has, thus, 114 dots. There are two main conclusions to be drawn from this purely non-parametric approach to the study of errors in decision making. First, errors occur mainly when incentives are small. The larger the stakes the less frequent are mistakes. In particular, the fraction of errors is higher than 40% almost uniquely for incentives below 0.4 (i.e., when the belief on either state being the true one is lower than 0.7). Second, errors increase more than linearly when incentives fall.

Overall, although our theory is unable to capture the behavior we observe in the laboratory perfectly, we consider the experimental results very encouraging. The theory organizes the data remarkably well. All the comparative statics go in the right direction. And crucially, our key results on aggregate cascades have strong predictive power in the lab. While aggregate down cascades do not occur there is a strong tendency for aggregate up cascades to

⁵We define the marginal incentive to take the optimal action as the absolute level of the difference between the beliefs that the state is good or bad. In other words, if the belief b that the state is good is higher than $1/2$, the marginal incentive to invest is defined as $2b - 1$; if the belief is lower than $1/2$, then the incentive not to invest is $1 - 2b$.

	Treatment A		Treatment B	
	$\sigma_i = 0$	$\sigma_i = 1$	$\sigma_i = 0$	$\sigma_i = 1$
$T_i = 0$	0.48	0.32	0.36	0.76
$T_i = 1$	0.26	0.52	0.04	0.90
$T_i = 2$	0.04	0.72	0.44	0.96

Table 5: Marginal incentives for playing optimally given beliefs.

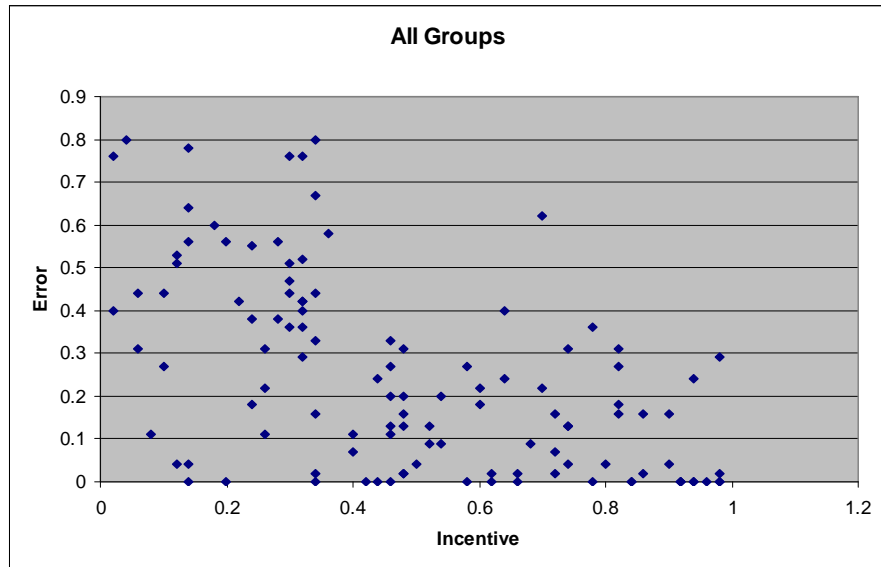


Figure 1: Observed error rates as a function of imputed incentives.

form. While not offering a full-fledged empirical analysis of our model, in our view, this experimental evidence lends it credibility and should, in fact, encourage further investigations.

6 Discussion

We have introduced a new model of information cascades. The crucial difference between our model and those already in the literature is that only one action taken by agents is observable to others. When it is their turn to make a binary decision, agents simply receive aggregate information about how many others before them took the observable action. We argue that this setup is in many cases realistic: for example, when entrepreneurs seek investors they will typically not inform them about how many others have turned them down before, but, surely, they will mention who else decided previously to invest in their project. This asymmetry in observability, which in many cases arises naturally, dramatically affects all equilibria in such games. Most importantly, there can be no “down cascades:” if an action is unobservable, there can never be an information cascade where agents take this action.

We have provided a first experimental test of our model. The theory organizes the data remarkably well. Deviations from optimal play mainly occur where deviations are not very costly. As a consequence, all major comparative statics are as predicted. And while we do observe up cascades, down cascades occur very rarely.

Our result has important implications. In particular, it implies that a new, good project (e.g., a technological innovation, a new product or service, a new medical treatment) will not be neglected for ever simply because there is lack of interest at the beginning. Sooner or later (i.e., as soon as people start receiving good information on it) the new project will start diffusing. Or, at least, a lack of initial interest will not represent a barrier to future adoption because of informational considerations. Our study has also an important consequence for applications where a third party (e.g., an health agency) can decide which information it is to release. If, because of externalities or other considerations, the agency wants to avoid a particular type of cascade, say on adopting a new medical treatment, then it should *withhold* information, and not publish the number of doctors who have decided to adopt the new treatment.

References

- [1] Anderson, L., and C. Holt (1997) Information Cascades in the Laboratory. *American Economic Review* 87 (5), 847-862.
- [2] Apestegua, J., S. Huck, and J. Oechssler (2007) Imitation: Theory and Experimental Evidence. *Journal of Economic Theory* forthcoming.
- [3] Banerjee, A.V. (1992) A Simple Model of Herd Behavior. *Quarterly Journal of Economics* 107, 797-817.
- [4] Banerjee, A.V., and D. Fudenberg (2004) Word-of-Mouth Learning. *Games and Economic Behavior* 46, 1-23.
- [5] Bikhchandani, S., D. Hirshleifer, and I. Welch (1992) A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* 110, 992-1026.
- [6] Çelen, B., and S. Kariv (2004) Observational Learning under Imperfect Information. *Games and Economic Behavior* 47, 72-86.
- [7] Dominitz, J., and A.A. Hung (2004) Homogeneous Actions and Heterogeneous Beliefs: Experimental Evidence on the Formation of Information Cascades. *mimeo*.
- [8] Kübler, D., and G. Weizsacker (2004) Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory. *Review of Economic Studies* 71, 425-441.
- [9] Kübler, D., and G. Weizsacker (2005) Are Longer Cascades More Stable? *Journal of the European Economic Association* 3, 330-339.
- [10] Larson, N. (2006) In with the New Or Out with the Old? Bottlenecks in Social Learning. *mimeo*
- [11] Smith, L., and P. Sørensen (1998) Rational Social Learning with Random Sampling. *mimeo*.
- [12] Sufi, A. (2007) Information Asymmetry and Financing Arrangements: Evidence from Syndicated Loans. *Journal of Finance* 62, 629-668.
- [13] Weizsacker, G. (2006) Do We Follow Others When We Should? A Model-Free Meta Analysis of Social Learning Experiments. *mimeo*.

7 Appendix

A Proof of Lemma 6

The proposition is equivalent to saying that in any equilibrium $\mathcal{I}(t_i, \sigma_i) \leq \mathcal{I}(t_i + 1, \sigma_i)$ for any $t_i = 0, 1, 2, \dots$ and both $\sigma_i = 0$ and $\sigma_i = 1$. Because of expected payoff maximization, this inequality holds if, whenever $\Pr(\omega = 1 | T_i = t, \sigma_i) \geq \frac{1}{2}$, we have $\Pr(\omega = 1 | T_i = t + 1, \sigma_i) > \frac{1}{2}$.

There are four relevant possibilities:

1. $\Pr(\omega = 1 | T_i = t, \sigma_i = 0) > \frac{1}{2}$ and $\Pr(\omega = 1 | T_i = t, \sigma_i = 1) > \frac{1}{2}$
2. $\Pr(\omega = 1 | T_i = t, \sigma_i = 0) < \frac{1}{2}$ and $\Pr(\omega = 1 | T_i = t, \sigma_i = 1) > \frac{1}{2}$
3. $\Pr(\omega = 1 | T_i = t, \sigma_i = 0) = \frac{1}{2}$ and $\Pr(\omega = 1 | T_i = t, \sigma_i = 1) > \frac{1}{2}$
4. $\Pr(\omega = 1 | T_i = t, \sigma_i = 0) < \frac{1}{2}$ and $\Pr(\omega = 1 | T_i = t, \sigma_i = 1) = \frac{1}{2}$

Case 1 is the case of an informational cascade. In such a case,

$$\Pr(\omega = 1 | T_i = t, \sigma_i) = \Pr(\omega = 1 | T_i = t + 1, \sigma_i)$$

for both σ_i , and therefore the proposition obviously holds.

Now let us consider Case 2. In this case we want to show that $\Pr(\omega = 1 | T_i = t + 1, \sigma_i = 1) > \frac{1}{2}$ (while nothing must be shown for the case of a bad signal). Suppose not, i.e., suppose $\Pr(\omega = 1 | T_i = t + 1, \sigma_i = 1) \leq \frac{1}{2}$. Let us consider, first, the case of the strict inequality.

By Bayes' rule,

$$\begin{aligned} & \Pr(\omega = 1 | T_i = t + 1, \sigma_i = 1) = \\ &= \frac{\Pr(T_i = t + 1 | \omega = 1, \sigma_i = 1) \Pr(\omega = 1 | \sigma_i = 1)}{\Pr(T_i = t + 1 | \omega = 1, \sigma_i = 1) \Pr(\omega = 1 | \sigma_i = 1) + \Pr(T_i = t + 1 | \omega = 0, \sigma_i = 1) \Pr(\omega = 0 | \sigma_i = 1)}. \end{aligned}$$

As we suppose that this is strictly smaller than $\frac{1}{2}$ we know that

$$\frac{\Pr(T_i = t + 1 | \omega = 1, \sigma_i = 1)}{\Pr(T_i = t + 1 | \omega = 0, \sigma_i = 1)} < \frac{\Pr(\omega = 0 | \sigma_i = 1)}{\Pr(\omega = 1 | \sigma_i = 1)},$$

which is equivalent to

$$\frac{\Pr(T_i = t + 1 | \omega = 1)}{\Pr(T_i = t + 1 | \omega = 0)} < \frac{\Pr(\omega = 0 | \sigma_i = 1)}{\Pr(\omega = 1 | \sigma_i = 1)}.$$

By the law of total probabilities,

$$\Pr(T_i = t + 1 | \omega = 1) \tag{2}$$

$$= \Pr(T_i = t + 1 | \omega = 1, T_{i-1} = t) \Pr(T_{i-1} = t | \omega = 1) \tag{3}$$

$$+ \Pr(T_i = t + 1 | \omega = 1, T_{i-1} = t + 1) \Pr(T_{i-1} = t + 1 | \omega = 1) \tag{4}$$

$$= q \Pr(T_{i-1} = t | \omega = 1) + \Pr(T_{i-1} = t + 1 | \omega = 1). \tag{5}$$

Notice that the last equality comes from the fact that we are analyzing Case 3 and that we are assuming (by contradiction) no investment after observing $t + 1$.

Now the decision problem of agent $i - 1$ is identical to the one of agent i . So, by applying recursively the same law, we obtain:

$$\begin{aligned}
& \Pr(T_i = t + 1 \mid \omega = 1) \\
&= q \Pr(T_{i-1} = t \mid \omega = 1) + \Pr(T_{i-1} = t + 1 \mid \omega = 1, \sigma_i = 1) \\
&= q \Pr(T_{i-1} = t \mid \omega = 1) + [q \Pr(T_{i-2} = t \mid \omega = 1) + \Pr(T_{i-2} = t + 1 \mid \omega = 1)] \\
&\quad + q \Pr(T_{i-1} = t \mid \omega = 1) + q \Pr(T_{i-2} = t \mid \omega = 1) + [q \Pr(T_{i-3} = t \mid \omega = 1) \\
&\quad\quad + \Pr(T_{i-3} = t + 1 \mid \omega = 1)] + \dots \\
&= q \Pr(T_{i-1} = t \mid \omega = 1) + q \Pr(T_{i-2} = t \mid \omega = 1) + q \Pr(T_{i-3} = t \mid \omega = 1) \\
&\quad + \dots + q \Pr(T_{i-m} = t \mid \omega = 1)
\end{aligned}$$

for some m (note that m depends on the value of i : indeed, for any value of i there is an m such that $\Pr(T_{i-m} = t + 1 \mid \omega = 1) = 0$). Similarly, conditioning on $\omega = 0$,

$$\begin{aligned}
& \Pr(T_i = t + 1 \mid \omega = 0) \\
&= (1 - q) \Pr(T_{i-1} = t \mid \omega = 0) + (1 - q) \Pr(T_{i-2} = t \mid \omega = 0) \\
&\quad + \dots + (1 - q) \Pr(T_{i-m} = t \mid \omega = 0).
\end{aligned}$$

Some algebraic computations show that for any pair of terms in the two expressions above, the following inequality holds:

$$\frac{\Pr(T_{i-j} = t \mid \omega = 1)}{\Pr(T_{i-j} = t \mid \omega = 0)} \geq \frac{\Pr(T_i = t \mid \omega = 1)}{\Pr(T_i = t \mid \omega = 0)}.$$

Since we know that $\Pr(\omega = 1 \mid T_i = t, \sigma_i = 1) > \frac{1}{2}$ and, therefore,

$$\frac{\Pr(T_i = t \mid \omega = 1)}{\Pr(T_i = t \mid \omega = 0)} > \frac{\Pr(\omega = 0 \mid \sigma_i = 1)}{\Pr(\omega = 1 \mid \sigma_i = 1)}$$

simple algebra shows that

$$\frac{\Pr(T_i = t + 1 \mid \omega = 1)}{\Pr(T_i = t + 1 \mid \omega = 0)} > \frac{\Pr(\omega = 0 \mid \sigma_i = 1)}{\Pr(\omega = 1 \mid \sigma_i = 1)},$$

a contradiction.

Note that the same proof holds true when, by contradiction, we assume that

$$\Pr(\omega = 1 \mid T_i = t + 1, \sigma_i = 1) = \frac{1}{2}.$$

The only difference is that in such a case

$$\begin{aligned} & \Pr(T_i = t + 1 \mid \omega = 1) \\ &= q \Pr(T_{i-1} = t \mid \omega = 1) + s \Pr(T_{i-1} = t + 1 \mid \omega = 1), \end{aligned}$$

where s represents the probability by which an agent receiving the good signal decided not to invest. This change does not affect the above inequalities.

Finally, note that the proofs for Case 3 (for both the good and the bad signal) and Case 4 are identical to Case 2 just described, with the exception that in Case 3,

$$\Pr(T_i = t + 1 \mid \omega = 1, T_{i-1} = t) = q + (1 - q)u,$$

and

$$\Pr(T_i = t + 1 \mid \omega = 0, T_{i-1} = t) = qu + (1 - q),$$

where u is the probability of investment by an agent receiving a bad signal; similarly, in Case 4,

$$\Pr(T_i = t + 1 \mid \omega = 1, T_{i-1} = t) = qu$$

and

$$\Pr(T_i = t + 1 \mid \omega = 0, T_{i-1} = t) = (1 - q)u.$$

■

B Instructions

Welcome to our experiment!

Please be quiet during the entire experiment. Do not talk to your neighbours and do not try to look at their screens. Simply concentrate on what you have to do. If you have a question, please raise your hand. We will come to you and answer it privately.

You are participating in an economics experiment in which you interact with two other participants for 15 rounds. There are more participants in this room, but you will interact with only two of them.

Depending on your choices, the other two participants' choices and some luck you can earn a considerable amount of money. You will receive the money immediately after the experiment. Notice that all participants have the same instructions.

The experiment

What you have to do

You will have to decide whether you want to invest in a project or not. The project may be good or not and we will give you some useful information about how the chances are. Additionally, you will also know something about what the other two participants decided to do.

What determines whether the investment is good or not

The computer will decide randomly whether in a given round the investment is good or not. The two possibilities are equally likely. This is equivalent to say that the computer will choose whether the investment is good or not by tossing of a coin.

Note that if in a given round the investment is good, it is good for all three participants. Similarly, if it is bad, it is bad for all three of you.

What you earn if you decide to invest or not

In real life, if you choose a good investment the prize is that you enjoy a good return. And if it is bad, the cost is that you have spent money on something that was not profitable. In our experiment, if the project is good and you choose to invest we give you £3 for the smart decision. If, instead, you decide not to invest, then you get nothing. Similarly, if the project is bad and you decide not to invest, we pay you £3 for the smart decision. And if you decide to invest, we pay you nothing.

How you can make your decision

As we said the computer will choose one of the two projects. Of course, we will not tell you which one has been chosen. But we will give you a piece of information.

If the project is good, then the computer will draw a ball from an urn containing 70 green and 30 red balls. If it is bad, it will draw a ball from an urn containing 70 red and 30 green balls. You (and only you) will be told the colour of this ball. Clearly if the ball is green, it is more likely (but not sure)

that the project is good. If it red, it is more likely (but not sure) that it is bad.

Note that the computer will draw a different ball for each participant. It will choose a ball for you and then replace it in the urn. Then a ball for another participant and then replace it. And so on. Therefore, it is well possible that you receive a green ball and another participant a red one, and vice versa.

This is not the only information that you will get. In the next paragraph you will discover why.

When you make your decision

You and the other two participants will make the decision to invest or not in the project in sequence. Therefore, you may be the first, or the second, or the third. Your position in the sequence is assigned to you randomly by the computer. Whether you will be first, or second, or third is equally likely. However, we will not tell you your position. But we will tell you the number of people who have decided to invest before you.

Let us briefly look at the different possibilities that can arise. You might see that two others have decided to invest in which case you know that you must be the last to make a decision. If you see that just one other participant has invested before you the situation is less clear. Obviously, you are not the first in the sequence. You might be the second and the first might have decided to invest. But you might also be the last with one of your predecessors having decided to invest and the other having decided to pass the opportunity. Finally, you might observe that none of the others has decided to invest so far. In that case you might be the first in the sequence. But you might also be second and the first passed the opportunity, or you might be the third and both others decided not to invest.

The procedure

At the beginning of each round the computer will decide whether the project is good or bad. Moreover, it will draw a ball for you and one for each other participant. And it will decide the sequence in which the three of you decide.

Then it comes to your decision. But notice instead of telling you your ball colour and the number of people who have decided to invest before you, we will do something different. We will ask you to make your decision for each

possible case. We will ask you to make a decision to invest or not depending on the number of people who have already decided to invest and on the colour of your ball. Specifically, you will see a table like this:

	Green Ball	Red Ball
Nobody has invested before you		
One other has invested before you		
Two others have invested before you		

For each possible combination you will have to decide whether you invest or not. Of course, when we compute your payoff, we will take into account only the decision corresponding to the actual situation (that will be revealed to you afterwards). The other five decisions will not be taken into account. Therefore, you can decide for each case as you if knew the ball colour and the actual number of people who decided to enter before you.

Is all this clear? If not, do not worry, here are two examples.

Example 1

Suppose you make the following decisions:

	Green Ball	Red Ball
Nobody has bought before you	NO	NO
One other has bought before you	NO	INVEST
Two others have bought before you	NO	INVEST

At the end of the round the computer tells you that actually in this round you received the red ball. Moreover, that you were third and another participant decided to invest before you. This is equivalent to the case “One other has invested before you”. Therefore, we will compute your payoff considering only your decision to INVEST for the case “One other has INVESTED before you”. All the other decisions are irrelevant. Hence, if the project is good, we will pay you £3.

Example 2

Suppose you make the following decisions:

	Green Ball	Red Ball
Nobody has bought before you	NO	NO
One other has bought before you	INVEST	INVEST
Two others have bought before you	INVEST	INVEST

At the end of the round the computer tells you that actually in this round you received the green ball and you were first. This is equivalent to the case

“Nobody has invested before you”. Therefore, we will compute your payoff considering only your decision NOT to INVEST for the case “Nobody has invested before you”. All the other decisions are irrelevant.

Procedures for each round

Remember that the experiment is organized into different rounds and that within each round you will have to make six investment decisions. So, now it is time to summarize what happens within each round.

1. At the beginning of each round the computer randomly chooses whether the project is good or bad. The project is the same for all participants. But you will not be told which project has been chosen.
2. The computer draws a ball from an urn for each participant. The ball is drawn and then replaced, so that the total number of balls in the urn is always the same before the computer makes another draw. The computer draws a ball for participant one and puts it back into the urn. Then it does the same for participant 2. And then for participant
3. If the project is good, the computer draws the ball from an urn containing 70 green and 30 red balls. If it is bad, from an urn containing 70 red and 30 green balls. You will be the only one knowing your ball colour.
4. The computer decides randomly which participant is first, who is second and who is third.
5. You will make your decisions for the six cases illustrated above. Of course, given that you do not know in which case you actually are, you will have to think of the best solution for each of the six cases in which you may be.

Once the round is over, you will be informed of your ball colour. We will also tell how many other participants decided to invest before you. And of course we will tell you whether the project was indeed good or bad and how much you earned.

Then, we will repeat the same procedure for the second round at the beginning of which the computer will choose again a project and so on. Note that at the beginning of each round the computer chooses the project always

with an equal chance of being good or bad, independently of what was chosen in previous rounds. We will repeat the same procedures for altogether 15 rounds.

Final payment

For the simple fact that you showed up in time for the experiment you earn £4. The rest of the payment depends on how you perform. The computer will randomly choose one round out of the first 5 rounds, one among the 6th through the 10th and one among the 11th through the 15th. Your payment will depend on how you performed in the selected rounds. We will sum up your payoffs in these three rounds. Your final payment will be equal to this amount plus the £4 for showing up.