

Self-correcting Information Cascades¹

Jacob K. Goeree, Thomas R. Palfrey, Brian W. Rogers, and Richard D. McKelvey²

November 2003 (revised April 15, 2004)

Abstract

In laboratory experiments, information cascades are ephemeral phenomena, collapsing soon after they form, and then reforming again. These formation/collapse/reformation cycles occur frequently and repeatedly. Cascades may be reversed (collapse followed by a cascade on a different state) and more often than not, such a reversal is self-correcting: the cascade switches from the incorrect to the correct state. Past experimental work focused on relatively short horizons, where these interesting dynamics are rarely observed. We present experiments with a longer horizon, and also investigate the effect of signal informativeness. We propose a theoretical model, based on quantal response equilibrium, where temporary and self-correcting cascades arise as equilibrium phenomena. The model also predicts the systematic differences we observe experimentally in the dynamics, as a function of signal informativeness. We extend the basic model to include a parameter measuring base rate neglect and find it to be a statistically significant factor in the dynamics, resulting in somewhat faster rates of social learning.

JEL classification numbers: C92, D82, D83

Key words: social learning, information cascades, laboratory experiments

¹Financial support from the National Science Foundation (SBR-0098400 and SES-0079301) and the Alfred P. Sloan Foundation is gratefully acknowledged. The theory and experimental design was partially completed, and pilot experiments were conducted in collaboration with Richard McKelvey, who died in April 2002. He is not responsible for any errors in the paper. We acknowledge helpful comments from Bogacen Celen, Terry Sovinsky, seminar participants at UCLA, NYU, Harvard University, Universitat Autònoma de Barcelona, University of Edinburgh, Washington University, the 2003 annual meeting of ESA in Pittsburgh, the 2003 Malaga Workshop on Social Choice and Welfare Economics, the 2003 SAET meetings in Rhodos, and the 2003 ESSET meetings in Gerzensee. We thank Iva Rashkova for research assistance.

²Goeree: CREED, Faculteit voor Economie en Econometrie, Universiteit van Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands. Palfrey and Rogers: Division of Humanities and Social Sciences, California Institute of Technology, Mail code 228-77, Pasadena, CA 91125, USA. McKelvey: Deceased.

“Those who err in mind will know the truth,” Isaiah 29:24¹

1. Introduction

In an information cascade, a sequence of imperfectly informed decision makers each of whom observes all previous decisions, quickly reach a point at which they rationally ignore their private information. Hence, after a few decisions, learning ceases as subsequent decision makers infer nothing new from observing any of the actions. Information cascades are predicted to occur despite the wealth of information available and despite the common interest of all decision makers. This result, if robust to variations in the basic model, has obvious and pernicious implications for economic welfare, and raises problematic issues for various applications of mass information aggregation, such as stock market bubbles and crashes, bank runs, technology adoption, mass hysteria, and political campaigns.

In this paper, we reconsider the canonical model of information cascades for which some laboratory data (from short sequences) exist, see Anderson and Holt (1997) henceforth denoted by AH.² There are two equally likely states of nature, two signals, two actions, and T decision makers. Nature moves first and chooses a state, and then reveals to each decision maker a private signal about the state. The probability a decision maker receives a correct signal is $q > 1/2$ in both states of the world. Decision makers choose sequentially, with each decision maker observing all previous actions (and her private signal). A decision maker receives a payoff of 1 if she chooses the correct action and 0 otherwise. In this environment, learning never progresses very far in a Nash or Sequential Equilibrium. In fact, regardless of T , the beliefs of all decision makers before considering their private information (or, equivalently, the beliefs of an external observer) are confined to a narrow interval centered at $1/2$.

In previous experimental data, however, there are numerous and repeated action choices that are inconsistent with Nash equilibrium given the realized signals, by nearly all subjects.

¹New American Standard Bible (1962).

²Subsequent studies by Anderson (2001), Hung and Plott (2001), Domowitz and Hung (2003), Nöth and Weber (2003), and others replicate the AH findings and include some additional treatments.

	$q = 5/9$	$q = 5/9$	$q = 6/9$	$q = 6/9$
	$T = 20$	$T = 40$	$T = 20$	$T = 40$
# sequences	116	56	90	60
# sequences with pure cascades	5	0	12	8
# sequences without cascades	0	0	0	0
# sequences with broken cascades	111	56	78	52

Table 1. *Percentages of (broken) cascades in our data.*

AH observe that in their experiment with $q = 2/3$ and $T = 6$, more than 25% of the time subjects make a choice against the cascade after receiving a contradictory signal. And nearly 5% of subjects who receive a signal consistent with the cascade choose the opposite action. In the experiments reported below we vary the signal precision, $q = 5/9$ and $q = 6/9$, and the number of decision makers, $T = 20$ and $T = 40$. With this many decision makers we should observe cascades arising in 100% of the sequences according to the theoretical model of Bikhchandani, Hirschleifer, and Welch (1992). However, with $T = 40$, for instance, a cascade arises and persists in only 8 out of 116 sequences ($< 7\%$).

Table 1 gives an indication of a few ways the standard theory misses badly. At a minimum, a reasonable theory should explain two systematic features of the data. First, off-the-Nash-equilibrium-path actions occur with significant probability. The theory as it stands does not place adequate restrictions off the equilibrium path. Second, deviations from equilibrium are systematic, indicating that such behavior is informative! Why? Because going off the equilibrium path (i.e., choosing an action opposite to the cascade) happens much more frequently if the player received a signal contradicting the cascade choices, see Table 2. Indeed, when a break occurs, the observed frequency with which the received signal was contradictory is 84%.³ This should come as no surprise as a deviation following a confirmatory signal is a worse deviation (e.g., in terms of expected payoffs, and also intuitively) than a deviation following a contradictory signal.

³When averaged over the four treatments. In the $(q = \frac{5}{9}, T = 20)$, $(q = \frac{5}{9}, T = 40)$, $(q = \frac{6}{9}, T = 20)$, and $(q = \frac{6}{9}, T = 40)$ treatments the numbers are 87%, 78%, 87%, and 82% respectively.

		T=20		T=40		
q=5/9	Decision \ Signal	Confirming	Contrary	Decision \ Signal	Confirming	Contrary
	Don't Break	633	348	Don't Break	555	291
	Break	41	283	Break	86	302
q=6/9	Decision \ Signal	Confirming	Contrary	Decision \ Signal	Confirming	Contrary
	Don't Break	575	304	Don't Break	1150	605
	Break	26	175	Break	30	138

Table 2. *Frequency of confirmatory/contrary signals when cascades are (not) broken.*

An alternative approach to the perfect Bayesian equilibrium is to consider models which admit a random component to behavior. The introduction of a random component ensures that all paths can be reached with positive probability, so Bayes' rule places restrictions on future rational inferences and behavior when a deviation from a cascade occurs. We consider such a model, quantal response equilibrium (QRE), where deviations from optimal play occur according to a statistical process and players take these deviations into account when making inferences and decisions. In a QRE, deviations or mistakes are payoff dependent in the sense that the likelihood of a mistake is inversely related to its cost.⁴ We demonstrate that QRE predicts the temporary and self-correcting nature of cascades and also predicts the systematic differences we observe experimentally in the dynamics, as a function of signal informativeness.

The remainder of the paper is organized as follows. Section 2 presents the basic model and theoretical results. Section 3 describes the experimental design. Section 4 contains an analysis of the data. Section 5 presents an econometric analysis of the basic model and develops extensions to better explain the data. Section 6 discusses efficiency properties and section 7 concludes. Appendix A contains proofs and Appendix B contains a program for estimation and estimation results.

⁴We only consider *monotone* quantal response equilibrium, where choice probabilities are monotone in expected utilities, see McKelvey and Palfrey (1995, 1998).

2. The Basic Model

There is a finite set $\mathcal{T} = \{1, 2, \dots, T\}$ of agents who sequentially choose between one of two alternatives, A and B . For each $t \in \mathcal{T}$ let $c_t \in \{A, B\}$ denote agent t 's choice. One of the alternatives is selected by nature as “correct,” and an agent receives a payoff of 1 only when she selects this alternative, otherwise she gets 0. The correct alternative (state of the world) is denoted by $\omega \in \{A, B\}$, which is unknown to the agents who have common prior beliefs that $\omega = A$ or $\omega = B$ with probability $\frac{1}{2}$. Further, they receive conditionally independent private signals s_t regarding the better alternative. If $\omega = A$ then $s_t = a$ with probability $q > 1/2$ and $s_t = b$ with probability $1 - q$. Likewise, when $\omega = B$, $s_t = b$ with probability q and $s_t = a$ with probability $1 - q$.

We will be concerned with the evolution of agents' beliefs, and how these beliefs co-evolve with actions. Agent t observes the actions of all her predecessors, but not their types. Thus a history H_t for agent t is simply a sequence $\{c_1, \dots, c_{t-1}\}$ of choices by agents $1, \dots, t - 1$, with $H_1 = \emptyset$. Agents care about the history only to the extent that it is informative about which alternative is correct. So let $p_t \equiv P(\omega = A|H_t)$ denote the (common knowledge) posterior belief that A is correct given the choice history H_t , with $p_1 \equiv \frac{1}{2}$, the initial prior. We first determine agent t 's private posterior beliefs given the public beliefs p_t and given her signal s_t . Applying Bayes' rule shows that if $s_t = a$, agent t believes that alternative A is correct with probability

$$\pi_t^a(p_t) \equiv P(\omega = A|H_t, s_t = a) = \frac{q p_t}{q p_t + (1 - q)(1 - p_t)}. \quad (2.1)$$

Likewise,

$$\pi_t^b(p_t) \equiv P(\omega = A|H_t, s_t = b) = \frac{(1 - q)p_t}{(1 - q)p_t + q(1 - p_t)} \quad (2.2)$$

is the probability with which agent t believes that A is correct if her private signal is $s_t = b$. A direct computation verifies that $\pi_t^a(p_t) > p_t > \pi_t^b(p_t)$ for all $0 < p_t < 1$. In other words, for any interior public belief an agent believes more strongly that $\omega = A$ after observing an a signal than after observing a b signal.

2.1. Nash Equilibrium

Following Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992) we first discuss optimal behavior under the assumption that full rationality is common knowledge. Given that each agent's private information is of the same precision, and the initial prior puts equal mass on both states, indifference occurs with positive probability resulting in a multiplicity of sequential equilibria.⁵ This multiplicity is potentially relevant for interpreting data from information cascade experiments, since the restrictions on action sequences are minimal. Indeed, below we illustrate how any action sequence is consistent with some sequential equilibrium for some sequence of signals. For field data, where signals cannot be directly observed, this means there are essentially no restrictions imposed by the Nash equilibrium. Furthermore, implications about behavior off the equilibrium path are quite ambiguous.

As an example to see that any action sequence is consistent with equilibrium, suppose the following sequence of actions is observed in the first four periods of an information cascade game: $\{A, A, A, B\}$. What (outsider) posterior beliefs are consistent with these actions if we assume they are generated by equilibrium behavior? There are three restrictions derived from equilibrium behavior that drive possible beliefs. The first is that player 1 must have observed signal a . Next the second and third players must have been using strategies that are uninformative (follow player 1 regardless of signal, which is a weak best response), otherwise choosing action B could not have been optimal for player 4. Third, since the intervening A choices by players 2 and 3 were uninformative, player 4 must have observed a b signal. Notice that extending this argument implies that, for an outsider who cannot observe the private information of the players, *any sequence of actions is consistent with some equilibrium for some realization of signals*. Fortunately, with experimental data, the outside observer (i.e., the experimenter) has the luxury of observing both signals and actions, and hence can place some restrictions on the data, although these restrictions are limited.⁶

⁵This multiplicity is non-generic and occurs because of the symmetric information structure (uniform prior, symmetric signal technology).

⁶The above argument also holds for Perfect Bayesian Equilibrium and Sequential Equilibrium since these

The multiplicity of equilibria disappears when indifferent players follow their signal with non-zero probability, no matter how small. In this case, the “pure cascade” Nash equilibrium identified by Bikhchandani, Hirshleifer and Welch (1992) is the only equilibrium.⁷

The pure cascade Nash equilibrium works as follows.⁸ The first agent chooses A if $s_1 = a$, and chooses B if $s_1 = b$, so that her choice perfectly reveals her signal. If the second agent’s signal agrees with the first agent’s choice, the second agent chooses the same alternative, which is strictly optimal. On the other hand, if the second agent’s signal disagrees with the first agent’s choice, the second agent is indifferent, as she effectively has a sample of one a and one b . Rather than making a specific assumption, suppose she follows her signal with some probability $\beta \geq \frac{1}{2}$. The third agent faces two possible situations: (i) the choices of the first two agents coincide, or (ii) the first two choices differ. In case (i), it is strictly optimal for the third agent to make the same choice as her predecessors, even if her signal is contrary. Thus her choice imparts no information to her successors, resulting in the onset of a cascade. The fourth agent is then in the same situation as the third, and so also makes the same choice, a process which continues indefinitely. In case (ii), however, the choices of the first two agents reveal that they have received one a signal and one b signal, leaving the third agent in effectively the same position as the first. Her prior (before considering her private information) is therefore $\frac{1}{2}$, so that her signal completely determines her choice. The fourth agent would then be in the same situation as the second agent described above, et cetera.

One quantity of interest is the probability that “correct” and “incorrect” cascades have formed after a particular number of choices. After the first two choices, the probabilities of a correct cascade, no cascade, and an incorrect cascade are

$$q(1 - \beta + \beta q), 2\beta q(1 - q), (1 - q)(1 - \beta q),$$

refinements do not rule out indifferent players always following the cascade.

⁷In fact, the trembling-hand perfect equilibrium selects a unique equilibrium in which indifferent players follow their signal with probability 1.

⁸As we will see, most data are not consistent with this “pure cascade” Nash equilibrium. In fact, most sequences of choices observed in the laboratory are not consistent with *any* Nash equilibrium.

respectively. More generally, after $2t$ choices, these probabilities are

$$q(1 - \beta + \beta q) \left(\frac{1 - (2\beta q(1 - q))^t}{1 - 2\beta q(1 - q)} \right), (2\beta q(1 - q))^t, (1 - q)(1 - \beta q) \left(\frac{1 - (2\beta q(1 - q))^t}{1 - 2\beta q(1 - q)} \right).$$

Taking limits as t approaches infinity yields the long run probabilities of the three regimes. First note that the probability of not being in a cascade vanishes as t grows. The probability of eventually reaching a correct cascade is $\frac{q(1-\beta+\beta q)}{1-2\beta q(1-q)}$, and the complementary probability of eventually reaching an incorrect cascade is $\frac{(1-q)(1-\beta q)}{1-2\beta q(1-q)}$.⁹ Once a cascade has formed, all choices occur independently of private information, and hence public beliefs remain unchanged. The points at which public beliefs settle are the posteriors that obtain after two consecutive choices for the same alternative, beginning with uninformative priors.

2.2. Quantal Response Equilibrium

We now describe the logit quantal response equilibrium (QRE) of the model described above. In the logit QRE, each individual t privately observes a payoff disturbance for each choice, denoted ϵ_t^A and ϵ_t^B . The payoff-relevant information for agent t is summarized by the difference $\epsilon_t \equiv \epsilon_t^A - \epsilon_t^B$. Denote agent t 's type by $\theta_t = (s_t, \epsilon_t)$. The logit specification assumes that the ϵ_t are independent and obey a logistic distribution with parameter λ .¹⁰ The disturbance, ϵ_t , can be interpreted in several different ways. For example, it could represent a stochastic part of decision making due to bounded rationality, or it could be an individual-specific preference shock that occurs for other reasons. Irrespective of the interpretation of the noise, the resulting logit choice model implies that the stronger the belief that A is the correct state, the more likely action A is chosen. The logit QRE model assumes that the distribution of the payoff disturbances is common knowledge.¹¹ The logit QRE is calculated as the sequential equilibrium of the resulting game of incomplete information, where each player observes only her own type $\theta = (s, \epsilon)$.

⁹Thus as q increases from $\frac{1}{2}$ to 1, the probability of eventually reaching a good cascade grows from $\frac{1}{2}$ to 1.

¹⁰This arises, when ϵ_t^A and ϵ_t^B are *i.i.d.* extreme-value distributed.

¹¹In general, the distributions of payoff disturbances in a logit QRE need not be the same for every decision maker, but these distributional differences would be assumed to be common knowledge.

It is straightforward to characterize the optimal decision of agent t given her type θ_t and the history H_t (which determines public beliefs p_t). The expected payoff of choosing A is $\pi_t^{s_t}(p_t) + \epsilon_t$, and that of selecting alternative B is $1 - \pi_t^{s_t}(p_t)$. Thus given agent t 's signal, the probability of choosing A is given by¹²

$$P(c_t = A|H_t, s_t) = P(\epsilon_t > 1 - 2\pi_t^{s_t}(p_t)) = \frac{1}{1 + \exp(\lambda(1 - 2\pi_t^{s_t}(p_t)))}, \quad (2.3)$$

and B is chosen with complementary probability $P(c_t = B|H_t, s_t) = 1 - P(c_t = A|H_t, s_t)$. When $\lambda \rightarrow \infty$ choices are fully rational in the sense that they do not depend on the private realizations ϵ_t and are determined solely by beliefs about the correct alternative. It is easy to show that the logit QRE converges to the pure cascade Nash equilibrium (with $\beta = \frac{1}{2}$) in this limit. On the other hand, as λ approaches 0 choices are independent of beliefs and become purely random.¹³

The dynamics of posterior beliefs depend on λ . To derive the evolution of the public belief that A is correct, note that given p_t there are exactly two values that $p_{t+1} = P(\omega = A|H_t, c_t)$ can take depending on whether c_t is A or B . These are denoted p_t^+ and p_t^- respectively. The computation of the posterior probabilities p_t^+ and p_t^- given p_t is carried out by agents who do not know the true state, and so cannot condition their beliefs on that event. In contrast, the transition probabilities of going from p_t to p_t^+ or p_t^- (i.e., of a choice for A or B) depend on the objective probabilities of a and b signals as dictated by the true state. Thus when computing these transition probabilities, it is necessary to condition on the true state. Conditional on $\omega = A$, the transition probabilities are:

$$\begin{aligned} T_t^{\omega=A} &= P(c_t = A|H_t, \omega = A) \\ &= P(c_t = A|H_t, s_t = a)P(s_t = a|\omega = A) + P(c_t = A|H_t, s_t = b)P(s_t = b|\omega = A) \\ &= \frac{q}{1 + \exp(\lambda(1 - 2\pi_t^a(p_t)))} + \frac{1 - q}{1 + \exp(\lambda(1 - 2\pi_t^b(p_t)))}, \end{aligned}$$

¹²Note that indifference occurs with probability zero under the logit specification, and hence plays no role.

¹³For any $\lambda \in (0, \infty)$, an agent chooses equi-probably when indifferent. Thus the logit QRE naturally selects exactly one of the continuum of Nash equilibria as λ diverges to infinity.

with the probability of a B choice given by $1 - P_t^A$. Similarly, conditional on $\omega = B$, the probability agent t chooses A is

$$T_t^{\omega=B} = \frac{1 - q}{1 + \exp(\lambda(1 - 2\pi_t^a(p_t)))} + \frac{q}{1 + \exp(\lambda(1 - 2\pi_t^b(p_t)))}.$$

Using Bayes' rule, we now obtain the two values that p_{t+1} may take as

$$p_t^+ \equiv P(\omega = A | H_t, c_t = A) = \frac{p_t T_t^{\omega=A}}{p_t T_t^{\omega=A} + (1 - p_t) T_t^{\omega=B}}, \quad (2.4)$$

and

$$p_t^- \equiv P(\omega = A | H_t, c_t = B) = \frac{p_t(1 - T_t^{\omega=A})}{p_t(1 - T_t^{\omega=A}) + (1 - p_t)(1 - T_t^{\omega=B})}. \quad (2.5)$$

These expressions can be used to derive the following properties of the belief dynamics (see Appendix A for proofs), where without loss of generality we assume the true state is $\omega = A$.

Proposition 1. *In the logit QRE:*

- (i) *Beliefs are interior: $p_t \in (0, 1)$ for all $t \in \mathcal{T}$.*
- (ii) *Actions are informative: $p_t^- < p_t < p_t^+$ for all $t \in \mathcal{T}$.*
- (iii) *Beliefs about the true state rise on average: $E(p_{t+1} | p_t, \omega = A) > p_t$ for all $t, t + 1 \in \mathcal{T}$.*
- (iv) *Beliefs converge to the truth: conditional on $\omega = A$, $\lim_{t \rightarrow \infty} p_t = 1$ almost surely.*

2.3. Testable Restrictions on the Data

We formalize different kinds of cascade-like behavior.¹⁴ A *pure A (B) cascade* is said to form at time $t \leq T$ if in the $t - 1$ periods before the number of A (B) choices exceed the number

¹⁴One might argue that we should use the term “herd” instead of cascade, since cascade refers to belief dynamics, while “herds” refer to choice dynamics. In the context of quantal response equilibrium, this distinction is artificial, since neither herds nor cascades ever arise in a logit equilibrium.

of B (A) choices by 2 for the first time, and *all* choices from t to T are A (B) choices. Thus, for example, if $T = 6$ and the sequence of choices is $\{A, B, A, A, A, A\}$, then we say a pure A cascade forms at $t = 5$. In periods 5 and 6, we say the decision makers are in a pure A cascade. Note that any pure cascade beginning at time t , will have length $T - t + 1$.

A *temporary A (B) cascade* or *A (B) craze*¹⁵ is said to form at time $t \leq T$ if in the $t - 1$ periods before (but not in the $t - 2$ periods before) the number of theoretically informative A (B) choices¹⁶ exceed the number of theoretically informative B (A) choices by 2 and some decision maker τ , with $t \leq \tau \leq T$, makes a contrary choice.¹⁷ The number of periods decision makers follow the craze, $\tau - t$, defines the *length* of a craze. Note that in the sequence of decisions $\{A, A, B\}$ we say that an A craze of length zero occurs at $t = 3$.

Temporary cascades are particularly interesting because subsequent play of the game is off the equilibrium path. Moreover, if the sequence is long enough it is possible for a new cascade to form after a temporary cascade has broken. Following AH, we define a simple counting procedure to classify sequences of decisions and determine whether a new cascade has formed. This *ad hoc* counting rule roughly corresponds to Bayesian updating when the probability that indifferent subjects follow their signals, β , equals the probability that subjects who break cascades hold contrary signals.¹⁸ Under the counting rule, every A decision when *not* in a cascade increases the count by 1 and every B decision when *not* in a cascade decreases the count by 1. Recall that we enter the first cascade of a sequence when the count reaches 2 or -2 . Then the decisions during the cascade do not change the count, until there is an action that goes against the cascade, which decreases the count to 1 if it was an A cascade or increases the count to -1 if it was a B cascade. The count continues to change in this way, until the count reaches either 2 or -2 again, and then we are in a new cascade, which we call a *secondary* cascade.

¹⁵According to the Oxford English Dictionary (1980), a craze is defined as a “great but often short-lived enthusiasm for something.”

¹⁶Choices made during a (temporary) cascade are called theoretically uninformative.

¹⁷These definitions extend in a natural way to more complex environments.

¹⁸These conditions are closely approximated in our data, where we find 85% of indifferent subjects go with their signals and 84% of cascade breakers received contrary signals.

We distinguish three different kinds of secondary cascades. One possibility is that actions cascade on the same state as the previous cascade: a *repeat cascade*. The other possibility is that the actions cascade on a different state: a *reverse cascade*. If a cascade reverses from the incorrect state to the correct state, it is called a *self-correcting cascade*.

The logit equilibrium implies several properties of the length and frequency of different kinds of cascades, and how this depends on our two main treatment parameters, q and T .¹⁹

Properties. *In the logit QRE:*

- (P1) The probability a pure cascade breaks is increasing in T and decreasing in q .
- (P2) For any q the probability of a pure cascade goes to 0 as T gets large.
- (P3) The number of cascades is increasing in T and decreasing in q .
- (P4) The length of cascades is increasing in T and q .
- (P5) Incorrect cascades are shorter than correct cascades.
- (P6) Incorrect cascades reverse (a self-correcting cascade) more frequently than correct cascades reverse, for any $T \geq 6$.²⁰
- (P7) Correct cascades repeat more frequently than incorrect cascades repeat, for any $T \geq 4$.²¹
- (P8) Later cascades are more likely to be correct than earlier ones.
- (P9) The *ex ante* (i.e., before player t has drawn a private signal) probability of a correct decision is increasing in both t and q . An interim version of this statement is true, but

¹⁹These properties are stated informally here: they have been verified by extensive simulations and in some cases can be proved formally.

²⁰Six periods are required for a reverse cascade, at least two periods to start the first cascade, and then at least four periods to reverse itself. For example, $\{A, A, B, B, B, B\}$ is the shortest possible sequence for a reverse from an A cascade to a B cascade. If $T < 6$ then there is not enough time for a reverse cascade to occur, illustrating the necessity of conducting experiments with sufficiently *long* sequences.

²¹Repeat cascades require at least 4 periods. For example, $\{A, A, B, A\}$ is the shortest possible sequence for a repeated A cascade.

only conditional on receiving an incorrect signal.²²

- (P10) More informative signals lead to faster learning. That is, the rate of increase in the *ex ante* probability of a correct decision as t increases is increasing in q . An interim version of this statement is true, but only conditional on receiving an incorrect signal.
- (P11) The probability of a correct decision is higher for a correct signal than for an incorrect signal.

While the last three properties are also true for the initial few decisions in the Nash equilibrium model, the effects go away quickly with longer sequences.²³

3. Experimental Design

In order to assess the quality of these models in the cascade setting, and to gain insights into how the basic model might be improved, we conducted a set of experiments implementing the situation described above under controlled laboratory conditions. These experiments were conducted at the Social Sciences Experimental Laboratory (SSEL) at Caltech and the California Social Sciences Experimental Laboratory (CASSEL) at UCLA between September 2002 and May 2003. The subjects included students from these two institutions who had not previously participated in a cascade experiment.²⁴

We employ a 2×2 design, where we consider two values of both the signal quality q and the number of individuals T . Specifically, q takes values $5/9$ and $6/9$, and T takes values 20 and 40. The number of games in each session is denoted M . Table 3 summarizes the experimental sessions.

²²It is *not* true conditional on receiving a correct signal. To see this, note that the interim probability of a correct decision at $t = 1$ with a correct signal is 1 as it is optimal to follow one's signal. In later periods, it is less than 1, because of the probability of a cascade on the wrong state.

²³An exception is the second part of property 8. In the perfect Nash equilibrium, the probability of a correct decision is approximately equal to the probability of ending up in a correct cascade, which quickly approaches $q^2/(q^2 + (1 - q)^2)$, which is increasing in q .

²⁴There was one subject who had previously participated in a related pilot experiment.

Session	T	q	M	Subject Pool
03/14/03A	20	5/9	30	Caltech
09/26/02B	20	5/9	30	Caltech
09/19/02A	20	5/9	26	Caltech
04/03/03AB	20	5/9	30	UCLA
04/14/03A	20	6/9	30	UCLA
04/14/03C	20	6/9	30	UCLA
04/14/03E	20	6/9	30	UCLA
05/05/03D	40	5/9	17	UCLA
05/05/03F	40	5/9	19	UCLA
05/05/03G	40	5/9	20	UCLA
04/16/03B	40	6/9	20	UCLA
04/21/03C	40	6/9	20	UCLA
04/21/03E	40	6/9	20	UCLA

Table 3. *Experimental sessions.*

In each session, a randomly chosen subject was selected to be the “monitor” and the remaining subjects were randomly assigned to computer terminals in the laboratory. All interaction among subjects took place through the computers; no other communication was permitted. Instructions were given with a voiced-over Powerpoint presentation in order to minimize variations across sessions.²⁵ After logging in, the subjects were taken slowly through a practice trial (for which they were not paid) in order to illustrate how the software worked, and to give them a chance to become familiar with the process before the paid portion of the experiment started.

Before each trial, the computer screen displayed two urns. For the $q = 5/9$ treatment, one urn contained 5 blue balls and 4 red balls and the other contained 4 blue balls and 5 red balls. For the $q = 6/9$ treatment, one urn contained 6 blue balls and 3 red balls and the other contained 3 blue balls and 6 red balls. The monitor was responsible for rolling a die at the beginning of each game to randomly choose one of the urns with equal probabilities. This process, and the instructions to the monitor (but not the outcome of the roll) was done publicly. At this point, the subjects

²⁵See <http://www.hss.caltech.edu/~rogers/exp/index6.html> for the instructions.

saw only one urn on the computer screen, with all nine balls colored gray, so that they could not tell which urn had been selected. Each subject then independently selected one ball from the urn on their screen to have its color revealed. Then, in a random sequence, subjects sequentially guessed an urn. During this process, each guess was displayed on all subjects' screens in real time as it was made, so each subject knew the exact sequence of guesses of all previous subjects. After all subjects had made a choice, the correct urn was revealed and subjects recorded their payoffs accordingly. Subjects were paid \$1.00 for each correct choice and \$0.10 for each incorrect choice. Subjects were required to record all this information on a record sheet, as it appeared on their screen. Due to time constraints, the number of matches (sequences of T decisions) was $M = 30$ in each $T = 20$ session and $M = 20$ in each $T = 40$ session.²⁶ After the final game, payoffs from all games were summed and added to a show-up payment, and subjects were then paid privately in cash before leaving the laboratory.

4. Results I: Cascades and Off-the-Equilibrium-Path Behavior

In this section, we provide some descriptive aggregate summary information about the extent of cascade formation, off-the-equilibrium-path behavior, and the number and lengths of cascades of different kinds. We also compare these aggregate features across our four treatments, and compare them to the shorter (6-period) cascade experiment reported in AH.

4.1. Infrequency of Pure Cascades and Frequency of Crazes

In AH's experiment with only $T = 6$ decision makers, all cascades were necessarily very short making it difficult to sort out pure cascades from crazes. In contrast, our experiments investigated sequences of $T = 20$ and $T = 40$ decision makers, allowing for the first time an opportunity to observe long cascades and the length distribution of crazes. As Table 4 shows

²⁶A few sessions contained fewer sequences due to technical problems.

	Our Data				AH Data	HP Data	Nash		QRE - BRF	
	q = 5/9		q = 6/9		q = 6/9	q = 6/9	q = 5/9	q = 6/9	q = 5/9	q = 6/9
	N = 20	N = 40	N = 20	N = 40	N = 6	N = 10				
	M = 116	M = 56	M = 90	M = 60	M = 45	M = 89				
First 6	36	36	50	53	64	70	98	99	29	45
First 10	15	14	31	42		62	100	100	11	29
First 20	4	2	13	32			100	100	1	15
First 40		0		13			100	100	0	7

Table 4. *Percentages of pure cascades by treatment.*

very clearly, pure cascades essentially do not happen in the longer trials. The cascades that persisted in the AH experiments simply appear to be pure cascades, a likely artifact of the short horizon. Our numbers are comparable to those of AH when we consider only the first six decision makers in our sequences. These numbers are given in the row marked "First 6" in Table 4. In contrast, we observe pure cascades in only 17 out of 206 sequences with $T = 20$ decision makers, and only 8 of 116 sequences with $T = 40$ decision makers.

The final columns of Table 4 give the predicted frequency of pure cascades according to the Nash equilibrium (and the QRE-BRF model, which we will discuss later). In general, the equilibrium probability of a pure cascade with T decision makers is $1 - (2\beta q(1 - q))^{T/2}$, with $\beta = 0.85$ the fraction of indifferent subjects who follow their signals. The data contradicts this in three ways. First, the data exhibit far fewer pure cascades than theory predicts. Second, according to theory, the frequency of pure cascades increases with T but it goes the other way in the data. Third, the frequency of pure cascades in the data is steeply increasing in q , while the Nash equilibrium predicts almost no effect. In our data, pure cascades occurred nearly five times as often in the $q = 6/9$ treatment than when $q = 5/9$ (20/150 compared to 5/172).²⁷

In contrast to pure cascades, crazes are quite common in all treatments. Table 5 shows the frequency of temporary cascades in our data. The rows and columns mirror Table 4, but the

²⁷Further evidence indicates this continues to increase with q . In a single additional session with $q = 3/4$ and $T = 20$, we observed pure cascades in 28/30 sequences.

	Our Data				AH Data	HP Data	Nash		QRE - BRF	
	q = 5/9		q = 6/9		q = 6/9	q = 6/9	q = 5/9	q = 6/9	q = 5/9	q = 6/9
	N = 20	N = 40	N = 20	N = 40	N = 6	N = 10				
	M = 116	M = 56	M = 90	M = 60	M = 45	M = 89				
First 6	58	55	42	38	27	24	0	0	62	49
First 10	84	84	69	55		38	0	0	88	70
First 20	96	98	87	68			0	0	99	86
First 40		100		87			0	0	100	93

Table 5. Percentages of temporary cascades by treatment.

entries now indicate the proportion of sequences in a given treatment that exhibit at least one temporary cascade that falls apart. Clearly, for large T , essentially all cascades we observe are temporary. Even with the short horizon of the AH experiment, crazes occur 27% of the time.

4.2. Number and Lengths of Crazes

For larger T , we generally observe *multiple* crazes along a single sequence. On average, there are three or more crazes per sequence in all treatments, see Table 6. Furthermore, the number of crazes rises with the sequence length, T , and falls with the signal precision, q , in the sense of first-degree stochastic dominance, see the top panel of Figure 1. This figure also shows the Nash prediction of exactly 1 cascade per sequence, independent of q and T , and the predictions of the QRE-BRF model discussed below.

The average length of crazes also varies by treatment. Recall that in the Nash model, cascades can only begin after an even number of choices. Moreover, for t even, the probability a cascade forms after $t+2$ choices conditional on one not having yet formed after t choices is $1 - 2\beta q(1-q)$, with $\beta = 0.85$ see section 2.1. Finally, in the Nash model, once a cascade forms it persists through period T . The predicted length distributions of crazes can be calculated easily, see Table 6 and the bottom panel of Figure 1. The Nash predictions are virtually independent of q and contrast

		q = 5/9		q = 6/9	
		N = 20	N = 40	N = 20	N = 40
		M = 116	M = 56	M = 90	M = 60
average number cascades	Our Data	3.47	7.54	2.99	3.73
	QRE - BRF	3.85	7.31	2.81	4.19
	Nash	1.00	1.00	1.00	1.00
average length cascades	Our Data	2.43	2.00	3.27	7.83
	QRE - BRF	1.59	2.04	3.79	6.50
	Nash	17.32	37.25	17.43	37.44

Table 6. *Number and lengths of cascades by treatment.*

sharply with the observed (average) lengths. In contrast, the QRE-BRF model does reasonably well and captures the comparative static effects predicted by the different treatments.

4.3. Off-the-Equilibrium-Path Behavior

The previous subsections show that the vast majority (92%) of cascades are *temporary* and short in duration, and nearly all (90%) sequences of length $T = 20$ and $T = 40$ exhibit *multiple* cascades. One immediate conclusion is that there are many choices off the (Nash) equilibrium path. Table 2 in the Introduction characterizes a subset of these choices for the different treatments as a function of the signals of the deviating decision maker. The table shows the behavior of what we call *cascade breakers*, since these are all terminal decisions of a temporary cascade.

In the continuation play following such actions, choices can rationally depend on private signals again if subsequent players believe with sufficiently high probability²⁸ that cascade breakers received contrary signals. Under these beliefs, a subject who observes a deviation from the cascade and also receives a contrary signal would rationally follow the previous subject's deviation,

²⁸The critical probability-level depends on the behavior of indifferent agents. In the data, indifferent agents follow their signal with probability 0.85 and cascade breakers received contrary signals with probability 0.84. These numbers imply that immediately after a break, subjects with signals contrary to the recent cascade are (close to) indifferent and it is rational for them to follow their signals. Those with signals confirming the recent cascade should rationally follow their signals as well and re-start the cascade. In our data, 75% of the choices following a cascade break are in line with the signals received.




Figure 1: The left panels depict the observed distributions of the number of cascades (top) and of cascade lengths (bottom), color coded by treatment: dark (light) gray lines correspond to $q = 5/9$ ($q = 6/9$) and they are solid (broken) for $T = 40$ ($T = 20$). The right panels show predictions of the Nash and QRE-BRF models. In the top right panel, the solid line that jumps to 100% at 1 corresponds to Nash predictions and the other lines the QRE-BRF predictions. In the bottom right panel, the lines that jump to 100% at 18 and 38 correspond to Nash predictions and the others to QRE-BRF predictions.

		T=20			T=40		
q=5/9	Decision \ Signal	Confirming	Contrary	Decision \ Signal	Confirming	Contrary	
	Confirming	42.8%	20.9%	Confirming	45.4%	20.3%	
	Contrary	3.3%	33.0%	Contrary	8.4%	25.9%	
		# obs = 306			# obs = 379		
q=6/9	Decision \ Signal	Confirming	Contrary	Decision \ Signal	Confirming	Contrary	
	Confirming	48.4%	14.7%	Confirming	58.2%	26.7%	
	Contrary	6.3%	30.5%	Contrary	2.4%	12.7%	
		# obs = 190			# obs = 165		

Table 7. Choice percentages confirming/contradicting the recent cascade after a break.

whereas she would follow the cascade if she receives a signal consistent with the cascade (as should future subjects, regardless of their signal). This observation allows us to classify the behavior of the person following the deviator, since we observe the signal they draw.²⁹ For example, in the sequence $\{A, A, B, B\}$, player 3 is a deviator, and player 4 is on some equilibrium path of the continuation game if and only if he observed a b signal. A player who observes a signal consistent with the recent cascade must follow the cascade, a prediction that is borne out by our data: only 10% of these players are *secondary deviators* who follow the recent break. More generally, choices directly after a break are in line with the received signals more than 75% of the time, indicating that play can be on the equilibrium path of the continuation game after a break occurs. Table 7 gives a complete breakdown of the choices directly following a cascade break, by treatment.

The two key conclusions of the analysis above is that play off the equilibrium path occurs frequently, and that play off the equilibrium path is informative. The second of these observations has been made in AH, but the first observation, indicating that the standard cascade theory is completely contradicted by the data, was underplayed in AH, as this could easily be missed in short sequences. While it would in principle be impossible to design an experiment to observe

²⁹Of course subjects in the experiment cannot observe the signal of the deviator and the follower, and so cannot rule anything out.

		q = 5/9		q = 6/9	
		T = 20	T = 40	T = 20	T = 40
		M = 116	M = 56	M = 90	M = 60
average number repeat cascades	Our Data	2.14	5.98	1.69	2.60
	QRE - BRF	2.40	5.31	1.56	2.86
	Nash	0.00	0.00	0.00	0.00
average number reversed cascades	Our Data	0.34	0.55	0.30	0.13
	QRE - BRF	0.46	0.98	0.25	0.34
	Nash	0.00	0.00	0.00	0.00

Table 8. *Frequency of repeated and reversed cascades by treatment.*

an infinite sequence of identical choices, it is quite simple to ask whether the data does or does not contradict the predictions of cascade models.

4.4. Repeated and Reversed Cascades: Self Correction

In past information aggregation experiments, with short sequences of few decision makers, there are many fewer observations off the equilibrium path. Since this off-path behavior is central to the dynamic properties of QRE (where such behavior is actually *not* off-path) and to the resulting convergence of beliefs, our design, with much longer sequences, should allow us to better observe the kinds of complex dynamics predicted by the theory, particularly the phenomenon of self correction.

Table 8 shows the average number of repeated and reversed cascades per sequence, by treatment. Theoretical expectations according to the Nash and QRE-BRF models are also given. The latter model predicts the observed number of reversed and repeated cascades remarkably well, while the Nash equilibrium erroneously predicts these types of cascades do not exist.

Table 9 shows how frequently correct and incorrect crazes repeat or reverse themselves.³⁰ Averaging over the four treatments shows that when a correct cascade breaks, it reverses to an

³⁰The percentages listed are based on the total number of correct (or incorrect) crazes minus the number of final correct (or incorrect) crazes, since the latter cannot repeat or reverse themselves because of the finite time horizon.

		T = 20			T = 40		
q = 5/9	From\To	Correct	Incorrect	From\To	Correct	Incorrect	
	Correct	92.7%	7.3%	Correct	93.6%	6.4%	
	Incorrect	22.7%	77.3%	Incorrect	11.0%	89.0%	
		# obs =		# obs =			
q = 6/9	From\To	Correct	Incorrect	From\To	Correct	Incorrect	
	Correct	91.4%	8.6%	Correct	98.7%	1.3%	
	Incorrect	30.5%	69.5%	Incorrect	20.0%	80.0%	
		# obs =		# obs =			

Table 9. *Transitions between correct and incorrect cascades in our data.*

incorrect one in less than 6% of all cases. In contrast, an incorrect cascade that breaks leads to a self-corrected cascade in more than 21% of all cases. Table 9 also lists the initial, final, and total number of correct and incorrect crazes by treatment. Notice that the fraction of correct crazes is always higher among the final cascades than among the initial cascades, confirming the predictions of Proposition 1.

4.5. Summary of Results

Here we summarize our findings by relating them to the properties of the logit QRE discussed in section 2.3.

- Properties P1 and P2: *The occurrence of pure cascades decreases with T and increases with q .* The effect of T is obvious from comparing the different rows in Table 4. Both for $q = 5/9$ and $q = 6/9$, the percentages of pure cascades fall quickly with each successive row. Comparing columns 1 and 3 and columns 2 and 4 in Table 4 shows the effect of signal informativeness.

- Property P3: *The number of cascades increases with T and decreases with q .* Longer sequences have more cascades because they allow for more cycles of formation and collapse, see Table 6 and Figure 1. This is in contrast to the Nash equilibrium, where repeat cascades cannot arise and the number of cascades is 1 in all treatments. These effects are barely noticeable in short sequences: AH's experiment averaged slightly more than 1 cascade per sequence.
- Property P4: *Cascades lengths increase with T for $q = 6/9$ and increase with q .*³¹ These comparative static effects are clear from Table 6 and Figure 1. Again, these differences are barely noticeable in past experiments because the sequences were so short. Note that the effect of T is not borne out by the $q = 5/9$ data, where the distributions of cascade lengths are very similar for the $T = 20$ and $T = 40$ treatments.³²
- Property P5: *Correct crazes last longer on average.* The observed average lengths of (correct, incorrect) crazes in the different treatments are: (2.55, 2.24) for $q = 5/9$ and $T = 20$, (2.08, 1.91) for $q = 5/9$ and $T = 40$, (3.42, 2.85) for $q = 6/9$ and $T = 20$, and (8.31, 5.50) for $q = 6/9$ and $T = 40$. So, this is found in every treatment.
- Properties P6 and P7: *Reverse cascades are usually corrected cascades.* See Table 9. Across the four treatments, the probability that a reversed cascade is a self-corrected cascade is 63% (even though there are many more correct than incorrect crazes to reverse from). It is this feature of the dynamics that produces the asymptotic full information aggregation result of Proposition 1.
- Property P8: *Later cascades are correct more frequently than earlier cascades.* See Table

³¹The first result is due to two different effects that lead to censoring. First, and most obvious, if T is short then some cascades that would have lasted longer are interrupted at T . This censoring effect is further compounded by the fact that the probability of collapse is decreasing in the duration of the cascade. That is, the probability of a collapse in period $t + s$, given the cascade started in period t is decreasing in s . This second effect is predicted by QRE but not by Nash. These two effects combined result in a fat tail of the length distribution and in a mass of cascades at $T - 2$.

³²This may be due to subject pool effects, since the ($q = 5/9, T = 20$) treatment was the only one that used mostly Caltech students.

9, which lists the number of initial and final cascades that are correct or incorrect.

5. Results II: Estimation

We start by describing the estimation procedure for the basic logit QRE model. The only parameter being estimated is the slope of the logit response curve, which in the context of these games can be interpreted as a proxy for rationality, experience, and task performance skill. In subsequent subsections, we jointly estimate logit and other parameters, using standard maximum likelihood estimation. For comparability, we choose to normalize payoffs in all experiments to equal 1 if a subject guesses the state correctly and 0 otherwise.

Since subjects' choice behavior will depend on λ , the public belief follows a stochastic process that depends on λ . The evolution of the public belief can be solved recursively (see equations (2.4) and (2.5)), so implicitly we can write $p_t(c_1, \dots, c_{t-1}|\lambda)$. Given $\{\lambda, s_t, (c_1, \dots, c_{t-1})\}$, the probability of observing player t choose A is:

$$P(c_t = A|\lambda, s_t, c_1, \dots, c_{t-1}) = \frac{1}{1 + \exp(\lambda(1 - 2\pi_t^{s_t}(p_t(c_1, \dots, c_{t-1}|\lambda))))},$$

and $P(c_t = B|\lambda, s_t, c_1, \dots, c_{t-1}) = 1 - P(c_t = A|\lambda, s_t, c_1, \dots, c_{t-1})$. Therefore, the likelihood of a particular sequence of choices, $c = (c_1, \dots, c_T)$, given the sequence of signals is simply:

$$l(c|\lambda) = \prod_{t=1}^T P(c_t|\lambda, s_t, c_1, \dots, c_{t-1}).$$

Finally, assuming independence across sequences, the likelihood of observing a set of M sequences $\{c^1, \dots, c^M\}$ is just:

$$L(c^1, \dots, c^M|\lambda) = \prod_{m=1}^M l(c^m|\lambda).$$

The estimation results for the logit QRE model are given in Table B1 of Appendix B, which also contains a detailed estimation program written in GAUSS. The λ estimates for the four

treatments are quite stable and the pooled estimate is close to that estimated from the AH data. Notice that the estimated value of λ for the ($q = 5/9$, $T = 20$) treatment is somewhat greater than the other three treatments. We attribute this to subject pool effects, since that treatment was the only one that used mostly Caltech students.

5.1. Alternative Models: the Base Rate Fallacy

In their seminal article, Kahneman and Tversky (1973) present experimental evidence showing that individuals' behavior is often at odds with Bayesian updating. A particularly prevalent judgement bias is the Base Rate Fallacy (BRF), or as Camerer (1995, pp. 597-601) more accurately calls it, "base rate neglect". In the context of our social learning model, the base rate fallacy amounts to the assumption that agents weight their own signal more than they should relative to the public prior. We formalize this idea as a non-Bayesian updating process in which a private signal is counted as α signals, where $\alpha \in (0, \infty)$.³³ Rational agents correspond to $\alpha = 1$, while agents have progressively more severe base-rate fallacies as α increases above 1.³⁴

We assume that decision makers do not overweight the actions of others, but treat them without bias, so the public belief process unfolds as in the Nash equilibrium. The updating rules in (2.1) and (2.2) now become

$$\pi_t^a(p_t|\alpha) = \frac{q^\alpha p_t}{q^\alpha p_t + (1-q)^\alpha(1-p_t)}. \quad (5.1)$$

and

$$\pi_t^b(p_t|\alpha) = \frac{(1-q)^\alpha p_t}{(1-q)^\alpha p_t + q^\alpha(1-p_t)} \quad (5.2)$$

respectively.³⁵

³³This could also be interpreted as a parametric model of "overconfidence" bias in the sense of Griffin and Tversky (1992). See also Kariv (2003) and Nöth and Weber (2002).

³⁴Values of $\alpha < 1$ correspond to under-weighting the signal, or "conservatism" bias, as discussed in Edwards (1968) and Camerer (1995, pp. 601-2). Although this latter kind of bias has less support in the experimental literature, it is sufficiently plausible that we choose not to assume it away.

³⁵From these equations, it is easy to see that for $\alpha > 1$ the learning process is faster as agents' choices depend more on their own signals, in the sense that the expected change in posterior is greater.

Before we report estimation results, let us provide some intuition for why this alternative model may better describe some features of the data. First, when $\alpha = 1$, QRE predicts that indifferent agents randomize uniformly, while the data show that 85% of the indifferent subjects follow their signals. This prediction is in line with $\alpha > 1$. Furthermore, when $\alpha > 1$, cascades take longer to start. For example, after two A choices the third decision maker need not choose A if she sufficiently overweighs her b signal. The base rate fallacy therefore provides one possible explanation for the prevalence of length zero crazes in our data set (see Figure 1).

The estimation results for the QRE-BRF model are reported in the second panel of Table B1. For all treatments, the BRF parameter, α , is significantly greater than 1. To test for significance we can simply compare the loglikelihood of the QRE-BRF model to that of the constrained model (with $\alpha = 1$) in the top panel. Obviously, the BRF parameter is highly significant.³⁶ Furthermore, the constrained model yields a significantly (at the 0.01 level) higher estimate of λ for all treatments.

5.2. Alternative Models: Non-Rational Expectations

It is possible that players update incorrectly because they do not have rational expectations. The QRE model implicitly assumes that λ is common knowledge. In particular, if players believed other players' λ were lower than it truly was, this could lead to a phenomenon that would be qualitatively similar to base rate neglect. Accordingly, we consider a model that allows for separate belief and action precision parameters, as proposed by Weizsäcker (2003). These different parameters are labelled λ_a (action lambda) and λ_b (belief lambda). That is, players choice probabilities follow the logit choice function with parameter λ_a but they believe that other players' choice probabilities follow a logit choice function with parameter λ_b . We call this the *non-rational expectations model*, or QRNE model.

³⁶For the pooled data the difference in loglikelihoods is nearly 200. A simple t -test also rejects the hypothesis that $\alpha = 1$, with a t -statistic of 14.6. Tests conducted for the AH data also reject the constrained model, with a slightly lower estimate of α .

The estimation results for the QRNE model are also given in Table B1 in Appendix B. While the two-parameter model performs significantly better than the QRE model, the increase in likelihood is not overwhelming. This is especially true when the model is combined with the base rate fallacy. In this case, the action and belief λ are virtually identical when estimated from the pooled data, and the increase in likelihood barely significant. A similar conclusion holds for the AH data, indicating the assumption of rational expectations is (approximately) valid.

5.3. Alternative Models: Cognitive Heterogeneity

As a final model, suppose there exists heterogeneity in the cognitive sophistication across players. For example, some players may behave completely randomly, while other players optimize against such behavior. Following Camerer, Chong, and Ho (2003), we extend this to allow for multiple levels of sophistication.³⁷ Specifically, level 0 players are random, level 1 believe all the other players are level 0, level 2 players believe all others are a mixture of level 0 and level 1, and so forth. The proportion of level k players is given by a Poisson distribution with parameter τ . Players are assumed to have truncated rational expectations, i.e. level k players believe all other players are a mixture of levels less than k with the mixture distribution given by the Poisson distribution, truncated at k . This is called the *cognitive hierarchy (CH) model*, see Camerer, Chong, and Ho (2003).

The presence of level 0 players will lead higher-level players to implicitly discount the information contained in the choices of their predecessors. This way the CH model can pick up some of the same features of the data as QRE. Furthermore, like QRE, the CH model is “complete” in the sense that it is consistent with any sequence of choices and signals. Hence we can obtain estimates of the τ parameter, without using QRE, see Table B1 in Appendix B.

The CH model is also estimated together with QRE to allow for comparison with QRE. All three models are then re-estimated including the BRF parameter, also for purposes of model

³⁷Stahl and Wilson (1995) explored a related model with levels of sophistication to study behavior in experimental games, but that model was different from the one considered here. See Camerer, Chong, and Ho (2003) for a discussion of the differences between the two models.

comparison. Note that the estimates for the combined QRE-BRF-CH model are stable across data sets and generally result in the highest likelihood. All three are significant factors, based on likelihood ratio tests, and leaving out any one of these factors changes the magnitudes of the other estimates.³⁸

5.4. Summary of Estimation Results

Table 10 presents the pooled estimates for all models. The estimation confirms our intuition that these models are alternative good explanations. However, the models are really quite different conceptually, and the BRF model is clearly a significant factor even if other behavioral factors are also present. The QRE-BRF model is simple and intuitively appealing, which is why we used it for simulation and comparisons with data (see Tables 4-8 and Figures 1-4). The QRE-CH-BRF model results in a slightly higher likelihood, but the model is conceptually harder and the effects on the descriptive statistics reported in the previous section are negligible.

6. Results III: Efficiency

We consider both informational efficiency and allocative efficiency. Regarding informational efficiency the relevant questions are: How well is the information from private signals aggregated? How close to 1 is the public belief about to the correct alternative? How does this vary with our treatment variables, q and T ? In contrast, allocative efficiency concerns a different set of questions: How frequently are actions correct? How does this change over time? And how does this change as a function of the informativeness of the signal?

³⁸The only unusual finding is that the estimate for τ is larger in magnitude than has been typically found in other settings. Camerer, Chong, and Ho (2003) report estimates in the range of 1.5 to 2.5, while our estimate in the combined model is 2.9 (with a standard error of 0.10). This appears to be due to an interaction between τ , λ , and α . The estimate of τ in the pure CH model is 1.9, and its estimate in the CH-QRE model (without BRF) is 2.5. Combining QRE and CH also leads to substantially larger estimates of λ . The reason for this is that both are rationality parameters that substitute for each other. The 0 types in the CH model absorb a lot of the randomness in the QRE model. In other words, the random behavior that can only be explained by 0 types in the CH model, is also explained by quantal response randomness. Hence we find relatively low values of either parameter if the models are estimated separately, but both increase significantly when the models are combined.

Model	λ	α	τ	λ_A	λ_B	logL
QRE	6.12 (0.14)					-3650
QRE-BRF	4.23 (0.11)	2.46 (0.10)				-3466
QNRE				6.32 (0.14)	4.48 (0.28)	-3636
QNRE-BRF		2.59 (0.12)		4.09 (0.12)	4.92 (0.33)	-3462
CH			1.91 (0.02)			-3648
QRE-CH	13.12 (0.75)		2.54 (0.08)			-3486
QRE-CH-BRF	7.69 (0.50)	1.81 (0.08)	2.90 (0.10)			-3411

Table 10. *Comparison of model estimates with our pooled data.*

6.1. Informational Efficiency

As shown in the theoretical section of the paper, in a QRE the public belief about the correct alternative converges to 1 with probability 1 as T approaches infinity. This process converges more slowly for the $q = 5/9$ treatments than for the $q = 6/9$ treatments. Of course, in any finite sequence, information cannot possibly reveal the correct alternative, because of noise in the signal generation process. Moreover, this noise in signal generation is compounded by strategic considerations that affect the social learning process.

We have three hypotheses about informational efficiency:

- H1. For each q , the public belief about the correct alternative is closer to 1 in the final period of the $T = 40$ treatments than in the $T = 20$ treatments.
- H2. For each t , the public belief about the correct alternative is closer to 1 in the $q = 6/9$ treatments than in the $q = 5/9$ treatments.
- H3. For all treatments, the average public belief about the correct alternative rises with t .

Since we do not observe beliefs directly, we use the theoretical QRE-BRF model together




Figure 2: *Estimated beliefs using the QRE-BRF model for all sequences in one of the ($q = 5/9$, $T = 20$) sessions.*

with the observed data to obtain estimated public belief paths.³⁹ This is done for each sequence. Using the pooled estimates $\lambda = 4.23$ and $\alpha = 2.46$, each sequence of action choices implies a unique public belief. This is illustrated in Figure 2, which shows the belief paths for all sequences in one of the $q = 6/9$ and $T = 20$ sessions. Here the horizontal axis represents the sequence of decisions, and the vertical axis the belief about the correct alternative. Each upward tick in the belief paths corresponds to a correct choice and each downward tick to an incorrect choice. Theoretically, for long enough sequences, the belief paths for almost all sequences should converge to 1.

Hypotheses 1-3 can be examined in several ways. The simplest way is to just average the public belief about the correct alternative across all sequences for a given treatment. This produces the four curves in the left panel of Figure 3. The middle and right panels depict simulated average beliefs using the QRE-BRF model and Nash model respectively. The curves are obviously consistent with the theoretical hypotheses. The right most panel shows that the difference between the two $q = 6/9$ treatments is caused by the particular signals drawn in these treatments. The comparison between the different q treatments is admittedly a weak test

³⁹Domowitz and Hung (2003) recently reported a social learning experiment using a belief elicitation procedure. We did not elicit beliefs for several reasons. For example, it introduces incentive problems, as noted by the authors.




Figure 3: *Estimated public beliefs about the true state by treatment (coded as in Figure 1). In the left panel, estimated beliefs are based on observed signals and decisions. The middle panel is based on the average of 100 QRE-BRF simulations of decisions, always using the same sequence of signals as in the experiment. The right panel shows the accuracy of Nash equilibrium decisions, using the same sequence of signals as in the experiment.*

since the paths are constructed using the theoretical model. That is, even if the sequences of signals and decisions were exactly the same for all sequences in $q = 6/9$ and $q = 5/9$ session, the $q = 6/9$ curves necessarily would lie strictly above the $q = 5/9$ curves. That said, the ordering also reflects a salient difference between our $q = 5/9$ and $q = 6/9$ data, namely that cascades fall apart more quickly, and are more often incorrect in the $q = 5/9$ data than in the $q = 6/9$ data (see Tables 5-8 of the previous section).

However, that the curves are increasing in t is not an artifact of the construction. That simply reflects the fact that there are more good cascades and fewer bad cascades toward the end of a session than toward the beginning. To summarize, we find strong support for H1 and H3 and somewhat weaker support for H2.

6.2. Allocative Efficiency

Allocative efficiency is quite a different story from informational efficiency, for at least two reasons. First, in contrast to beliefs, allocative efficiency is directly measured (by the proportion

of correct decisions), since both the state and the action of each individual is observed in the data. Second, full allocative efficiency will *not* be theoretically achieved in a quantal response equilibrium, even for arbitrarily large T .⁴⁰ After beliefs have converged to the true state, decisions keep fluctuating because of the stochastic nature of QRE, resulting in efficiency losses.⁴¹ At one extreme, when λ is close to 0, choice behavior is random and decisions are correct 50% the time, regardless of history, signal, or q . We take this as a plausible lower bound for allocative efficiency.⁴²

The hypotheses we have regarding allocative efficiency are based on Properties 9-11 of the logit QRE listed in section 2.3. First, average allocative efficiency will increase over time because expected beliefs converge monotonically to the true state. Also, allocative efficiency should be positively affected by signal informativeness, in three ways. There is the direct effect that more good signals are received with a higher q , but there are two indirect effects as well: with more informative signals, social learning is faster because actions are more informative, and conditional on being in a cascade, the cascade is more likely to be correct.⁴³ Because of these two indirect effects, there should be a difference in allocative efficiency in the different q treatments *controlling for the signals subjects receive*. Summarizing:

H4. The probability of a correct choice is increasing in t .⁴⁴

H5. The probability of a correct choice is higher for a correct than for an incorrect signal.

H6. Controlling for signal correctness, the probability of a correct choice is increasing in q .

⁴⁰If λ increased without bound as T increased, then full allocative efficiency may be possible, but here we are only considering QRE models with constant precision.

⁴¹The base rate neglect, at least as we have modelled it, is another source of inefficiency. Again, this contrasts with informational efficiency, where base rate neglect *speeds up* the learning process.

⁴²This is the lower bound in an aggregate analysis that looks at average efficiency over many sequences. The theoretical lower bound for any particular sequence is even lower, since it is possible for every action in a sequence to be incorrect. In fact this happens in two of our sequences, where a pure cascade on the wrong state starts at the very beginning.

⁴³A third minor effect going in the same direction is that with a higher q the posterior beliefs are, on average, further from $\frac{1}{2}$, so the expected payoff difference between a correct and incorrect action is generally increasing in q .

⁴⁴As pointed out in the theory section, this is true *ex ante* and conditional on receiving an incorrect signal, but not conditional on receiving a correct signal.

H7. The rate of increase of the probability of a correct choice as t increases is increasing in q .

The panels of Figure 4 show the time-dependence of decision accuracy for different treatments. The top row is ($q = 6=9$, $T = 20$), the second row is ($q = 6=9$, $T = 40$) the third row is ($q = 5=9$, $T = 20$) and the bottom row is ($q = 5=9$, $T = 40$). In each graph, the thick solid black line shows decision accuracy for all signals, the dashed red line for correct signals, and the thin blue line for incorrect signals. The panels in the left column gives Nash predictions, the middle column the actual decisions made in the experiment, and the right column QRE-BRF simulations, all based on the actual signals from the experiment. The graph shows moving averages: a point at time t represents the average decision accuracy between $t - 2$ and $t + 2$ for $3 \leq t \leq T - 2$.

In the Nash equilibrium, decision accuracy becomes independent of signals very quickly, reflecting the formation of cascades that never break. Decision accuracy for correct signals declines for a few rounds and then levels off. Decision accuracy for incorrect signals increases for a few rounds and then levels off. Unconditional decision accuracy increases for only the first few periods, before levelling out, because nearly all cascades are formed in the first 6 rounds, and never break.

This contrasts sharply with the dynamics in the actual data, and in the QRE-BRF simulations, where unconditional decision accuracy continues to increase, rather than stopping after the first four periods or so. Furthermore, there is a strong signal dependence that persists throughout the experiment, as the sequences pass through cycles of temporary cascades, which break and re-form. The decision accuracy for incorrect signals is always less than that for correct signals in both the actual data and the QRE-BRF simulations, in contrast to the Nash model. For incorrect signals (see the top-right panel) there is a clear and strong upward trend in decision accuracy, indicating that information continues to be aggregated throughout the experiment, resulting in later decision makers with incorrect signals making correct choices more often. There is only a very small, early, downward trend for decision makers with correct signals, because of the possibility of being in a wrong cascade. This levels out or even reverses sign later, because it is being offset by the fact that later cascades are more likely to be correct than earlier ones,

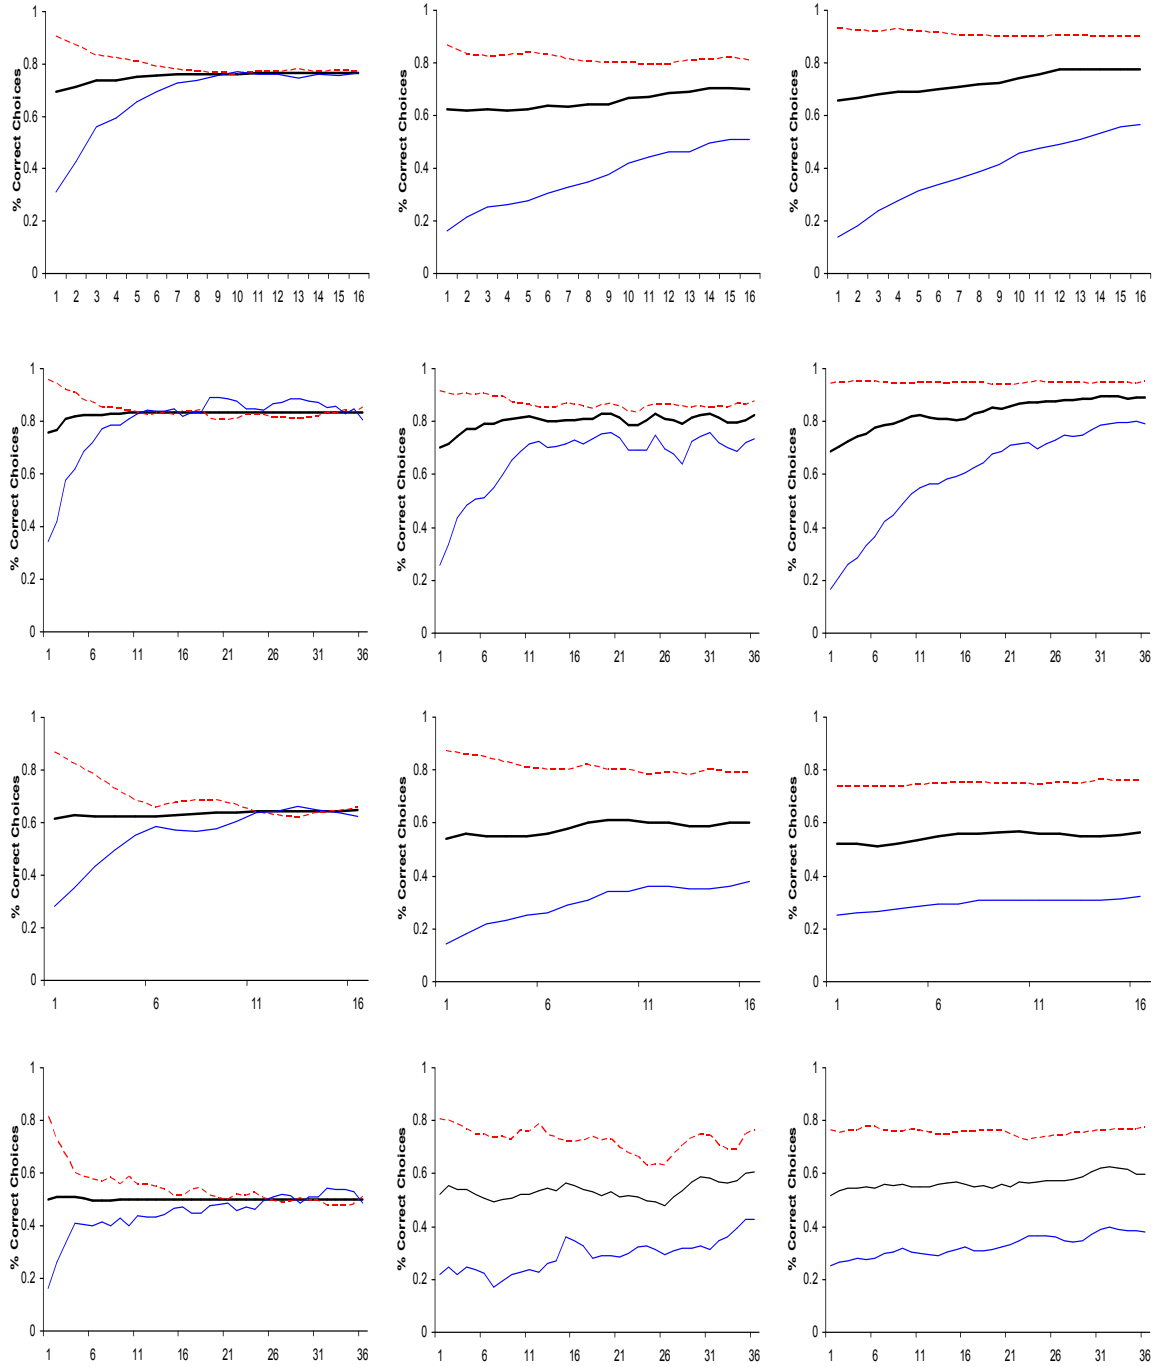


Figure 4: *Decision accuracy along the sequence of decision makers by treatment: ($q = 6/9, T = 20$) top row, ($q = 6/9, T = 40$) second row, ($q = 5/9, T = 20$) third row, and ($q = 5/9, T = 40$) bottom row. In each graph, the thick solid black line shows decision accuracy for all signals, the dashed red line for correct signals, and the thin blue line for incorrect signals. The left column gives Nash predictions, the middle column data, and the right column QRE-BRF predictions. The graphs show moving averages: a point at time t represents the average decision accuracy between $t - 2$ and $t + 2$ for $3 \leq t \leq T - 2$.*

Dependent Variable Correct Choice	Data	Simulation 1	Simulation 2
Constant	-1.57 (0.30)	-2.58 (0.33)	-3.27 (0.32)
q	1.26 (0.51)	2.52 (0.55)	4.51 (0.53)
t	0.021 (0.0025)	0.033 (0.0026)	0.019 (0.0025)
q * t	0.017 (0.0029)	0.016 (0.0031)	0.012 (0.0031)
signal * t	-0.037 (0.0029)	-0.049 (0.0031)	-0.027 (0.0030)
signal	1.76 (0.055)	2.23 (0.059)	1.71 (0.057)
match	0.0047 (0.0021)	0.0077 (0.0022)	-0.0015 (0.0006)
# obs	8760	8760	8760
logL	-4620	-4021	-4178

Table 11. *Probit estimation of the effects of q and t on efficiency.*

because of self-correction. The net effect is almost neutral, as reflected by the curves in the top-middle panel. The decision accuracy for incorrect signals is always less than that for correct signals in both the actual data and the QRE-BRF simulations, in contrast to the Nash model.

For a more formal test of H4-H7, we conduct a Probit regression with six independent explanatory variables: t , q , $q * t$, signal, signal * t , match. Signal is dummy variable that takes on the value of 1 if the signal is correct. The variable $q * t$ is an interaction of signal informativeness and time period,⁴⁵ which, according to H4, should be positive. The variable signal * t is an interaction between time and signal correctness. From H1, the effect of t on decision accuracy should be positive only for incorrect signals, with a possible small negative effect for correct signals. Match is a variable that is included to control for possible experience effects. Notice that we do not include T in the regression, because the theory does not predict any effect except through the variable t .

Table 11 shows the estimated coefficients in column 2, with standard errors in parentheses. All coefficients have the expected sign and are highly significant. These results deserve closer inspection for at least two reasons. First, the regression is not based on any kind of structural model of decision making. Second, there are obvious dependencies in the data, and un-modelled

⁴⁵Here $q * t$ equals 0 if $q = 5/9$ and $q * t$ equals t if $q = 6/9$.

sources of error, including quantal response errors and variation in signal sequences.

To check the robustness of our findings and to check it against the theoretical model, we generated two simulated data sets based on the QRE-BRF model, using the pooled estimates $\lambda = 4.23$ and $\alpha = 2.46$. The first of these simulations uses the same signal sequences as in the laboratory experiment but decisions are generated by the QRE-BRF model. The second simulation uses a completely new draw of signal sequences. The Probit estimations based on the simulated data sets are reported in columns 3 and 4 of Table 11. While there are some small differences in magnitude, all coefficients of theoretical interest are significant with the correct sign.⁴⁶ Note that the log-likelihoods for the simulated data are higher than for the real data. This is likely caused by the fact that the simulations assume homogeneous agents, while we would expect some heterogeneity to be present in the laboratory data.

Also of interest is the question of whether decision accuracy is improved or hampered by the stochastic choice and the base-rate neglect. On the positive side, information is better, but on the negative side, subjects are making decision errors and are also subject to decision errors. From the figures, we can see that the negative effects dominate in the early periods, while the positive effects of information aggregation dominate in later periods. In the very long run (T going to infinity) beliefs in the QRE-BRF model converge to the true state, so that private beliefs and public beliefs will be the same, independent of the signal. Hence the BRF part of the model plays no role in the limit, nor does signal informativeness. Moreover, Using our pooled estimate of $\lambda = 4.23$, decision accuracy converges asymptotically to 0.99.

7. Conclusion

This paper reports the results of an information cascade experiment with two novel features:

⁴⁶The only notable difference is the experience variable, which is not significant in the simulation using a new batch of signal sequences, suggesting that its significance was spurious, due to more favorable order of signals in later matches. (Indeed, there is no reason that experience should have had a significant effect in the first simulation.) In any case, the magnitude of the experience effects, to the extent they may possibly not be spurious, is negligible.

longer sequences of decisions and systematic variation of signal informativeness. According to standard game theory, neither of these treatments should be interesting, and neither should produce significantly different results. We find, however, that both of these treatment effects are strong and significant, with important implications for social learning, information aggregation, and allocative efficiency.

The longer sequences have several effects. First, they have fewer permanent cascades, more temporary cascades, more repeated cascades, more reversed cascades, and more self-corrected cascades. In contrast, standard theory predicts that longer sequences will have more permanent cascades, and that temporary, repeated, reversed, and self-corrected cascades never occur. The observed dynamics are highly dependent on signal informativeness, again in contrast to Nash predictions. Relatively uninformative signals lead to less stable dynamics, in the sense that cascades are much shorter, more frequent, and reverse more often. These subtle but important features of the dynamics are impossible to detect in the short sequences employed in previous experiments (Anderson and Holt, 1997).

To explain the observed features of the dynamics and the dependence on signal informativeness, we consider the logit quantal response equilibrium (QRE). In addition, we apply QRE as a structural model to estimate base rate neglect and to test for heterogeneity in levels of rationality. We find both to be significant factors in observed behavior. In particular, subjects tend to overweight their signals, or, alternatively, underweight the public prior generated by past publicly-observed choices (base rate neglect).

Our experimental results confirm a basic property of the QRE with profound implications in this context: deviations happen and their likelihood is inversely related to their cost. This property implies that cascade breakers more often than not hold contrary signals, and, hence, that deviations from cascades are highly informative. Learning continues in a QRE even after a cascade forms or breaks, and temporary, repeated, reversed, and self-correcting cascades arise as equilibrium phenomena. While standard cascade theory predicts that learning ceases after a few initial decisions, our data show that information is continuously being aggregated, providing

evidence for the QRE prediction that ultimately the truth will prevail.

A. Appendix: Proof of Proposition 1

Proofs of (i) and (ii): The proof of (i) is by induction. Recall that $p_1 = \frac{1}{2}$, so we only need to show that $0 < p_t < 1$ implies $0 < p_t^- < p_t < p_t^+ < 1$. Equation (2.4) can be expanded as

$$p_t^+ = \frac{qp_t(1 - F_\lambda(1 - 2\pi_t^a)) + (1 - q)p_t(1 - F_\lambda(1 - 2\pi_t^b))}{(qp_t + (1 - q)(1 - p_t))(1 - F_\lambda(1 - 2\pi_t^a)) + ((1 - q)p_t + q(1 - p_t))(1 - F_\lambda(1 - 2\pi_t^b))},$$

with $1 > \pi_t^a > \pi_t^b > 0$ defined in (2.1) and (2.2), and $F_\lambda(x) = 1/(1 + \exp(-\lambda x))$ the logistic distribution with parameter λ and support $(-\infty, \infty)$. Since $\frac{1}{2} < q < 1$ and $0 < p_t < 1$ by assumption, the denominator exceeds the numerator: $p_t^+ < 1$. A direct computation shows

$$p_t^+ - p_t = \frac{p_t(1 - p_t)(2q - 1)(F_\lambda(1 - 2\pi_t^b) - F_\lambda(1 - 2\pi_t^a))}{(qp_t + (1 - q)(1 - p_t))(1 - F_\lambda(1 - 2\pi_t^a)) + ((1 - q)p_t + q(1 - p_t))(1 - F_\lambda(1 - 2\pi_t^b))},$$

which is strictly positive because $\pi_t^a > \pi_t^b$. The proof that $0 < p_t^- < p_t$ is similar. *Q.E.D.*

Proofs of (iii) and (iv): Let $\ell_t = (1 - p_t)/p_t$ denote the likelihood ratio that A is correct. For all $t \in \mathcal{T}$ we have

$$E(\ell_{t+1} | \omega = A, \ell_t) = \ell_t,$$

i.e. the likelihood ratio constitutes a martingale, a basic property of Bayesian updating. Note that p_t is a strictly convex transformation of the likelihood ratio ($p_t = (\ell_t + 1)^{-1}$), so

$$E(p_{t+1} | \omega = A, p_t) = E((\ell_{t+1} + 1)^{-1} | \omega = A, \ell_t) > (E(\ell_{t+1} + 1 | \omega = A, \ell_t))^{-1} = p_t,$$

by Jensen's inequality and the fact that $\ell_t^+ \neq \ell_t^-$, see (ii). We sketch the proof of (iv), see Goeree, Palfrey, and Rogers (2003) for details. First, limit points of the stochastic belief process $\{p_t\}_{t=1,2,\dots}$ have to be invariant under the belief updating process. But (ii) implies that $p_{t+1} \neq p_t$ when $p_t \neq \{0, 1\}$, so the only invariant points are 0 and 1. Next, the Martingale Convergence Theorem implies that ℓ_t converges almost surely to a limit random variable ℓ_∞ with finite expectation. Hence, $\ell_\infty < \infty$ with probability one, which implies that $p_\infty > 0$ with probability one and p_t thus converges to 1 almost surely. *Q.E.D.*

B. Appendix: Estimation Program and Results

Below we assume the experimental data are stored in an $MT \times 2$ matrix called "data"; every T rows correspond to a single sequence, or run, with a total of M runs, the first column contains subjects' signals and the second column subjects' choices. The coding is as follows: A choices and a signals are labelled by a 1 and B choices and b signals by a 0. The outcome of the procedure is the log-likelihood for a single treatment (i.e. with a fixed precision, q , and fixed length, T) although it is easy to adapt the procedure to deal with pooled data.⁴⁷

```

PROC loglikelihood( $\lambda$ );
LOCAL logL,signal,choice,m,t,p, $\pi^a$ , $\pi^b$ ,P(A|a),P(A|b),P(B|a),P(B|b), $p^+$ , $p^-$ ;
  logL=0; m=1;
  DO WHILE m<=M;
    p=1/2; t=1;
    DO WHILE t<=T;
       $\pi^a$ =qp/(qp+(1-q)(1-p));
       $\pi^b$ =(1-q)p/((1-q)p+q(1-p));
      P(A|a)=1/(1+exp( $\lambda$ (1-2 $\pi^a$ ))); P(B|a)=1-P(A|a);
      P(A|b)=1/(1+exp( $\lambda$ (1-2 $\pi^b$ ))); P(B|b)=1-P(A|b);
       $p^+$ =(pqP(A|a)+p(1-q)P(A|b))/((pq+(1-p)(1-q))P(A|a)+(p(1-q)+(1-p)q)P(A|b));
       $p^-$ =(pqP(B|a)+p(1-q)P(B|b))/((pq+(1-p)(1-q))P(B|a)+(p(1-q)+(1-p)q)P(B|b));
      signal=data[(m-1)T+t,1]; choice=data[(m-1)T+t,2];
      IF signal==1 AND choice==1; p= $p^+$ ; logL=logL+ln(P(A|a)); ENDIF;
      IF signal==0 AND choice==1; p= $p^+$ ; logL=logL+ln(P(A|b)); ENDIF;
      IF signal==1 AND choice==0; p= $p^-$ ; logL=logL+ln(P(B|a)); ENDIF;
      IF signal==0 AND choice==0; p= $p^-$ ; logL=logL+ln(P(B|b)); ENDIF;
      t=t+1;
    ENDO;
    m=m+1;
  ENDO;
  RETP(logL);
ENDP;

```

⁴⁷The procedure is simple because information cascade experiments concern individual decision-making environments, not games, so there is no need to solve fixed-point equations to compute the QRE.

	Our Data				AH Data	
	p = 5/9		p = 6/9		Pooled	p = 6/9 T = 6
	T = 20	T = 40	T = 20	T = 40		
# obs	2320	2240	1800	2400	8760	270
QRE						
λ	11.36 (0.42)	7.19 (0.32)	4.38 (0.18)	4.69 (0.19)	6.12 (0.14)	6.62 (0.72)
logL	-981.0	-1181.4	-682.0	-634.0	-3650.3	-79.0
QRE-BRF						
α	2.33 (0.18)	2.97 (0.36)	2.01 (0.16)	1.67 (0.16)	2.46 (0.10)	1.51 (0.19)
λ	7.07 (0.45)	3.68 (0.32)	3.47 (0.16)	4.09 (0.18)	4.23 (0.11)	5.90 (0.76)
logL	-930.7	-1147.6	-653.0	-622.5	-3466.0	-74.5
QRNE						
λ_A	14.45 (0.62)	9.82 (0.49)	5.16 (0.23)	4.74 (0.18)	6.32 (0.14)	7.93 (0.92)
λ_B	4.07 (0.37)	1.86 (0.18)	1.86 (0.18)	3.45 (0.33)	4.48 (0.28)	3.78 (0.66)
logL	-947.7	-1156.3	-660.8	-627.9	-3636.6	-74.7
QRNE-BRF						
α	3.24 (0.34)	2.64 (0.41)	1.82 (0.24)	1.54 (0.16)	2.59 (0.12)	1.75 (0.23)
λ_A	5.43 (0.44)	4.06 (0.51)	3.65 (0.27)	4.19 (0.20)	4.09 (0.12)	5.35 (0.83)
λ_B	12.56 (1.87)	3.25 (0.47)	2.93 (0.52)	3.40 (0.34)	4.92 (0.33)	15.68 (10.58)
logL	-925.6	-1147.1	-652.5	-620.5	-3462.8	-73.3
CH						
τ	1.67 (0.06)	1.24 (0.04)	1.96 (0.04)	2.82 (0.03)	1.91 (0.02)	2.20 (0.22)
logL	-964.0	-1180.4	-694.3	-656.6	-3648.1	-77.1
QRE-CH						
τ	2.00 (0.11)	1.67 (0.14)	2.52 (0.20)	3.63 (0.23)	2.54 (0.08)	2.44 (0.25)
λ	26.45 (3.31)	16.99 (2.33)	7.07 (0.86)	6.23 (0.51)	13.12 (0.75)	28.34 (14.16)
logL	-940.7	-1162.1	-672.3	-632.2	-3486.3	-74.3
QRE-CH-BRF						
α	1.91 (0.16)	2.67 (0.27)	1.90 (0.16)	1.50 (0.15)	1.81 (0.08)	1.36 (0.32)
τ	2.56 (0.23)	3.23 (0.73)	3.70 (0.73)	3.80 (0.28)	2.90 (0.10)	3.54 (2.28)
λ	12.77 (1.80)	4.50 (0.73)	3.97 (0.47)	5.21 (0.45)	7.69 (0.50)	7.47 (3.96)
logL	-911.9	-1144.3	-652.0	-616.1	-3411.3	-73.9

Table B1. Parameter estimates for the different models with standard errors in parentheses.

References

- Anderson, L. (2001) "Payoff Effects in Information Cascade Experiments," *Economic Inquiry*, **39**:609-15.
- Anderson, L. and C. Holt. (1997) "Information Cascades in the Laboratory," *American Economic Review*, **87**: 797-817.
- Banerjee, A. (1992) "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, **107**, August, 797-817.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, **100**, October, 992-1026.
- Camerer, C. (1995) "Individual Decision Making," in *The Handbook of Experimental Economics*, ed. J. Kagel and A. Roth, Princeton University Press: Princeton, 587-703.
- Camerer, C., T. Ho, and J. Chong (2004) "A Cognitive Hierarchy Theory of Games," Working Paper, California Institute of Technology, Pasadena, CA.
- Domowitz, J. and A. Hung. (2003) "Homogeneous Actions and Heterogeneous Beliefs: Experimental Evidence on the Formation of Information Cascades." Carnegie Mellon University, Working Paper.
- Edwards, W. (1968). "Conservatism in Human Information Processing," in *Formal Representation of Human Judgement*, ed. B. Kleinmuntz, New York: Wiley, 17-52.
- Goeree, J., T. Palfrey, and B. Rogers (2003), "Social Learning with Private and Common Values," *Social Science Working Paper 1187*, California Institute of Technology, Pasadena, CA.

- Griffin, D. and A. Tversky (1992) "The Weighting of Evidence and the Determinants of Confidence" *Cognitive Psychology*, **24**, 411-35.
- Kahneman, D. and A. Tversky (1973) "On the Psychology of Prediction," *Psychological Review*, **80**, 237-51.
- Kariv, S. (2003), "Overconfidence and Information Cascades," Draft, University of California, Berkeley, CA.
- McKelvey, R. and T. Palfrey (1995), "Quantal Response Equilibria in Normal Form Games," *Games and Economic Behavior*, 10(July):6-38.
- McKelvey, R. and T. Palfrey (1998), "Quantal Response Equilibria in Extensive Form Games," *Experimental Economics*, 1:9-41.
- Nöth, M. and M. Weber (2003) "Information Aggregation with Random Ordering: Cascades and Overconfidence," *Economic Journal*, **113**, 166-89.
- Plott, C. and A. Hung, "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions," *American Economic Review*, **91**: 1508-20.
- Stahl, D. and P. Wilson (1995), "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10(July):218-54.
- Weizsäcker, Georg (2003), "Ignoring the Rationality of Others: Evidence From Experimental Normal Form Games," *Games and Economic Behavior*, 44(July):145-71.