LEARNING AND STATUS IN SOCIAL NETWORKS

Thesis by

Brian W. Rogers

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2006

(Defended April 20, 2006)

Abstract

The patterns in which individuals interact with each other have important consequences for determining outcomes across a wide variety of contexts. One notable phenomenon that relies on these interactions is social learning. Social learning occurs when asymmetrically informed individuals have the opportunity to observe the choices of others and to incorporate this information when making their own choices. Under certain assumptions this process leads to information cascades in which the ability to learn from others ceases before any sound conclusion is reached. To the extent this theory is valid, the implications for information aggregation are quite negative. However, casual empiricism suggests that such an inefficiency is unlikely: many people making similar decisions over time are very unlikely to be continually wrong. The first two chapters address this discrepancy.

Experiments that implement a standard social learning paradigm are reported. A novel feature of the data is that we examine long sequences of decisions (up to forty) and study the effects of different signal qualities. In contrast to standard equilibrium predictions, a pattern of cascade formation, collapse, and re-formation is routinely observed. The main implication of these dynamics is that learning continues throughout the sequence of decisions, so that the truth is nearly revealed. Quantal Response Equilibrium, augmented to allow for non-rational belief updating in the form of base-rate neglect, explains nearly all the features of the data.

One assumption underlying most studies of observational learning is that the order in which individuals choose is exogenously fixed. In many applications, however, strategic considerations play a major role in determining the timing of decisions. To understand how timing issues impact the ability to learn from others' decisions, I study a model in which decision times are strategic variables and individuals have heterogeneous signal qualities. The main finding is that with two players, the player with better information announces first in (the unique) equilibrium. Consequently,

both players make the same decision, as in a herd, but because of the sorting effect, the outcome is informationally efficient. It is also shown that the game ends immediately as the time interval vanishes. In comparison to the standard exogenous sequence assumption, welfare is always higher under strategic timing. When there are many players, a herd forms immediately, and it is always on the correct action because early decision makers have the best information, and their choices reveal the true state.

We next study a model that addresses how social networks form in strategic settings. Individuals allocate a budget of resources across others, creating directed, continuously valued links. In many applications, a (directed) link confers benefit to both individuals involved in the link. By separating this benefit flow into "giving" and "taking" components, we are able to study the implications for efficiency. The main finding is that inefficiencies at equilibrium are due only to the giving of benefit. We also relate network structures to underlying heterogeneity of players.

The final chapter analyzes large-scale social networks. The main question concerns how correlation patterns in links across individuals affect the diffusion of a virus or rumor. A surprise is that in all of the simulations considered, the Susceptible-Infected-Susceptible (SIS) model behaves identically on networks with varying correlation patterns. These are the first results in this context that isolate the role of correlation structures.

Contents

List of Figures

List of Tables

Chapter 1

Summary

Social interactions are a fundamental ingredient of everyday life, affecting many of the decisions we make, ranging from whom to vote for in political elections, to what kind of personal computer to buy, from how to invest financial assets, to what job opportunities to explore. Despite the prevalence and importance of such examples, until recently the vast majority of economics research—with important exceptions—ignored most aspects of social interactions and, consequently, their possible implications for outcomes in these settings. This thesis, presented in the following chapters, addresses a series of questions relating to the economic incentives and consequences of social interactions.

The work contained herein adds to the understanding of social interactions from a variety of perspectives. There are at least two dimensions along which the chapters can be categorized. The first dimension is methodological. This thesis takes the perspective that, when feasible, the most powerful way to analyze a problem is through the combination of rigorous formal modeling and sound empirical investigation. The next chapter's analysis is based on data collected from a series of laboratory experiments. Experimental economics is quickly becoming one of the major paradigms of modern research, and for good reason: laboratory experiments give researchers the ability to control the environment in which decision makers operate in ways that are very useful, and usually impossible in the field.

The current application to social learning is a good example. As described below, we will study a situation in which individuals observe a number of others making a choice before having to make a similar choice themselves. It is thus important to be able to assess what information individuals have when making their choices, including their prior beliefs over an uncertain outcome, and which other individuals they observed, as well as the information obtained by those other individuals at the times of their decisions, and so forth. The use of laboratory experiments allows researchers to obtain very high controls over these quantities, and thereby facilitates the study of how people learn from the observations of others.

After Chapter 2, this thesis takes a formal mathematical approach to analyzing social interactions in Chapters 3 and 4. This work is quite complimentary with the experimental portion of the thesis, and together they provide some powerful conclusions. Whereas the experimental data helps answer the positive questions of "How do individuals behave, and what are the consequences for their welfare and the long run behavior of their beliefs?," the formal analysis is better suited to identifying situations in which welfare tends to be higher and seeks to provide ways in which to align the incentives of individuals so that they will find it optimal to make the choices that are socially welfare-improving.

The final chapter uses computer simulations to explore a diffusion model on large social networks. The particular question has proven very difficult to analyze analytically, and the scale of the problem is prohibitive to running experiments. It is in these situations that simulations can reveal regularities that would otherwise be elusive.

One of the reasons that the combination of experiments, simulations, and mathematical modeling is useful is that the approaches are complimentary. In addition to the fact that the methodologies are better suited to addressing different kinds of questions, a good theory is useful in designing experiments that will be the most informative, and good experiments give insights into the determinants of human behavior that should be incorporated into theoretical models. This interplay is perhaps best exemplified in Chapter 2 in which a long-run theoretical prediction motivated the design of the experiments. The same theory also makes short-run predictions, which are testable using our data. To the extent that the basic theory can not explain the data, additional behavioral factors are identified that, when incorporated, are able to better explain the data, and therefore generate a more adequate theory of behavior. Incorporating these additional motivations into the theory requires a re-assessment of its predictions, which is carried out to interesting conclusions.

The second dimension that these chapters can be categorized along is substantive. That is, the chapters address fundamentally different issues related to the phenomena of social interactions. The next two chapters are concerned with the idea of *social learning*, in which individuals learn through the observation of others' decisions before making their own decision. The last two chapters deal with *social networks* in which the main objects of study are the patterns in which individuals are connected to each other and the consequences this structure has for the benefits individuals can obtain from each other. The remainder of this summary explains in more detail the goals of the subsequent chapters and how they are related to each other.

Clearly, one function of social interactions that could have economic consequences is the exchange of ideas and information. The flow of information across individuals can have large effects in many contexts, including fashion, fads, purchases of consumer goods, product differentiation, and investments. In each of these examples, the ability to observe others making a particular decision can affect the optimal decision of the observer. The reason is very simple: in a common value framework, where individuals all share the same interests but face aggregate uncertainty, observing another individual's choice communicates what action that individual viewed as the optimal choice. If others have the same preferences (or preferences that differ in certain known ways), then the fact that one person found a particular choice optimal is valuable information to others. For this reason, when making decisions individuals should pay attention to how others have behaved in similar situations and to what information those individuals based their decisions on.

A first attempt at modeling the dynamics of choices and beliefs in such a context works as follows (see Bikchandani et al. (1992)). There is a random state variable, the outcome of which is unobserved, but there is a common prior distribution shared by all. Individuals all have the same preferences, which are such that the optimal action depends on the realization of the state variable. Each individual receives an informative signal about the state variable and must make a decision from a specified set of possibilities in an exogenously given sequence. The social learning element is that each individual observes the entire history of decisions that precedes them (but not the signals of those individuals).

Under certain symmetry assumptions, the equilibrium outcomes of this game have a very negative information aggregation property. In particular, there is no equilibrium in which anyone learns more about the outcome of the state variable than the amount of information contained in any two private signals. Thus even in arbitrarily large populations, with many people observing the actions of many others, hardly any learning takes place. In other words, the mechanism of social learning described by this model is extraordinarily inefficient.

The next two chapters aim to identify mechanisms through which this negative result may be alleviated or even reversed. The data presented in Chapter 2 are completely inconsistent with standard equilibrium predictions. The proposed explanation relies on Quantal Response Equilibrium (see McKelvey and Palfrey (1995, 1998)). In QRE, players systematically deviate from best responses by playing suboptimal strategies with positive probabilities that are decreasing in the expected cost of the deviation. The model is closed by assuming that the statistical choice frequencies generated by the deviations are correctly predicted by other players so that the model has the same equilibrium flavor as Nash equilibrium. QRE both explains many of the features of our data and makes a long-run prediction about social learning: that eventually the sequence of choices will perfectly reveal the outcome of the state variable. The fact that our data is consistent with the short-term implications of QRE provides strong evidence that the actual dynamics of this social learning paradigm imply a positive information aggregation result in large populations.

The main message of Chapter 2 is that the random element to choices inherent in actual human decision making provides a mechanism for much more information to be transmitted than under the classical assumption of purely optimal behavior. Chapter 3 describes a second mechanism that generates a higher level of information aggregation than the standard model. The idea is to use time as a screening device so that when an individual makes a decision, the timing of the announcement, in addition to the content, transmits useful information to observers. There is good reason to believe that the timing of decisions is informative. In many settings, the time at which an agent makes her decision affects her payoff, and so modeling the timing of the decision as a choice variable allows individuals to respond to differences in their beliefs about the state variable through adjusting the time at which they decide. The benefit of delaying a decision is the possibility of observing another's decision, which can be informative. The cost of delay essentially comes from explicit temporal preferences so that making a correct decision is more valuable the sooner it is made. Once individuals are thought of as having control over the timing of their decisions, one should expect them to balance these effects optimally.

In the case where some individuals are better informed than others, the endogenous timing of choices reveals important information about the quality of information possessed by decision makers. When there are exactly two players, a unique equilibrium is identified in which the better informed individual announces first, and the other individual mimics her decision. In other words, a herd forms in which the decisions of both individuals are identical. But in contrast to the classical model, which assumes a fixed order of decisions, the herd here is informationally efficient in the sense that the decisions coincide with those of the first-best outcome, in which individuals pool their information before announcing. Still, the outcome is not perfectly efficient, because time costs are incurred in order to identify the optimal decisions. However, even when the delay costs are taken into account, average welfare is higher in the case where individuals choose the time of their decisions. A second result pertaining to large populations shows that with sufficiently many players, the outcome of the state variable is learned very quickly, and an arbitrarily large proportion of the population is able to obtain this information prior to announcing. Thus, the main point from Chapter 3 is that the timing mechanism provides a means to improve the information aggregation properties of social learning.

It is important to understand that the results from Chapters 2 and 3 depend on the fact that everyone sees the entire history of decisions that have occurred prior to their own announcements. When this assumption is relaxed, the results can change dramatically. See Choi, Gale and Kariv (2005), Gale and Kariv (2003), and Celen and Gale (2004) for models that address this question in a number of interesting ways. One way to think about the information structure of the social learning paradigm is through the notion of a social network. In this framework, individuals are identified as nodes in the network, where a link from one node to another represents the idea that the individual represented by the first node can observe the decision of the second individual. Since every individual is capable of observing the actions of every other individual, one way to represent the standard social learning model is through a complete network–one in which every possible link is present. Of course, this is a very special kind of network. In general, only a subset of links would exist in the network, representing the notion that some individuals are not in contact with each other and can not observe each other's actions, while others are in contact.

The concept of modeling observation patterns with network models is new in economics, although related ideas have been present in the sociology literature for decades. Moreover, we seek not just an understanding of how a particular network structure affects the outcomes of a particular interaction, but also an understanding of what kinds of networks structures will form in different kinds of settings. That is, one would like to be able to model the choices of what links to form from an understanding of the relative costs and benefits of maintaining different sets of links. Since a general treatment of the incentives facing agents in networking situations has not been developed, Chapter 4 takes a step towards understanding what kinds of communication structures are likely to form among individuals strategically choosing whom to interact with.

Since we are interested in modeling networking choices across a range of applications—including, but not limited to, social learning—Chapters 4 and 5 consider social networks from two different points of view. Chapter 4 explores a model where the benefits individuals derive from their relationships to others are described abstractly. This approach has several advantages. First, the relative simplicity of the model allows for sharp analytical results characterizing the efficiency of equilibrium networks. This is an important question because identifying situations in which equilibrium choices tend to produce socially inferior outcomes allows one to focus on situations that could benefit from intervention. Second, the utility specification is general enough to capture features shared by many social networking applications. The main ideas are (i) that maintaining connections has (an opportunity) cost, (ii) that individuals benefit from those to whom they connect, (iii) that some individuals are inherently more valuable, and (iv) that being well-connected is another source of value.

Thinking about the social learning model, it is clear that observation is not necessarily a symmetric relationship. That is, one person could observe another without the reverse relationship being possible. More generally, it is possible for one individual to receive benefits from another without any benefits flowing in the opposite direction. The model in Chapter 4 analyzes a setting in which this directed nature of relationships is a key part of the analysis. Two cases are considered. In the first case, individuals decide whom to derive benefits from, and in the second case, they decide whom to pass benefits to. The main result is that, while in the first case the equilibrium networks are necessarily efficient for the group, in the second case they are essentially never efficient. In this sense, the source of inefficiency in contexts where benefits flow in both directions is isolated as being due to benefits flowing in one particular direction.

Whereas the incentive structure and its implied information requirements are best suited for small populations, Chapter 5 explicitly considers the role of large-scale social networks. The networks studied are chosen because they exhibit the full spectrum of stylized facts that have been shown to characterize social networks. One of the primary roles of social networks is to provide a means to pass things–information, goods, ideas, behaviors–among a large group of people. However, the same networks also serve to allow disease and viruses to propagate. These networks typically have interesting correlation patterns in the placement of links between individuals, and we aim to study the effect of those patterns on diffusion. Although there is intuitive reason to expect such patterns to matter, the surprising result is that there is little if any evidence to suggest the correlation patterns have an effect.

Although the issues explored in the following chapters are clearly distinct, they are also inter-related. There are a great many aspects of social interactions to be studied, and this thesis represents my enquiries into some of the issues that I have found most interesting.

Chapter 2

Self-Correcting Information Cascades

2.1 Summary

This chapter reproduces the work of "Self-Correcting Information Cascades," a paper written jointly with Jacob K. Goeree, Thomas R. Palfrey, and Richard D. McKelvey. The part of the work conducted by the present author cuts across all aspects of the paper and may be summarized as follows. The idea of conducting the experiments reported in this chapter is due to McKelvey and Palfrey. The present author helped identify optimal parameters under which the theoretical framework would be implemented in the laboratory. Second, he shared in the work of actually running the experiments at CASSEL (UCLA) and SSEL (Caltech). Next, the theoretical background outlined in this chapter is based largely on "Social Learning with Private and Common Values" by Goeree, Palfrey, and Rogers (2006). The authors split, more or less equally, the ideas on how to analyze the data, what body of theoretical models to consider, and the mathematical programs to estimate the models considered.

2.2 Introduction

An information cascade arises when a sequence of imperfectly informed decision makers, each of whom observes all previous decisions, has reached a point after which all future decision makers will rationally ignore their private information. Hence, learning ceases as subsequent decision makers infer nothing new from observing any of the actions. Information cascades are predicted to occur, possibly after very few decisions, despite the wealth of information available and despite the common interest of all decision makers (Banerjee 1992, Bikhchandani et al. 1992). 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 bank runs, technology adoption, mass hysteria, and political campaigns.

We conducted laboratory experiments with very long sequences of decision makers in canonical social learning environments. The data is examined and analyzed through the lens of quantal response equilibrium (QRE), which makes systematic predictions about the long-run dynamics of choice behavior, beliefs, and efficiency. Some of these predictions are essentially the opposite of Nash equilibrium. Because of the complicated dynamics implied by QRE, a careful test of many of these properties of QRE demands the observation of long sequences. In addition, we vary the informativeness of individuals' signals, which systematically affects the observable properties of QRE dynamics.

The QRE approach to the analysis of data enables two additional innovations. First, using a Logit equilibrium error structure we are able to structurally estimate a parametric model of base-rate neglect and a cognitive hierarchy model of strategic sophistication. The existence and magnitude of judgement fallacies in these environments have important systematic implications about efficiency and dynamics. Second, this structural estimation approach yields estimates of the entire trajectory of public beliefs, for each sequence in the experiment. That is, the belief dynamics can be estimated indirectly without eliciting beliefs from the decision makers.

We use the simplest possible social learning environment in our experiment because QRE makes especially crisp predictions in these environments, enabling relatively straightforward tests of the predictions while at the same time simplifying the structural estimation procedure. 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 that 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 equilibrium. In fact, regardless of T, the equilibrium beliefs of all decision makers are confined to an interval centered around 1/2.

The need for an alternative theory of behavior in these environments is obvious from looking at data from short decision sequences, such as those reported in Anderson and Holt (AH) (1997).[1] In that experiment, cascades are observed; however some action choices are inconsistent with Nash equilibrium given the realized signals, and many subjects exhibit such behavior. For example, Anderson and Holt (1997) 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. Such deviations become even more pronounced in the experiments reported below where 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

[1] Indeed AH use a recursive version of Logit equilibrium to describe their data.

	q = 5/9 T = 20	q = 5/9 T = 40	q = 6/9 T = 20	q = 6/9 T = 40
number of sequences	116	56	90	60
percentage with pure cascades	4.3%	0.0%	13.3%	13.3%
percentage without cascades	0.0%	0.0%	0.0%	0.0%
percentage with broken cascades	95.7%	100.0%	86.7%	86.7%

Table 2.1: *Percentages of (broken) cascades in our data.*

observe cascades arising in 100% of the sequences according to the theoretical model of Bikhchandani et al. (1992). However, with $T = 40$, for instance, a cascade arises and persists in only 8 out of 116 sequences ($< 7\%$).

Table 2.1 gives an illustration of a few ways that the standard theory fares badly. At a minimum, a plausible 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 2.2. Indeed, when a break occurs, the observed frequency with which the received signal was contradictory is 83%.[2] 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.

The introduction of a random component in QRE ensures that all paths can be reached with positive probability, so Bayes' rule places restrictions on future rational

[2]When 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		
	Decision \ Signal	Confirming	Contrary	Decision \ Signal	Confirming	Contrary
q=5/9	Don't Break	48.5%	26.7%	Don't Break	45.0%	23.5%
	Break	3.1%	21.7%	Break	7.0%	24.5%
	# obs = 1305			# obs = 1234		
	Decision \ Signal	Confirming	Contrary	Decision \ Signal	Confirming	Contrary
q=6/9	Don't Break	53.2%	28.2%	Don't Break	59.8%	31.4%
	Break	2.4%	16.2%	Break	1.6%	7.2%
	# obs = 1080			# obs = 1923		

Table 2.2: *Frequency of confirmatory/contrary signals when cascades are (not) broken.*

inferences and behavior when a deviation from a cascade occurs. Deviations from optimal play occur according to a statistical process and players take these deviations into account when making inferences and decisions. Moreover, deviations or mistakes are payoff-dependent in the sense that the likelihood of a mistake is inversely related to its cost.[3]

In this chapter, we demonstrate that QRE predicts the temporary and self-correcting nature of cascades and also predicts several features of the long-run dynamics, as a function of signal informativeness. QRE predicts that with an infinite horizon the true state will be revealed with probability one, i.e., learning is complete. While no finite experiment can formally test this prediction, our ability to structurally estimate public beliefs with QRE allows us to draw inferences about the rate at which beliefs are converging to full revelation.[4]

Following the pioneering paper of Anderson and Holt (1997), there have been a number of studies exploring different questions related to information cascades. Hung

[3]We only consider *monotone* quantal response equilibrium, where choice probabilities are monotone in expected utilities, see McKelvey and Palfrey (1995, 1998) and Goeree, Holt, and Palfrey (2005).

[4]Longer sequences of decisions could possibly be obtained from an Internet experiment where agents are successively invited to participate (see Drehmann, Oechsler, and Roider, 2005).

and Plott (2001) replicate the original findings and also explore information aggregation in a voting mechanism. Çelen and Kariv (2004) differentiate between information cascades and herds. Huck and Oechssler (2000), Dominitz and Hung (2004), Nöth and Weber (2003), and Oberhammer and Stiehler (2003) explore whether decision makers are following Bayes' rule in their updating process and find evidence that Bayes' rule is systematically violated. Some of the other extensions involve cascades in networks (Choi et al., 2005), the effect of advice (Çelen et al., 2005), costly signals (Kübler and Weizsäcker, 2004), and herd behavior in stock markets (Cipriani and Guarino, 2005; Drehmann et al., 2005). The negative relationship between the duration of a cascade and the probability of collapse is demonstrated in Kübler and Weizsäcker (2005) across several different studies and is consistent with our own findings and with the predictions of the QRE model.

The remainder of the chapter is organized as follows. Section 2.3 describes the basic model and presents the main theoretical properties of QRE dynamics, which deliver hypotheses that are directly testable with data from our experiment. Section 2.4 explains the experimental design. Section 2.5 contains a descriptive analysis of the data, focusing on cascade dynamics and choice behavior. Section 2.6 presents an econometric analysis of the basic model and extensions that better explain the data. Section 2.7 discusses the belief dynamics implied by the structural estimation and the resulting efficiency properties of the data. Appendix A contains proofs and Appendix B contains the basic estimation program.

2.3 The Basic Model

There is a finite set $\mathcal{T} = \{1, 2, \ldots, T\}$ of agents who sequentially choose between one of two alternatives, A and B. Agent t chooses at time t, and 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 (or state of the world), denoted by $\omega \in \{A, B\}$, 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 $q \in (\frac{1}{2}, 1)$ 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 of her predecessors, but not their signals. Thus a history H_t for agent t is simply a sequence $\{c_1, \ldots, c_{t-1}\}$ of choices by agents $1, \cdots, t-1$, 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

$$\pi_t^a(p_t) \equiv P(\omega = A|H_t, s_t = a) = \frac{q\, p_t}{q\, p_t + (1-q)(1-p_t)}. \tag{2.3.1}$$

Likewise,

$$\pi_t^b(p_t) \equiv P(\omega = A|H_t, s_t = b) = \frac{(1-q)p_t}{(1-q)p_t + q(1-p_t)} \tag{2.3.2}$$

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.

2.3.1 Nash Equilibrium

We first discuss the dynamics of beliefs and choice behavior in a Bayesian Nash equilibrium. The unique trembling hand perfect equilibrium of the game, identified by Bikhchandani et al. (1992), involves rapid convergence to an information cascade.

This pure cascade Nash equilibrium works as follows.[5] 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. For comparison to the Quantal Response Equilibrium discussed next, we assume that the second agent randomizes uniformly when indifferent.[6] 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 posterior (before considering her private information) is $p_3 = \frac{1}{2}$, so that her signal completely determines her choice. The fourth agent would then be in the same situation as the second agent described above, and so forth. Thus a cascade begins after some even number of agents have chosen and $|\#A - \#B| = 2$, where $\#A$ is

[5]As we will see, almost all choice sequences in our laboratory data are inconsistent with the behavior implied by this Nash equilibrium.

[6]This randomization holds in any Logit QRE. There are other Nash equilibria where players randomize with different probabilities when indifferent, but none of these equilibria are trembling hand perfect. In the unique trembling-hand perfect equilibrium indifferent players follow their signal with probability 1. For details see Goeree et al. (2006).

the number of decision makers who have chosen A and $\#B$ is the number of decision makers who have chosen B.

One quantity of interest is the probability that "correct" and "incorrect" cascades will form in equilibrium. First, the probability of being in neither cascade vanishes rapidly as t grows. The probability of eventually reaching a correct cascade is $\frac{q(1+q)}{2-2q(1-q)}$, and the complementary probability of eventually reaching an incorrect cascade is $\frac{(q-2)(q-1)}{2-2q(1-q)}$.[7,8] 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 prior.

2.3.2 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.[9,10] 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

[7] After the first two choices, the probabilities of the three regimes, correct cascade, no cascade yet, or incorrect cascade, are: $\frac{1}{2}q(1+q)$, $q(1-q)$, and $\frac{1}{2}(q-2)(q-1)$, respectively. More generally, after $2t$ choices, these probabilities are $\frac{1}{2}q(1+q)\left(\frac{1-(q(1-q))^t}{1-q(1-q)}\right)$, $(q(1-q))^t$, $\frac{1}{2}(q-2)(q-1)\left(\frac{1-(q(1-q))^t}{1-q(1-q)}\right)$. Taking limits as t approaches infinity yields the long-run probabilities of the three regimes.

[8] Thus as q increases from $\frac{1}{2}$ to 1, the probability of landing in a good cascade grows from $\frac{1}{2}$ to 1.

[9] This arises when ϵ_t^A and ϵ_t^B are *i.i.d.* extreme-value distributed.

[10] The properties derived in this section hold for all atomless error distributions that have full support over the interval $[-1, 1]$. The logit specification is convenient because its behavior is determined by a single parameter with a natural "rationality" interpretation.

model implies that the stronger the belief that A is correct, 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 θ_t.

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 by[12]

$$P(c_t = A|H_t, s_t) \;=\; P(\epsilon_t > 1 - 2\pi_t^{s_t}(p_t))) \;=\; \frac{1}{1 + \exp(\lambda(1 - 2\pi_t^{s_t}(p_t)))}, \qquad (2.3.3)$$

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 $\lambda \to \infty$ 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 in which indifferent subjects randomize uniformly.[13] On the other hand, as λ approaches 0 choices are independent of beliefs and become purely random.

The belief dynamics also 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

[11] In 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.

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

[13] This is because for any $\lambda \in (0, \infty)$, an agent chooses equi-probably when indifferent.

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:

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

with the probability of a B choice given by $1 - T_t^{\omega=A}$. Similarly, conditional on $\omega = B$, the probability that 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

$$p_t^+ \equiv P(\omega = A|H_t, c_t = A) = \frac{p_t T_t^{\omega=A}}{p_t T_t^{\omega=A} + (1-p_t)T_t^{\omega=B}}, \tag{2.3.4}$$

and

$$p_t^- \equiv P(\omega = A|H_t, c_t = B) = \frac{p_t(1 - T_t^{\omega=A})}{p_t(1 - T_t^{\omega=A}) + (1-p_t)(1 - T_t^{\omega=B})}. \tag{2.3.5}$$

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. *For all $\lambda > 0$ there is a unique logit QRE with the following properties:*

(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 $\omega = A$, $\lim_{t\to\infty} p_t = 1$ almost surely.*

2.3.3 Classification of Cascades Observed in the Laboratory

We distinguish several kinds of cascade-like behavior.[14] A *pure A (B) cascade* is said to form at time $t \leq T$ if after period $t - 1$ the number of A (B) choices exceeds the number 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 that 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 after period $t - 1$ (but not after period $t - 2$) the number of theoretically informative A (B) choices[16] exceeds the number of theoretically informative B (A) choices by 2 and

[14]One might argue for using 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 can last forever. All choices occur with positive probability at every point in time, and learning never ceases.

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

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

some decision maker τ, with $t \leq \tau \leq T$, makes a contrary choice.[17] The number of periods decision makers follow the cascade, $\tau - t$, defines its *length*. Thus in the sequence of decisions $\{A, A, B\}$ we say that an A cascade of length zero occurs at $t = 3$.

Temporary cascades are particularly interesting because subsequent play of the game is off the Nash 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 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.

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*. A *self-correcting cascade* is a cascade that reverses from the incorrect state to the correct state.

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

[18]These 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.

2.3.4 Hypotheses

A wide range of observable implications follow from the theoretical results about the logit equilibrium in these dynamic games of incomplete information. We distinguish four categories of hypotheses depending on their object: cascade length and frequency, self-correction of cascades, efficiency of decisions, and belief dynamics. Most of these hypotheses are in the form of the comparative statics with respect to the two main treatment parameters, q and T.

Testable implications of the logit QRE. *For all $\lambda > 0$ observed behavior in the unique logit QRE will have the following properties:*

1. **Cascades: Frequency and length**[19]

(C1) For any q and sufficiently large T, the probability of observing a pure cascade is decreasing in T, converging to 0 in the limit. For any q and $T > 2$, the probability of observing a temporary cascade is increasing in T, converging to 1 in the limit.

(C2) For any T, the probability of a pure cascade is increasing in q.

(C3) For any q, the expected number of cascades is increasing in T.

(C4) For sufficiently large T, the expected number of cascades is decreasing in q.

(C5) The probability that a cascade, which has already lasted k periods, will break in the next period is decreasing in k.

(C6) For any q, the average length of cascades is increasing in T.

(C7) For any T, the average length of cascades is increasing q.

[19]Several of these hypotheses are only sensible if T is sufficiently large. At least 2 periods are required for any cascade to form, and at least 6 periods are required to observe a cascade and its reversal. For example, $\{A, A, B, B, B, B\}$ is the shortest possible sequence for a reverse from an A cascade to a B cascade, and $\{A, A, B, A\}$ is the shortest possible sequence for a repeated A cascade.

2. **Self-Correction**

(SC1) Incorrect cascades are shorter on average than correct cascades.

(SC2) Incorrect cascades are more likely to reverse than correct cascades (self-correction).

(SC3) Correct cascades are more likely to repeat than incorrect cascades.

(SC4) Later cascades are more likely to be correct than earlier ones.

(SC5) A decision maker with a contradictory signal is more likely to break a cascade than a decision maker with a confirmatory signal.

3. **Efficiency: The probability of correct decisions**

(E1) The *ex ante* (i.e., before decision maker t has drawn a private signal) probability of a correct decision is increasing in t. An interim version of this statement is true, but only conditional on receiving an incorrect signal.[20]

(E2) The probability of a correct decision is higher for a correct than for an incorrect signal.

(E3) The probability of a correct choice is increasing in q.

4. **Beliefs: Informational efficiency**

(B1) For each q, on average the public belief on the true state is closer to 1 in the final period of the $T = 40$ treatments than in the $T = 20$ treatments.

(B2) For each t, on average the public belief on the true state is closer to 1 in the $q = 6/9$ treatments than in the $q = 5/9$ treatments.

[20]It is *not* true conditional on receiving a correct signal. To see this, note that the interim probability of a correct decision at time $t = 1$ with a correct signal approaches 1 as λ diverges as it is optimal to follow one's signal. In later periods it is bounded away from 1 because of the probability of a cascade on the wrong state.

(B3) For all treatments, on average the public belief on the true state is increasing in t.

These hypotheses follow from a few basic properties implied either by QRE or by the informative signal process itself. We list them below,[21] and refer to them in the ensuing discussion that explains the intuition of the hypotheses. For any positive value of λ:

1. There is a (positive) lower bound on the probability a decision maker chooses either decision, because payoffs are bounded. This lower bound is independent of beliefs.
2. The higher the public belief on a state, the greater the probability the decision maker will choose the optimal action for that state.
3. If a decision maker breaks a cascade, he is much more likely to have a contradictory signal than a confirmatory signal.
4. The higher is q, the more likely it is that any given cascade will be correct.
5. In a correct cascade, confirmatory signals are more likely than contradictory signals.
6. In an incorrect cascade, confirmatory signals are less likely than contradictory signals.
7. When an action is taken at time t, the public belief on the corresponding state increases. That change in public belief is an increasing function of q.
8. The higher the public belief on the true state, the higher the probability the decision maker receives a signal favoring that state.

[21] The proofs are straightforward and are omitted.

9. The expected change in beliefs in the true state from t to $t+1$ is always positive.

Hypothesis (C1) (which applies to $T > 2$) follows from (1) which implies that the probability a cascade breaks in any round is strictly positive. Hypothesis (C2) follows from (3),(4), and (5). Hypothesis (C3) follows because the probability that a first cascade has formed is increasing in T, the probability that a cascade has formed and broken is increasing in T, the probability that a cascade has formed and broken and then another one has formed is increasing in T, and so forth.

Hypothesis (C4) is more complicated and can only be proved for T sufficiently large. For example, if $T = 2$, then the expected number of cascades is simply the probability that exactly one cascade occurs, which is the probability of two either correct or two incorrect signals, which is $q^2 + (1-q)^2$. This expression is increasing in q. The difficulty is that there are two opposing effects of increasing q. The probability of a cascade forming is increasing in q but the probability of a cascade breaking is decreasing in q. For sufficiently large T the latter effect dominates because decisions are more frequently in a cascade than not in a cascade. The higher is λ, the greater must be T for this to be true.

Hypothesis (C5) follows from (2) and (6). Hypotheses (C6) and (C7) follow from (2), (6), and (8) and the fact that the probability of a cascade breaking once you are in a cascade is decreasing in q. Hypothesis (SC1) follows from (3), (4), and (5). The logic behind the next two hypotheses about self-correction, (SC2) and (SC3), is fairly obvious. They follow from (7) and the fact that decision makers are more likely to receive correct than incorrect signals. Hypothesis (SC4) is a consequence of the self-correction process and follows from (2) and (8). Hypothesis (SC5) is equivalent to (3).

The efficiency hypotheses address the frequency of correct decisions. First, on average, efficiency will increase over time because expected public belief converges monotonically to the true state (the ex ante part of E1). Second, decision makers who receive a correct signal are obviously more likely to make the correct decision than decision makers with incorrect signals (E2), but this difference will decline over time, because the public belief on the true state converges to 1 (interim part of E1). Third, 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.[22] Because these three effects all go in the same direction, there should be a difference in efficiency in the different q treatments.

All the C, SC, and E hypotheses are tested with simple direct tests on sample means. However, (C3), (C4), (C6), and (C7) can be strengthened because the comparative statics results on length and frequency of cascades holds for *the entire distribution* of lengths and frequencies, not just the means.

Because beliefs, unlike actions, are not directly observable, we test the B hypotheses by estimating beliefs using our QRE structural estimation approach. Thus the analysis of beliefs in our data is quite different and depends on the estimation, so we discuss the results about beliefs later, after presenting the QRE estimates of the underlying parameters of the model.

While some of these properties are also true for the initial few decisions in the pure cascade Nash equilibrium, the effects vanish quickly with longer sequences. An exception is (E3). 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 $q^2/(q^2 + (1 - q)^2)$ and rises with q.

[22] Another 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.

For the purposes of evaluating and explaining empirical data, it is clear that any reasonable model must allow for the possibility of noise or trembles. However, it is important to note that the structure of the noise in QRE is essential to almost all of the qualitative predictions we derive. A model with payoff-uncorrelated trembles would generate very different predictions which, as we show below, would not be consistent with the data. In such a model, if the sequence (A, A, B) is observed, then for small tremble rates, the only way to rationalize the deviating B choice is through a tremble, which means that the decision imparts no information about the private signal, since the tremble rate is independent of beliefs. Thus, cascades will form and persist in such a model (modulo occasional trembles) just as in the standard Nash equilibrium. Because the standard information cascade logic applies to this setting, no information is learned from observing long decision histories, so that posterior beliefs are constrained to an interval around the prior. For instance, when $q = 6/9$ and there is a tremble rate of $\epsilon = 0.05$, terminal beliefs are 0.22 and 0.78. At these beliefs, subjects who get a signal contrary to the cascade have posteriors of 0.63 and 0.37, respectively, so that breaking the cascade with a contrary signal results in expected losses of 41%.

2.4 Experimental Design

The two innovations of our experimental design are the use of much longer choice sequences and the use of different signal precisions. These innovations allow us to assess the predictions of the logit QRE model in ways that are not possible with past designs and to gain insights into how the basic models might be improved.

The experiments reported here 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 sub-

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 2.3: *Experimental sessions.*

jects included students from these two institutions who had not previously participated in a cascade experiment.[23]

The experiments employ a 2×2 design, where we use 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 experimental session is denoted M. Table 2.3 summarizes the design.[24]

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

[23]There was one subject who had previously participated in a related pilot experiment.

[24]The design is not balanced with respect to subject pool because Caltech's laboratory has a maximum capacity of 32 subjects. In the estimations reported below we checked for subject pool effects but found no major differences in parameter estimates.

[25]See www.hss.caltech.edu/~rogers/exp/ for the instructions.

subjects were taken slowly through a practice match (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 commenced.

Before each match, 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), were done publicly. At this point, the subjects 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.

Before discussing the results, we comment briefly on the choice of parameters in the experimental design. Regarding T, and M, it is clear that more observations give a more complete picture of the dynamics and allow better statistical tests. There is

[26] A few sessions contained fewer sequences due to technical problems; see Table 2.3.

a tradeoff given time constraints for experimental sessions, in that longer sequences (higher T) require fewer repetitions (lower M). Given that variation in T is central to our design, we decided that the larger value $T = 40$ was sufficient to study choice and belief dynamics, while allowing time for enough repetitions (20 or 30) to accurately estimate the theoretical models and do statistical tests. With respect to the signal precision q, we wanted values that generated enough difference in behavior to be readily observed. However, too much separation in the values would make the data less interesting, since if q is close to one, nearly all subjects receive signals that indicate the true state, so that we would expect the vast majority of sequences of choices to cascade instantly (ie, when the true state is A, for all T decisions to be a). On the other hand, if q is close to one-half, posteriors will not move far away from the neutral prior, and we would expect the choice data to be nearly uniformly random. Last, their is the consideration that, in order to minimize differences in presentation across treatments, we wanted the denominator–which corresponds to the number of balls in each jar in the experimental screen–to be constant across values of q. Since balancing these effects is a non-trivial problem, we chose the experimental values through a comprehensive study of simulations of data generated according to QRE. Although the choice is nonetheless somewhat arbitrary ex ante, the results below show that these aims were met (see, e.g., Tables 2.4 and 2.5, as well as Figure 2.1, and many of the following tables and figures for evidence that the data is non-degenerate and significantly different across treatments).

2.5 Results I: Cascades, Self-Correction, and Efficiency

In this section, we examine the aggregate properties of our data. The analysis is focused by the hypotheses in the previous section about cascade frequency and length,

	Our Data				AH Data	Nash		QRE - BRF			
	q = 5/9		q = 6/9		q = 6/9	q = 5/9	q = 6/9	q = 5/9		q = 6/9	
	T = 20	T = 40	T = 20	T = 40	T = 6			T = 20	T = 40	T = 20	T = 40
	M = 116	M = 56	M = 90	M = 60	M = 45						
First 6	36%	36%	50%	53%	64%	98%	99%	27%	26%	44%	44%
First 10	15%	14%	31%	42%		100%	100%	9%	11%	28%	27%
First 20	4%	2%	13%	32%		100%	100%	1%	1%	15%	14%
First 40		0%		13%		100%	100%		0%		7%

Table 2.4: *Percentages of pure cascades by treatment.*

self-correction of cascades, and efficiency of decisions.

2.5.1 Infrequency of Pure Cascades; Frequency of Temporary Cascades

In AH's experiment with only $T = 6$ decision makers, all cascades were necessarily very short. 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, the length distribution of temporary cascades, and the self-correcting property. As Table 2.4 clearly demonstrates, pure cascades essentially do not happen in the longer matches. 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 2.4. 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 2.4 give the predicted frequency of pure cascades according to the Nash equilibrium (and out of sample predictions from the QRE-BRF model, which we explain and discuss in a later section). The Nash equilibrium

	Our Data				AH Data	Nash		QRE - BRF			
	q = 5/9		q = 6/9		q = 6/9	q = 5/9	q = 6/9	q = 5/9		q = 6/9	
	T = 20	T = 40	T = 20	T = 40	T = 6			T = 20	T = 40	T = 20	T = 40
	M = 116	M = 56	M = 90	M = 60	M = 45						
First 6	58%	55%	42%	38%	27%	0%	0%	64%	62%	50%	48%
First 10	84%	84%	69%	55%		0%	0%	88%	87%	71%	71%
First 20	96%	98%	87%	68%		0%	0%	99%	99%	85%	86%
First 40		100%		87%		0%	0%		100%		93%

Table 2.5: *Percentages of temporary cascades by treatment.*

probability of a pure cascade with T decision makers is $1 - (q(1-q))^{T/2}$.

The data contradict the Nash predictions in three ways. First, there are far fewer pure cascades than theory predicts. Second, there are far fewer than were observed in past experiments with very short decision sequences. According to theory, the frequency of pure cascades should increase with T but in fact the data show the opposite. 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, temporary cascades are common in all treatments. Table 2.5 shows the frequency of temporary cascades in our data. The rows and columns mirror Table 2.4, but the 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. With the short horizon of the AH experiment, temporary cascades occur only in about one-fourth of the sequences.

[27]Further 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.

		q = 5/9		q = 6/9	
		T = 20	T = 40	T = 20	T = 40
		M = 116	M = 56	M = 90	M = 60
average number cascades	Our Data	3.47	7.54	2.99	3.73
	QRE - BRF	3.88	7.30	2.88	4.21
	Nash	1.00	1.00	1.00	1.00
average length cascades	Our Data	2.43	2.00	3.27	7.83
	QRE - BRF	1.54	2.26	4.15	6.39
	Nash	17.32	37.25	17.43	37.44

Table 2.6: *Number and lengths of cascades by treatment.*

2.5.2 Number and Lengths of Temporary Cascades

With larger T, we almost always observe *multiple* temporary cascades along a single sequence. Table 2.6 (top) displays the average number of cascades in each treatment. The number of temporary cascades rises with the sequence length, T, and falls with the signal precision, q, not only on average, but also in the sense of first order stochastic dominance; see the top panel of Figure 2.1. This evidence supports hypotheses (C3) and (C4). The table and the figure also show the Nash prediction of exactly 1 cascade per sequence, independent of q and T, and out-of-sample predictions generated by the QRE-BRF model (discussed later).

Figure 2.2 graphs, separately for each treatment, the empirical probability of collapse as a function of the duration of the cascade: i.e., the probability of a collapse in period $t+s$, given the cascade started in period t. This probability is sharply decreasing in s. In other words, longer cascades are more stable (Kübler and Weizsäcker, 2005), which is predicted by QRE but is not true in the Nash equilibrium. This finding supports hypothesis (C5).

The average length of temporary cascades for each treatment is displayed in Table 2.6 (bottom), and the complete distributions of length are shown in Figure 2.1 (bot-

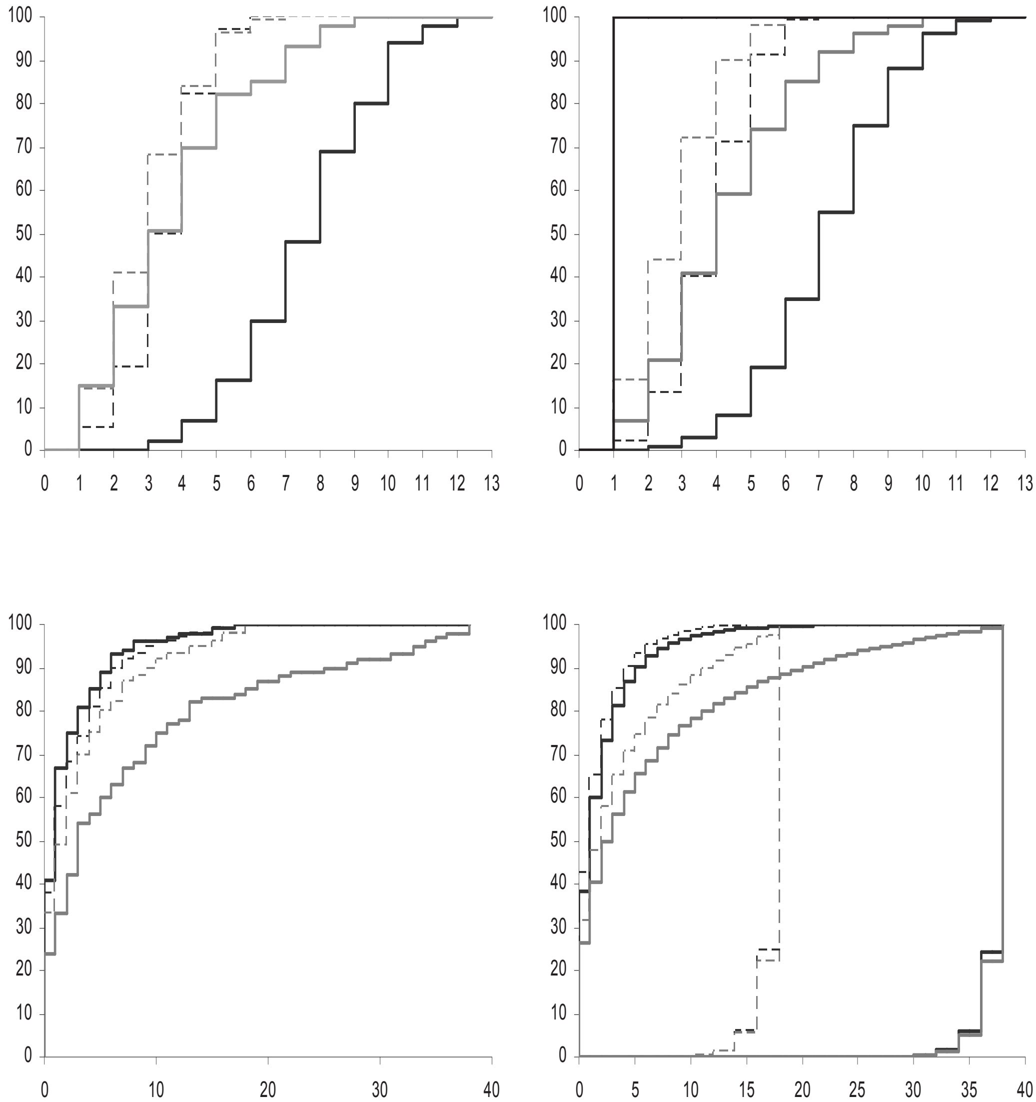

Figure 2.1: *The left panels depict the observed distributions of the number of cascades (top) and of cascade lengths (bottom), color-coded by treatment: dark (light) gray lines correspond to* $q = 5/9$ *(*$q = 6/9$*) and they are solid (broken) for* $T = 40$ *(*$T = 20$*). The right panels show predictions 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* $T - 2$ *correspond to Nash predictions and the others to QRE-BRF predictions.*

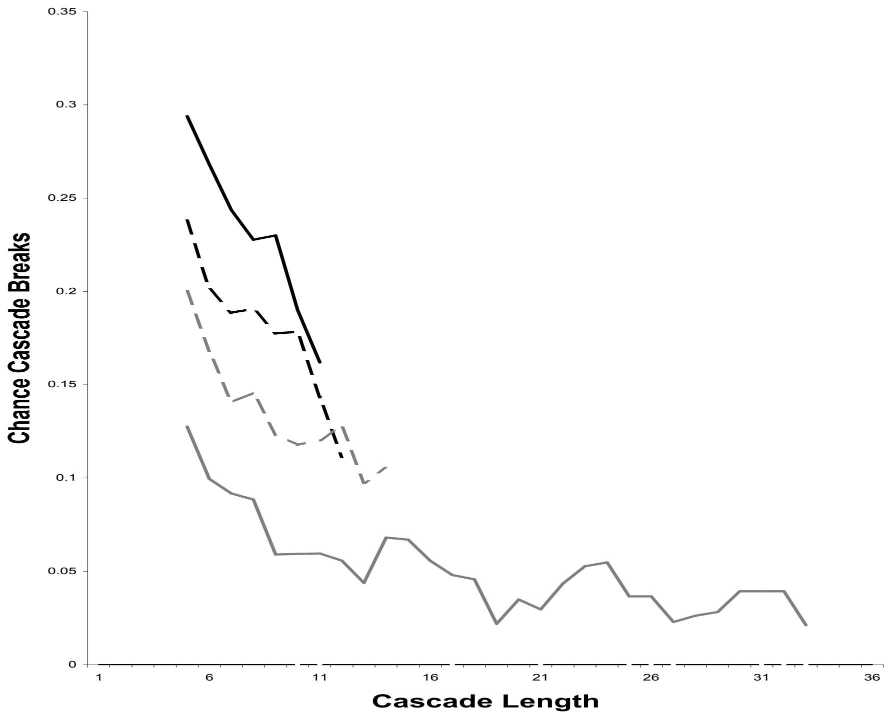

Figure 2.2: *Chance of cascade breaking as a function of cascade length. The lines show 5-period moving averages of the probability of a break in each of the treatments (color-coded as in Figure 2.1).*

tom).[28] Average length of temporary cascades rises with the sequence length, T when $q = 6/9$ but not when $q = 5/9$, with the difference insignificant at the 5% level in the latter case. Average length also rises with the signal precision, q, for both $T = 20$ and $T = 40$. Thus we find strong support for (C7) but only weak support for (C6). A comparison of the entire distribution of lengths is given in the top panel of Figure 2.1. The table and the figure also show the Nash prediction of exactly 1 cascade per

[28] To compute the Nash predictions for cascade lengths, recall that a cascade can only begin after an even number of choices. For t even, the probability that a cascade forms after $t + 2$ choices conditional on one not having yet formed after t choices is $1 - 2\beta q(1 - q)$, where β is the probability that an indifferent subject follows her signal. In the data, $\beta = 0.85$. Since upon forming, a cascade persists through period T, the predicted length distributions of temporary cascades can be calculated easily.

sequence, independent of q and T, and out-of-sample predictions generated by the QRE-BRF model (discussed later).

2.5.3 Off-the-Equilibrium-Path Behavior

Given that the vast majority (92%) of cascades are *temporary* and short in duration, and nearly all (90%) sequences in our data exhibit *multiple* cascades, an immediate conclusion is that there are many choices off the (Nash) equilibrium path. Table 2.2 in the Introduction characterizes a subset of these choices for the different treatments as a function of the deviating decision maker's signal. The table shows the behavior of what we call *cascade breakers*, since these are all terminal decisions of a temporary cascade.

Over all treatments, cascades were broken a total of 1081 times. These contrary actions were *five times more likely* to be taken by subjects with contradictory signals than with confirmatory signals (898 compared with 183). In fact, if we compare the *rates* of breaking cascades for decision makers with contradictory versus confirmatory signals, the difference is even starker (37% compared to 6%). This supports hypothesis (SC5).

The behavior of decision makers immediately following a cascade breaker also plays a critical role in the dynamics. Because the first break is so informative, a second break moves beliefs close to .5, essentially eliminating the trend in beliefs that had developed during the cascade.

As expected, the probability of a second break by the next decision maker is sharply increased. Approximately 75% of the decision makers immediately following a cascade break follow their signals. A player who observes a signal consistent with the recent cascade of course should rationally follow the cascade, a prediction that is borne out by our data: 90% of these decision makers follow the action corresponding

	T=20			T=40		
	Decision \ Signal	Confirming	Contrary	Decision \ Signal	Confirming	Contrary
q=5/9	Confirming	42.8%	20.9%	Confirming	45.4%	20.3%
	Contrary	3.3%	33.0%	Contrary	8.4%	25.9%
	# obs = 306			# obs = 379		
	Decision \ Signal	Confirming	Contrary	Decision \ Signal	Confirming	Contrary
q=6/9	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 2.7: *Percentages of choices confirming/contradicting the recent cascade after a break.*

to the recently broken cascade. Only 10% are *secondary deviators* who follow the recent break. Thus, they behave roughly the same as they would have if the cascade had never been broken. Those who received contradictory signals behaved much differently. Well over half (56%) of the decision makers with contradictory signals are secondary deviators. Pooling over all treatments, they outnumber the secondary deviators with confirmatory signals by a factor of five to one (277 compared to 58). Table 2.7 gives a complete breakdown of the choices directly following a cascade break, by treatment.

The two key conclusions of this subsection are that play off the equilibrium path *occurs frequently* and is *highly informative*, setting the stage for self-correction. As a result, we will find that the long-run implications of the standard theory are completely contradicted by the data.

2.5.4 Repeated and Reversed Cascades: Self Correction

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

		q = 5/9		q = 6/9	
		T = 20	T = 40	T = 20	T = 40
		M = 116	M = 56	M = 90	M = 60
	Our Data	2.14	5.98	1.69	2.60
average number repeat cascades	QRE - BRF	2.39	5.37	1.62	2.91
	Nash	0.00	0.00	0.00	0.00
	Our Data	0.34	0.55	0.30	0.13
average number reversed cascades	QRE - BRF	0.50	0.93	0.26	0.30
	Nash	0.00	0.00	0.00	0.00

Table 2.8: *Frequency of repeated and reversed cascades by treatment.*

experimental design with much longer sequences allows us to better observe the kinds of complex dynamics predicted by the theory, in particular the phenomenon of self correction.

Table 2.8 shows the average number of repeated and reversed cascades per sequence, by treatment, and also gives theoretical expectations according to the Nash and out-of-sample QRE-BRF predictions (explained later). While such cascades are not possible in the Nash equilibrium, the latter model predicts the observed number of reversed and repeated cascades remarkably well.

Table 2.9 shows how frequently correct and incorrect cascades repeat or reverse themselves.[29] The number of repeat cascades is increasing in T and decreasing in q, which is consistent with the QRE model.

Averaging over the four treatments shows that when a correct cascade breaks, it reverses to an incorrect one in approximately 6% of all cases (39/637). In contrast, an incorrect cascade that breaks leads to a self-corrected cascade in more than 21% of all cases (66/369). This confirms hypotheses (SC2) and (SC3).

Table 2.9 also lists the initial, final, and total number of correct and incorrect

[29]The percentages listed ignore terminal cascades, since they can neither repeat nor reverse, by definition.

	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%
# correct	total = 252	start = 65	final = 78	total = 237	start = 29	final = 34
# incorrect	total = 151	start = 51	final = 38	total = 185	start = 27	final = 22
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%
# correct	total = 197	start = 66	final = 71	total = 186	start = 48	final = 52
# incorrect	total = 72	start = 24	final = 19	total = 38	start = 12	final = 8

Table 2.9: *Transitions between correct and incorrect cascades in our data.*

cascades by treatment. In all four treatments, the fraction of incorrect cascades is always lower among the final cascades compared with the initial cascades. Overall, initial cascades were incorrect nearly 35% of the time (114/322) and final cascades were incorrect only 27% of the time (87/322). This supports hypothesis (SC4).

2.5.5 Efficiency

How frequently are actions correct? How does this change over time? And how does this change as a function of signal informativeness? These questions can be directly answered in our data by checking the proportion of correct decisions, since both the state and the action of each individual are observed in the data.

There are two important observations to note before delving into the analysis of the efficiency results. First, the probability of a correct decision, and the way that probability changes over time, will be much different for decision makers who received correct versus incorrect signals. Decision makers with incorrect signals will do badly at the beginning, but will do increasingly well over time. Decision makers with correct signals will do very well at the beginning (perfectly in the Nash equilibrium), but will do worse for a while until the public belief gets close enough to 1. Second, overall efficiency is extremely sensitive to the specific sequence of signals individuals receive and also (in a quantal response equilibrium) to the specific action choices. Since in 20 rounds there are over one million possible signal sequences (and many more signal-action sequences), our experimental data represent only a small fraction of the possible sequences. Therefore, there is a lot of sample variation.

Figure 2.3 shows the time-dependence of decision accuracy, by treatment. The middle column displays the actual data, averaged across all experimental sequences, with the four rows each corresponding to a 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 the fraction of correct choices for all signals; the dashed (upper) and thinner (lower) lines display the fraction of correct choices for correct and incorrect signals, respectively.

It is useful to contrast the data with the efficiency predictions of Nash equilibrium, which are displayed in the left column of the same graph, again based on the actual signal draws in the experiment. In the Nash equilibrium, decision accuracy quickly becomes independent of signals, reflecting the formation of pure cascades where all learning stops and all future decisions are the same.[30] The decision accuracy for (in)correct signals (rises) falls for a few rounds and then levels off. As a result, the unconditional decision accuracy increases for only a short amount of time as nearly all cascades are formed in the first five periods and never break. This contrasts sharply with the dynamics in the actual data, where unconditional decision accuracy continues to rise as the sequence of decision makers passes through cycles of temporary cascades that break and re-form.

[30] As an illustration of the sample variation induced by the specific sequence of signal draws, decision makers in the $q = 5/9$, $T = 20$ treatment by chance drew many more correct signals in the early rounds than did decision makers in the $q = 5/9$, $T = 40$ treatment. This is most easily seen by comparing the Nash predictions of decision accuracy for the two treatments.

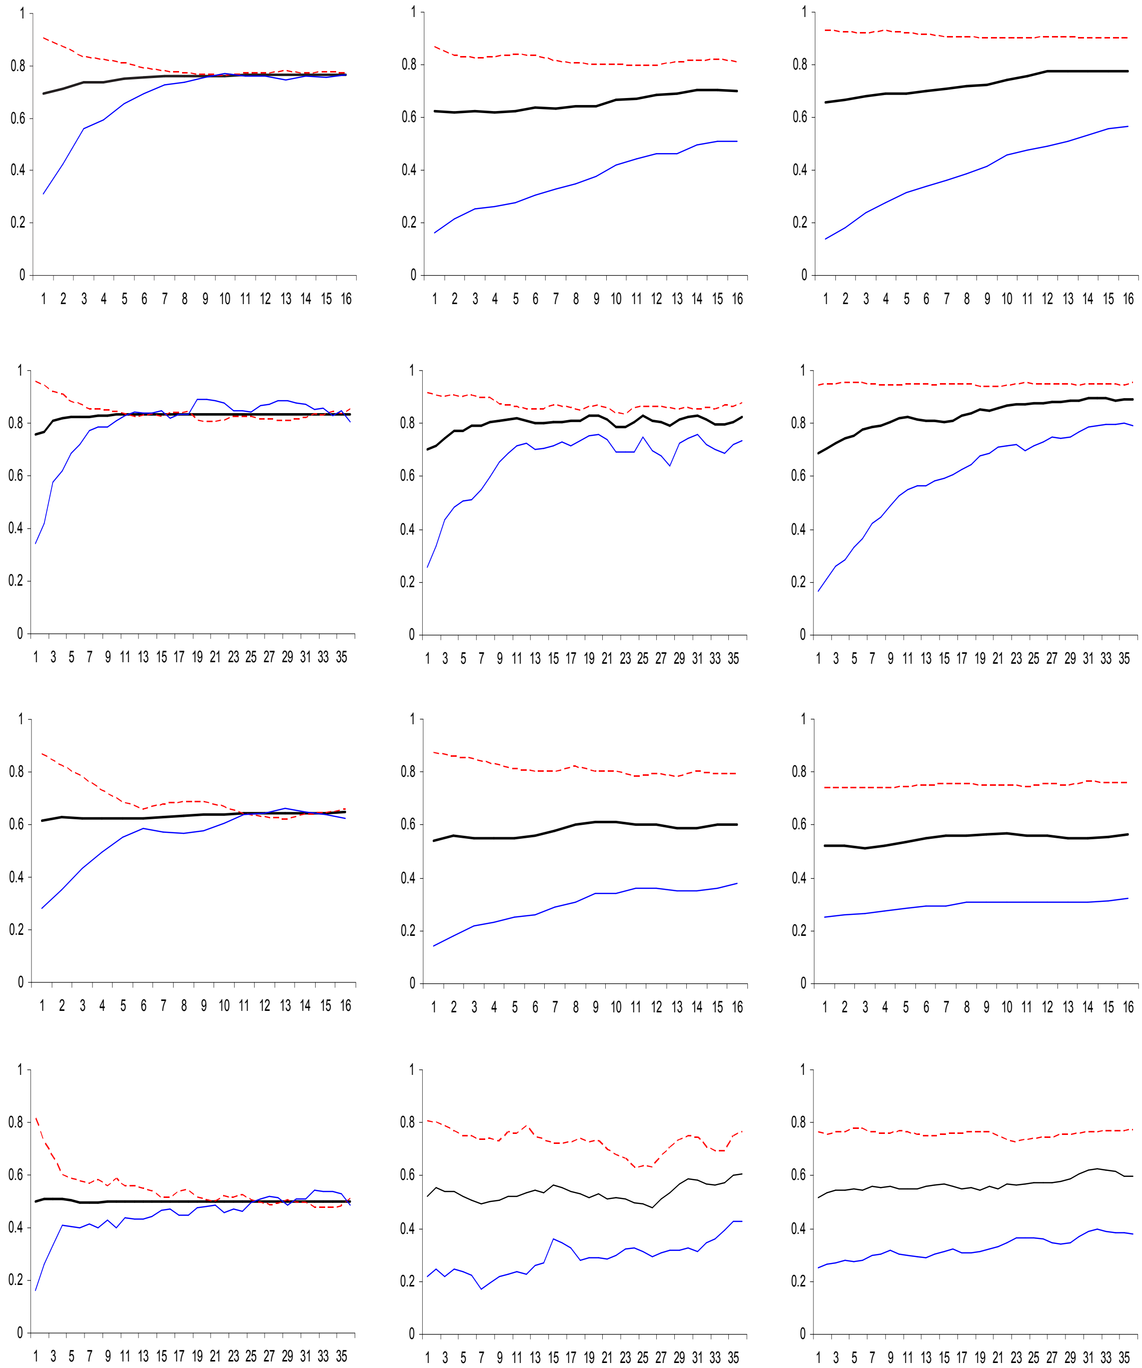

Figure 2.3: *Decision accuracy along the sequence of decision makers by treatment: ($q = 6/9$, $T = 20$) top row, ($q = 6/9$, $T = 40$) second row, ($q = 5/9$, $T = 20$) third row, and ($q = 5/9$, $T = 40$) bottom row. In each graph, the thick solid black line shows the fraction of correct choices for all signals, the dashed red line for correct signals, and the thin blue line for incorrect signals. The lines show moving averages: a point at time t represents average decision accuracy between $t-2$ and $t+2$ for $3 \leq t \leq T-2$. The left column gives Nash predictions, the middle column data, and the right column QRE-BRF simulations, all based on the actual signals used in the experiment.*

There is a strong signal dependence that persists throughout the experiment. The decision accuracy for incorrect signals is always less than for correct signals in the actual data *in every round t*, providing strong support for (E2). For incorrect signals, there is a clear and persistent upward trend in decision accuracy (due to information aggregation)and there is a small, early downward trend for decision makers with correct signals, as hypothesized (the interim part of E1). For decision makers with correct signals this levels off and even reverses sign later because later cascades are more likely to be correct due to the phenomenon of self-correction.

For a more formal statistical test of hypotheses E1-E3, Table 2.10 shows the results of a Probit regression with six independent explanatory variables: t, q, $q * t$, signal, signal$* t$, match. Signal is a 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,[31] which, according to hypothesis H4, should be positive. The variable (signal$* t$) is an interaction between time and signal correctness. From hypothesis 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.[32] Notice that we do not include T in the regression because the theory does not predict any effect except through the variable t.

The second column of Table 2.10 shows the estimated coefficients with standard errors in parentheses. All coefficients have the expected sign and are statistically 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 sources of error, including quantal response errors and variation in signal sequences. (The third and fourth

[31] Here $q * t$ equals 0 if $q = 5/9$ and $q * t$ equals t if $q = 6/9$.

[32] Match=1 corresponds to the first sequence in a session and ranges up to 20 or 30 depending on whether T equals 40 or 20, respectively.

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 2.10: *Probit estimation of the effects of q and t on efficiency.*

columns of the table are discussed later.)

Finally, it is natural to ask whether efficiency is higher under the QRE-BRF model than it is under the standard Nash model. Information is aggregated better under QRE-BRF (see Proposition 1) but decision-making is worse in this case as subjects are prone to errors. Figure 2.5 shows that efficiency levels are increasing with time under the QRE-BRF model throughout the duration of the experiment. In fact, in the long run as T grows large, beliefs in the QRE-BRF model converge to the true state so that private beliefs and public coincide, independent of signals. Using the pooled data to estimate the parameters of the QRE-BRF model, we can compute the asymptotic decision accuracy: 0.99, i.e., almost full allocative efficiency is achieved in this limit.

2.5.6 Summary of Results

Here we summarize our findings by relating them to the properties of the logit QRE discussed in Section 2.3.4.

- (C1) and (C2): *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 2.4. 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 2.4 shows the effect of signal informativeness.

- (C3 and C4): *The number of cascades increases with T and decreases with q.* See Table 2.6 and Figure 2.1. Longer sequences have more cascades because they allow for more cycles of formation and collapse. These effects are barely noticeable in short sequences: AH's experiment averaged slightly more than 1 cascade per sequence.

- (C5): *The probability of collapse sharply decreases as a function of the duration of the cascade.* See Figure 2.2.

- (C6 and C7): *Cascade lengths increase with T for $q = 6/9$ and increase with q.* The effect of T can be decomposed as follows. First, and most obvious, if T is short then some cascades that would have lasted longer are interrupted at T. Second, by (C5) longer cascades are less likely to break. The two effects combined result in a fat tail of the length distribution and in a mass of cascades at $T - 2$; see Table 2.6 and Figure 1. The effect of T is observed in the $q = 5/9$ data, where the distributions of cascade lengths are very similar for the $T = 20$ and $T = 40$ treatments.

- (SC1): *Correct cascades last longer on average.* The observed average lengths of (correct, incorrect) cascades 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 $T = 40$.

- (SC2) and (SC3): *Reverse cascades are usually self-correcting, and repeat cascades are usually correct.* See Table 2.9. Across the four treatments, the probability that a reversed cascade is self-correcting is 63% (even though there are many more correct than incorrect cascades to reverse from). It is this feature of the dynamics that produces the full information aggregation result of Proposition 1.

- (SC4): *Later cascades are correct more frequently than earlier ones.* See Table 2.9, which lists the number of (in)correct cascades among initial and final cascades.
- (SC5): *Cascades are almost always broken by decision makers with contradictory signals.* See Table 2.2.
- (E1): *Ex ante efficiency is increasing in t.* Efficiency is increasing in t, conditional on an incorrect signal. Efficiency is initially decreasing in t conditional on a correct signal, but this eventually reverses (see Figure 3).
- (E2): *Correct signals lead to more efficient decisions than do incorrect signals.* Again, see Figure 3.
- (E3): *More informative signals lead to more efficient decisions.* See Figure 2 and Table 2.10.

The final three hypotheses, (B1)-(B3), address the evolution of beliefs during a sequence and are discussed in section 2.7. The next section describes the QRE estimation and our base rate fallacy model.

2.6 Results II: Estimation

We start by describing the estimation procedure for the basic logit QRE model. The only parameter 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.[33]

Since subjects' choice behavior depends on λ, public beliefs follows a stochastic process that depends on λ. The evolution of the public belief can be solved recursively (see equations (2.3.4) and (2.3.5)), so implicitly we can write $p_t(c_1, \cdots, c_{t-1}|\lambda)$. Given $\{\lambda, s_t, (c_1, \cdots, c_{t-1})\}$, the probability of observing player t choose A is:

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

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

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

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

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

The estimation results for the logit QRE model are given in Table 2.11. A detailed estimation program written in GAUSS is contained in Appendix B. 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 ($q = 5/9$,

[33] Recall that in the experiment subjects received \$1 for a correct choice and \$0.10 for an incorrect choice. The difference of \$0.9 is normalized to 1 unit in the estimations. Without this normalization the estimates reported below would be multiplied by a factor of 1/.9.

$T = 20$) treatment is somewhat greater than the other three treatments. This may reflect a subject pool effect, since that treatment was the only one that used mostly Caltech students.

Since comparison with Nash equilibium does not provide a particularly informative benchmark for the logit QRE, the following three subsections consider extensions and alternatives to the basic model. This allows us to access the extent to which the choice behavior in our data is explained by quantal response-type decision errors as opposed to other sources, such as non-Bayesian updating and non-rational expectations.[34] Using parametric specifications we measure the extent of certain types of these biases in the data.

2.6.1 Incorporating 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. As noted in the introduction, there is considerable evidence in the literature on cascade experiments that players are non-Bayesian. We explore two of these here. First, 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 would imply that agents weight the public prior too little relative to their own signal. Because past experiments have been suggestive of these effects, we construct an analytical model of this and estimate it using the error structure of the Logit equilbrium.[35] We formalize this idea as a non-Bayesian updating process in which the private signal is counted by the decision maker as α signals, where $\alpha \in (0, \infty)$.[36] Rational agents correspond to

[34]Huck and Oechssler (2000) find strong evidence of violation of Bayesian updating in a similar context.

[35]Some error structure is required for the estimation because the α-BRF model is deterministic.

[36]This could also be loosely interpreted as a parametric model of "overconfidence" bias in the sense of Griffin and Tversky (1992). Kariv (2005) and Nöth and Weber (2003) use this terminology.

$\alpha = 1$, while agents have progressively more severe base-rate fallacies as α increases above 1.[37]

While agents over-weight their private signals we retain the assumption that they have rational expectations about others' behavior. This implicitly assumes that α is common knowledge (as well as λ). The updating rules in (2.3.1) and (2.3.2) now become

$$\pi_t^a(p_t|\alpha) = \frac{q^\alpha p_t}{q^\alpha p_t + (1-q)^\alpha(1-p_t)} \tag{2.6.6}$$

and

$$\pi_t^b(p_t|\alpha) = \frac{(1-q)^\alpha p_t}{(1-q)^\alpha p_t + q^\alpha(1-p_t)}, \tag{2.6.7}$$

respectively.[38]

The public belief, p_t, in equations (5.1) and (5.2) is derived recursively using (2.3)-(2.5). In particular, this means that subjects not only overweight signals, but also take into account that other subjects overweight signals too, and the public belief is updated accordingly. Thus, for $\alpha > 1$, the public belief is updated more quickly than in the pure Bayesian model.

There is good reason to think this model may better describe some features of the data. First, when $\alpha = 1$ QRE predicts that indifferent agents randomize uniformly. However in the data 85% of indifferent subjects follow their signals, which is consistent with $\alpha > 1$.[39] Second, when $\alpha > 1$, cascades take longer to start.[40] The base rate fallacy therefore provides one possible explanation for the prevalence of length zero

[37]Values 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.

[38]From 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.

[39]A subject is indifferent when the counting rule applied to previous decisions and the subject's private signal balances to zero.

[40]For example, after two A choices the third decision maker need not choose A if she sufficiently overweighs her b signal.

temporary cascades in our data set (see Figure 1).

The estimation results for the QRE-BRF model are reported in the second panel of Table 2.11. For all treatments, the BRF parameter, α, is significantly greater than 1.[41] To test for significance we can simply compare the log-likelihood 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.[42] Furthermore, the constrained model yields a significantly (at the 0.01 level) higher estimate of λ for all treatments.

There is at least one alternative interpretation to the finding that subjects respond too strongly to their signal. By doing so, they are giving better information to later decision makers, which increases efficiency and raises the expected utility of the other players in the game. Evidence from experiments on public goods and some game theory experiments suggest some degree of altruism by subjects. Conceivably, what we are calling a base rate neglect (or overweighting of signals) may simply be a manifestation of altruistic behavior. However, there is some counter evidence that suggests this is probably not the case. First, if altruism is the motivating force, one would expect higher estimates of α for $T = 40$ than for $T = 20$. This is not the case. Second, once would expect less overweighting of signals in later periods than in earlier periods. We tested for this and found no significant effect. Therefore, our interpretation is not that subjects are behaving altruistically, but rather the source of the distortion is a probability judgement fallacy.

2.6.2 Incorporating Non-Rational Expectations

Rather than simply over-weighting private information relative to the base rate (public belief), it is possible that players update incorrectly because they do not have

[41] Similar results are reported by Çelen and Kariv (2004).

[42] For the pooled data the difference in log-likelihoods 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.

rational expectations about the driving parameters of the model. The QRE model implicitly assumes that λ is constant across the population and common knowledge. In particular, if players believed other players' λ were lower than it truly was, then beliefs, and hence choice dynamics, would be qualitatively similar to those under a base rate neglect. The reason is that when choices are believed to be generated by a noisier process, players draw weaker inferences about predecessors' signals from observing their choices. 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.[43] We call this the *non-rational expectations model*, or QRNE model.

The estimation results for the QRNE model are also given in Table 2.11. While this two-parameter model performs significantly better than the QRE model, the increase in likelihood is smaller in magnitude than the increase of QRE-BRF relative to the simple QRE.

An advantage of using the QRE model is that we can explore the relative importance of different biases, by nesting them in the same model. In this case we can see whether the BRF bias is more or less important in our data compared to updating failures due to irrational expectations about other players' error probabilities. When BRF and QRNE are combined so that the model includes both sources of bias, the action and belief λ are virtually identical when estimated from the pooled data, and the increase in likelihood from the QRE-BRF model is barely significant. A similar conclusion holds for the AH data, indicating that the assumption of rational expectations ($\lambda_a = \lambda_b$) is (approximately) valid in both data sets, while $\alpha > 1$ indicates a robust effect of base rate neglect.

[43]See Kübler and Weizsäcker (2004) for a more extensive discussion of this model.

2.6.3 An Alternative Model: Cognitive Heterogeneity

It is instructive to consider other models with non-quantal response sources of noise that could also potentially explain our data. This helps to check the validity of our basic story for choice behavior, in light of the observation that the Nash equilibrium does not provide a way to challenge any of the predictions of QRE. One natural question to ask is where the source of scatter (error) in our data is really coming from. In QRE, it is assumed to come entirely from payoff-monotone choice errors, and this behavior is assumed to be homogeneous across the population. An alternative possibility is that this apparent noise in the data is due to some kind of underlying heterogeneity. We explore one possible model of heterogeneity in this section.

Although there are many options, a natural first step is to suppose that some players behave completely randomly, while other players optimize against such behavior. Camerer, Ho, and Chong (2003) extend this idea to allow for multiple levels of sophistication.[44] Specifically, level 0 players are random, and all other players use optimal strategies given their beliefs. Level 1 players 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 in the population is given by a Poisson distribution with parameter τ. That is, the probability of a level k player in the population, given the Poisson parameter τ, is equal to $\frac{\tau^k e^\tau}{k!}$. Thus, for example, if $\tau = 1.5$ then the distribution of types $0, 1, 2, 3, ...$ is equal to $(0.22, 0.33, 0.25, 0.125, ...)$. 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 their relative probabilities given by the true Poisson distribution. Thus, again using the example of $\tau = 1.5$, 22% of the players are simply randomizing, 33% are optimizing assuming they face only rational players, 25% are optimizing assuming they face a mixture of level 0 and level 1 in proportions

[44] Stahl and Wilson (1995) explored a related but different model with levels of sophistication to study behavior in experimental games. See Camerer, Ho, and Chong (2003) for a discussion of the differences between the two models.

equal to $\frac{2}{5}$ and $\frac{3}{5}$, and so forth. Therefore, assuming the model is correct, very high level types have very accurate beliefs about the distribution of types. This implies they also have accurate beliefs about the distribution of strategies in the population, and therefore they are almost optimizing. This is called the *cognitive hierarchy (CH) model*.

The presence of randomizing level 0 players will lead higher-level players to implicitly discount the information contained in the choices of their predecessors. In this way the CH model can pick up some of the same features of the data as QRE. To see this, it is instructive to look at exactly what the behavior of the lowest three types are. Level 0's of course are just random. Level 1's simply follow their own signal, since they assume there is no useful information in the observations of previous decision makers (they are believed to be totally random). Level 2's optimize against a mixture of such players, so they simply act as if each previous decision is a noisy (but informative) signal about the signals of earlier decision makers. Again using the example of $\tau = 1.5$, if the second mover is a level 2 player and observes the first player choose A, he believes that the first mover received an A signal with probability $\frac{4}{5}$ and a B signal with probability $\frac{1}{5}$. Thus, such a player's posterior on state A will be less than q. That is, level 2's have dampened updating, but also note that level 2's will reach a point quickly where they no longer follow their own signal. In the example above, they will act exactly like a player following the Nash equilibrium and will herd after one of the decisions has been chosen two more times than the other decision. (This is independent of q.) Furthermore, like QRE, CH is "complete" in the sense that it is consistent with any sequence of choices and signals. Hence we can obtain maximum likelihood estimates of the parameter τ via the same methodology, without using QRE; see Table 2.11.

We also estimate CH together with QRE to allow for further comparison with QRE. To do so, we suppose that each agent is assigned a level k in the hierarchy, as in CH, but quantal responds to her beliefs, as in QRE. Thus CH-QRE is a model parameterized by (τ, λ), which are assumed to be common knowledge. All three models (CH, QRE, and CH-QRE) are then re-estimated with the inclusion of the BRF parameter α to allow for the possibility of over- or under-weighting of private information in each case; see Table 2.11. 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.

A surprising finding is that the estimate for τ is larger in magnitude than has been typically found in other settings. Camerer, Ho, and Chong (2004) 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.

2.6.4 Implications of Estimates for the Data

The QRE-BRF model is simple and intuitively appealing and we use it to create simulated data for comparisons with the actual data.[45] For each of the four treatments,

[45]The QRE-CH-BRF model would have been an alternative model for simulation, but the additional randomness of 0-level types would have necessitated many more simulated sequences. Because the fit improvement over QRE-BRF is negligible, we decided to use the simpler QRE-BRF model for our simulations.

	Our Data					AH Data
	p = 5/9		p = 6/9		Pooled	p = 6/9
	T = 20	T = 40	T = 20	T = 40		T = 6
# 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	-3636.1	-77.1
CH-BRF						
τ	1.98 (0.08)	1.40 (0.07)	2.03 (0.08)	2.87 (0.08)	2.06 (0.04)	2.45 (0.27)
α	1.49 (0.10)	0.97 (0.12)	1.54 (0.06)	1.65 (0.07)	1.34 (0.04)	0.99 (0.07)
logL	-925.7	-1167.2	-683.1	-650.0	-3537.5	-74.3
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 2.11: *Parameter estimates for the different models with standard errors in parentheses.*

we used the data from the *other* three treatments to obtain out-of-sample estimates for λ and α. We then applied the out-of-sample estimates to the signals realized in the experiment to obtain simulated choices for the treatment. Based on this simulated data set we computed descriptive statistics about the numbers, lengths, and types of cascades: pure and temporary, repeated and reversed, self-correcting, etc. These are reported in the right two columns of Tables 2.4 and 2.5 (pure and temporary cascades, respectively), and the second and fifth rows of Tables 2.6 and 2.8 (numbers/lengths of cascades and reversals, respectively). Because the simulations were constructed using out-of-sample estimates of λ and α, they represent out-of-sample predictions of the properties of cascades in our data, which makes a comparison to the actual data meaningful. Indeed, the match with the actual data is quite remarkable.

We are also able to construct out-of-sample simulated efficiency dynamics in the the same way for each of the four treatments, again using the actual sample draws. These are displayed in the four charts in the right-hand column of Figure 2.3. Again, it reproduced the patterns observed in the data.

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 2.10. 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

[46]The 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.

likely caused by the fact that the simulations assume homogeneous agents, while we would expect some heterogeneity to be present in the laboratory data.

2.7 Results III: Estimated Belief Trajectories

We use belief estimates generated from the QRE-BRF model to examine both the the informational efficiency and to address hypotheses about the evolution of beliefs (B1-B3). How well is the information from private signals aggregated? How high is the public belief on the correct alternative after a sequence of decisions? How does this vary with our treatment variables, q and T?

2.7.1 Informational Efficiency

As shown in Proposition 1, in a QRE the public belief about the correct alternative increases on average with t and converges to 1 as T approaches infinity. The convergence is slower 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 a combination of noise in the signal generation process and noise in the decision making process. Moreover, this noise in signal generation is compounded by strategic considerations that affect the social learning process.

Although we do not observe beliefs directly, we can use the theoretical QRE-BRF model together with the observed choice data to obtain estimated public belief paths.[47] This is done for every sequence in the experiment. 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 4, which shows the belief paths for all sequences in one of the $q = 6/9$ and $T = 20$ sessions. The belief trajectories for other sessions exhibit

[47] Domowitz and Hung (2003) recently reported a social learning experiment using a belief elicitation procedure.

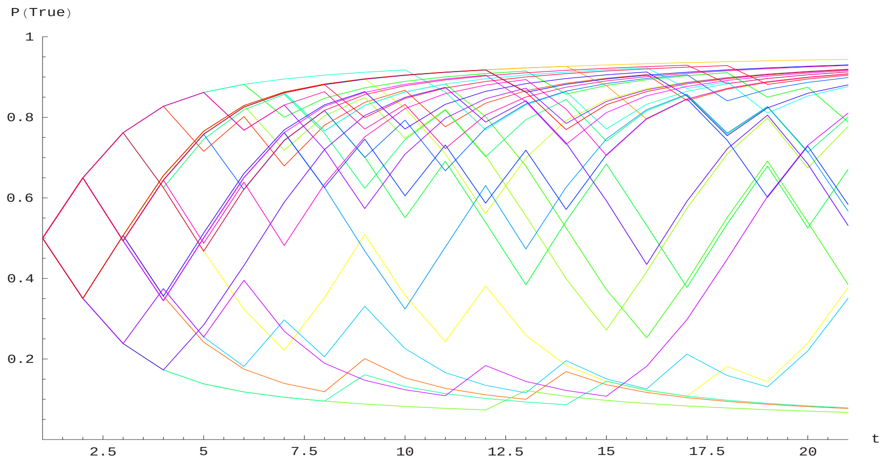

Figure 2.4: *Estimated beliefs using the QRE-BRF model for all sequences in one of the* $(q = 5/9, T = 20)$ *sessions.*

similar features. Here the horizontal axis represents the sequence of decisions, and the vertical axis represents the belief about the correct alternative. Each upward tick in the belief paths corresponds to a correct choice and each downward tick corresponds to an incorrect choice. Theoretically, for long enough sequences, the belief paths for almost all sequences should converge to 1.

The simplest way to test Hypotheses (B1)-(B3) is to 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 2.5. 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.[48]

The comparison between the different q treatments is a weak test 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

[48]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.

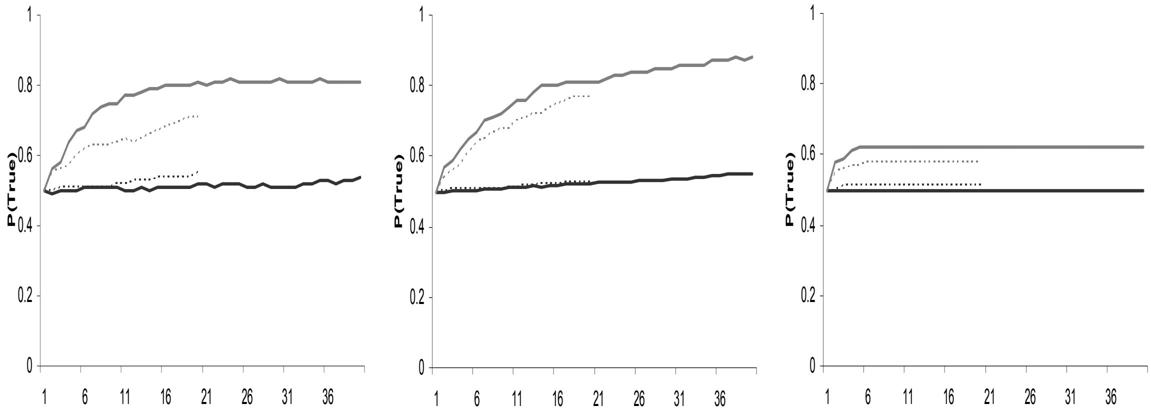

Figure 2.5: *Estimated public beliefs about the true state by treatment (coded as in Figure 2.1). In the left panel, estimated beliefs are based on observed signals and decisions. The middle panel is based on the average of 100 QRE-BRF simulations of decisions, always using the same sequence of signals as in the experiment. The right panel shows estimated beliefs implied by Nash decisions based on the sequence of signals employed in the experiment.*

the $q = 5/9$ data than in the $q = 6/9$ data (see Tables 2.5-2.8 of the previous section).

However, that the curves are increasing in t is *not* an artifact of the construction, but simply reflects the fact that there are more good cascades and fewer bad cascades toward the end of a session than toward the beginning. In summary, we find strong support for hypotheses (B1), (B2), and (B3), and somewhat weaker support for hypothesis (B4).

2.8 Conclusion

This chapter reports the results of an information cascade experiment with two novel features: 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 efficiency.

The longer sequences have several interesting features. First, there is almost a complete absence of pure cascades and a proliferation of temporary cascades, including many repeated, reversed, and self-correcting cascades. Standard theory predicts that longer sequences will have more permanent cascades and that temporary, repeated, reversed, and self-corrected cascades never occur. 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.

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.

Our experimental results confirm a wide range of hypotheses about the number and frequency of different kinds of cascades, efficiency, and belief dynamics. Most of these hypotheses follow logically from the informativeness of signals and a basic property of the QRE: deviations from rationality occur and their likelihood is inversely related to their cost. In the context of information cascades, 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 for long enough sequences public beliefs would be approximately correct.

Chapter 3

The Strategic Timing of Social Learning

3.1 Introduction

When individuals share a common objective but are faced with limited information, they may base their decisions partly on the observable actions of others. Empirical evidence to this effect is provided in the previous chapter. There are also examples from the field. For instance, portfolio managers tend to bias their investment decisions towards the decisions of their colleagues.[1] In seminal papers, Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) explore settings where each of a set of ex-ante identical agents receive private signals and publicly guess the outcome of a state variable in a predetermined sequence. The common prediction of these models is that at some point a predominance of choices for a particular action overwhelms any privately held information. At this point all remaining decision makers also take the same action, even if it is contradicted by their signals.

It is this phenomenon that has come to be known as an *information cascade.*

[1]See, for example, Scharfstein and Stein (1990).

A cascade entails an informational inefficiency in that once decisions do not reflect signals, there is no mechanism by which private information can be distributed to other agents and aggregated, and the process of social learning halts. Thus even in large economies where there is an abundance of distributed information, with substantial probability most individuals make suboptimal decisions.

In this standard setup, the order in which agents make decisions is fixed exogenously. Yet there may be interesting phenomena that drive the observed timing of decisions in the field. One factor is that as the number of announcements grows, inferences based on these decisions become more accurate on average. Consequently, individuals who decide later in the sequence expect to do better.[2] On the other hand, there is typically an explicit cost to delay. For instance, a financial manager making an investment decision pays the cost of holding a suboptimal portfolio while collecting the information to make her decision.[3] The choice of which candidate to endorse in a national election also has an important timing element. As more endorsements are made public, the likely winner becomes known with better certainty, but endorsing a candidate too late in the campaign may pay smaller benefits. Consider also the decision by a group of firms about whether to adopt a new production technology.[4] The profitability of the technology is uncertain. Each firm would like to observe the adoption by other firms before making its decision, as this provides valuable information. Yet delaying the decision is costly since there is a possibility that the firm is using a suboptimal technology in the interim. Thus there is an inherent tension between deciding early and late that must be resolved. At equilibrium, agents balance this tradeoff optimally, realizing that others are simultaneously making the same calculations. This interaction is modeled below by allowing timing decisions to be strategic variables in a social learning game.

[2] Gale (1996) explains this logic in some detail, and also provides an overview of related phenomena from the herding literature.

[3] Dasgupta (2000) examines this setting in more detail.

[4] See Kapur (1994) for an analysis of technology diffusion.

The model presented below may be summarized as follows. Agents are asymmetrically informed about a binary state variable. Signals are conditionally independent and take values on the unit interval, with a monotonic likelihood ratio. Agents must publicly guess the outcome of the state variable at a time of their choosing. The payoff is positive only when the guess is correct, and future payoffs are discounted at a common rate so that waiting is costly. We characterize equilibrium outcomes in this setting and the implications for efficiency.

Notice that, as opposed to the canonical model, the accuracy of private information is heterogeneous. Signals vary in quality, ranging from uninformative to fully revealing the state variable.[5] The combination of strategic timing and differential signal quality drives the intuition of the analysis below. To illustrate, consider the situation of a worker seeking employment, where job offers are publicly observable. It seems likely that some potential hirers would have better information regarding the worker's quality than others. For instance, some firms may gain more information from the interview or may personally know those who wrote recommendation letters. Further, firms with better information could use this difference to gain a strategic advantage over other firms in the hiring process, e.g., by making quicker decisions. Another example derives from bank runs, as experienced in the United States before the institution of federal deposit insurance.[6] Under this interpretation of the model, the timing of deposit withdrawals is related to an individual's private information. Those with the strongest beliefs that the bank is insolvent should be the first to act. But taking this action is costly if the bank does not fail, so there is an incentive to delay in order to observe the others' decisions.

This chapter provides an explanation of how rational agents with differing information qualities behave when they strategically choose the timing of their decisions.

[5] Thus signals are unbounded in the sense of Milgrom (1979). Smith and Sorenson (2000) contains a thorough analysis of social learning in the case of unbounded signals.

[6] Caplin and Leahy (1994) studies a model of market dynamics with an application to bank runs.

Since utility depends only on an agent's own action and the state of the world, all strategic interactions are due solely to an informational externality. We show that in two player games there is a unique equilibrium. It exhibits monotonic "sorting," i.e., agents with better information decide earlier, and those with lower quality signals behave more patiently and incur more delay cost in order to learn from others. In this model, continuously varying signal qualities are important for the existence of a symmetric equilibrium in pure strategies, since signal quality differentiates agents and breaks indifferences about the timing decision that would otherwise occur. As the time grid becomes increasingly fine, the equilibria approach a limit where all action occurs immediately, and delay costs vanish. Since the sorting improves as the time intervals shrink, the agent with the lower quality signal learns from the better-informed agent with probability converging to one, and the equilibrium outcome approaches the first best. However, this limit outcome can only be supported as an equilibrium in continuous time under a particular "tie-breaking" assumption regarding simultaneous decisions. Further, since the equilibria in the discrete games are unique, there is no other approachable equilibrium in the continuous game.[7]

The main emphasis of this chapter concerns how the welfare properties of social learning are affected by the strategic timing of decisions relative to exogenous sequencing, for which welfare implications are very negative. Despite the equilibrium phenomenon of herding, information is aggregated efficiently in the present model due to the sorting property when there are only two agents. Yet these outcomes are not fully efficient because of the delay cost implicit in the sorting. We find that efficiency increases monotonically as the time intervals shrink, and, as mentioned above, approaches the first best outcome in the limit. For any finite time interval, aggregate welfare is higher in the endogenous timing model than in the standard exogenous

[7]Several authors have provided general existence theorems for discontinuous games, including Reny (1999), Jackson, Simon, Swinkels, Zame (2002), Dasgupta and Maskin (1986a,b), and Milgrom and Weber (1985). Despite the fact that none of these results are readily applicable to the present model, we are also able to find asymmetric equilibria, although they are not perfect.

sequence model. In games with many players full efficiency is approximated when the time intervals are short. The reason is that with unbounded signals, there are always some signals that produce a dominant strategy to announce immediately, and with many players some individuals receive such signals and reveal the true state to others with high probability. Thus the message of this chapter is positive in the sense that we view the assumptions of this model as more appropriate in many settings, and welfare is unambiguously higher in this setting.

The rest of the chapter is organized as follows. The next section describes in more detail how the present model relates to existing literature and highlights this chapter's contributions. Section 3.3 presents the model for the case of two players. Section 3.4 contains the equilibrium characterization and the convergence result that as time intervals vanish the game ends immediately. Section 3.5 then generalizes the model to allow for a large number of players and analyzes equilibrium outcomes in that context. The comparison of information aggregation and efficiency to the case of exogenous ordering of decisions is handled in Section 3.6. Some extensions and limitations of the model are discussed in Section 3.7. Finally, Section 3.8 concludes. Proofs of some main results and technical details are included in Appendix C.

3.2 Related Literature

A number of authors, including Chamley and Gale (1994), Gul and Lundholm (1995), Chamley (2004), and Zhang (1997), have also analyzed models of social learning in which timing decisions are strategic variables. Interestingly, the results of these models are mixed and, in contrast to the exogenous sequencing models, depend crucially on modeling assumptions about the nature of time, uncertainty, and the richness of the type and action spaces. Consequently, it is difficult to draw unambiguous conclusions regarding welfare implications.

Chamley and Gale (1994) consider an investment model where there is only one action, but infinite delay is not a dominated strategy. The most important difference from the current chapter is that there is no notion of differentiated signal quality. Because of this, social learning does not lead to welfare gains in their model. The reason is that the only symmetric equilibrium involves mixing probabilities that cause agents to be indifferent between announcing and delaying, so that informational gains are exactly balanced by delay costs in equilibrium. One reason for the welfare improvement in the present chapter is that the timing of announcements reveals information about signal qualities, which can not happen in Chamley and Gale (1994). A common finding of these models, however, is that the game ends quickly as the time interval shrinks.

The model of Gul and Lundholm (1995) operates in continuous time and uses a continuous action space, and again signals are not differentiated by quality. The focus of that study is on "clustering," the tendency of announcements to be closer together than if they were made independently. Under their utility specification they find that social welfare is lower than in the standard exogenous sequencing model, which is the opposite result of the present chapter. As before, time can not serve as a screening device for worse informed agents to learn. A second effect is that, since time is continuous, the relevant comparison to the exogenous sequencing case involves no discounting for later announcers. This suggests one sense in which the modeling of time as a discrete quantity may be more appropriate in social learning models.

Building on the work of Chamley and Gale (1994), Chamley (2004) incorporates differential signal quality into the model, which generates results more similar to this model: (i) as the time intervals vanish, the game ends immediately, and (ii) the true state is asymptotically revealed as the number of players diverges. Interestingly, however, he finds a multiplicity of equilibria with widely varying information aggregation properties. This contrasts with the present model for which we find a unique equilibrium[8] with unambiguous welfare implications. For instance, a rush in which all agents announce immediately constitutes an equilibrium in Chamley (2004) but not in the present chapter. This contrast is due to the fact that there are two possible announcements in the present context, so that agents with neutral beliefs will find it optimal to delay given that others are announcing. Another difference resulting from the number of possible announcements is that if ever there is a period in which no agent makes an announcement, the game effectively ends.[9]

The most closely related work to the present model is Zhang (1996). The signal structure there is more general than the one defined below, with the important difference that signals there are bounded, whereas in this model there is a possibility of being arbitrarily well informed. Action takes place in continuous time. In the unique symmetric equilibrium there is delay, followed by an immediate cascade based on the (single) highest quality signal. This result, however, depends crucially on an additional source of uncertainty that has no counterpart here.[10] The present model also has a unique equilibrium, without restricting attention to symmetric strategy profiles, which does not converge to the continuous time result found by Zhang. Further, with many players herding is inefficient in Zhang's model whereas the opposite is true here. These differences relate to the presence of unbounded signals in my model. With many players, unbounded signals guarantee that some agents announce early on, and the information conveyed from their choices reveals the state variable with high probability. As a result, the asymptotic welfare properties are opposite in the two models. Despite the similarity of these models, Chamley (2004) argues for the

[8]This is without restricting attention to symmetric equilibria a priori, as do Chamley and Gale (1994) and Gul and Lundholm (1995).

[9]This property is due to the fact that there is only one kind of announcement, and so a period with no announcements reveals the worst possible news about the profitability of announcing. As such, the equilibrium analysis in Chamley (2004) can make use of a "two-step" property that effectively allows the payoff from announcing to be compared with the payoff from delaying exactly one period, rather than having to compute the continuation value of delay as a function of others' strategies over all future periods.

[10]Zhang's model includes an *ex ante* unknown cost of adopting either action that is publicly revealed at the time of the first announcement.

importance of analyzing a discrete time version of Zhang's model.[11] In fact, Chamley (2004) conjectures that there may be multiple equilibria in such a game. This chapter suggests that this is not true.

One contribution of this chapter is that, when combined with other results from the previous literature, one now obtains a rather complete picture of how various modeling assumptions about information and timing translate into different results. With respect to the information structure, two general points can be made. First, when signals are unbounded, the implications for welfare in large economies are positive. With arbitrarily accurate signals and many individuals, there will always be some individuals who are sufficiently sure about the true state that it can not benefit them to delay, and their decisions collectively reveal the state to all other individuals, who then have nothing else to gain and so announce correctly immediately following the first announcers. In contrast, when signals are bounded the welfare results remain negative: inefficient herds may still arise under strategic timing since the option value of delay must balance the cost. Second, when signal quality is homogenous, symmetric equilibria typically involve mixed strategies in order to balance the marginal benefits and costs of delay to maintain indifference, whereas with a continuum of signal qualities the model may admit symmetric pure strategy equilibria where agents with better qualities announce earlier.

With respect to the modeling of time, one finding is that equilibria in discrete time games may not converge to the equilibria of corresponding continuous time games. For instance, a result that is somewhat robust in the literature is that as discrete time intervals vanish, all action takes place instantly in the limit, whereas such equilibria are not easily supported in continuous time.

[11]See footnote 13 in that paper.

3.3 Model

There is a countable set of dates $T^r = \{0, r, 2r, \ldots\} \cup \{\infty\}$, with the interval between consecutive dates given by $r > 0$, at which agents in a set $N = \{1, 2\}$ make decisions.[12] The payoff-relevant states of the world are given by $\omega \in \Omega = \{A, B\}$. Agents have uniform prior beliefs regarding the true state, so that $\Pr(A) = \Pr(B) = 1/2$. The model is one of incomplete information. Agents receive signals $p_i \in [0, 1]$, $i = 1, 2$.[13] Conditional on ω, the p_i's are independently and identically distributed with *pdf*

$$f(p|\omega) = \begin{cases} f(p|A) = 2p & , \text{ if } \omega = A \\ f(p|B) = 2(1-p) & , \text{ if } \omega = B \end{cases} \quad \text{with support } [0,1].$$

Thus low realizations of p are evidence of state B while high realizations are evidence of A. In particular, a signal of $p = 1/2$ is uninformative, while signals of $p = 0$ or $p = 1$ perfectly reveal the true state, so that the signals are unbounded.[14] The unconditional distribution of types is given by

$$f(p) = \frac{1}{2}\left(f(p|A) + f(p|B)\right),$$

which is uniform on the unit interval. The most important feature of the signal structure is the monotone likelihood ratio. More general structures are discussed in Section 3.7. All elements of the model, except the realizations of the signals and the state variable, are common knowledge among the agents.

Each agent must guess the state of the world and so chooses a $c_i \in C = \{a, b\}$

---

[12] The notation T is used in place of T^r when no confusion arises.

[13] We use the phrases "agent with signal p" and "agent of type p" interchangeably for convenience, even though the language "signal" is more accurate.

[14] Smith and Sorenson (2000) analyze social learning in a general framework with unbounded signals and exogenous ordering of decisions.

and a time $t_i \in T$ at which to make her announcement. Denote by $x_i^t \in X \equiv C \cup \{0\}$ agent i's action at date t, with the interpretation that $x_i^t = 0$ records the event that i did not make an announcement at t, and let $x^t = (x_1^t, x_2^t)$ be the profile of actions at time t. A history h^t at time t specifies whether or not each agent has made an announcement and, if so, what the announcement was and the time at which it was made. Thus we can write $h^t = \cup_{\tau \in T : \tau < t} x^\tau$, with the convention that $h^0 = \emptyset$.[15] A strategy for i specifies an action to take at every history. Let $\mathcal{P}(Z)$ denote the set of probability distributions over a set Z. A *behavioral strategy* σ_i for i maps signals and histories into probability distributions over X, specifying the probabilities that i announces a, announces b, and delays, respectively.[16] Denote by $\sigma = (\sigma_1, \sigma_2)$ the profile of strategies.

Agents receive a payoff of one if they guess correctly and zero otherwise. A common instantaneous discount factor $\gamma > 0$ accounts for temporal preferences, so that the present value of guessing correctly at time t is $\exp(-\gamma t)$. Denote the game so defined by G^r.

Given the instantaneous discount rate γ and interval length r, one immediately infers that $\delta \equiv \exp(-\gamma r)$ is the effective period-to-period discount factor. Treating temporal preferences as fixed and exogenous, there is a one-to-one correspondence between δ and r, such that $\delta \to 1$ as $r \to 0$. In discussing the results below, we make use of this equivalence.

Agents use Bayes' Law to update their beliefs. Let $\pi_i(p_i, t, h^t, \sigma)$ represent i's belief that the true state is A at time t conditional on i's type p_i and on the history h^t, given the strategy profile σ. Thus after observing her type and before the game

[15] To formally express the fact that each agent can make only one announcement, write $\mathcal{H}^t$ for the set of possible histories at time t and let $\mathcal{H} = \cup_{t \in T} \mathcal{H}^t$ denote the space of all histories. We require that for all $h^t \in \mathcal{H}$, $x_i^\tau \neq 0 \Rightarrow x_i^{\tau'} = 0$ for all $\tau < \tau' < t$.

[16] The restriction that only one announcement may be made can be written as follows: for all $p_i \in [0, 1]$ and all $h^t \in \mathcal{H}$, $\sigma_i(p_i, h^t) = (0, 0, 1)$ whenever there exists $\tau < t$ such that $x_i^\tau \neq 0$.

starts, agent i's posterior belief is given by

$$\pi_i(p_i, 0, \emptyset, \sigma) = \frac{f(p_i|A)P(A)}{f(p_i|A)P(A) + f(p_i|B)P(B)} = p_i.$$

We use perfect Bayes Nash equilibrium (PBNE) as the solution concept.[17] In the next section we find a unique PBNE. In fact, the game is dominance solvable, so that the equilibrium strategy profile is determined through iterative elimination of strictly dominated strategies.

Although the set of strategies is quite large, a few observations greatly restrict the set of possible equilibrium strategies. First, note that if the other agent announces first, the unique best response is to make a decision immediately thereafter since there is no possible benefit in delaying further, and delay is costly. Thus after one agent announces, the other agent announces as quickly as possible, payoffs are realized, and the game effectively ends. Second, note that if an agent announces before her opponent, she should always announce a if and only if $\pi_i(p_i, t, h^t, \sigma) \geq 1/2$.[18] Given these two simple observations, the equilibrium strategies for i can then be fully described by a mapping $s_i : [0, 1] \to \mathcal{P}(T)$ that gives the distribution of stopping times at which an agent will announce as a function of her type p_i provided the other agent has not yet announced. We will be primarily concerned with pure strategies, which in this reduced form may be expressed by a measurable function $s_i : [0, 1] \to T$. That is, i is willing to wait only until time $s_i(p_i)$ to make her announcement and will announce earlier if and only if her opponent announces first. Clearly, any equilibrium strategy must satisfy $s_i(0) = s_i(1) = 0$, as an agent who knows the true state with certainty

[17]Let $g_i(\cdot|p_i, t, h^t, \sigma)$ denote i's probability assessment of j's type given her own type, the time, the current history, and strategy profile σ. A PBNE is characterized by a strategy profile σ and a belief profile g such that for $i = 1, 2$, $j \neq i$, (i) σ_i specifies a best response to σ_j at every history given beliefs g_i, and (ii) g_i is consistent with Bayes' Law at every history that is reached with positive probability under σ.

[18]The results are not sensitive to how we specify an agent's decision when she is indifferent between the two choices, as this occurs with zero probability.

has nothing to gain by delaying and would still incur the cost, i.e., she has a strictly dominant strategy.

In what follows we will be predominantly concerned with symmetric equilibria. In such cases, subscripts on strategies and beliefs are dropped when no confusion arises. In addition, the equilibrium strategies we find treat the two states of the world symmetrically. Formally, we have the following:

Definition 1 *A strategy s_i for i is information-symmetric if for all $0 \leq p \leq 1/2$, $s_i(p) = s_i(1-p)$.*

Thus information-symmetric strategies are symmetric mappings about the uninformative signal ($p = 1/2$), which means that we can restrict attention to the *quality* of an agent's type, as measured by $q = |p - 1/2|$. Note that when agents use information-symmetric strategies, beliefs about the true state $\pi_i(p_i, t, h^t, s)$ change only when an announcement is made. To see this, note that whenever a type $p < 1/2$ plans to announce at a time t, there is a corresponding type $1-p$ who also plans to announce at t. Thus agents learn nothing about the likelihood of the true state through the passage of time, unless an announcement is made.[19]

We refer to PBNE that are both symmetric and information-symmetric as SISPBNE. Since under information-symmetric strategies we can restrict attention to types $p \leq 1/2$, a SISPBNE in pure strategies (pure SISPBNE) is fully described by a function $s^* : [0, 1/2] \rightarrow T$. So s^* gives the planned stopping time for each signal quality an agent may have.

3.4 Equilibrium Analysis

The equilibrium strategies below satisfy the following:

[19] See Appendix C for a formal statement and proof of this important fact.

Definition 2 *Given $\gamma > 0$, fix $r > 0$ and an information-symmetric pure strategy s_i for i in G^r. s_i is characterized by cutpoints if there exist a number $k < \infty$ and numbers $0 = p_0 < p_1 < \cdots < p_k = 1/2$ such that if $p_{l-1} < p < p_l$ then $s_i(p) = (l-1)r$, $l = 1, \ldots, k$.*

Notice that any strategy characterized by cutpoints satisfies two properties. First, there can be no "empty dates," meaning that for any $l < k$, there exists a set of positive measure of types for which $s(p) = lr$. Second, a monotonicity condition is satisfied, in the sense that as an agent's signal quality q increases, her stopping time (weakly) decreases. We can now characterize the equilibria of G^r. The following proposition shows that G^r has a unique equilibrium and that it is symmetric, information-symmetric, in pure strategies, and characterized by cutpoints.

Proposition 1 *Fix $r > 0$ and $\gamma > 0$. G^r has an (essentially) unique PBNE s^*. Moreover, s^* is symmetric, information-symmetric, in pure strategies, and characterized by cutpoints.*

Proof. See Appendix C. ■

The intuition for Proposition 3 is as follows. That there exists a symmetric and information-symmetric equilibrium results from the ex-ante symmetry of the game. Uniqueness is obtained by iteratively eliminating strictly dominated strategies in the following way. At time zero, there is a group of types of each agent that is sufficiently well informed that waiting can not possibly be beneficial–the cost of delaying is too high even if by waiting one period they could learn the outcome of the state variable. Given that these types must announce at time zero in any equilibrium, it turns out there is enough expected benefit from waiting to induce all less well-informed agents to delay. This is shown through iteratively eliminating the strategies "announce immediately with positive probability" for types who are increasingly well informed, starting from the uninformative type ($p = 1/2$), until only the types who had the

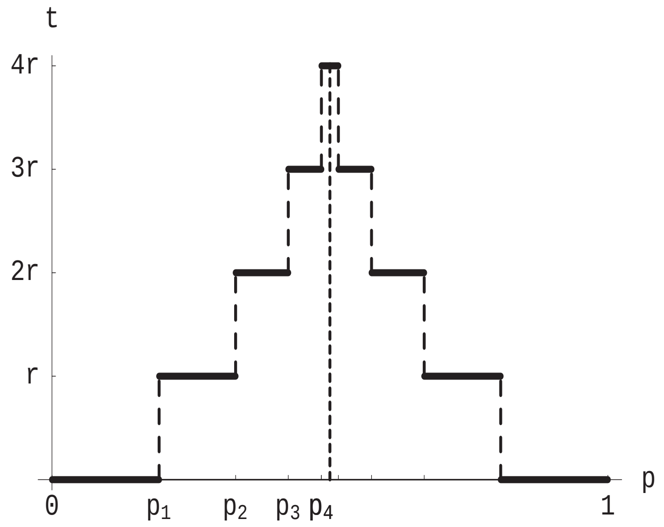

Figure 3.1: *The equilibrium strategy $s^*(p)$ as a function of type p for $r = 0.045$ and $\gamma = 1$.*

dominant strategy to announce remain. Then at the next date, if neither agent has announced, it is common knowledge that both agents are not so well informed that they dropped out at the previous date. The same logic is applied again, so that of those types remaining, some agents can not possibly gain enough from waiting to offset the cost, etc. This construction results in a unique equilibrium. Figure 3.1 depicts a typical equilibrium strategy. In this case, the maximum length of the game in equilibrium is $4r$.

Proposition 3 provides a strong result characterizing the equilibrium of G^r. Note that for any fixed temporal preference γ, as r decreases, the cost of delaying one period decreases, and so agents become more willing to delay on the margin. This means that as r tends to zero, the cutpoints depicted in Figure 3.1 shift continuously towards the boundaries of the unit interval, additional cutpoints enter from the midpoint $p = 1/2$, and the number of potential periods grows without bound. Since the marginal cost of delaying one period approaches zero as time progresses, one could imagine that for sufficiently small r, the game could last an infinite number of periods with positive probability. The next result shows that this is in fact not possible, i.e., that G^r must end in finite time for every positive r.

Proposition 2 *For every $\gamma > 0$ and $r > 0$, G^r ends in finite time. That is, there exists $k < \infty$ such that $s^*(p) < kr$ for all $p \in [0,1]$ in the unique equilibrium s^*.*

Proof. See Appendix C. ■

The proof is established as follows. First from the proof of Proposition 3 (see Appendix C), we have an equation that implicitly defines the cutpoint p_l as a function of the previous cutpoint p_{l-1} for any $\delta < 1$. This equation may be solved to provide an explicit function $p_l = p_l(p_{l-1}; \delta)$, which is graphed in Figure 3.2. The properties of this function can be used to derive a positive minimum distance between any consecutive cutpoints p_l and p_{l-1}, i.e., a minimum interval length $d(\delta)$ of types who plan to announce at time $(l-1)r$. Graphically, the interval length $p_l(p_{l-1}; \delta) - p_{l-1}$ is the vertical distance between the bold curve $p_l(p_{l-1}; \delta)$ and the dashed 45^o line at the point $p = p_{l-1}$. The cutpoint function takes valid values $p_l \leq 1/2$ up to the point $p_{l-1} = v$, where the function crosses the value $\frac{1}{2}$. Thus the red line segment depicts the smallest possible interval length d for a given δ, which is always positive. Given this bound, the maximal number of equilibrium intervals, which corresponds to the maximum number of periods the game can last, is $1/2d$.

One consequence of Proposition 3 is that unless the agents happen to announce simultaneously, they will necessarily make the same decision. To see why this is so it is enough to consider the decision problem facing an agent whose opponent has already announced. Given the sorting nature of equilibrium, the first announcer has strictly higher signal quality, so the posterior beliefs of the second announcer are dominated by the information from the decision of the first announcer, independent of which state the second announcer's private information indicates. That is, in equilibrium, $\pi_i(p_i, t, h^t, s^*) \geq 1/2$ if $x_j^\tau = a$ for some $\tau < t$. Thus the second announcer finds it optimal to make the same announcement as the first announcer. The only way in

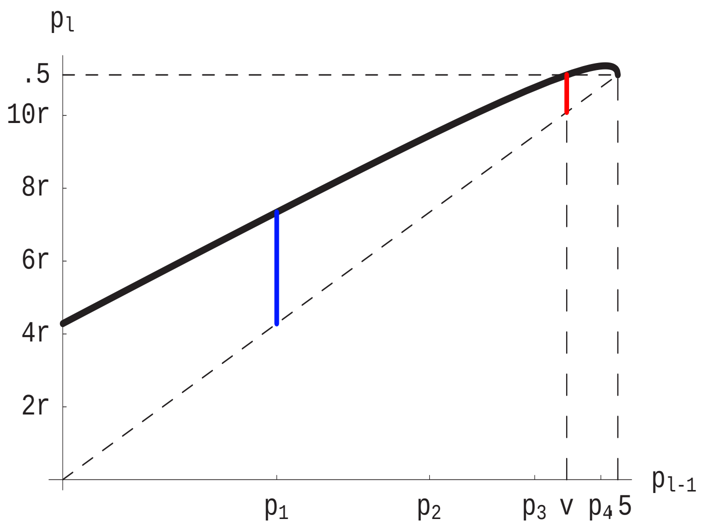

Figure 3.2: *The equilibrium cutpoint function* $p_l(p_{l-1}; 0.956)$.

which the announcements of the two agents can differ is if they decide simultaneously, in which case they have approximately the same quality of information, and neither agent learns from the announcement of the other.

Having characterized the equilibrium, we now turn attention to analyzing the comparative statics as r tends to zero. By the results of the previous section, we know that the equilibrium strategy s^* partitions the set of types $p \leq 1/2$ into intervals, where each interval of types is willing to wait until the same time to announce, conditional on the other agent having not previously announced. For a fixed δ, call the number of equilibrium cutpoints $k(\delta)$. Thus the game necessarily ends by time $rk(\delta)$, since at each date there must be some interval of types planning to announce. By Proposition 4, $k(\delta) < \infty$ for all $\delta < 1$. Note however that $k(\delta)$ diverges to infinity as δ approaches unity, meaning that as the time intervals become vanishingly small, the number of periods that the game potentially lasts becomes infinite. The relative rates at which r and k change determines the limiting behavior of equilibrium timing.

The next result shows that as r becomes vanishingly small, the game ends immediately.

Theorem 1 *For every $\gamma > 0$, G^r ends immediately as the time interval vanishes. That is,* $\lim_{r\to 0} rk(\delta) = 0$.

Proof. See Appendix C. ■

Theorem 1 is established by deriving an upper bound $\bar{k}(\delta)$ on the number of equilibrium cutpoints and showing that $r\bar{k}(\delta)$ limits to zero. The proof is essentially an application of L'Hopital's rule; it is constructing the upper bound $\bar{k}(\delta)$ that requires work. As in the proof of Proposition 4, the properties of the function $p_l(p_{l-1};\delta)$ (see (C.4) in Appendix C) are exploited. Using a linear approximation to $p_l(p_{l-1};\delta)$ instead of the actual function, one can define alternative equilibrium cutpoints z_l with the property that $z_l \leq p_l$ for all $l \leq k$. Because of the linearity, the recurrence relation that defines the $\{z_l\}$ can be solved. Setting this expression equal to 1/2 defines $\bar{k}(\delta)$, the number of cutpoints that result by using the approximation to $p_l(p_{l-1};\delta)$, which is an upper bound on the actual number of cutpoints.

3.5 Many Players

In this section we expand the model to allow for an arbitrary finite number of agents. Let the set of agents be $N = \{1, 2, \ldots, n\}$, $n < \infty$. Denote these games by G_n^r, where all other aspects of the game are unchanged. The details of this extended model are difficult to work out precisely, so we are not able to explicitly characterize the equilibria as in the case of two players. The main difficulty is that with many players, there are many subgames that must be considered at all dates after $t = 0$. This problem is avoided when $n = 2$ because there is only one interesting history to consider at any date, namely, that both players have not yet announced. In contrast, there are many nondegenerate histories with several players.

When there are more players, there is more information that may be learned from the announcements of others. Thus one would expect individuals to behave more patiently, that is, to wait longer in equilibrium, all else equal. While this is in fact true, we can show that as the number of players grows large, the game necessarily ends quickly. The intuition is that no matter how many players there are, there is always a positive mass of types with good information who cannot be induced to delay $t = 0$, since even if they could learn the true state with certainty by delaying one period, the cost of delay (given $r > 0$) outweighs this benefit. As n grows, the probability of at least one player having such a signal approaches one. Thus as the number of players becomes large, the announcements of these players reveal the true state with near certainty after the first date $t = 0$. Let $L(\mathbf{p})$ denote the maximum realized length of G_n^r in any equilibrium given the n-vector of signals $\mathbf{p}= (p_1, p_2, \ldots, p_n)$. Our next result states that for any positive r, the probability that the game ends after one period converges to one as n diverges.

Proposition 3 *Fix $r > 0$ and $\gamma > 0$. For every $\epsilon > 0$ there exists $\bar{n}$ such that $\Pr(L(\mathbf{p}) = r) > 1 - \epsilon$ for all $n > \bar{n}$.*

Proof. By definition of G_n^r, the maximal continuation value of delay at $t = 0$ is $\delta = exp(-\gamma r) < 1$. Thus for all $p < 1 - \delta$ and $p > \delta$ announcing at 0 is a strictly dominant strategy for every n. Thus for each n there will be some cutpoint $1 - \delta \leq p_1^n \leq 1/2$ such that all $p < p_1^n$ announce at $t = 0$. Assume that $\omega = B$ (the case of $\omega = A$ is parallel). Note that $\beta \equiv \Pr(p_i < p_1^n | \omega = B) > \Pr(p_i > 1 - p_1^n | \omega = B)$, since in state B low signals are more likely. Let $n_B = \#\{i \,|\, p_i < p_1^n\}$ be the number of agents who announce b at 0 as a function of $\mathbf{p}$. By a law of large numbers, for every $\epsilon > 0$, $\lim_{n\to\infty} \Pr(|n_b/n - \beta| < \epsilon) = 1$. It follows that for every $\epsilon > 0$, $\lim_{n\to\infty} \Pr(\tilde{\pi}_i(p_i, x^0) < \epsilon) = 1$, where $\tilde{\pi}_i(p_i, x^0)$ represents the posterior obtained from signal p_i and actions x^0. Thus as $n \to \infty$ beliefs converge to the true state after time 0, and so the marginal benefit of delay at the next date r approaches 0 since there is nothing more to be learned. Therefore all remaining types find it optimal to announce at $t = r$, since the cost of delay is strictly positive. ■

Since Proposition 3 is true for all $r > 0$, we can also state as a corollary that as r goes to zero, the probability that the game ends almost immediately is arbitrarily close to 1 for sufficiently large n. Note also that herding is observed in the equilibrium outcomes of G_n^r for large n. In particular, all agents who announce at time 0 do so independently from each other (so that most but not all of these agents announce correctly), but then all agents who delayed at 0 learn the true state and make the same optimal decision at time r. In contrast to the standard cascade model, where the order of announcements is fixed, there is no informational inefficiency associated with the herding. In fact, the opposite is true: all agents in the herd are announcing correctly with arbitrarily high probability. Observe also that as r tends to zero, the proportion of agents who decide at $t = 0$ vanishes, and these agents announce correctly with probability converging to 1. Thus full efficiency is approximated in large economies for small r. These observations are summarized in the following:

Proposition 4 *Fix $\gamma > 0$. For every $\bar{t} > 0$, $\epsilon > 0$, and $\eta < 1$, there exists an $r > 0$ and $\bar{n}$ that satisfy*

(i) delay efficiency: $\Pr(L(\mathbf{p}) < t) > 1 - \epsilon$, *and*

(ii) informational efficiency: $\Pr(\frac{\#\{i \,|\, x_i = \omega\}}{n} > \eta) > 1 - \epsilon$ *for all $n > \bar{n}$.*

Proof. By Proposition 3 for any $r < \bar{t}$ there exists $\bar{n}$ to satisfy the first claim. For any r, the first cutpoint p_1 approaches $\delta = \exp(-\gamma r)$ as n diverges. Therefore, we can choose r small enough and $\bar{n}$ big enough so that the individuals announcing at time zero (i) are an arbitrarily small proportion of individuals and (ii) are arbitrarily well informed. This implies that an arbitrarily high proportion of individuals announce correctly with arbitrarily high probability, satisiying the second claim. ■

Given γ, for any n there is a maximum amount of time $\bar{r}_n$ that a player with an uninformative signal is willing to wait to see the others' announcements. When $r = \bar{r}_n$ agents with every other signal strictly prefer to announce immediately, rather than waiting until time r. When $n = 2$, an uninformed player compares the expected utility of announcing at zero, which is $1/2$, to the expected utility of waiting, in which case he sees the other player's announcement, resulting in an expected utility of $(3/4)\delta$, so that $\bar{r}_2 = (1/\gamma)\ln(3/2)$. By the argument above, as n grows, the expected utility of delay when $r = \bar{r}_n$ converges up to δ, so that $\bar{r}_n$ approaches $(1/\gamma)\ln(2)$. It is in this sense that players become more patient as n increases. However since the amount of time agents are willing to wait does not increase too dramatically with n, it may be possible to bound the length of G_n^r for arbitrary values of n, although we do not have this result.

3.6 Information Aggregation and Efficiency

The model aims to provide an explanation of how rational agents with heterogeneous signal qualities will use time as a screening device to learn about others' private information. We have determined that when agents choose the timing of their decisions strategically, they order themselves in a manner such that those with the best information decide first, and those with worse information are willing to wait in order to learn from early decision makers. Thus it is natural to measure the extent to which information is aggregated more efficiently in the endogenous timing model relative to the standard social learning setup, where an exogenous order is assigned to the players.

For tractability we consider the limiting case $\delta \to 1$ with two players. We first examine the welfare properties of the standard exogenous sequencing model as a benchmark for comparison to the present model. To do so, we suppose that before the game starts, and independently of the signal realizations, agents are randomly assigned a position in the queue, so that each agent is equally likely to be the first or second announcer.

Let Cor_i denote the event that i announces the correct state, i.e., $x_i = \omega$. The probability that the first agent guesses correctly is given by

$$\Pr(Cor_1) = \Pr(Cor_1|A) = \Pr(p_1 \geq 1/2|A) = \int_{1/2}^{1} f(p|A)dp = 3/4.$$

Given the optimal decision rule of player 1, the probability that player 2 announces correctly is 13/16.[20] Note that this is consistent with the observation made above that agents who decide later in the sequence do better on average, in the sense that they announce correctly with higher probability. But deciding later is also costly. To make the fairest comparison with the strategic timing setting, we assume that the value of announcing correctly is 1 for player 1 and $\delta < 1$ for player 2. Thus being the second announcer is preferred if and only if $(13/16)\delta > 3/4$, or for fixed discount rate γ, if $r < (1/\gamma)\ln(13/12)$. The ex ante expected sum of utilities in the exogenous sequencing model is

$$U^{ex}(\delta) = 3/4 + (13/16)\delta.$$

The welfare properties under strategic timing also depend on δ, but in a more complicated way. Specifically, as δ decreases from 1, two things happen: (i) the probability that the players announce simultaneously, and hence no learning occurs, increases (from 0) and (ii) the game typically lasts longer, so greater delay costs are incurred. The exact welfare properties are difficult to work out in closed form for $\delta < 1$, so we consider the limiting case of welfare as $\delta \to 1$, since in this case we can

[20] $\Pr(Cor_2) = \Pr(Cor_2|A) = \Pr(x_2 = a|A, x_1 = a)\Pr(x_1 = a|A) + \Pr(x_2 = a|A, x_1 = b)\Pr(x_1 = b|A) = \Pr(p_2 > 1/4|A)\Pr(p_1 > 1/2|A) + \Pr(p_2 > 3/4|A)\Pr(p_1 < 1/2|A) = (15/16)(3/4) + (7/16)(1/4) = 13/16$.

easily derive explicit comparisons between the two models.

The probability that the first announcer (i.e., the one with higher signal quality) is correct can be worked out as $\Pr(Cor_1) = \Pr(Cor_1|A) = \Pr(p_1 \geq 1/2|A, |p_1 - 1/2| \geq |p_2 - 1/2|) = 5/6$. Since the second decision maker always makes the same announcement as the first, $\Pr(Cor_2) = 5/6$ as well. Thus we find that the difference in social welfare from endogenous timing is $(5/3 - (3/4 + 13/16))$ when there is no discounting, or almost 7%. The intuition is that, conditional on receiving a relatively uninformative signal, the exogenous sequencing forces a player to announce first with probability 1/2, whereas when announcements are timed strategically, with high probability the poorly informed agent will learn from her opponent's signal and thus make a more informed announcement.

As δ decreases from 1, the expected welfare improvement from strategic timing increases so that the difference of 7% as $\delta \to 1$ provides a lower bound on the social utility differential between the models. The reason is that as δ decreases in the exogenous sequencing model, the second decision maker is penalized with certainty, whereas when announcement times are chosen strategically, each player always chooses whether or not to delay and is willing to delay only if she expects to benefit from doing so. For example, if both agents are sufficiently well informed, then both announce at time zero, and no delay costs are incurred. Notice that when δ is sufficiently small (less than 2/3), both agents will announce independently from each other at $t = 0$, regardless of their signals. Thus the ex ante probability that each is correct is 3/4, and therefore $U^{en}(\delta) = 3/2$. In this case, $U^{ex}(\delta) \leq 3/4 + (13/16)(\frac{2}{3}) = 31/24 \approx 1.3$, so that the increase in welfare from endogenous timing is at least 15%. Finally observe that for $\delta < 2/3$ the expected social welfare in the endogenous timing model is constant, whereas it continues to decrease in the exogenous sequencing model. The welfare gains provided by the opportunity to strategically time announcements always fall short of full efficiency for $\delta < 1$ since either a delay cost is incurred by at least one agent, or else both agents announce independently at time 0, in which case no social learning takes place. Yet the limit point of the equilibrium outcomes as δ tends to one is the first best outcome.

A different basis of comparison between the models is to ask how frequently both agents announce the *pooled information state*, that is, the optimal choice if all private information were aggregated prior to announcing. Clearly, this happens with probability one under strategic timing in the limit $\delta \to 1$, since the agent with the more informative signal always announces "first" in equilibrium. On the other hand, in the worst case scenario for the endogenous timing model, δ is small enough that all types announce at 0, no learning occurs, and the pooled information state is announced by both agents with probability $(3/4)^2 + (1/4)^2 = 5/8$. In the exogenous sequencing model, in contrast, this happens about 79% of the time independently of δ.[21] Even though whenever δ is small enough that both agents announce the pooled information state more frequently in the exogenous sequencing model, the effect of greater delay costs dominates this informational advantage, so that expected welfare remains higher under strategic timing.

3.7 Extensions

3.7.1 Continuous Time

Most of the results extend to a continuous time setting in the following way. Consider instead of G^r a game played in continuous time with "reaction lags." That is, players may make an announcement at any time $t \geq 0$, but there is a strictly positive finite time that it takes to observe, process, and react to an announcement by another player. By a slight abuse of notation, call this parameter $r > 0$. Thus if j makes an

[21] Note that the pooled information state is A if and only if $p_2 > 1 - p_1$. Thus one needs to compute $\Pr(p_1 > 1/2, p_2 > 1, 4 \,|\, p_2 > 1 - p_1) + \Pr(p_1 < 1/2, p_2 < 3/4 \,|\, p_2 < 1 - p_1)$ conditional on ω.

announcement at some time t_j, i can not incorporate that information into her own decision until time $t_j + r$. If i announces between t_j and $t_j + r$, it must be because she had already planned to do so, for she did not have time yet to react to j's announcement. In such a case i makes her decision independently of j's announcement. Denote by Γ^r the continuous time game so defined with reaction lag r.

The next lemma states that in any equilibrium of Γ^r, agents will announce only at the discrete times T^r that are allowed in the corresponding discrete time game G^r.

Lemma 1 *Given $\gamma > 0$, fix $r > 0$ and consider any PBNE σ^* of Γ^r. Then T^r has full measure under $\sigma_i^*(p)$ for all $p \in [0, 1]$, $i = 1, 2$.*

Proof. We claim that in any PBNE of Γ^r, agents will make decisions only at times kr, $k = 0, 1, \ldots$ with probability one. Consider times $\tau \in (0, r)$. Any strategy that places positive probability on announcing at such times is strictly dominated by the strategy that is identical to σ_i^* except for shifting that mass to time $t = 0$, since conditional on being correct, the payoff at 0 is strictly greater, and, by hypothesis, no information can be acquired by any $\tau < r$. Therefore beliefs must not change before time r. Thus if an agent does not make an announcement at $t = 0$, she will wait at least until $t = r$. Next consider times $\tau' \in (r, 2r)$. Announcing at such times with positive probability is strictly dominated by announcing at r instead. Given that agents will not announce at times $\tau \in (0, r)$, no information will arrive in $(r, 2r)$, implying that beliefs are again constant on this interval, and conditional on being correct, the payoff at r is strictly greater than at any time after r. Thus if an agent has not announced at r, she will wait at least until $2r$. The claim follows in this fashion by induction on k. Thus in Γ^r agents will announce only at times that are allowed by G^r. ■

Lemma 1 establishes a connection between the games G^r and Γ^r, which can be extended to show that the equilibrium outcomes in discrete and continuous time coincide for any $r > 0$.

3.7.2 Costly Information

Thus far we have presumed that agents are endowed with private information and that the quality of the information varies randomly across agents. Further, we assumed that agents use this information strategically in choosing when to make an announcement and have observed that, in equilibrium, those with better information do better on average. It is then natural to ask by what means agents acquire their information.

In this section, we suppose that information is available to the agents, but that there is a cost to gathering it and that better information is more costly. One way to model this is to introduce an information source that sells signals to players. Suppose that before playing the "announcement game" G^r, players have the opportunity to simultaneously and privately purchase information. Any player who chooses not to buy information would begin the game with the initial prior $P(A)$. The information source offers a menu of contracts, selling signals of higher qualities at higher prices. Recall that $q = |p - 1/2| \in [0, 1/2]$ denotes the quality of a signal p. The menu takes the form of a function $c(q)$ with $c(0) = 0$, $c' > 0$ and $c'' \geq 0$. Conditional on purchasing a quality q, the signal indicates the correct state with probability $q + 1/2$. Thus signals of highest quality, $q = 1/2$, perfectly reveal the true state ω while signals of lowest quality, $q = 0$, which reveal no information, are equally likely to be correct or incorrect.

The value of signal quality is determined by the equilibrium welfare properties of the model. In particular, we must consider the expected utility of a player with signal quality q_i. Again consider the limiting case $\delta \to 1$ since it is the most tractable for this analysis, and the welfare properties of G^r approximate this limit for small r. In this case, the two players always make the same announcement, and this announcement is based solely on the better of the two pieces of information. Thus players' welfare is determined by the quantity $\max\{q_1, q_2\}$. Signal quality can be viewed as a pure public good since both players benefit equally from either player purchasing a high quality signal. Given signal qualities, expected utilities are given by $u_i(q_1, q_2) = 1/2 + \max\{q_1, q_2\} - c(q_i)$.

If $c(q_i) = q_i$ it is easy to check that there is a continuum of asymmetric pure strategy Nash equilibria of the purchasing game, where one agent does not purchase information and the other agent purchases an arbitrary signal quality. There is also a unique symmetric Nash equilibrium, where both agents purchase no information. Note that the strategy of not purchasing ($q_i = 0$) is weakly dominant, and that it is a strict best response to any positive purchase by the other agent.

Allowing more general cost functions reduces the multiplicity problem and alleviates the perverse result that in the symmetric equilibrium, no information is purchased. In particular, assume $c'(0) < 1$, $c'' > 0$, and $c'(1/2) \geq 2$. Then there are exactly two asymmetric equilibria, where one agent chooses $q_i = 0$ and the other chooses $q_j = q^*$ given by $c'(q^*) = 1$. There is also a symmetric equilibrium where each agent mixes between $q_i = 0$ and $q_i = q^*$ with probabilities $c(q^*)/q^*$ and $1 - c(q^*)/q^*$, respectively. Thus the probability of purchasing information decreases with the cost of the optimal signal $c(q^*)$ and increases with the optimal quality q^*. In all these equilibria information is under-provided with probability one relative to the socially optimal arrangement where one agent purchases no information and the other purchases the quality characterized by $c'(q^o) = 2$.

3.7.3 Risk Aversion

We have assumed throughout the discussion that players are risk neutral. One interesting extension is to explore the implications of more general risk preferences. In the risk neutral case, the expected utility of announcing at time zero with belief π is $u(\pi) = \max\{\pi, 1-\pi\}$. A risk averse agent has a utility function satisfying $\tilde{u}(\pi) < \max\{\pi, 1-\pi\}$ for $\pi \in (0,1)$, since the certainty equivalent to the lottery π is less than it's expected value. We would like to be able to describe the qualitative effects that this has on the model's equilibrium predictions. Consider the case $n = 2$ for simplicity. The primary effect is that risk averse agents behave more patiently–they are willing to pay a greater cost in order to reduce uncertainty about the state variable. One can verify this directly by transforming the utility functions in the equation that define the equilibrium cutpoints (see (C.2) in Appendix C) and re-solving. Upon doing this, the result is that the cutpoints shift towards the boundaries so that the maximum number of periods the game may last $\bar{k}(\delta)$ (weakly) increases. That is, the length of the game typically increases with risk aversion. In a particularly extreme form of risk aversion, $\tilde{u}(1/2) \to 0$. Yet even in this case, it can be verified that the game will end in finite time for any $\delta < 1$. Thus the main results do not depend on assuming a particular form of risk preferences. I have not verified that as δ tends to one that the game ends immediately for arbitrary utility functions, but I conjecture that this is true.

3.7.4 Other Utility Functions

The game we have analyzed has a specific payoff structure. One may wish to extend the model by allowing for more general utility functions. There are several interesting directions this could take. The first is to relax the maintained assumption of common interest. That is, one could assume that in addition to the common value component, there is an idiosyncratic private value of choosing either alternative. Indeed most of economic theory presumes that choices reflect private preferences, rather than an attempt to identify a common objective. In a model of exogenous sequencing, Goeree, Palfrey, and Rogers (2006) explore such a setting and find that the true state becomes known with certainty as the number of players diverges. The effect of introducing private values is that the information from observing a decision is reduced, since the decision may reflect an idiosyncratic taste, rather than a signal about the common value aspect. Thus the conjecture is that agents would behave less patiently, since the expected benefit of delay is decreased.

A second possibility is to dispense with the pure information externalities environment. Note that doing so would complicate strategic aspects of the game, since it would introduce signalling issues. That is, when a player cares about how her announcement will effect the decisions of those who announce later, her optimal decision rule is much more difficult to characterize in general. One could assume that the utility of a player depends on the true state and the *group* decision (as defined by an appropriate voting rule), as opposed to the private decision.[22] Such a model could, as a special case, maintain the common interest assumption and capture situations such as committee decisions and jury verdicts. In this case, the motivation to announce quickly given a strong signal could arise either from the desire to signal this fact to those who follow or, as is the case above, from an explicit cost to waiting (or a combination of the two). Extending the model in this direction seems particularly interesting.

One could also imagine settings where individual interests are not perfectly aligned. This could arise in reputational models of information cascades,[23] or in competitive environments, such as a patent race. The strategic timing aspects of these situations are more difficult to analyze, as there exists a tension between announcing correctly

[22] See Hung and Plott (2001) for an experimental study of such an environment.
[23] See, for instance, Scharfstein and Stein (1990) or Zwiebel (1992).

and quickly, and giving away information to others, which could be costly.

3.7.5 Relaxing Assumptions of the Model

The model presented above is somewhat restrictive in a number of ways. However, many of the results presented would still obtain under more general specifications. For instance, the assumption that the initial priors place equal weight on both states is not essential. Suppose instead that $P(A) \in (0, 1)$. Then the signal that results in a flat posterior is $p = 1 - P(A)$, rather than $p = 1/2$. The important part about a player's signal is her resulting posterior beliefs about the true state, which derive partially from her beliefs about the other players' information. Thus this generalization does not alter the fact that the game ends in finite time, and converges to zero length as the time interval vanishes. Note, though, that the definition of information-symmetry would have to be suitably modified. One natural way to do this is to phrase the definition in terms of posteriors rather than signals. In addition, the signal structure could also be suitably generalized. The essential feature of the conditional signal distributions, in addition to having convex support and well-behaved densities, is the monotone likelihood ratio property. That is, we require that higher types yield higher posterior beliefs for state A. It would be possible to prove analogues of the main results under less restrictive assumptions. Note, however, that an explicit welfare comparison between endogenous timing and exogenous sequencing models, as in Section 3.6, would be considerably more difficult.

3.8 Conclusion

This chapter provides a framework in which it is convenient to analyze the properties of social learning where agents with differing signal qualities strategically choose the time at which to make their decisions. One nice feature of the model is a particularly simple and strong equilibrium characterization for the case of two players. The uniqueness of equilibria allows for unambiguous welfare comparisons to the standard model of exogenous sequencing. The primary finding is that expected social welfare is higher when timing is endogenous and that this difference increases with the length of time between consecutive decisions. Another important finding is that, in sharp contrast to the standard model, information is used efficiently. There is no possibility of getting "stuck" in a situation where players may be choosing a suboptimal action despite the prevalence of useful information. Moreover this finding is not shared by similar models. For instance, in Zhang (1997), after the first agent has made an announcement, all agents decide immediately with the same announcement, so that interesting dynamics of social learning may not be apparent.

Because of the necessity to allow instantaneous reactions in a continuous time model, the results of the discrete time model presented here may be more appropriate where there are institutional features that create lags. We may conclude that although a herd forms under strategic timing, the optimal timing decisions of agents are such that the herd is based on the *best* pieces of information in the economy, rather than on a random selection from the private information. This suggests that herds may be more often "right" than would be suspected from the model of exogenous sequencing. A final conclusion of the model may be drawn from the result that the information problem is solved asymptotically as n grows. Thus for large economies, one would expect that after an initial period where the best informed agents announce, at some point announcements occur very fast. This feature of the model may help to explain the dynamics of phenomena such as bank runs and currency crises.

Chapter 4

A Strategic Model of Network Status

4.1 Introduction

People devote much of their time to maintaining relationships with each other. In addition to having certain qualities independent of their social position, some aspects of what a person has to offer depend on who their acquaintances are. For instance, some people are more likely to pass on job offers because of their professional contacts, and others are likely to introduce two friends with common interests because they know many similar people. The process of network formation and its impact on the flow of information and services across individuals or organizations has broad implications. The ease with which one finds information on the world wide web, the spread of infectious and electronic viruses, (un)employment dynamics, international trade, and co-authorship structures, to name just a few examples, have all been studied and shown to crucially depend on network structures.[1] From individuals' viewpoints,

[1]See, for instance, Myers and Shultz (1951), Rees and Shultz (1970) for early studies of employment contacts, Calvó-Armengol and Jackson (2001, 2004) for models of employment dynamics, Adamic (1999) for analysis of www data collected by James Pitkow that demonstrate short path

these relationships affect the range of available opportunities, the amount of information to which they have access, or their relative influence in the system. From a social perspective, the structure of relationships affects overall system performance, such as how quickly and widely information is transmitted, or how efficiently resources are allocated. Thus it is important to understand the incentives facing agents in these situations in order to gain insight into what kinds of structures are likely to form and how these relate to structures that are optimal for the group.

Much has been learned through the study of strategic network formation,[2] yet nearly all of the formal modeling of network formation treats links as binary quantities, that is, links are either present or absent.[3] While this adds tractability to the analysis, it is clearly a simplification of the kinds of links that are observed in many applications of interest. Relationships such as friendships, trade partnerships, and research collaborations are characterized not only by the presence or absence of pairwise interactions, but also by the intensity, frequency, or reliability with which they occur.

The importance of considering links with different strengths originates in the work of Granovetter (1973, 1983). Allowing for the possibility of variable-strength links adds two important ingredients: (i) the ability to study finer details of individual linking decisions and how they adjust to (small) changes in the underlying characteristics of others, and (ii) the possibility to analyze macro-level consequences of the tradeoffs involved in adjusting resources across an actor's maintained links, includ-

lengths, Auerbach et al. (1984) for an analysis of contacts between homosexuals with AIDS, and Casella and Rauch (2001) for an anlysis of international trade.

[2]Jackson and Wolinsky (1996) and Bala and Goyal (2000) are seminal contributions on which much of this literature is based. See Jackson (2003) for a thorough review.

[3]A recent exception is Bloch and Dutta (2005), who derive partial characterizations for a variety of cases and find that optimal linking decisions are either all-or-nothing or involve equal spread. Goyal (2005) discusses implications of allowing for more general link qualities. Goyal, Konovalov, and Morága (2005) consider a special form of variable strength links in the context of joint R&D investments among firms and derive a number of interesting predictions in that setting. Calvo-Armengol and Jackson (2001) model employment dynamics in a network setting that allows for very general information structures that can be interpreted as a weighted graph.

ing what kinds of configurations are stable and efficient for the group in this richer environment.

In settings where network structures are important, some kind of service or information is transferred across actors. When two actors are connected, they each typically receive benefits from the other (though possibly in an asymmetric way). Thus an actor's investment in a link affects the direct flow of benefits not only to himself but also to the other actor involved in the relationship. In the analysis below, we separate these components, or directions, of benefit flow into "taking" (T) and "giving" (G) and analyze these extreme cases. Separating these effects into the taking and giving components proves useful for analyzing the efficiency properties of equilibrium networks. An important lesson is that the taking component is efficient, whereas the giving component need not be. A number of additional questions arise: How should people allocate their time devoted to others, and how do their optimal linking choices change with a change in an individual's intrinsic quality? How are values affected by an increase in linking capacity?

Even though many applications have elements of both, there are some contexts that correspond closely to one or the other cases of unidirectional flow analyzed here, as the following examples argue.

The giving model is relevant where the quantification of reputation is important. Examples include measuring the imputation of trustfulness in online P2P networks, the influence of scientific journals, and determining popularity in social friendship networks. In all these cases agents choose whom to "nominate" or promote by their actions. Calculations similar to the framework presented here have indeed been utilized in some of these contexts.[4] In the case of P2P networks, which are becoming increasingly prevalent in online applications, a prevalent issue is how to trust certain

[4]Beth, Borcherding, and Klein (1994) and Kamvar, Schlosser, and Garcia-Molina (2003) examine the case of trust networks and have results that are suggestive of the model proposed here. Palacio-Huerta and Volij (2004) formally study the problem of journal rankings.

network nodes. A system of recommendations provides one answer to this problem. The model predicts what kinds of recommendation patterns are likely to emerge when nodes make their decisions in order to increase their own level of trust, and as we will see, suggests that these decisions are typically subotimal as far as maximizing the reputation of the group.

The taking model can be viewed as a strategic setting for the model of interpersonal influence proposed by Friedkin (1991, 1998) and Friedkin and Johnsen (1990). Influence is often a desired quality, and so it is natural to think of a model of influence in a strategic setting. In this context the intrinsic values represent the quality of information they possess. Agents choose how to spend their effort gaining access to the value of others. The fact that some people are more active, or better able, to get value from others will be modeled through allowing some agents to maintain a richer set of connections than others. Individuals who have much asked of them are central. The analysis will tie the intrinsic quality and networking ability of the individuals to the idea of centrality. As will be shown, the equilibrium choices in this setting will necessarily be optimal for the group.

To understand the way in which benefits accrue to agents through the network, consider four points. First, each individual has an *intrinsic value.* Second, each individual allocates a resource budget across links to others. Third, the benefit that one individual receives from another is the product of the *total value* of the other agent and the strength of the relevant link. Thus, more valuable agents and stronger links confer greater benefit. Fourth, the total value of each agent is the sum of the intrinsic value and the benefits from all of the connections to other agents. Thus value is achieved both through intrinsic value and through high quality connections to other high value agents.

To motivate this way of constructing utility, think of attempting to quantify scholarly potential among academic economists. One measure of quality is represented by the CV, which contains, e.g., a publication record. This corresponds to the intrinsic quality of individuals. Another important consideration, though, is the content of recommendation letters for hiring or promotion decisions, which represent the choice variables in the model. Clearly, recommendations from different kinds of individuals should be treated differently. In particular, recommendations from people who are "highly regarded" should contribute more to an individual's status. In the model, being highly regarded comes from a combination of intrinsic value (i.e., good publications) and strong recommendations from other highly regarded individuals.

The benefit from a link depends on the value of the agent linked to, and the value of the agent depends on its set of links and the values of its neighbors, and so forth. Consequently, given linking choices, the equations that determine network values must be solved simultaneously in order to compute utilities. Importantly, this is true despite the fact that an agent's utility depends explicitly on only those links that directly involve the agent. The implicit dependence on the remaining links enters through the values of those to whom the agent is directly connected. One implication of this interdependence is the presence of "feedback effects," whereby the benefits associated with a particular link are counted many times. In the previous example, one can think of the feedback effects in the following terms: if individual i recommends individual j, then j's status increases and (s)he is more likely to get a good job, whereupon any recommendation or help that j provides to i becomes more valuable.

Previous work has limited the extent to which link externalities are accounted for. Consider the "connections model" of Jackson and Wolinsky (1996), where the utility one agent derives from another is a function of the length of the shortest path between the agents. There is no added value in having multiple paths (of any length) between agents, even though redundancy may be an important consideration. In contrast, in the present model a small change in the intensity of a link not only has a positive affect on the agents involved in the link, but also has an indirect positive affect on every other agent who is in some way connected to them. A second feature of the connections model is that the potential value of connecting to an agent is independent of the network structure. However in some applications one might desire, for instance, that this benefit depend on the degrees (number, or weight, of connections) of the nodes involved. Unlike the connections model, this model considers such factors in the utility calculation by assuming the recursive definition of values. Thus this work is complimentary to much of the work that has stemmed from variations on the connections model.

Another contribution is that heterogeneity among agents is allowed along two dimensions. Agents may differ both in the measure of intrinsic quality or skill, and also in the amount of resources they can expend forming links.[5] Such heterogeneity is both natural and important for the model. Some individuals have more social interest or ability, and naturally form stronger links to others, while others are more valuable contacts for reasons independent of their connections. These features of the model admit a more interesting analysis of the strategic considerations involved in link formation. For instance, when exogenous qualities are homogenous, the efficient networks are the same under giving and taking behavior, but the distribution of value can be very different. Second, when linking budgets are homogenous, giving behavior is also efficient: budget differences are required to generate the inefficiency mentioned above.

The rest of the chapter proceeds as follows. The next section, Section 4.2 presents a simple example illustrating the intuition for the model. Section 4.3, interprets the network values and connects them with various concepts in related fields. Section 4.4 formally presents the framework for modeling the network game. The main body of

[5]In the binary link case, Galeotti (2004), extending the work of Bala and Goyal (2005), allows for similar kinds of heterogeneity. He finds that such heterogeneity is important in determining equilibrium network configurations, but interestingly, the results differ from the ones presented here.

results, characterizing which network structures are outcomes of Nash equilibria, and which maximize aggregate utility, is contained in Section 4.5. The findings are related to other work through a series of examples in Section 4.6, where the implications of heterogeneity are emphasized. Finally, Section 4.7 concludes. Some technical results and extensions are included in Appendix D.

4.2 A Preview of Network Effects

In this section we present a simple example to illustrate how utility accrues in the model. First we briefly introduce some notation that formalizes the model and allows us to relate this model to concepts in social networks below.

Consider a simple example with three actors, $N = \{1, 2, 3\}$. Each actor has a common intrinsic value denoted α_i and a budget of resources β_i to allocate across links to others, $i \in N$. For simplicity make the normalization $\alpha_i = 1$ for each i. The list of budgets is given by $\beta = (0.5, 0.1, 0.01)$, so that, for instance, agent 1 has five times the linking resources of agent 2. Figure 4.1 depicts the Nah equilibrium and socially optimal networks for two cases: when linking decisions represent the *giving* of benefits (Model G), and when they represent the *taking* of benefits (Model T). Observe that the equilibrium network in Model G is inefficient in the sense that the sum of utilities, $U = u_1 + u_2 + u_3$, is not maximized. In particular, agents invest too heavily in others with large linking budgets relative to the efficient network.[6] In contrast, the equilibrium linking choices are efficient in Model T, a result that holds true in general. Thus by separating the benefit transfers into the giving and taking components, we show that the source of inefficiency is due only to the giving aspect.

Each arrow represents a direct link, with its label denoting the amount of resources invested by the agent initiating the link. We will denote the resources i devotes to the

[6] More specifically, agents 2 and 3 invest too heavily in 1, and 1 invests too heavily in agent 2.

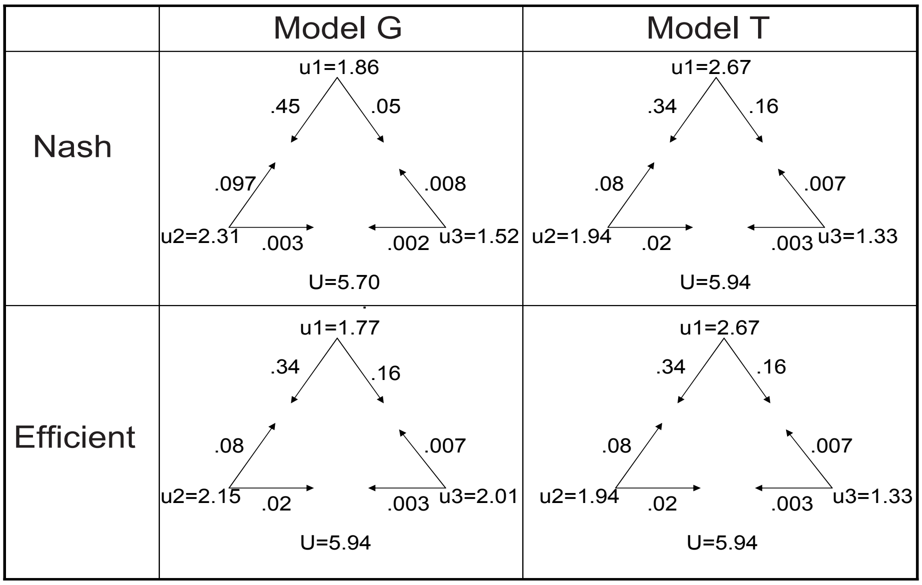

Figure 4.1: *Equilibrium and efficient networks in Models T and G.*

link with j by ϕ_{ij}, so that, for instance, $\phi_{12} = 0.34$ in both networks depicted for Model T. It will also be convenient to collect these choices in a matrix $\Phi = [\phi_{ij}]$. It should also be mentioned at this point that link strengths are derived from the investment choices via a technology function f, so that the strength of the link from i to j is $f(\phi_{ij})$. In this example we use $f(x) = \sqrt{x}$. To illustrate the calculation of utility, consider agent 1in the taking case. We compute $u_1 = \alpha_1 + f(\phi_{12})u_2 + f(\phi_{13})u_3$. Notice that, as mentioned above, computing the utilities involves solving a set of equations in u.

4.3 Interpretation of Network Value: Network Centrality, Interpersonal Influence, and Interdependent Utilities

Although the computation of values is simple, some of the implications deserve comment. In this section we provide some interpretation of the model and relate it to a number of ideas proposed in other contexts.

The idea that value is determined in part by the value of one's neighbors creates the necessity for values to be determined simultaneously, since each individual's value may depend, through paths in the networks, on the values of all other agents. Consider the following illustration. If an agent i connects to another agent j, i derives benefit proportional to j's utility from this link. Agent j's utility is determined in part by how well-connected she is. Thus if j strengthens a connection to a third agent k, j's value and, hence, i's utility increase as well. Furthermore, if agent k increases her value by strengthening one of her connections to yet another agent, this also has a positive externality on i's utility through an increase in agent j's value, since j's neighbor k became more valuable. Of course, this effect is in addition to effects on i through other paths. This logic shows that an increase in the quality of any link increases the utilities of all agents who are not completely disconnected from the augmented link, since an increase in any link quality brings the whole world closer together, benefiting everyone. In this sense the model may be said to take full account of network externalities, since a change in the quality of a link that does not involve a particular agent still has an affect on her utility.

This idea has motivated a literature in the study of social networks that attempts to measure the importance, status, influence, or centrality of a node in a network. The earliest notions of centrality in the social network literature worked on binary networks and did not take all these considerations into account. For instance, one simple measure of centrality is a node's degree, the number (or total weight) of incoming paths. While degree centrality is a useful concept, and certainly captures important features of a node's position, it ignores the origin of the incoming paths.[7] In an effort to allow for such a calculation, Katz (1953) developed the first notion that resembles the calculation of values in this model. If Φ is a binary network of relationships, the measure in Katz (1953) is a scaled version of $(I - \delta\Phi)^{-1}\alpha$, where α is a vector of ones, analogous to homogeneous intrinsic values, and $\delta < 1$ is an attenuation factor. Such an attenuation factor is necessary when dealing with binary networks rather than the weighted graphs studied here or else the contributions from longer paths cause the status measure to diverge to infinity. The fact that the links take fractional values in the present chapter subsumes the role of the attenuation factor. The main innovation of Katz's approach to centrality is that it allows all paths in the network to be considered, rather than calculating the centrality score based on some subset of links, e.g., the set of links pointing to a given node in the case of degree centrality.

Building on Katz's work, Hubbell (1965) introduced the measure $(I - f(\Phi))^{-1}\alpha$, where again, α is assumed to be a vector of ones, but which operates directly on a weighted graph $f(\Phi)$.[8] As Hubbell notes, this formulation mathematically resembles the input-output model of interlinked sectors in an open economy due to Leontief (1951).[9] Hubbell's work introduces a measure of centrality that very closely resembles the formulation of utility here. While Katz's goal was to derive a measure of status, Hubbell aimed his work at clique identification by analyzing the properties of $(I - f(\Phi))^{-1}$ to identify closely connected subgroups. While such an analysis is beyond

[7]Similar comments can be made about many other notions of centrality, including closeness, betweenness and eccentricity, all of which do not take into account a node's importance as a function of its position in the network when computing its contribution to the centrality scores of others. Wasserman and Faust (1994) (ch. 5) provide an excellent review of this literature, and Freeman (1979) discusses their differences in applicability.

[8]Here denote a weighted matrix as $f(\Phi)$ for consistency with the model below.

[9]Some of the tools useful for analyzing these models originated from the study of Leontief's model and in general rely on results about linear systems, a good review of which is Gale (1960).

the scope of this chapter, an interesting question is what kind of clique structures form in the strategic setting studied here.

Related measures have gained popularity through the influential work of Bonacich (1972, 1987, 2005) and others. Bonacich's measure that is most relevant to this model, in that it operates on weighted and asymmetric graphs, is essentially the eigenvector corresponding to the largest eigenvalue of the network $f(\Phi)$. A number of similar indices have been proposed by various authors; Poulin, Boily, and Mâsse (2000) provide a useful review and comparison. Interestingly, many of the measures can be calculated in a number of ways, including solving a linear system of equations, computing an infinite sum of status contributions along paths of different lengths, and a factor analytic approach, thus providing various ways to interpret the centrality scores.

One way to understand our formulation of utility here is to view it as the outcome of a repeated process of interactions occurring along the links of the network. This is the approach taken by Friedkin (1991, 1998) and Friedkin and Johnsen (1990) in studies of interpersonal influence. Under this approach, α is a vector of exogenous qualities, which are shared across links during an infinite sequence of time periods. At date t, the interaction is described by $u^t = \alpha + f(\Phi)u^{t-1}$, with $u^0 = \alpha$. At each date, individuals receive a contribution from their own exogenous quality and contributions proportional to their links to others and others' acquired utility at the time.[10] Under our assumptions, the limit outcome of this process is $u^\infty = (I - f(\Phi))^{-1}\alpha$. As this is the expression for utility functions in the model, utility can be viewed as the long-run result of repeated interactions through the network.

This formulation makes it obvious that the utility vector, or limit outcome of the

[10] In this intertemporal formulation, it is not clear that the Nash equilibrium of the game with utility functions u is the most compelling stability notion, since repeated game considerations may become important. While such an analysis is outside the scope of this chapter, the game corresponding to utilities u is relevant if players are patient, the choices must be made once and for all, and interactions take place repeatedly thereafter.

dynamic process, dominates the vector of exogenous quality α. In the context of interpersonal influence, Friedkin (1991) refers to this effect as polarization, meaning the tendency of opinions to drift towards a more extreme position due to repeated social interactions (rather than to a separation of opinions within the group). The mathematical reason for this is that the value in α is counted multiple times in the computation of the utility vector. For our purposes, such a formulation captures synergies of interactions. For instance, research collaborations and joint R&D projects are often thought to create output beyond what the parties would generate independently. Katz (1953) rejected this kind of effect and proposed instead to normalize the status measure explicitly so that his measure has a more relative interpretation.

The fact that the model here produces Hubbell's status score as utilities provides an interesting re-interpretation of status. That is, in this model agents seek to maximize their status when forming their links, as measured in the way suggested by sociologists. Since the analysis here is strategic, one of the benefits of this chapter is to understand what kinds of networks form when individuals are seeking explicitly to maximize their status in the group. The results on efficiency then highlight when such networks will and will not be efficient, in terms of maximizing the overall status of the group. This chapter is not the first to make an economic connection to the centrality literature. A very different connection is contained in Ballester, Calvó-Armengol, and Zenou (2005), in which the equilibrium action is proportional to Bonacich centrality, taking the network as given. Another paper, de Martí Beltran (2005), is complimentary to this model in that it asks what the optimal way is to distribute intrinsic value among the agents as a function of an exogenously given network.

It is also interesting to note that the model is mathematically quite similar to some of the work on interdependent utilities, as seen in Bergstrom (1989, 1999) and Bramoullé (2001).[11] Under that interpretation, the links in the network represent pa-

[11] Other contributions in this field include Hori (1997) and Shinotsuka (2003).

rameters in utility functions that describe how individuals' consumption utility affects the happiness of others. Again the analysis here is very different principally because that work always takes the network structure as given and addresses interesting questions that do not arise in the present chapter such as the effect of "negative links." On the other hand, this chapter takes individual attributes as given and asks which network structures form Nash equilibrium outcomes and which structures maximize aggregate utility in the group.

4.4 A Model of Link Formation

There is a finite set of agents $N = \{1, \ldots, n\}$, $n \geq 2$, endowed with publicly observed intrinsic qualities $\alpha = (\alpha_1, \ldots, \alpha_n) \in \mathbb{R}^n_+$ and budgets $\beta = (\beta_1, \ldots, \beta_n)$, with $0 < \beta_i < 1$ for each $i \in N$. The agents are players in a *network formation game* with strategies $\phi_i = (\phi_{i1}, \ldots, \phi_{in})$ that satisfy $\phi_{ii} = 0$, $\phi_{ij} \geq 0$ for all $i, j \in N$ and $\sum_j \phi_{ij} \leq \beta_i$ for all $i \in N$. The *intensity* of the directed link ij is given by $f(\phi_{ij})$, where $f(\cdot)$ is increasing and satisfies $f(0) = 0$.

We first distinguish two cases that relate to how link choices ϕ_{ij} translate into link intensities.

Assumption 1 *(Concave specification) The function $f : [0, 1] \to \mathbb{R}$ is continuously differentiable, strictly increasing, strictly concave, $f(0) = 0$, $\lim_{x \to 0} f'(x) = \infty$, and satisfies $(n-1)f(\frac{\max_i \beta_i}{n-1}) < 1$.*

The concave specification reflects the notion of diminishing returns to investment in a link. It is convenient mathematically because it guarantees interior solutions for equilibrium and efficient network structures. Essentially, analogous with the Inada conditions on a production function, the first order conditions characterizing optimal tradeoffs across links are very simple to write down under Assumption 1. An alternative assumption is the following.

Assumption 2 *(Linear specification) The function $f(\cdot)$ is the identity mapping.*

Hence under the linear specification the intensity of link ij is simply ϕ_{ij}.

The bulk of the analysis hinges on which direction benefits flow through the network. In Model T (for "take") agent i setting $\phi_{ij} > 0$ confers benefits to i from j. One can think of i placing phone calls to others in proportion to the $\phi_{ij}s$ asking for information. Given the linking choices and intrinsic values of others, agent i chooses the relative amounts of benefits to receive from others from her budget of resources. The opposite case is handled in Model G (for "give"). There, agent i setting $\phi_{ij} > 0$ confers benefits to j from i. In this case one can think of ϕ_{ij} as representing a nomination of j by i, as in the case of journal article citations. In this case an agent chooses how to give benefit given the linking choices of all others.

The description of the network formation game is completed by formally defining how network structures generate utilities. To this end, collect the strategy profiles in the matrix $\Phi = [\phi_{ij}]$ and let $f(\Phi)$ denote the matrix with elements $f(\phi_{ij})$, i.e., $f(\Phi)$ is the network structure generated by strategy profile Φ. Total utility is the sum of intrinsic value and the value derived through sharing others' values via the network structure. Let v_i denote the (endogenously determined) component of i's utility derived from her relations in the network $f(\Phi)$, so that i's (total) utility is $u_i = \alpha_i + v_i$. In Model T,

$$v_i = \sum_j (\alpha_j + v_j) f(\phi_{ij}),$$

whereas in Model G,

$$v_i = \sum_j (\alpha_j + v_j) f(\phi_{ji}).$$

Notice that network values are linear combinations of the total utilities of other players, weighted by the intensity of the corresponding link (which varies according to the taking or giving component of the model). Thus, for instance, in Model T, $u_i = \alpha_i + v_i = \alpha_i + \sum_j u_j f(\phi_{ij})$. Collecting these equations in matrix notation, we obtain $u = \alpha + f(\Phi)u$, where $u = (u_a, \ldots, u_n)'$ and $\alpha = (\alpha_a, \ldots, \alpha_n)'$ are column vectors of utilities and intrinsic values, respectively. Solving for u yields $u = (I - f(\Phi))^{-1}\alpha$, provided $I - f(\Phi)$ is invertible. Letting $A = (I - f(\Phi))^{-1}$, we have $u = A\alpha$. A similar derivation for Model G produces $u = A'\alpha$. Section **??** contains a more thorough interpretation of the value computation and how it relates to various other models.

Notice that under either Assumption 1 or 2 the requirements on $f(\cdot)$ imply that $I - f(\Phi)$ satisfies the dominant diagonal condition. It is well-known[12] that in this case

$$A = \sum_{p=0}^{\infty} f(\Phi)^p.$$

That is, A can be written as the infinite sum of powers of the network structure matrix $f(\Phi)$. Each of these matrices, in turn, has a natural interpretation. $(f(\Phi)^p)_{ij}$ measures the total weight of all directed paths from i to j that have length p. Thus the matrix A represents the total weight of *all* paths from i to j. Because individual links take values strictly less than unity, longer paths have less weight. This is why the dominant diagonal condition is sufficient for the above sum to converge.

4.5 Results

This section contains the main results of the analysis. In particular, we characterize the equilibrium and efficient networks in both models for the concave and linear cases. Following the literature on strategic network formation, we restrict attention to pure strategy Nash equilibria. The concave specification in Models T and G is presented first, followed by a discussion of the linear case.

[12] See, for example, Gale (1960) for a proof, and other related facts.

4.5.1 Concave Specification

Most of the networks we focus on for the concave specification are *interior* networks, in which all links have positive value: $\phi_{ij} > 0$ for all $i, j \in N$. One can think of these networks as a generalization of the complete network from the binary link setting.

The results for the concave case rely on the following simple lemma, which describes how changes in individual link choices translate into changes in the components of utility.

Lemma 1 $\frac{\partial a_{ij}}{\partial \phi_{kl}} = f'(\phi_{kl}) a_{ik} a_{lj}$.

Proof. See Appendix D. ■

Since utilities are linear combinations of intrinsic qualities, with weights given by the matrix A, Lemma 1 shows how changes in strategies produce changes in utilities, in a form that depends only on the entries of A. In particular, elements of A are strictly increasing in Φ for interior networks.

We proceed with two propositions for Model G, characterizing the Nash networks and the efficient networks, respectively.

Proposition 1 *(Nash networks in Model G) In Model G under the concave specification, there exist multiple Nash equilibria, at least one of which is interior. Interior equilibria satisfy the conditions* $\sum_j \phi_{ij} = \beta_i$ *for all* $i \in N$*, and*

$$f'(\phi_{ij}) a_{ji} = f'(\phi_{ij'}) a_{j'i}$$

for all distinct $i, j, j' \in N$*. Non-interior equilibria have the form that a partition of players exists such that there are no positive links between players in different partitions, and links within a partition constitute an interior equilibrium restricted to that group.*

Proof. See Appendix D. ■

Proposition 2 *(Efficient networks in Model G) In Model G under the concave specification, a socially efficient network exists, is interior, and satisfies the conditions* $\sum_j \phi_{ij} = \beta_i$ *for all* $i \in N$*, and*

$$f'(\phi_{ij}) \sum_k a_{jk} = f'(\phi_{ij'}) \sum_k a_{j'k}$$

for all distinct $i, j, j' \in N$*.*

Proof. See Appendix D. ■

Notice that there are many non-interior Nash networks in Model G, for instance, the empty network, which corresponds to a partition of singletons. The reason is that if an individual i has no paths pointing to him, i.e., if $\phi_{ji} = 0$ for all j, then i is indifferent over his entire strategy space. More generally, $\phi_{ij} = 0$ in equilibrium implies that $a_{ji} = 0$, i.e., that there are no paths of any positive weight from j to i. However, once the assumption of interiority is made, this indifference disappears, and the model makes a sharp prediction.

The characterizations in Propositions 1 and 2 shed light on the (possible) differences between Nash and efficient networks in Model G. Efficiency requires a *weakly regular* network in the following sense: the ratio of marginal link choices to agents j and j', $\frac{f'(\phi_{ij})}{f'(\phi_{ij'})}$, is constant across individuals i. When budgets are homogeneous, this implies that *strong regularity* networks, where $\phi_{ij} = \frac{\beta}{n-1}$ for all $i, j \in N$, are efficient. But in this case, strongly regular networks also form a Nash equilibrium. Interestingly, there may also exist asymmetric interior equilibria even in symmetric environments. Whether or not this occurs depends on the particular shape of f. If it is sufficiently concave, then only the regular network is Nash. The difference between equilibrium and efficient networks arises under heterogeneous budgets because for efficiency, linking choices are determined by total connectivities $\sum_k a_{jk}$, whereas the equilibrium choices depend on individual connectivities a_{ji}.

The next result describes the relationship between Nash equilibria and efficient networks in Model G.[13]

Theorem 1 *Assume $N \geq 3$. In Model G under the concave specification, there is an efficient Nash equilibrium if and only if $\beta_i = \bar{\beta}$ for all $i \in N$.*

Proof. See Appendix D. ■

Recalling the example presented in the introduction, the differences in budgets are necessary (and sufficient) to generate the inefficiency possible under giving behavior.

The possibility of tension between equilibrium and efficiency, however, is not present in Model T. The next result shows that Nash and efficient networks coincide for all parameters under Model T.[14]

Proposition 3 *In Model T under the concave specification, all Nash networks and socially efficient networks are interior. They satisfy the conditions $\sum_j \phi_{ij} = \beta_i$ for all $i \in N$, and*

$$f'(\phi_{ij})u_j = f'(\phi_{ij'})u_{j'}$$

for all distinct $i, j, j' \in N$.

Proof. See Appendix D. ■

The condition in Proposition 3 has a particularly simple structure. Individuals allocate their budgets so as to balance the marginal values of their *direct* links, treating the utility of others as fixed. Model T produces the unusual result that equilibrium networks are necessarily stable. The intuition for this fact can be summarized as

[13] The proof of the "only if" part is not yet complete, so this statement is still a conjecture.

[14] This remark and the discussion after the proposition assume that the first order conditions in Proposition 3 are uniquely satisfied. When this is true, the Nash and efficient networks must coincide. While this has not been proven, it is true in all examples I have computed.

follows. Each individual i makes her linking choices in order to maximize her derived network value. This means sending more effort to individuals with higher values, but sending some effort to everyone because of the concavity of f. Now consider any other individual j linking to i. j also makes her decisions selfishly and benefits from having high value individuals to link to. It turns out that i's optimal decision is also optimal from j's (and every other individual's) point of view. Put differently, if j was given the power to choose ϕ_i as well as ϕ_j, j's choices for i would be the choices that i finds optimal himself. This can be verified directly by considering the conditions $\frac{\partial u_j}{\partial \phi_{ik}} = \frac{\partial u_j}{\partial \phi_{ik'}}$, which generate the same characterization as i's optimization problem from Proposition 3.

This argument of course fails for Model G. Under giving behavior, individuals control their utility only indirectly, in the sense that their linking choices confer direct benefits to others, rather than to themselves. Thus, to optimize, an individual i sends more effort to those agents who have the strongest paths back to i. When other individuals make their linking choices to i, they care not what i's network value is, as in Model T, but instead how strong the paths are from i back to themselves. Thus the linking choices of i are not in general optimal from any other agent's perspective, and this causes the tension between efficiency and equilibrium considerations. Another observation that this logic highlights is that the first order conditions in Model G for stability and efficiency do not depend on utilities or intrinsic values, but only on the network, as measured by the entries in A. We have the following.

Corollary 1 *In Model G under the concave specification, both the efficient and Nash networks are independent of the intrinsic values α.*

Recall that in the example from the introduction, the aggregate utilities at efficient networks were the same under taking and giving. The next result shows that this property depends on having homogeneous intrinsic values.

Corollary 2 *When $\alpha_i = \bar{\alpha}$ for all $i \in N$, the efficient networks in Model T and Model G coincide. Moreover, aggregate utility is the same across models at the efficient solution.*

Proof. See Appendix D. ■

However, there is not a clear relationship between the distribution of utility. Under the taking model, utilities are ranked the same way as budgets, but this need not be true under giving. Recall the example from the introduction where agent 2, who has the middle-size linking budget, has the highest utility in the efficient network.

As mentioned above, the reason for this is that individuals care only about the network structure when making their linking decisions and not about the intrinsic values of the agents they link to. In Model T, however, changing α changes utilities both through the direct effect and also through an equilibrium effect, whereby optimal decisions change in response to the exogenous change.

4.5.2 Linear Specification

It is clear that the concavity of f is crucial in the preceeding analysis. The concave specification makes the analysis easier, since optimal networks are mainly interior, and the first order conditions simplify. This section considers the case of linear returns to investment. The results here are less complete, since we are obliged to restrict attention to special cases of the model's parameters. Nonetheless, it is clear that the shape of f can change the properties of the model substantially.

We begin with a definition of paired networks.

Definition **1** *A network $f(\Phi)$ is said to be paired if for every i $\phi_{ij} = B_i$ for some $j \neq i$ and for every i (except one agent in the case that n is odd), $\phi_{ij} = B_i$ implies $\phi_{ji} = B_j$. The network $f(\Phi)$ is assortatively paired if, in addition, players can be* *labeled so that* $\alpha_1 \geq \cdots \geq \alpha_n$*, and* $\phi_{ij} = B_i$ *implies* $j = i + 1$ *when* i *is odd and* $j = i - 1$ *when* i *is even, for* $i = 1, 2, \ldots, n - 1$*.*

As before, none of the results for Model G depend on the intrinsic values α. We show that under homogeneous budgets, the efficient networks are those that exhaust all the budgets. There must be one equilibrium that is efficient, but there are necessarily inefficient equilibria as well.

Proposition 4 *Assume* $\beta_i = \beta$ *for all* $i \in N$*. In Model G, the efficient networks are those networks for which* $\sum_j \phi_{ij} = \beta_i$ *for all* $i \in N$*. In addition, there are both efficient and inefficient Nash equilibria.*

Proof. See Appendix D. ■

That there are inefficient equilibria is obvious, as the empty network always forms an equilibrium, although it is in weakly dominated strategies. An efficient equilibrium exists in this case since the regular network is both efficient and Nash. The next proposition handles the case where the budgets can be strictly ordered. In this case we can label agents so that $\beta_1 > \cdots > \beta_n$.[15]

Proposition 5 *Assume* $\beta_1 > \cdots > \beta_n$*. In Model G, all paired networks are Nash equilibria. The unique efficient network is assortatively paired.*

Proof. See Appendix D. ■

Thus even in the case of unequal budgets, with linear investment costs there is an efficient equilibrium under giving behavior. We now turn to Model T.

Proposition 6 *Assume* $\alpha_i = \alpha$ *and* $\beta_i = \beta$ *for all* $i \in N$*. The Nash and efficient networks coincide in Model T and are those networks for which* $\sum_j \phi_{ij} = \beta$ *for all* $i \in N$*.*

[15] This is without generality, once we have assumed the generic condition of distinct budgets. Recall that the results are not sensitive to α.

Proof. ■

We can also say something about when either the intrinsic qualities or linking budgets are held constant across individuals, but not both. Notice that generically (i.e., for almost all configurations of α and β, and almost all feasible Φ) utilities will be strictly ordered. In this case, the model necessarily generates a star-like network at equilibrium and at efficient outcomes, since it is strictly optimal for each agent to take everything from the highest-utility agent (except that agent, who chooses the next-highest-utility agent). Under special assumptions, we can identify which agent must have the highest utility, and thus become the center of the star. In particular, when the linking budgets are held constant across agents and there is a strict ordering over intrinsic values, that each agent devotes her entire budget to the unique agent with maximal intrinsic value (and the maximal agent chooses the next highest value). Similarly, a parallel result holds when the intrinsic values are the same and the linking budgets are ranked.

4.6 Network Structures and Heterogeneity

This section presents some examples to illustrate how the forces of the model operate. In particular, we are interested in identifying how particular configurations of modeling parameters map into stylized network architectures that have attracted attention in previous work. In a setting that incorporates variable link intensities, the set of networks is considerably more complex than in the binary link case.

4.6.1 Binary-Link Models

The most closely related work originates in Bala and Goyal (2000), which considers a binary link model close to Model T.[16] When there is no decay and players are identical, Bala and Goyal (2000) show that strict Nash networks are either empty or "wheels."[17] In the presence of decay, a partial characterization shows that strict Nash networks are connected or empty, and include wheel, star, and complete networks, depending on the parameters.[18] For a range of parameters, strict Nash networks are efficient.

Galeotti (2003) extends this analysis to allow for individual heterogeneity of costs and values, which have natural counterparts here in terms of link budgets and intrinsic values. Under the no-decay assumption, strict Nash networks in this model include empty networks, wheels, and center-sponsored stars. When cost heterogeneity takes its most general form, almost any minimally connected network is strict Nash. When decay is allowed, strict Nash networks are difficult to characterize but include almost all minimally connected graphs.

Since the networks identified in the previous section are mostly interior networks,[19] there is not a direct correspondence between the network structures that emerge in this model and the networks studied in binary link models. However, much of the flavor of the networks in the binary case can be seen in this model. Moreover, this framework provides a way to adjust continuously between structures that are reminiscent of prominent binary link networks, such as the complete network and

---

[16]These authors call this "one-way flow," but do not consider the other direction of flow corresponding to Model G. Instead, they also consider "two-way flow," in which a directed link confers benefits to both agents involved. Values are not computed in the recursive way considered here, but depend only on the network only to the extent that it partitions nodes into components.

[17]A wheel is a strongly connected network in which each player maintains exactly one link, that is, players can be numbered so that i links to $i+1$ for all $i = 1, \dots, n-1$, and n links to 1.

[18]Without decay, Bala and Goyal (2000) prove their results for very general payoff functions, but when decay is present, as in the weighted networks in the present model, tractability requires using a linear form as in the connections model.

[19]The only exception is the multiplicity of Nash networks in Model G, many of which are non-interior.

star network, and to tie those adjustments to underlying heterogeneity of individuals. In fact, the networks often focused on in the binary link case are highly stylized and do not match well with the large complex networks that are often cited as motivation or applications. The additional flexibility of this model allows it to predict much more complex relationships among agents that may potentially come much closer to the kinds of networks found in the field.

4.6.2 Empty and Complete Networks

In binary link models, when costs are sufficiently high, empty networks are both efficient and stable. The counterpart to high costs in this context is small linking budgets, in which case all feasible networks converge to an empty network as budgets vanish. In addition, the empty network is Nash in Model G for budgets of every size. The reason for this difference is that the benefit flow in previous work corresponds instead to the taking behavior of Model T.

At the other extreme, provided there is some decay (to allow for networks that are not minimally connected), complete networks emerge for sufficiently small costs. In the present context most Nash networks, and all efficient networks, are interior, which is one form of "completeness," in that every link has positive strength. However, a stronger analogue of the complete binary network is the regular network with "equal spread," in which $\phi_{ij} = \frac{\beta_i}{n-1}$ for all $i, j \in N$. In this case, every link is given equal weight, so that all pairs of agents are directly connected and the network is symmetric in a rather strong sense.

This arises in Model T when both budgets and intrinsic qualities are homogeneous. The underlying symmetry of agents implies that the equilibrium and efficient network is symmetric in this way. However, it is also possible for a regular network to occur in the presence of budget and quality differences across agents, provided that they balance each other in exactly the right way. That is, if an agent has a high intrinsic value, she must have a correspondingly low linking budget in order to induce the other agents to treat her the same as the others. For example, consider 3 nodes with intrinsic values $\alpha = (3, 2, 2)$. When budgets are $\beta = (0.015, 0.1, 0.1)$, and the investment function is $f(x) = \sqrt{x}$, each agent splits her linking budget equally among the other two. However, it is possible that one agent's intrinsic value is sufficiently high that no matter how low her corresponding budget, agents always devote a majority of their effort to the high-value agent. This is the case in the above example if α_1 is increased to 4. The reason is that one agent can become arbitrarily more valuable than the others even on her own without any connections to the value of others.

Model G can also generate regular networks. In fact, it is in some sense biased in that direction because of the fact that the optimal network structures do not depend on intrinsic qualities. Thus, in contrast to Model T, differences in qualities do not translate into differences in linking choices. The requirement for strong regularity in Model G is identical budgets. Recall that this is the same condition that identifies when equilibrium networks are efficient. Hence under giving behavior, regular networks are the only equilibrium networks that are efficient. But there can be asymmetric equilibria even with identical budgets (and intrinsic values). For instance, when $f(x) = \delta x^\lambda$, and λ is sufficiently close to one, a 3-agent equilibrium exists in which two agents devote a vast majority of their budget to the third agent, who splits her own budget equally. This resembles a center-sponsored star.

4.6.3 Wheel and Star Networks

Although the wheel architecture is common in the binary link case, neither Model T nor G produces networks with wheel structures. One reason is the presence of decay considered here, which makes the conditions for generating wheel networks more restrictive in binary link worlds. With decay, the heterogeneity of costs and values must take specific forms to generate a wheel. These forms do not occur in this context. In Model T, the value of linking to any particular agent is the same for everyone; so if it is optimal for one agent to devote a disproportionate share of his resources to a particular agent, then it is optimal for others to do so as well. In Model G, an agent can have linking values that differ across others because it might have stronger paths to some agents compared to others. Nonetheless, there is not a sufficient degree of heterogeneity generated to obtain a wheel-like network in equilibrium.

A well-known network architecture that occurs robustly in "two-way flow" models is the star network, in which one node, the center, is connected to each of the other nodes, and there are no other connections.[20] In the case of directed networks, stars can be either center-sponsored or periphery-sponsored (or a combination of the two), depending on which agents maintain the links. The star network is indicative of the role of "hubs" in real networks, where a few nodes have many connections. The prevalence of hubs has important implications for network performance, mainly through their role in decreasing the distances between nodes in the network.

Interestingly, the star network is not prominent in previous work on one-way flow models, but it is easy to generate star-like networks in Models T and G. For instance, if in either model there is a single agent distinguished by a larger linking budget, then all other agents link more strongly to the distinguished agent, who then takes on a central role. In Model T this happens because the larger budget allows greater access to the value of others, which increases the value of linking to the agent with the large budget. In Model G, the agent with the larger budget creates stronger paths through the network to all other agents, which, on the margin, increases the incentive to link to the distinguished agent. As the common budget level of the peripheral agents

[20] Jackson and Wolinsky (1996) and Bala and Goyal (2000) contain early theoretical treatments, and Goeree, Riedl, and Ule (2005) extend the analysis and report on experiments. They find that heterogeneity plays a major role in helping star networks to form empirically.

vanishes, the prediction becomes exactly the center-sponsored star. On the other hand, holding the budgets of the peripheral agents fixed, as the budget of the center agent grows, peripheral agents devote larger and larger shares of their budgets to the center, as in the periphery-sponsored star. As mentioned above, star-like networks can occur in Model G even in symmetric environments.

A similar effect occurs in Model T when a single agent has a larger intrinsic value. The effects on the linking decisions of others are the same: as the distinguished individual becomes increasingly valuable, the remaining agents find it optimal to devote more of their budget to the center. This situation is more closely analogous to the periphery-sponsored star, since in this case the links pointing out from the center do not increase in strength. The only changes are in the proportion of the budget that periphery agents allocate to the center.

One implication of the model is that increasing either intrinsic value or the linking budget is beneficial to the agent. In Model T, having a low intrinsic quality decreases the incentives for others to link to the agent. But having a large linking budget can more than compensate for what is lacked in intrinsic value. That is, an agent with low intrinsic value can be the center of a star provided it has a large enough linking budget. In this sense, investing in becoming "well-connected" can make up for a low value and create enough value through paths to others to be a focal node in the network.

With regards to efficiency, the results of the previous section are easy to apply. In Model T the star network, when it forms, is always efficient. In Model G, it is never efficient because stars form only when budgets are not homogeneous.

The predictions of the model, while being focused for any particular set of parameters, are very flexible. In particular, many of the prominent networks from the binary case have natural analogues in this environment. Moreover, we can relate these very different outcomes to each other through the role of the underlying value and budget parameters in a way that adjusts smoothly between these predictions.

4.6.4 Fixed Costs, Interiority, and Binary Links

With the exception of some Nash equilibria in Model G, all of the stable and efficient networks identified are interior. In other words, a fairly robust feature of this analysis is that individuals tend to be connected, at least to some extent, to all other individuals in the system. Such a structure may be descriptive of relatively small groups, such as colleagues in an academic department, where people know each other. However, in large systems, such as the national collection of researchers in a field, this prediction is clearly at odds with the observation that large networks tend to be very sparse–each individual is connected to only a small proportion of the system.

One way to account for these facts is to introduce a fixed cost of initiating a link into the model. This is a natural feature to model since we may think that initiating a relationship takes some amount of resources, and only after that expenditure can benefits be realized. For instance, there may be search costs of finding an individual's webpage, or researchers may have to spend some time coordinating their ideas before becoming productive. Another example is that people may need to invest effort in becoming compatible, such as adopting to new software choices, or learning a new language.

To model this idea formally, define a new link technology function $\bar{f}$ from the original function f by $\bar{f}(0) = 0$ and $\bar{f}(x) = c + f(x)$ for $x > 0$, where the fixed cost is $c > 0$. If the fixed cost c is set at an appropriate level relative to the marginal costs, then it would be possible to obtain interior networks for small group sizes n, and to obtain increasingly sparser networks as n grows large. The reason is that the presence of a finite fixed cost sets an upper bound on the number of connections each individual can have. Thus it would be possible–and potentially optimal–to connect to everyone in small groups, whereas this would not be possible in large groups.

An additional advantage of allowing for fixed costs is that it allows a cleaner comparison of the present model with models of binary links. In fact, the model allowing for fixed costs nests, in a sense, both the model presented above and a purely binary link model. To see this, note that as the $c \to 0$ fixed costs vanish and we are left with the original model. On the other hand, with $c > 0$, as the variable costs f' diverge, the model approaches a binary link model in that, while it may be useful to invest a small amount in a link, further strengthening the link becomes prohibitively costly. In other words, optimal decisions in such a model would take the form of nearly homogenous link strengths across β_i/c individuals. An open question with fixed costs is what kind of network structures will form, and in particular, will they be similar to those of binary link models? It should be noted that, even in the limit where the present model approximated a binary link model, the utility specification is different from previous models, and differences in predictions can be expected as a result.

4.7 Conclusion

This is among the first studies of strategic network formation that incorporate links of variable strength. In addition, the computation of network externalities is handled in a way that considers all possible paths in the network, resulting in a formulation that creates a bridge to the sociology literature on centrality in networks. A novel feature of this work is that both equilibrium and efficient networks display interesting heterogeneity in link strength, which other models of variable strength links do not predict. Moreover, variations in link strengths are related to underlying characteristics of individuals in intuitive ways. The model thus provides a better understanding of the relationship between the underlying heterogeneity of individuals and the kinds of network structures that are likely to form.

While in some contexts maintaining a link confers benefits to both parties, there are also situations where the flow of benefits is, at least primarily, in only one direction. By separating the flow of benefits into taking and giving components, we are able to derive new insights into the question of efficiency regarding stable network structures. In particular, the source of inefficiency is identified with giving behavior, in which a link investment by an individual confers benefits to the individual who is linked to. However, the inefficiency is only present when there is heterogeneity across the linking budgets of individuals; when this is not the case, both taking and giving behavior generate efficient choices at equilibrium.

The role of budget constraints in the model is to bound link investments. One way to make the level of investments endogenous is to replace the budget constraints with a (convex) cost function. In this way one would be able to model not only the relative allocation of link investments but also the total amount. A conjecture is that the efficiency result of the taking model would no longer hold, since each link investment generates a positive externality on others which is not accounted for in the individual's decision. Thus, while the relative investments would continue to be socially optimal, the total amount of investments would be inefficiently low.

Another way of making the amount of link investments endogenous is by allowing the resource budget to be strategically spent between self-investment and link investment. The simplest way of accommodating this consideration is by allowing self-linking in the form of $\phi_{ii} > 0$. Extra assumptions would have to be made to ensure that A is well-defined, but otherwise the efficiency result of Model T would remain intact. Another possibility is to allow individuals to allocate some of their resource budget to increasing α_i. Although it is less clear, one may still think the efficiency result will continue to hold in Model T. Enriching the model along these lines is a promising direction of future work.

There are a number of interesting directions for future work. First, binary-link studies have analyzed one-way flow (corresponding to Model T) and two-way flow (where both parties benefit from the link) and have shown that different networks arise at equilibrium in the two cases. Instead, we focus on two different kinds of one-way flow: taking and giving. The model here can be easily extended to allow for two-way flow.

Second, there are a number of models in the literature that predict different kinds of networks corresponding to different kinds of heterogeneity in the population. It will be important to relate the theoretical analysis to empirical studies to begin to determine what kinds of network benefits are more appropriate for modeling. A good example of such work is Goyal, van der Leij, and Moraga-González (2003), who analyze a network of economists linked by joint research papers. Laboratory experiments provide another useful tool for empirical examination of models, and experiments incorporating variable link strengths promise to add insight into the analysis of strong and weak links.

Last, the model could be made more general in terms of the separation of utility flow. Combining the roles of give and take could take a number of forms. One idea is to propose a utility model that is a convex combination of the utilities in the giving and taking cases. The interpretation is that a fixed percentage of each investment ϕ_{ij} is spent benefitting i, and the remainder is spent benefitting j. Whenever there is a non-trivial weight placed on the giving component, inefficiencies will remain in equilibrium. Another idea is to allow giving and taking choices to be made separately and simultaneously. A natural way to bound taking investments is by requiring them to be no larger than giving investments, and assuming that giving is costly. The utility calculations could be carried out similarly by assuming that the total flow from i to j is the sum of what i gives to j and what j takes from i.

Chapter 5

The Role of Clustering in Diffusion on Networks

5.1 Introduction

One of the primary reasons that it is important to study social networks is that they are the systems through which many kinds of things are distributed. Applications are as diverse as contagion of electronic viruses over computer networks, spreading of rumor across citizens of a nation, or adoption of a new behavior among consumers. This chapter considers diffusion of something—a disease or opinion, for instance—through a network of individuals. The goal is to gain a deeper understanding of how the characteristics of a network affect the properties of a diffusion process across its nodes.

Individuals are connected to each other through a series of bilateral relationships. We represent the agents as nodes in a network, and their relationships through links. The diffusion process we consider is based on the Susceptible-Infected-Susceptible (SIS) model common in epidemiology (see Bailey (1975)). There are two states, susceptible and infected, with the former being the default. At the outset, a small proportion of nodes are randomly selected to be infected. In each discrete period that follows, the states of individuals are randomly re-assigned according to the following rule. Susceptible individuals are infected with a probability that depends on the number of neighbors they have in the network and how many of them are currently infected. Each infected individual recovers with some probability. We study the long-run predictions of this process as a function of the correlation structure of the social network.

Since the goal is to understand diffusion in large systems, it is important to study networks that have realistic properties. For this reason we analyze network structures that have features that have been shown to be exhibited by large-scale networks. We focus on two aspects in particular.

The first is concerned with the degree of a node, or the number of links it is involved in. The distribution of degree has been claimed to obey a power law in many applications (see, for example Barabási and Albert (2000)). While this does appear to be generally true in the upper tail of the degree distribution, the rest of the distribution may not be scale-free (see Jackson and Rogers (2006) and the references therein for details). The family of degree distributions we examine includes both scale-free and random networks at the extreme, and a range of distributions in between that fit observed networks remarkably well.

The second aspect is concerned with the correlation structure of the links in the network. In the random networks first studied by Erdös and Rényi (1960), each link is formed with some probability that is independent of the presence or absence of other links. On the other hand, nearly all social networks exhibit correlations (see Newman (2003, 2004) for detailed discussion and findings). One way to measure this is through clustering coefficients, which measure the tendency of two friends of a given node to themselves be connected. Social networks tend to have much higher clustering coefficients than random networks with the same numbers of nodes and links.

The model of Jackson and Rogers (2006) generates networks that are particularly well suited to our analysis for two reasons. First, the networks are realistic. The degree distributions provide the best match to observed networks of any model that we are aware of. Furthermore, the model generates networks with high clustering coefficients, in contrast to other random network models, but as observed in the field. Second, we are able to generate networks that have identical degree distributions but different clustering coefficients, allowing for the first time the possibility of isolating the role of clustering in the diffusion dynamics.

To illustrate the ideas behind the networks we will study in this chapter, consider Figure 5.1. Nodes in corresponding positions each have the same degree in the two networks, so the degree distributions are clearly identical. Yet the network structures are clearly different. Notice that there are no triads in the right-hand network, so that the clustering coefficient is zero. The left-hand network, however, contains two triads out of four triples with at least two edges, so that the total clustering coefficient is one-half. The larger networks analyzed below have essentially this same structure: similar degree distributions with different levels of clustering.

We are concerned with three aspects of the diffusion process. First, it has been shown in other contexts that there is a minimum necessary spreading rate in order to achieve a long-run positive rate of infection. How does this cutoff rate vary with the clustering coefficient? Second, in the case where there is long-run infection, how does the steady state level of infection vary with the clustering coefficient? Last, does the clustering coefficient affect the rate of convergence to the steady state?

There are a number of other studies that have been interested in these properties of diffusion in the context of the SIS model and its variants. Lopez-Pintado (2004) and Pastor-Satorras and Vespignani derive results relating the diffusion process to differences in the degree distribution. Obtaining those results, however, requires

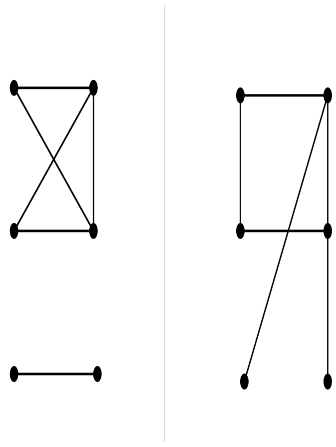

Figure 5.1: *The two networks have identical degree distributions, but different correlation structures. The left-hand network has a total clustering of one-half, and the right-hand network has a total clustering of zero.*

assuming that the link formation process is independent, so that it is not possible to study the role of clustering. Jackson and Yariv (2006) also study a diffusion process that is more incentive-based than the SIS model studied here, but again, the results are in reference to changes in the degree distribution. Eguiluz and Klemm (2002) is the only other work of which we are aware that studies the role of clustering for diffusion. However, that model considers a more restrictive form of the SIS model and is built on a different model of network formation, and our results are not directly comparable.

Holding the degree distribution fixed, we had the following conjectures about how increasing the clustering coefficient affects diffusion dynamics. First, the minimal spreading rate necessary for long-run infection will be lower. Second, the steady-state infection level will be lower. Third, convergence to the steady state will be slower. The reasoning for these conjectures are as follows. When clustering is higher, there is more overlap in the neighborhoods of linked nodes. Thus the total number of nodes reachable in k links from a given node tends to be smaller. It should therefore take more time for initial infection to spread, leading to the third conjecture. Also with high clustering, there tend to be cliques in the population–groups of individuals with many connections among them. Thus some individuals could be in a position where they are consistently being re-infected, leading to the first conjecture. The second conjecture follows similarly, since with more overlap in connections, one might expect there to be larger portions of the population where the disease does not spread to.

The results confirm none of these conjectures. In fact, all of our simulations suggest that clustering has no effect whatsoever on the diffusion process. We formally describe the model in the next section. Then Section 5.3 reports on simulations designed to test our hypotheses. Section 5.4 concludes.

5.2 The Framework

To fix ideas, we will speak of the spreading of disease across people, but it should be understood that the model applies in other settings as well. This section describes the framework in which we will simulate diffusion of a contagious disease through a population. The next subsection describes the structure of the networks connecting individuals. The following subsection formally describes the diffusion mechanism.

5.2.1 Random Networks

We consider a class of networks generated by the model of Jackson and Rogers (2006). The model may be described as follows.

For any node $i \in N$, let $d_i(g) = |\{j \in N \,|\, g_{ji} = 1\}|$ denote the in-degree of i. Let $n_i(g) = \{j \in N \,|\, g_{ij} = 1\}$ denote i's neighborhood.

The network is generated as follows. Action takes place at a countable set of dates $t \in \{1, 2, \ldots\}$. At each time t a new node is added to the population. Let N_t denote the set of all agents or nodes present at time t. Denote by $g(t)$ the network consisting of the links formed on the nodes N_t at the end of time t. That is, $g(t)$ is an $N \times N$ matrix where entry $g(t)_{ij}$ indicates whether a directed link exists from node i to node j. $g(t)_{ij} = 1$ indicates the presence of a directed link while $g(t)_{ij} = 0$ indicates the absence of a directed link.

Links are formed as follows. Let us denote the new node born at time t by t. Upon birth, the node t identifies m_r nodes uniformly at random (without replacement) from N_{t-1}. We shall call these "parent" nodes. The new node forms a directed link to a given parent node with probability p_r.

Additionally, (regardless of whether the node forms a link to the parent) the node t meets other nodes in the parents' neighborhoods. The new node t finds m_n nodes through these network-based meetings (over all parents). Let p_n denote the probability that the new node obtains a positive utility from linking to a given node found through such a network-based meeting.

Let $r = \frac{p_r m_r}{p_n m_n}$ denote the ratio of number of links formed through the random meeting process to those through the network-based process. Denote the expected number of links formed by $m = p_r m_r + p_n m_n$. The following result appears in Jackson and Rogers (2006) and relies on a mean-field approximation to the network formation process. That is, one analyzes a continuous time process in which all changes happen deterministically at a rate equal to their mean rates under the original process. This defines a system of differential equations that can be solved to determine a steady state.

The limiting (as t grows) distribution of degree is

$$F(d) = 1 - \left(\frac{rm}{d + rm}\right)^{1+r}, \tag{5.2.1}$$

for $d \geq 0$.

Next, define the total clustering coefficient on a graph g by

$$C(g) = \frac{\sum_{i;j\neq i;k\neq j,i} \widehat{g}_{ij}\widehat{g}_{jk}\widehat{g}_{ik}}{\sum_{i;j\neq i;k\neq j,i} \widehat{g}_{ij}\widehat{g}_{jk}},$$

where $\widehat{g}_{ij} = \max g_{ij}, g_{ji}$. That is, the clustering coefficient $C(g)$ ignores the direction of links by considering the symmetric graph formed by including a link whenever at least one of the corresponding directed links is present.

Even though the model is constructed as a directed graph, the SIS model as implemented here operates on symmetric relations so that clustering measures that take direction into account are less relevant. An alternative way to define clustering is to do so locally for each node and then average across nodes. We thus define

$$C^{Avg}(g) = \frac{1}{n}\sum_{i} \frac{\sum_{j\neq i;k\neq j,i} \widehat{g}_{ij}\widehat{g}_{jk}\widehat{g}_{ik}}{\sum_{j\neq i;k\neq j,i} \widehat{g}_{ij}\widehat{g}_{ik}}.$$

In Section 5.3 we begin by showing that we can generate networks with the same degree distribution but different clustering measures. This is because the degree distribution depends only on m and r, but clustering is affected by changing p_r, which can be done without altering the degree distribution.

5.2.2 The Diffusion Process

As described above, we analyze an SIS model that is run on a network generated according to the above process. Even though the network is generated in a directed fashion, from this point on we want to consider only the symmetric network $\widehat{g}$. At each date, an infected node recovers with probability δ. Let $a(t)_i$ denote the number of i's neighbors who are infected at time t. Then if i is susceptible at t, i becomes infected with probability $\nu f(n(t)_i), a(t)_i)$.

In the simulations below we use to formulations of the diffusion function. The first is $f(n(t)_i), a(t)_i) = a(t)_i/n(t)_i$. The interpretation is that i interacts with exactly one of her neighbors at each date, so that f gives the probability of interacting with an infected neighbor. There second formulation is $f(n(t)_i), a(t)_i) = a(t)_i$. In this case, the interpretation is that individuals interact with all of their neighbors at each date, so that when ν is small, i becomes infected from each of her neighbors with an independent probability.

5.3 Simulations

We want to consider networks with the same degree distributions but different total and average clustering coefficients. The low clustering (LC) networks are generated according to the following parameters: $T = 250, m_r = 10, p_r = .199, m_n = 2, p_n = .983$. The high clustering (HC) networks use: $T = 250, m_r = 2, p_r = .992, m_n = 2, p_n = .983$. Thus, in both the LC and HC networks, $r = 1.01$ and $m = 3.95$, so that according to the result above, the networks should have the same distributions. However, the parameterizations were chosen so that they would have maximal separation in clustering coefficients. This was accomplished through results in Jackson and Rogers (2006) about limiting clustering coefficients, again based on a mean-field approximation. These results show that clustering increases with p_r provided that $r > 1$ and that the sensitivity is greatest when r and m are small (and $r > 1$).

Ten LC and ten HC networks were generated. To check that their degree distributions are the same, consider Figure 5.2. Shown are average complimentary cdf's for the two kinds of networks, where each point is the average proportion of nodes in the network with at least the degree plotted on the horizontal axis. The LC networks are shown with black dots and the HC networks with a solid curve. The distributions are shown on a log-log scale, where it is clear that the degree distributions are very similar. The slight difference in the upper tail is exaggerated in this scale, and is due to the finite size of the networks.

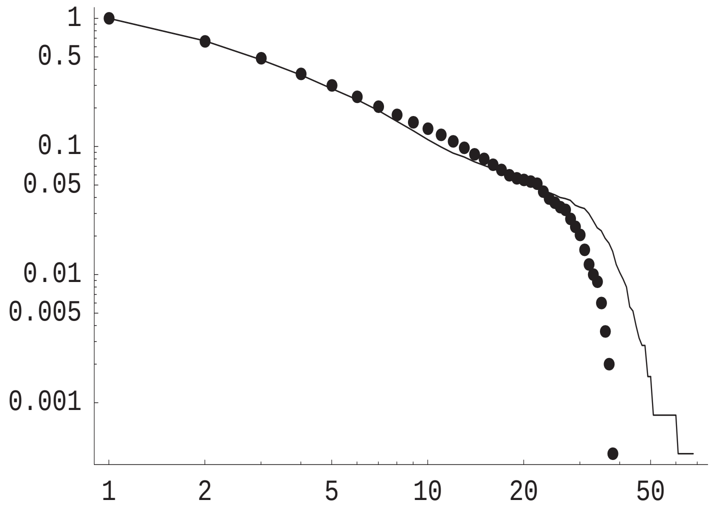

Figure 5.2: *Average degree distributions of LC (black dots) and HC (solid curve) networks. The complimentary cdf is plotted against degree. Each point is the average ccdf taken over the ten networks. The LC and HC networks have very similar degree distributions.*

We next verify that the total and average clustering measures are consistent within the category and systematically higher for the HC networks. This data for all twenty networks is shown in Table 5.1, and it is clear that the clustering is higher in the HC networks as intended.

Extensive simulations have confirmed that there is little if any evidence for an observable effect of clustering on the diffusion process using the LC and HC networks. That is, the dynamics of diffusion look identical in every way across the two different classes of network structures.

The following figures show some examples of SIS simulations on LC and HC

	Total Clustering									
	1	2	3	4	5	6	7	8	9	10
LC	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
HC	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
	Average Clustering									
	1	2	3	4	5	6	7	8	9	10
LC	0.11	0.13	0.10	0.13	0.12	0.12	0.12	0.11	0.13	0.11
HC	0.32	0.32	0.33	0.33	0.32	0.33	0.33	0.34	0.32	0.34

Table 5.1: *Total and average clustering measures for LC and HC networks verifying that the HC networks indeed have higher clustering.*

networks. In the first case we consider the first variation on f corresponding to the case where individuals interact with all of their neighbors each period, and in the second case we consider the f corresponding to the case where individuals interact with one neighbor per period. The first example in Figure 5.3 shows diffusion on both classes of networks converging to the same positive steady-state level of infection, at exactly the same rates. The data represent averages of 200 SIS simulations, where $\nu = 0.04$ and $\delta = 0.1$. Figure 5.4 tells a similar story, but for a parametrization that implies that the infection will die out. In this case, δ has been increased to 0.5, and the initial infection level is 50%, with the other parameters the same as in Figure 5.3. It appears that there may be a very slight effect in the direction that the HC networks may collapse to zero infection somewhat slower.

5.4 Conclusion

This chapter has presented a model that is useful for studying the effects of clustering measures on a dynamic diffusion process. This framework is particularly useful because the network structures studied are consistent on many dimensions with observed

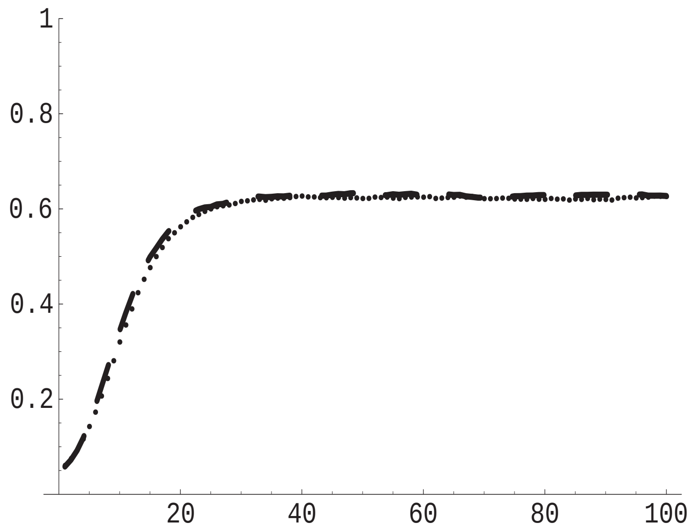

Figure 5.3: *SIS simulations where the proportion of the infected population is plotted against time. LC (dots) and HC (dashed curve) are essentially identical. The data represent averages of 200 simulations on randomly selected networks. Initial infection is 5%,* $\nu = 0.04$*, and* $\delta = 0.1$.

social networks. In addition to fitting data on degree distributions and clustering measures, the networks also fit other features of social networks, including diameter (maximum distance between any pair of nodes measured in the number of links required to go between them), correlation of degrees of linked nodes, and correlations between degree and local clustering. Importantly, it is possible to vary the clustering measures of the network while holding fixed the degree distribution, allowing one to isolate the possible effects of clustering.

Surprisingly, we find very little evidence that clustering affects diffusion on any dimension. The examples presented are, of course, not comprehensive. But similar results hold for different values of ν and δ, as well as for other specifications of the contagion function f. One might also conjecture that increasing the size of the networks may change the results presented here. They appear not to. We have also run simulations on networks with up to 2,000 nodes and have obtained results in line with those presented above. It should be observed that the network formation model we have used has only limited room to vary clustering while holding fixed the degree

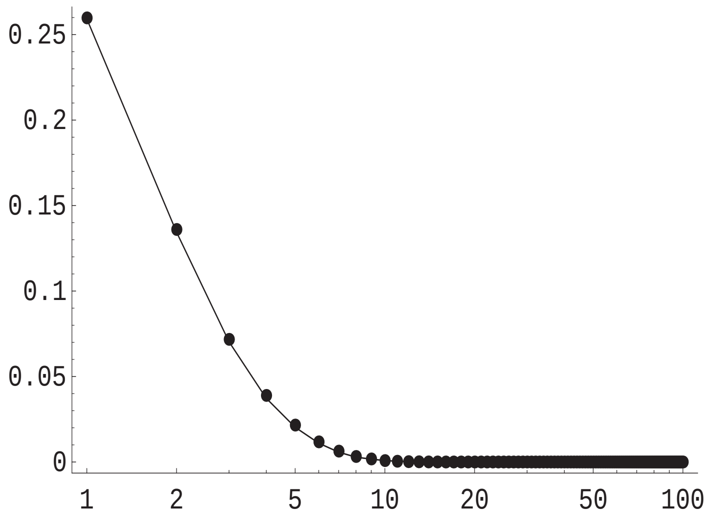

Figure 5.4: *SIS simulations where the proportion of the infected population is plotted against time, coded as in Figure 5.3. Now* $\delta = 0.5$ *and the initial infection is 50%, with the other parameters as in Figure 5.3. Time is plotted on a log scale to make the convergence more clear.*

distribution. This can be seen in Table 5.1, which shows that the clustering measures typically differ by about a factor of two, despite the effort to optimize this dispersion. Future research is needed to see if these results will hold up in the context of other kinds of random networks.

Of course, analytic results would be desirable here. The reason they do not exist is that directly analyzing the discrete dynamic system that governs link formation is a difficult problem. The complication is that the probability a given node attracts a new link depends on its current degree, so that there is a high degree of path-dependence in the stochastics. It is for this reason that researchers have frequently resorted to mean-field approximations. However, the mean-field approach necessarily ignores the correlation structures that are necessary for high clustering, so that the role of clustering can not be studied analytically using these techniques. We hope the simulations provided here are suggestive of deeper results concerning diffusion in complex social networks.

Chapter 6

Appendices

A Appendix to Chapter 2: 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 (2.3.4) 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 (2.3.1) and (2.3.2), 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 et al. (2006) for proof details, and Smith and Sorensen (2000) for a similar argument if there are continuous signals with unbounded beliefs. First, limit points of the stochastic belief process $\{p_t\}_{t=1,2,\cdots}$ 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 to Chapter 2: Estimation Program

In the GAUSS program below we assume that 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 contains 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.[1]

```
PROC loglikelihood(λ);
LOCAL logL,signal,choice,m,t,p,πᵃ,πᵇ,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;
     πᵃ=qp/(qp+(1-q)(1-p));
     πᵇ=(1-q)p/((1-q)p+q(1-p));
     P(A|a)=1/(1+exp(λ(1-2πᵃ))); P(B|a)=1-P(A|a);
     P(A|b)=1/(1+exp(λ(1-2πᵇ))); 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;
```

[1] The 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.

C Appendix to Chapter 3: Collected Proofs

C.1 Beliefs Change Only with Announcements

We claim that, given a signal p_i and information-symmetric strategies s, $\pi_i(p_i, t, h^t, s)$ is constant on all times t and histories h^t such that $x_j^\tau = 0$ for all $\tau < t$. Informally, an agent's posterior beliefs about the true state can change only when the other agent makes an announcement. Note, however, that as time passes, each agent *does* update her beliefs about her opponent's type, since given j's strategy s_j, the fact that the opponent has not yet announced restricts the support of $g(p_j|p_i)$.[2] In particular, let the updated distribution of j's type be $g(p_j|p_i, t, h^t, s)$. Thus:

$$g(p_j|p_i, t, h^t, s) = \frac{g(p_j|p_i)}{\int_{\mathcal{S}(t,h^t,s)} g(p_j|p_i)dp_j} \tag{C.1}$$

where $\mathcal{S}(t, h^t, s)$ denotes the support of $g(p_j|p_i, t, h^t, s)$. By information-symmetry of s_j the set $\mathcal{S}(t, h^t, s)$ is symmetric about 1/2. Thus it suffices to show that given a prior belief p_i, $\pi_i(p_i, t, h^t, s)$ is constant for all histories where neither agent has announced.

Also let $\tilde{\pi}(p_1, p_2) = \frac{p_1 p_2}{p_1 p_2 + (1-p_1)(1-p_2)}$ denote the posterior on A given two conditionally independent signals p_1 and p_2.

Using (C.1) we have:

$$\begin{aligned}
\pi_i(p_i, t, h^t, s) &= \int_{\mathcal{S}(t,h^t,s)} \tilde{\pi}(p_i, p_j) g(p_j|p_i, t, h^t, s)dp_j \\
&= \left(\int_{\mathcal{S}(t,h^t,s)} g(p_j|p_i)dp_j\right)^{-1} \int_{\mathcal{S}(t,h^t,s)} \tilde{\pi}(p_i, p_j) g(p_j|p_i)dp_j \\
&= 2p_i \frac{\int_{\mathcal{S}(t,h^t,s)} p_j dp_j}{2p_i \int_{\mathcal{S}(t,h^t,s)} p_j dp_j + 2(1-p_i)\int_{\mathcal{S}(t,h^t,s)}(1-p_j)dp_j}
\end{aligned}$$

[2] Recall that $g_i(\cdot|p_i, t, h^t, \sigma)$ denotes i's probability assessment of j's type given her own type, the time, the current history, and strategy profile σ, and that $g = (g_1, g_2)$.

$$
\begin{aligned}
&= 2p_i \frac{\int_{\mathcal{S}(t,h^t,s)} p_j dp_j}{2\left((2p_i - 1)\int_{\mathcal{S}(t,h^t,s)} p_j dp_j + (1-p_i)\int_{\mathcal{S}(t,h^t,s)} dp_j\right)} \\
&= p_i,
\end{aligned}
$$

where the last equality follows since:

$$
\begin{aligned}
\int_{\mathcal{S}(t,h^t,s)} p_j dp_j &= (2p_i - 1)\int_{\mathcal{S}(t,h^t,s)} p_j dp_j + (1-p_i)\int_{\mathcal{S}(t,h^t,s)} dp_j \\
(1-p_i)\int_{\mathcal{S}(t,h^t,s)} 2p_j dp_j &= (1-p_i)\int_{\mathcal{S}(t,h^t,s)} dp_j,
\end{aligned}
$$

which follows from the symmetry of $\mathcal{S}(t, h^t, s)$. Thus we have shown that when strategies are information-symmetric, agents learn nothing about the true state solely from the passage of time.

C.2 Main Results

Proposition 3 *Fix $r > 0$ and $\gamma > 0$. G^r has an (essentially) unique PBNE s^*. Moreover, s^* is symmetric, information-symmetric, in pure strategies, and characterized by cutpoints.*

Proof. Define $\tilde{\pi}(p_1, p_2) = \frac{p_1 p_2}{p_1 p_2 (1-p_1)(1-p_2)}$ to be the posterior belief that $\omega = A$ given two conditionally independent signals p_1 and p_2. Also define $f(p_j|p) = 2\left(pp_j + (1-p)(1-p_j)\right)$ to be i's probability assessment of p_j given type p.

Given r and γ set $\delta = \exp(-\gamma r)$. Let $p_0 = 0$ and for $l > 0$ define p_l implicitly by:

$$
1 - p_l = \frac{\delta}{\int_{p_{l-1}}^{1-p_{l-1}} f(p|p_l)dp}\left(\int_{p_{l-1}}^{1-p_l} [1 - \tilde{\pi}(p_l, p)] f(p|p_l)dp + \int_{1-p_l}^{1-p_{l-1}} \tilde{\pi}(p_l, p) f(p|p_l)dp\right) \tag{C.2}
$$

on the domain $(p_{l-1}, 1/2]$. It is straightforward to verify that the right-hand side of (C.2) evaluated at $p_l = p_{l-1}$ is less than $1 - p_{l-1}$ and has derivative strictly greater than -1 on $(p_{l-1}, 1/2]$, implying that there is at most one solution. Let k denote the greatest l for which (C.2) has a solution. We claim that for any $r > 0$ the $\{p_l\}_{l=1}^k$ defined recursively by (C.2) are the cutpoints of the unique perfect Bayes-Nash equilibrium of G^r. Note that by the symmetry of the game, whenever a type $p < 1/2$ has a strictly dominated strategy to announce at some t, announcing at t is also strictly dominated for type $1 - p$. Thus we focus mainly on types $p < 1/2$, although it should be understood that the behavior of types greater than $1/2$ is also being described.

We first prove that for all $l = 1, \ldots, k$, conditional on neither player having announced by time $(l-1)r$, all $p \in (p_{l-1}, p_l)$ (and all $p \in (1 - p_l, 1 - p_{l-1})$) have a strictly dominant strategy to announce immediately. To do so, we compare i's expected payoff of announcing at $(l-1)r$ to the maximal continuation value of delay, i.e., the continuation value that obtains when j plays the continuation strategy that is optimal from i's perspective. Specifically, this strategy is to announce at $(l-1)r$ if and only if $q_j > q_i$. Under this strategy, i learns all decision-relevant information immediately (at time $(l-1)r$) and is thus able to announce the pooled information state at the next time lr. Equation (C.2) compares the value of announcing at $(l-1)r$ to this maximal continuation value. As noted above, (C.2) has either zero or one solution on $p_l \in (p_{l-1}, 1/2]$. If it has no solution, then all types strictly prefer to announce at $(l-1)r$ since the payoff to announcing is greater than the maximal continuation payoff, and the game ends. If it has a solution p_l, then all types less than p_l have a strictly dominant strategy to announce at $(l-1)r$. Thus at $(l-1)r$, all types $p < p_{l-1}$ have already announced, and so they are eliminated from the support in (C.2).

We next prove that all types $\tilde{p} \in (p_l, 1/2]$ strictly prefer to delay at $(l-1)r$. To do so, we first show that for any such $\tilde{p}$, given that all $p \in (p_{l-1}, p_l)$ announce at $(l-1)r$, type $\tilde{p}$ strictly prefers to delay at $(l-1)r$ provided all $p' \in (\tilde{p}, 1/2]$ delay at $(l-1)r$. We do this by comparing the value of announcing at $(l-1)r$ to a lower bound on the continuation value of delaying. Specifically, we use the value at the strategy "delay at

$(l-1)r$, then announce at lr" when j uses a "worst case" strategy. This lower bound obtains when j plays the strategy that *only* types $p < p_l$ announce at $(l-1)r$. In this case, we show that the expected payoff from delaying one period and then announcing at lr (which is a lower bound on the continuation value of delay at $(l-1)r$) in this worst-case scenario is greater than the value of announcing at $(l-1)r$. Consider the following inequality that states this formally.

$$1 - \tilde{p} < \frac{\delta}{\int_{p_{l-1}}^{1-p_{l-1}} f(p|\tilde{p})dp} \left(\int_{p_{l-1}}^{1-p_l} [1 - \pi(\tilde{p}, p)] f(p|\tilde{p})dp + \int_{1-p_l}^{1-p_{l-1}} \pi(\tilde{p}, p) f(p|\tilde{p})dp \right) \tag{C.3}$$

It is straightforward to verify that (C.3) is satisfied for all $\tilde{p} \in (p_l, 1/2]$.

This implies that for type $\tilde{p} = 1/2$ announcing at $(l-1)r$ is strictly dominated. Then since for every strategy s_j of agent j, $Eu_i(t_i, s_j|p_i)$ is continuous in p_i, there exists $\epsilon > 0$ such that announcing at $(l-1)r$ is strictly dominated by the strategy "wait at $(l-1)r$, announce at lr" for all types $1/2 - \epsilon < p' < 1/2$. Let ϵ_1 be the supremum of all ϵ for which this holds. But then by (C.3) the fact that all types strictly above $1/2 - \epsilon_1$ delay implies $1/2 - \epsilon_1$ also delays. Then again by continuity, there exists a largest ϵ_2 such that all $1/2 - \epsilon_1 - \epsilon_2 < p'' < 1/2 - \epsilon_1$ delay. But then type $1/2 - \epsilon_1 - \epsilon_2$ has a strict preference to delay and we can take another open neighborhood to the left who strictly prefer to delay, etc. Define the monotonic sequence $p^\eta = 1/2 - \sum_{n=1}^{\eta} \epsilon_n$. Thus at each step, we know that all $p > p^\eta$ will wait. We claim that the p^η must converge to p_l. To show this, suppose instead that p^η converged to some $p^* > p_l$. Then for every $\epsilon > 0$ there exists an η such that $p^\eta - p^* < \epsilon$. Thus after a finite number η of rounds of iterative elimination of dominated strategies, all types $p \geq p^\eta$ wait at $(l-1)r$. Also note that by (C.3) if in addition all $p^* < p < p^\eta$ waited at $(l-1)r$ then p^* would have a strict preference to wait. But since $p^\eta - p^*$ is arbitrarily small, the probability that $p_j \in (p^*, p^\eta)$ is arbitrarily small, implying that for ϵ sufficiently small, player i with signal p^* strictly prefers to wait given only that types $p \geq p^\eta$ wait at $(l-1)r$.

To summarize, the $\{p_l\}$ defined by (C.2) define the cutpoints of the unique equilibrium of G^r. That is, at each date $(l-1)r$ all types $p < p_l$ have a strict preference to announce, and all types $p > p_l$ have a strict preference to delay. Type p_l is indifferent and thus may choose to delay with an arbitrary probability. But since the set of indifferent types is of measure zero, their choices do not affect the expected utilities of other types. Thus the game has an essentially unique equilibrium and is dominance solvable. ■

Proposition 4 *For every $\gamma > 0$ and $r > 0$, G^r ends in finite time. That is, there exists $k < \infty$ such that $s^*(p) < kr$ for all $p \in [0,1]$ in the unique equilibrium s^*.*

Proof. Given the information-symmetry of s^* from Proposition 3, we confine attention to types $p \leq 1/2$. We can solve (C.2) explicitly for $p_l = p_l(p_{l-1};\delta)$ to obtain the following:

$$p_l = 1/2\left(1 - 1/\delta(1-2p_{l-1}) + \sqrt{(1/\delta - 1)(1-2p_{l-1})}\sqrt{1/\delta(1-2p_{l-1}) + 3 - 2p_{l-1}}\right) \quad \text{(C.4)}$$

Thus (C.4) provides a way to recursively solve for the cutpoint p_l, given only the previous cutpoint and the discount factor δ. Let $v(\delta) = 1/2(3-2/\delta)$. It is straightforward to verify that $p_l(p_{l-1};\delta) > p_{l-1}$ on $[0,1/2)$, is strictly increasing and concave on $[0, v(\delta)]$, and satisfies $p_l(v(\delta);\delta) = 1/2$ and $p_l(0;\delta) > 1/\delta - 1$, so that $\inf_{p_{l-1}\leq v(\delta)} p_l(p_{l-1};\delta) - p_{l-1} = 1/\delta - 1$. Therefore, the number of cutpoints is bounded above by $\frac{\delta}{2(1-\delta)} < \infty$. ■

Before proving Theorem 1, we first prove a useful lemma.

Lemma 2 *Let $\bar{k}(\delta) = \frac{\ln\left(\frac{(1/\delta - 1 + m(\delta))(3-2/\delta(1+1/\delta - m(\delta)))}{2(3-2/\delta)(-1/\delta + m(\delta))}\right)}{-\ln(1/\delta + m(\delta))}$, where $m(\delta) = \sqrt{(1/\delta - 1)(1/\delta + 3)}$. For every $\delta < 1$, $k(\delta) \leq \bar{k}(\delta)$.*

Proof. Let $v(\delta) = \frac{3\delta-2}{2\delta}$. Consider the line $h(p;\delta)$ through $(0, p_l(0;\delta))$ and $(v(\delta), 1/2)$. By concavity of $p_l(\cdot;\delta)$, $h(p;\delta) - p_l(p;\delta) < 0$ on $(0, v(\delta))$. The vertical distance from the 45^o line to $h(p;\delta)$ is $d(p;\delta) = c(\delta)p + p_l(0;\delta)$, where $c(\delta) = \frac{1/\delta - m(\delta)}{3-2/\delta} - 1 \in (-1, 0)$. Consider the sequence $\{z_k\}$, where $z_{k+1} = z_k + d(z_k)$ and $z_0 = 0$. This sequence would describe the successive cutpoints if instead of using $p_l(p;\delta)$ we use the line $h(p;\delta)$ to compute cutpoints. Because $h(p;\delta)$ lies strictly below $p_l(p;\delta)$ on $(0, v(\delta))$, it follows that the first cutpoint z_k defined by $h(p;\delta)$ to lie above $1/2$ occurs no earlier than the first *true* cutpoint p_k, defined by $p_l(p;\delta)$ to do so. Using the linearity of $d(\cdot)$, one can solve the recurrence relation to obtain $z_k = \frac{p_l(0;\delta)}{c(\delta)}((c(\delta)+1)^k - 1)$. The claimed upper bound $\bar{k}(\delta)$ is derived by solving the equation $z_k = 1/2$ for k. ■

Theorem 1 *For every $\gamma > 0$, G^r ends immediately as the time interval vanishes. That is,* $\lim_{r\to 0} rk(\delta) = 0$.

Proof. Fix $\gamma > 0$. By Lemma 2, it is enough to show that $\lim_{\delta\to 1} \frac{-1}{\gamma}\ln(\delta)\bar{k}(\delta) = 0$. Applying L'Hopital's Rule and re-evaluating the limit proves the result. ■

D Appendix to Chapter 4: Collected Proofs

Here we collect proofs that were omitted from the text.

Lemma 1 $\frac{\partial a_{ij}}{\partial \phi_{kl}} = f'(\phi_{kl})a_{ik}a_{lj}$.

Proof. Differentiating $AA^{-1} = I$ yields:

$$\frac{\partial A}{\partial \phi_{kl}}A^{-1} + A\left(\frac{\partial A^{-1}}{\partial \phi_{kl}}\right) = 0.$$

Right-multiplying by A and rearranging produces:

$$\frac{\partial A}{\partial \phi_{kl}} = A\frac{\partial f(\Phi)}{\partial \phi_{kl}}A,$$

after substituting $I - f(\Phi) = A^{-1}$. The result is just the scalar form of this identity. ■

Proposition 1 *(Nash networks in Model G) In Model G under the concave specification, there exist multiple Nash equilibria, at least one of which is interior. Interior equilibria satisfy the conditions* $\sum_j \phi_{ij} = \beta_i$ *for all* $i \in N$*, and*

$$f'(\phi_{ij})a_{ji} = f'(\phi_{ij'})a_{j'i}$$

for all distinct $i, j, j' \in N$.

Proof. The maximization problem of i, taking the strategies of others as given, is:

$$\max_{\phi_i} \sum_k a_{ki}\alpha_k \; s.t. \; \sum_k \phi_{ik} \leq \beta_i.$$

Assume that Φ_{-i} is interior. Thus $a_{ji} > 0$ for all $j \neq i$, so i's utility is strictly increasing in each ϕ_{ij} (by Lemma 1). Thus the budget constraint must bind. The assumption that $\lim_{x \to 0} f(x) = \infty$ insures that any solution ϕ_i^* is interior as well. Therefore the necessary first order conditions are $\frac{\partial u_i}{\partial \phi_{ij}} = \frac{\partial u_i}{\partial \phi_{ij'}}$ for all $j, j' \neq i$. Since $\frac{\partial u}{\partial \phi_{ij}} = \frac{\partial A'}{\partial \phi_{ij}}\alpha$, it follows that:

$$\frac{\partial u}{\partial \phi_{ij}} = -A' \frac{\partial (A')^{-1}}{\partial \phi_{ij}} A'\alpha.$$

Using $A'\alpha = u$ and simplifying shows that:

$$\frac{\partial u}{\partial \phi_{ij}} = f'(\phi_{ij})u_i \begin{pmatrix} a_{j1} \\ \vdots \\ a_{jn} \end{pmatrix}.$$

The i^{th} component of the derivative is simply $f'(\phi_{ij})u_i a_{ji}$. Setting $\frac{\partial u_i}{\partial \phi_{ij}} = \frac{\partial u_i}{\partial \phi_{ij'}}$ for all $j, j' \in N$ produces the claimed expression for i. Existence of an interior equilibrium is established as follows. Define an ϵ-strategy for i as one in which $\phi_{ij} \geq \epsilon$ for all $j \neq i$, and $\sum_k \phi_{ik} = \beta_i$. So the space of profiles of ϵ-strategies is compact and convex. It can be shownthat there exists an $\epsilon > 0$ such that all best responses to ϵ-strategies are ϵ-strategies. An argument parallel to the existence proof in Proposition 3, using $A(\Phi) = (I - f(\Phi))^{-1}$ (which has strictly positive elements on ϵ-strategies) in place of $u(\Phi)$, and the first order conditions above instead of those from Model T, establishes an equilibrium in ϵ-strategies, which is interior.

For non-interior equilibria, notice that $a_{ji} > 0$ implies $\phi_{ij} > 0$, because $\lim_{x \to 0} f(x) = \infty$. This implies that if $a_{ij} > 0$, then for any two agents k, l that have a path to or from i or j, $a_{kl}, a_{lk} > 0$. Consider the finest partition such that there is no positive link across partition elements. The argument implies that within a partition element, links constitute an interior equilibrium restricted to that group. Finally, agents i, j in different elements of the partition, who have no paths between them ($a_{ij} = a_{ji} = 0$), do not have a unilateral profitable deviation of linking to one another. ■

Proposition 2 *(Efficient networks in Model G) In Model G under the concave specification, a socially efficient network exists, is interior, and satisfies the conditions* $\sum_j \phi_{ij} = \beta_i$ *for all* $i \in N$*, and*

$$f'(\phi_{ij}) \sum_k a_{jk} = f'(\phi_{ij'}) \sum_k a_{j'k}$$

for all distinct $i, j, j' \in N$.

Proof. An efficient network exists since the choice sets are compact and U is continuous in Φ. Since $U = \sum_k u_k$ is strictly increasing in Φ, which is implied by Lemma 1, each budget constraint must bind. Also, since $\lim_{x \to 0} f'(x) = \infty$, any optimum must be interior: $\phi_{ij} > 0$ for all i, j. Interiority of the solution implies necessary first order conditions are $\frac{\partial U}{\partial \phi_{ij}} = \frac{\partial U}{\partial \phi_{ij'}}$ for all $i, j, j' \in N$. From the proof of Proposition 1, we have:

$$\frac{\partial U}{\partial \phi_{ij}} = \sum_k \frac{\partial u_k}{\partial \phi_{ij}} = f'(\phi_{ij}) u_i \sum_k a_{jk}.$$

Substitution produces the claimed expression. ■

Theorem 1 *Assume $N \geq 3$. In Model G under the concave specification, there is an efficient Nash equilibrium if and only if $\beta_i = \bar{\beta}$ for all $i \in N$.*

Proof. (If) Assume $\beta_i = \bar{\beta}$ for all $i \in N$. We prove that the strongly regular network is both Nash and efficient. In this case, $a_{ji} = a_{j'i}$ for all distinct $j, j', i \in N$, and $a_{ii} = a_{jj}$ for all $i, j \in N$. Thus the first order conditions in Propositions 1 (Nash) and 2 (efficiency) are identical and reduce to $f'(\phi_{ij}) = f'(\phi_{ij'})$ for all i and all $j, j' \neq i$, which is satisfied at the strongly regular network. Verifying the second order conditions for a maximum is straight-forward.

(Sketch of only if) Next assume there exist i, j so that $\beta_i \neq \beta_j$. The only Nash equilibrium candidate for efficiency is an interior one, so it must satisfy the first order conditions in Proposition 1. $\beta_i \neq \beta_j$ implies that $a_{ji} \neq a_{j'i}$ for some j, j', i at an interior equilibrium, since a strongly regular network can not be an equilibrium. Consider agents with the largest ratio $\frac{a_{ji}}{a_{j'i}}$. Then $\frac{a_{ji}}{a_{j'i}} > \frac{\sum_k a_{jk}}{\sum_k a_{j'k}}$. This implies that i's Nash decision is inefficient, and we are done. ■

Proposition 3 *In Model T under the concave specification, all Nash networks and socially efficient networks are interior. They satisfy the conditions $\sum_j \phi_{ij} = \beta_i$ for all $i \in N$, and*

$$f'(\phi_{ij}) u_j = f'(\phi_{ij'}) u_{j'}$$

for all distinct $i, j, j' \in N$.

Proof. Existence of an efficient network is clear since U is continuous and the strategy sets are compact. Lemma 1 implies that each u_i, and so also $U = \sum_i u_i$, is strictly increasing in each ϕ_{ij}, so that all budget constraints must bind in a Nash or efficient network. Further, the assumption that $\lim_{x\to 0} f'(x) = \infty$ implies that any Nash or efficient network is interior. Therefore, maximizing U has necessary first order conditions $\frac{\partial U}{\partial \phi_{ij}} = \frac{\partial U}{\partial \phi_{ij'}}$ for all distinct $i, j, j' \in N$. We have

$$\frac{\partial U}{\partial \phi_{ij}} = \sum_k \frac{\partial u_k}{\partial \phi_{ij}} = f'(\phi_{ij}) u_j \sum_k a_{ki}.$$

Substituting into the first order conditions then produces the claimed conditions.

To show that a Nash equilibrium exists, let S_i denote i's strategy space, and set $S = \times S_i$. Define $BR_i(\bar{u}) : \mathbb{R}^n_{++} \to S_i$ as the unique ϕ_i that satisfies $\sum_j \phi_{ij} = \beta_i$ and $f'(\phi_{ij})\bar{u}_j = f'(\phi_{ij'})\bar{u}_{j'}$ for all $j, j' \neq i$, for given constants $\bar{u} > 0$. Set $BR(\bar{u}) = (BR_1(\bar{u}), \ldots, BR_n(\bar{u}))$. Let $u(\Phi) = (I - f(\Phi))^{-1}\alpha$ denote utilities as a function of choices. Then $Q : S \to S$ defined by $Q(\Phi) = BR(u(\Phi))$ has a fixed point by Brouwer's theorem, since S is convex and compact, and Q is continuous. Existence follows because a fixed point $\Phi^* = BR(u(\Phi^*))$ of Q is an equilibrium by construction. Necessary first order conditions for i's optimization problem, taking Φ_{-i} as given, are $\frac{\partial u_i}{\partial \phi_{ij}} = \frac{\partial u_i}{\partial \phi_{ij'}}$ for all distinct $j, j' \neq i$. We have:

$$\frac{\partial u}{\partial \phi_{ij}} = \frac{\partial A}{\partial \phi_{ij}}\alpha = A\frac{\partial f(\Phi)}{\partial \phi_{ij}}A\alpha \tag{D.5a}$$

$$= f'(\phi_i j)u_j \begin{pmatrix} a_{1i} \\ \vdots \\ a_{ni} \end{pmatrix}. \tag{D.5b}$$

Substitution into the first order conditions produces the same conditions as before, and we are done. ■

Corollary 2 *When $\alpha_i = \bar{\alpha}$ for all $i \in N$, the efficient networks in Model T and Model G coincide. Moreover, aggregate utility is the same across models at the efficient* *solution.*

Proof. Assume $\alpha_i = \bar{\alpha}$ for all $i \in N$. In Model T, $u_j = \bar{\alpha} \sum_k a_{jk}$. Therefore the first order conditions in Proposition 3 are identical to those in characterizing efficiency in Model G in Proposition 2. Thus the efficient networks coincide. In both Models, aggregate utility is proportional to the sum of all elements in A by the constant $\bar{\alpha}$ and so is the same across models. ■

Proposition 4 *Assume $\beta_i = \beta$ for all $i \in N$. In Model G, the efficient networks are those networks for which $\sum_j \phi_{ij} = \beta_i$ for all $i \in N$. In addition, there are both efficient and inefficient Nash equilibria.*

Proof. Since U is strictly increasing in each ϕ_{ij}, the condition that $\sum_j \phi_{ij} = \beta_i$ for all $i \in N$ is clearly necessary for efficiency. To show that it is also sufficient, write $U = \sum_i \sum_j a_{ji}\alpha_j = \sum_j \alpha_j \sum_i a_{ji}$. When the budget constraints are exhausted, the row sums of Φ are each β. This implies that the row sums of A are each $\frac{1}{1-\beta}$.[3] Thus $U = \frac{\sum_j \alpha_j}{1-\beta}$ for all such networks. Since U is constant on these networks, they must all be efficient. For the second claim, the empty network is Nash and is obviously not efficient. Last, we claim that the strongly regular network, in which $\phi_{ij} = \frac{\beta}{n-1}$ for all $i \neq j$, is Nash. When $N \backslash \{i\}$ set each link equal to $\frac{\beta}{n-1}$, then $a_{ki} = a_{k'i}$ for all $k, k' \neq i$. Thus $\frac{\partial a_{ji}}{\partial \phi_{ik}} = a_{ji}a_{ki}$ is constant across k. Thus a_{ji} is constant across all strategies of i that exhaust its budget. Since any other strategy decreases u_i, i has no profitable deviation from the strongly regular network. ■

Proposition 5 *Assume $\beta_1 > \cdots > \beta_n$. In Model G, all paired networks are Nash equilibria. The unique efficient network is assortatively paired.*

Proof. In a paired network, $a_{ji} = 0$ for all but at most one j. When there is a (unique) $a_{ji} > 0$, i's unique best response is to devote the entire linking budget to j.

[3] See Wilansky (1951) for a result that implies this fact.

If n is odd and there is an agent i such that $a_{ji} = 0$ for all $j \neq i$, then i is indifferent over all strategies. It can be verified by direct computation that the assortatively paired network generates maximal total utility out of all paired networks . ■

Proposition 6 *Assume $\alpha_i = \alpha$ and $\beta_i = \beta$ for all $i \in N$. The Nash and efficient networks coincide in Model T and are those networks for which $\sum_j \phi_{ij} = \beta$ for all $i \in N$.*

Proof. Under these assumptions, $u_i = \alpha \sum_j a_{ij}$. A necessary condition for either equilibrium or efficiency is that $\sum_j \phi_{ij} = \beta$ for all $i \in N$. In this case, the row sums of A are $\frac{1}{1-\beta}$ (see the proof of Proposition 4). Thus $u_i = \frac{\alpha}{1-\beta}$ for every network that exhausts budgets and is strictly smaller otherwise. This also proves that $U = \frac{n\alpha}{1-\beta}$ when budgets are exhausted and is strictly smaller otherwise. ■

References

Adamic, L.A. (1999) "The Small World Web," *Proceedings of the ECDL* vol 1696 of Lecture Notes in CS, 443-454.

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.

Auerbach, D., W. Darrow, H. Jaffe, and J. Curran (1984) "Cluster of cases of the acquired immune deficiency syndrome: patients linked by sexual contacts," *American Journal of Medicine* 76, 487-492.

Bailey, N.T.J. (1975) *The Mathematical Theory of Infectious Diseases*, Griffin: London.

Bala, V. and S. Goyal (2000) "A Non-cooperative Model of Network Formation," *Econometrica* 5, 205-228.

Ballester, G., A. Calvó-Armengol, and Y. Zenou (2005) "Who's Who in Networks. Wanted: The Key Player," mimeo.

Banerjee, A. (1992) "A Simple Model of Herd Behavior," *Quarterly Journal of Economics* 107, August, 797-817.

Barabási, A., R. Albert, and H. Jeong (2000), "Scale-Free Characteristics of random networks: the topology of the world-wide web," *Physica A* 281, 69-77.

Becker, G. (1974) "A Theory of Social Interaction," *Journal of Political Economy* 82(6), 1063-1093.

Bergstrom, T. (1989) "Puzzles: Love and Spaghetti, The Opportunity Cost of Virtue," *The Journal of Economic Perspectives* 3(2), 165-173.

Bergstrom, T. (1999) "Systems of Benevolent Utility Functions," *Journal of Public Economic Theory* 1(1), 71-100.

Bernardo, A. and I. Welch (2001) "On the Evolution of Overconfidence and Entrepreneurs," *Journal of Economics and Management Strategy* 10(3), 301-330.

Beth, T., M. Borcherding, and B. Klein (1994) "Valuation of Trust in Open Networks," *Proceedings of the European Symposium on Research in Computer Security,* LNCS 875, 3-18, Springer-Verlag: Brighton, UK.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy* 100 (October), 992-1026.

Bloch, F., and B. Dutta (2005) "Communication Networks with Endogenous Link Strength," mimeo: GREQAM.

Bonacich, P. (1972) "Factoring and Weighting Approaches to Status Scores and Clique Identification," *Journal of Mathematical Sociology* 2, 113-120.

Bonacich, P. (1987) "Power and Centrality: A Family of Measures," *The American Journal of Sociology* 92(5), 1170-1182.

Bonacich, P., and P. Lloyd (2006) "Eigenvector-like Measures of Centrality for Asymmetric Relations," mimeo.

Borgatti, S. (2002) "Centrality and Network Flow," mimeo.

Bramoulle, Y. (2001) "Interdependent Utilities, Preference Indeterminacy, and Social Networks," mimeo.

Braun, N. (1997) "A Rational Choice Model of Network Status," *Social Networks* 19, 129-142.

Calvó-Armengol, A., and M. Jackson (2001) "Networks in Labor Markets: Wage and Unemployment Dynamics and Inequality," *Journal of Economic Theory*, forthcoming, www.hss.caltech.edu/∼jacksonm/mobility.pdf.

Calvó-Armengol, A., and M. Jackson (2004) "The Effects of Social Networks on Employment and Inequality," *American Economic Review* 94(3): 426-454.

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 Model of Behavior in Games," *Quarterly Journal of Economics* 119, 861-898.

Caplin, A., and J. Leahy (1994) "Business as Usual, Market Crashes, and Wisdom After the Fact," *American Economic Review* 84, 548-565.

Casella, A., and J. Rauch (1999) "Anonymous Market and Group Ties in International Trade," mimeo, Columbia University and University of California at San Diego.

Çelen, B. and S. Kariv (2004), "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," *American Economic Review* 94, 484- 497.

Çelen, B., S. Kariv, and A. Schotter (2005), "An Experimental Test of Advice and Social Learning," Working Paper, NYU.

Chamley, C. (2004) "Delays and Equilibria with Large and Small Information in Social Learning," *European Economic Review* 48(3), 477-501.

Chamley, C., and D. Gale (1994) "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica* 62, 1065-1085.

Choi, S., D. Gale, and S. Kariv (2005), "Learning in Networks: An Experimental

Study," Working Paper, NYU.

Cipriani, M. and A. Guarino (2005) "Herd Behavior in a Laboratory Financial Market," *American Economic Review* 95, 1427-1443.

Dasgupta, A. (2000) "Social Learning with Payoff Complementarities," Working Paper: Yale University.

Dasgupta, P., and E. Maskin (1986a) "The Existence of Equilibrium in Discontinuous Games, I: Theory," *Review of Economic Studies* 53, 1-26.

Dasgupta, P., and E. Maskin (1986b) "The Existence of Equilibrium in Discontinuous Games, II: Applications," *Review of Economic Studies* 53, 27-42.

de Martí Beltran, J. (2005) "Group Influence: Efficient Allocations and Distributive Conflict," mimeo.

DeMarzo, P., D. Vayanos, and J. Zwiebel (2003) "Persuasion Bias, Social Influence, and Unidimensional Opinions," *The Quarterly Journal of Economics* August, 909-968.

Drehmann, M., J. Oechssler, and S. Roider (2005) "Herding and Contrarian Behavior in Financial Markets: An Internet Experiment," *American Economic Review* 95, 1403-1426.

Dominitz, J. and A. Hung (2004) "Homogeneous Actions and Heterogeneous Beliefs: Experimental Evidence on the Formation of Information Cascades," Working Paper, Carnegie Mellon University.

Edwards, W. (1968) "Conservatism in Human Information Processing," in *Formal Representation of Human Judgement*, ed. B. Kleinmuntz, New York: Wiley, 17-52.

Eguíluz, V., and K. Klemm (2002) "Epidemic Threshold in Structured Scale-Free

Networks," *Physical Review Letters* 88(10).

Erdös, P. and A. Rényi (1960) "On the Evolution of Random Graphs," Publication of the Mathematical Institute of the Hungarian Academy of Sciences, 5, 17-61.

Freeman, L. (1979) "Centrality in Social Networks: Conceptual Clarification," *Social Networks* 1, 215-239.

Friedkin, N. (1991) "Theoretical Foundations for Centrality measures," *The American Journal of Sociology* 96(6), 1478-1504.

Friedkin, N. (1998) "A Structural Theory of Social Influence," Cambridge: Cambridge University Press.

Friedkin, N., and E. Johnsen (1990) "Social Influence and Opinions," *Journal of Mathematical Sociology* 15(3-4), 193-205.

Gale, D. (1960) "The Theory of Linear Economic Models," New York: McGraw-Hill.

Gale, D. (1996) "What Have We Learned from Social Learning?" *European Economic Review* 40, 617-628.

Galeotti, A. (2004) "One-way Flow Networks: The Role of Heterogeneity," Tinbergen Institute Discussion Paper.

Goeree, J., C. Holt, and T. Palfrey (2005) "Regular Quantal Response Equilibrium," *Experimental Economics* 8, 347-367.

Goeree, J., T. Palfrey, and B. Rogers (2006) "Social Learning with Private and Common Values," *Economic Theory* 28(2), 245-264.

Goeree, J., A. Riedl, and A. Ule (2005) "In Search of Stars: Network Formation among Heterogeneous Agents," mimeo.

Goyal, S. (2005) "Strong and Weak Links," *Journal of the European Economic Association* 3(2-3): 608-616.

Goyal, S., A. Konovalov, and J. L. Moraga-González (2005) "Hybrid R&D," mimeo.

Goyal, S., M. van der Leij, and J. L. Moraga-González (2003) "Economics: An Emerging Small World?" mimeo: University of Essex.

Granovetter, M. (1973) "The Strength of Weak Ties," *American Journal of Sociology* 78: 1360-1380.

Granovetter, M. (1973) "The Strength of Weak Ties: A Network Theory Revisited," *Sociological Theory* 1, 201-233.

Griffin, D. and A. Tversky (1992) "The Weighting of Evidence and the Determinants of Confidence" *Cognitive Psychology* 24, 411-35.

Gul, F., and R. Lundholm (1995) "Endogenous Timing and the Clustering of Agents' Decisions," *Journal of Political Economy* 103, 1039-1066.

Hori H. (1997) "Dynamic Allocation in an Altruistic Overlapping Generations Economy," *Journal of Economic Theory* 73, 292:315.

Hubbell, C. (1965) "An Input-Output Approach to Clique Identification," *Sociometry* 28(4), 377-399.

Huck, S. and J. Oechssler (2000) "Informational Cascades in the Laboratory: Do they occur for the right reasons?" *Journal of Economic Psychology* 21(6), 661-671.

Hung, A. and C. Plott (2001) "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions," *American Economic Review* 91, 1508-20.

Jackson, M. (2003) "A Survey of Models of Network Formation: Stability and Efficiency," *Group Formation in Economics; Networks, Clubs and Coalitions,* ed. G. Demange and M. Wooders, Cambridge University Press: Cambridge.

Jackson, M. (2005) "The Economics of Social Networks," mimeo.

Jackson, M., and B. Rogers (2006) "Meeting Strangers and Friends of Friends: How Random are Social Networks?" Caltech Social Sciences Working Paper 1216.

Jackson, M., K. Simon, J. Swinkels, and W. Zame (2002) "Communication and Equilibrium in Discontinuous Games of Imcomplete Information," *Econometrica* 70(5), 1711-1740.

Jackson, M., and A. Wolinsky (1996) "A strategic model of social and economic networks," *Journal of Economic Theory* 71, 44-74.

Jackson, M., and L. Yariv (2006) "Diffusion on Social Networks," Caltech Social Sciences Working Paper 1251.

Kahneman, D. and A. Tversky (1973) "On the Psychology of Prediction," *Psychological Review* 80, 237-51.

Kamvar, S., M. Schlosser, and H. Garcia-Molina (2003) "The EigenTrust Algorithm for Reputation Management in P2P Networks," ACM.

Kapur, S. (1994) "Technological Diffusion with Social Learning," Working Paper: BirkBeck College.

Kariv, S. (2005) "Overconfidence and Information Cascades," Working Paper, University of California, Berkeley, CA.

Katz, L. (1953) "A New Status Index Derived from Sociometric Analysis," *Psychometrika* 18(1), 39-43.

Kolm, S. (1969) "The Optimal Production of Social Justice," in *Public Economics,* ed. J. Margolis and H. Guitton. London: MacMillan, 145-200.

Kübler, D. and G. Weizsäcker. (2004) "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory," *Review of Economic Studies* 71, 425-41.

Kübler, D. and G. Weizsäcker (2005) "Are Longer Cascades More Stable?" *Journal of the European Economic Association (Papers and Proceedings)* 3, 330-39.

Leontief, W. (1951) "Input-Output Economics," *Scientific American* 15-21.

Lopez-Pintado, D. (2004) "Diffusion in Complex Social Networks," Universidad de Alicante WP-AD 2004-33.

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.

McKenzie, L. (1960) "Matrices with Dominant Diagonals and Economic Theory," in *Mathematical Methods in the Social Sciences,* ed. K. Arrow, S. Karlin, and P. Suppes. Stanford: Stanford University Press, 47-62.

Milgrom, P. (1979) "A Convergence Theorem for Competitive Bidding with Differential Information," *Econometrica* 47(3), 679-688.

Milgrom, P., and R. Weber (1987) "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research* 10, 619-632.

Myers, C., and G. Shultz (1951) *The Dynamics of a Labor Market*, Prentice-Hall: NY.

Newman, M. (2003) "The structure and function of complex networks," *SIAM Review* 45, 167-256.

Newman, M. (2004) "Coauthorship networks and patterns of scientific collaboration," *Proceedings of the National Academy of Sciences* 101, 5200-5205.

Nikaido, H. (1968) "Convex Structures and Economic Theory," New York: Academic Press.

Nöth, M. and M. Weber (2003) "Information Aggregation with Random Ordering: Cascades and Overconfidence," *Economic Journal* 113, 166-89.

Oberhammer, C. and A. Stiehler (2003) "Does Cascade Behavior Reflect Bayesian Updating: An Experimental Study," Working Paper, Max Planck Institute of Economics.

Palacio-Huerta, I., and O. Volij (2004) "The Measurement of Intellectual Influence," *Econometrica,* 72(3): 963-977.

Pastor-Satorras, R., and A. Vespignani (2000) "Epidemic Spreading in Scale-Free Networks," *Physical Review Letters* 86(14), 3200-3203.

Poulin, R., M.-C. Boily, and B.R. Mâsse (2000) "Dynamical Systems to Define Centrality in Social Networks," *Social Networks* 22: 187-220.

Rees, A. (1966) "Information Networks in Labor Markets," *American Economic Review*, 56: 559-566.

Rees, A., and G. Shultz (1970) *Workers in an Urban Labor Market,* University of Chicago Press: Chicago.

Reny, P. (1999) "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games," *Econometrica* 67(5), 1029-1056.

Scharfstein, D., and J. Stein (1990) "Herd Behavior and Investment," *American Economic Review* 80, 465-479.

Shinotsuka, T. (2003) "Interdependent Utility Functions in an Intergenerational Context," mimeo.

Simon, L., and M. Stinchcombe (1989) "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica* 57(5), 1171-1214.

Smith, L., and P. Sørenson (2000) "Pathological Outcomes of Observational Learning," *Econometrica* 68(March), 371-398.

Smith, L., and P. Sørenson (2001) "Informational Herding and Optimal Experimentation," Working Paper, University of Michigan.

Stahl, D. and P. Wilson (1995), "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior* 10(July), 218-54.

Wasserman, S., and K. Faust (1994) *Social Network Analysis: Methods and Applications,* Cambridge University Press.

Weizsäcker, G. (2003) "Ignoring the Rationality of Others: Evidence From Experimental Normal Form Games," *Games and Economic Behavior* 44(July), 145-71.

Wilansky, A. (1951) "The Row-Sums of the Inverse Matrix," *The American Mathematical Monthly* 58(9), 614-615.

Zhang, J. (1997) "Strategic Delay and the Onset of Investment Cascades," *Rand Journal of Economics* 28, 188-205.

Zwiebel, J. (1992) "Corporate Conservatism, Herd Behavior and Relative Compensation," *Journal of Political Economy* 103, 1-25.