CHAPTER 3

Global Games: Theory and Applications

Stephen Morris and Hyun Song Shin

1. INTRODUCTION

Many economic problems are naturally modeled as a game of incomplete information, where a player's payoff depends on his own action, the actions of others, and some unknown economic fundamentals. For example, many accounts of currency attacks, bank runs, and liquidity crises give a central role to players' uncertainty about other players' actions. Because other players' actions in such situations are motivated by their beliefs, the decision maker must take account of the beliefs held by other players. We know from the classic contribution of Harsanyi (1967–1968) that rational behavior in such environments not only depends on economic agents' beliefs about economic fundamentals, but also depends on beliefs of higher-order – i.e., players' beliefs about other players' beliefs, players' beliefs about other players' beliefs about other players' beliefs, and so on. Indeed, Mertens and Zamir (1985) have shown how one can give a complete description of the "type" of a player in an incomplete information game in terms of a full hierarchy of beliefs at all levels.

In principle, optimal strategic behavior should be analyzed in the space of all possible infinite hierarchies of beliefs; however, such analysis is highly complex for players and analysts alike and is likely to prove intractable in general. It is therefore useful to identify strategic environments with incomplete information that are rich enough to capture the important role of higher-order beliefs in economic settings, but simple enough to allow tractable analysis. Global games, first studied by Carlsson and van Damme (1993a), represent one such environment. Uncertain economic fundamentals are summarized by a state θ and each player observes a different signal of the state with a small amount of noise. Assuming that the noise technology is common knowledge among the players, each player's signal generates beliefs about fundamentals, beliefs about other players' beliefs about fundamentals, and so on. Our purpose in this paper is to describe how such models work, how global game reasoning can be applied to economic problems, and how this analysis relates to more general analysis of higher-order beliefs in strategic settings.

One theme that emerges is that taking higher-order beliefs seriously does not require extremely sophisticated reasoning on the part of players. In Section 2, we present a benchmark result for binary action continuum player games with strategic complementarities where each player has the same payoff function. In a global games setting, there is a unique equilibrium where each player chooses the action that is a best response to a uniform belief over the proportion of his opponents choosing each action. Thus, when faced with some information concerning the underlying state of the world, the prescription for each player is to hypothesize that the proportion of other players who will opt for a particular action is a random variable that is uniformly distributed over the unit interval and choose the best action under these circumstances. We dub such beliefs (and the actions that they elicit) as being *Laplacian*, following Laplace's (1824) suggestion that one should apply a uniform prior to unknown events from the "principle of insufficient reason."

A striking feature of this conclusion is that it reconciles Harsanyi's fully rational view of optimal behavior in incomplete information settings with the dissenting view of Kadane and Larkey (1982) and others that rational behavior in games should imply only that each player chooses an optimal action in the light of his subjective beliefs about others' behavior, without deducing his subjective beliefs as part of the theory. If we let those subjective beliefs be the agnostic Laplacian prior, then there is no contradiction with Harsanyi's view that players should *deduce* rational beliefs about others' behavior in incomplete information settings.

The importance of such analysis is not that we have an adequate account of the subtle reasoning undertaken by the players in the game; it clearly does not do justice to the reasoning inherent in the Harsanyi program. Rather, its importance lies in the fact that we have access to a form of short-cut, or heuristic device, that allows the economist to identify the actual *outcomes* in such games, and thereby open up the possibility of systematic analysis of economic questions that may otherwise appear to be intractable.

One instance of this can be found in the debate concerning self-fulfilling beliefs and multiple equilibria. If one set of beliefs motivates actions that bring about the state of affairs envisaged in those beliefs, while another set of self-fulfilling beliefs bring about quite different outcomes, then there is an apparent indeterminacy in the theory. In both cases, the beliefs are logically coherent, consistent with the known features of the economy, and are borne out by subsequent events. However, we do not have any guidance on which outcome will transpire without an account of how the initial beliefs are determined. We have argued elsewhere (Morris and Shin, 2000) that the apparent indeterminacy of beliefs in many models with multiple equilibria can be seen as the consequence of two modeling assumptions introduced to simplify the theory. First, the economic fundamentals are assumed to be common knowledge. Second, economic agents are assumed to be certain about others' behavior in equilibrium. Both assumptions are made for the sake of tractability, but they do much more besides.

They allow agents' actions and beliefs to be perfectly coordinated in a way that invites multiplicity of equilibria. In contrast, global games allow theorists to model information in a more realistic way, and thereby escape this strait-jacket. More importantly, through the heuristic device of Laplacian actions, global games allow modelers to pin down which set of self-fulfilling beliefs will prevail in equilibrium.

As well as any theoretical satisfaction at identifying a unique outcome in a game, there are more substantial issues at stake. Global games allow us to capture the idea that economic agents may be pushed into taking a particular action because of their belief that others are taking such actions. Thus, inefficient outcomes may be forced on the agents by the external circumstances even though they would all be better off if everyone refrained from such actions. Bank runs and financial crises are prime examples of such cases. We can draw the important distinction between whether there can be inefficient equilibrium outcomes and whether there is a unique outcome in equilibrium. Global games, therefore, are of more than purely theoretical interest. They allow more enlightened debate on substantial economic questions. In Section 2.3, we discuss applications that model economic problems using global games.

Global games open up other interesting avenues of investigation. One of them is the importance of public information in contexts where there is an element of coordination between the players. There is plentiful anecdotal evidence from a variety of contexts that public information has an apparently disproportionate impact relative to private information. Financial markets apparently "overreact" to announcements from central bankers that merely state the obvious, or reaffirm widely known policy stances. But a closer look at this phenomenon with the benefit of the insights given by global games makes such instances less mysterious. If market participants are concerned about the reaction of other participants to the news, the public nature of the news conveys more information than simply the "face value" of the announcement. It conveys important strategic information on the likely beliefs of other market participants. In this case, the "overreaction" would be entirely rational and determined by the type of equilibrium logic inherent in a game of incomplete information. In Section 3, these issues are developed more systematically.

Global games can be seen as a particular instance of equilibrium selection though perturbations. The set of perturbations is especially rich because it turns out that they allow for a rich structure of higher-order beliefs. In Section 4, we delve somewhat deeper into the properties of general global games – not merely those whose action sets are binary. We discuss how global games are related to other notions of equilibrium refinements and what is the nature of the perturbation implicit in global games. The general framework allows us to disentangle two properties of global games. The first property is that a unique outcome is selected in the game. A second, more subtle, question is how such a unique outcome depends on the underlying information structure and the noise in the players' signals. Although in some cases the outcome is sensitive to the details of the information structure, there are cases where a particular outcome is selected and where this outcome turns out to be robust to the form of the noise in the players' signals. The theory of "robustness to incomplete information" as developed by Kajii and Morris (1997) holds the key to this property. We also discuss a larger theoretical literature on higher-order beliefs and the relation to global games.

In Section 5, we show how recent work on local interaction games and dynamic games with payoff shocks use a similar logic to global games in reaching unique predictions.

2. SYMMETRIC BINARY ACTION GLOBAL GAMES

2.1. Linear Example

Let us begin with the following example taken from Carlsson and van Damme (1993a). Two players are deciding whether to invest. There is a safe action (not invest); there is a risky action (invest) that gives a higher payoff if the other player invests. Payoffs are given in Table 3.1:

Table 3.1. *Payoffs of leading example*

	Invest	NotInvest
Invest	θ, θ	$\theta - 1, 0$
NotInvest	$0, \theta - 1$	$0, 0$

(2.1)

If there was complete information about θ, there would be three cases to consider:

- If $\theta > 1$, each player has a dominant strategy to invest.
- If $\theta \in [0, 1]$, there are two pure strategy Nash equilibria: both invest and both not invest.
- If $\theta < 0$, each player has a dominant strategy not to invest.

But there is incomplete information about θ. Player i observes a private signal $x_i = \theta + \varepsilon_i$. Each ε_i is independently normally distributed with mean 0 and standard deviation σ. We assume that θ is randomly drawn from the real line, with each realization equally likely. This implies that a player observing signal x considers θ to be distributed normally with mean x and standard deviation σ. This in turn implies that he thinks his opponent's signal x' is normally distributed with mean x and standard deviation $\sqrt{2}\sigma$. The assumption that θ is uniformly distributed on the real line is nonstandard, but presents no technical difficulties. Such "improper priors" (with an infinite mass) are well behaved, as long as we are concerned only with conditional beliefs. See Hartigan (1983) for a discussion of improper priors. We will also see later that an improper

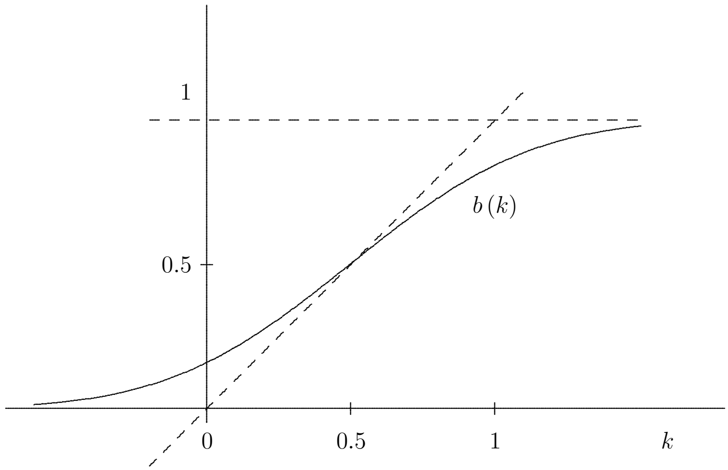

Figure 3.1. Function $b(k)$.

prior can be seen as a limiting case *either* as the prior distribution of θ becomes diffuse *or* as the standard deviation of the noise σ becomes small.

A strategy is a function specifying an action for each possible private signal; a natural kind of strategy we might consider is one where a player takes the risky action only if he observes a private signal above some cutoff point, k:

$$s(x) = \begin{cases} \text{Invest}, & \text{if} \quad x > k \\ \text{NotInvest}, & \text{if} \quad x \leq k. \end{cases}$$

We will refer to this strategy as the switching strategy around k. Now suppose that a player observed signal x and thought that his opponent was following such a "switching" strategy with cutoff point k. His expectation of θ will be x. He will assign probability $\Phi\big(1/\sqrt{2}\sigma(k - x)\big)$ to his opponent observing a signal less than k [where $\Phi(\cdot)$ is the c.d.f. of the standard normal distribution]. In particular, if he has observed a signal equal to the cutoff point of his opponent ($x = k$), he will assign probability $\frac{1}{2}$ to his opponent investing. Thus, there will be an equilibrium where both players follow switching strategies with cutoff $\frac{1}{2}$.

In fact, a switching strategy with cutoff $\frac{1}{2}$ is the unique strategy surviving iterated deletion of strictly interim-dominated strategies. To see why,[1] first define $b(k)$ to be the unique value of x solving the equation

$$x - \Phi\left(\frac{k - x}{\sqrt{2}\sigma}\right) = 0. \tag{2.2}$$

The function $b(\cdot)$ is plotted in Figure 3.1. There is a unique such value because the left-hand side is strictly increasing in x and strictly decreasing in k. These

[1] An alternative argument follows Milgrom and Roberts (1990): if a symmetric game with strategic complementarities has a unique symmetric Nash equilibrium, then the strategy played in that unique Nash equilibrium is also the unique strategy surviving iterated deletion of strictly dominated strategies.

properties also imply that $b(\cdot)$ is strictly increasing. So, if your opponent is following a switching strategy with cutoff k, your best response is to follow a switching strategy with cutoff $b(k)$. We will argue that if a strategy s survives n rounds of iterated deletion of strictly dominated strategies, then

$$s(x) = \begin{cases} \text{Invest}, & \text{if} \quad x > b^{n-1}(1) \\ \text{NotInvest}, & \text{if} \quad x < b^{n-1}(0). \end{cases} \tag{2.3}$$

We argue the second clause by induction (the argument for the first clause is symmetric). The claim is true for $n = 1$, because as we noted previously, NotInvest is a dominant strategy if the expected value of θ is less than 0. Now, suppose the claim is true for arbitrary n. If a player knew that his opponent would choose action NotInvest if he had observed a signal less than $b^{n-1}(1)$, his best response would always be to choose action NotInvest if his signal was less than $b(b^{n-1}(1))$. Because $b(\cdot)$ is strictly increasing and has a unique fixed point at $\frac{1}{2}$, $b^n(0)$ and $b^n(1)$ both tend to $\frac{1}{2}$ as $n \to \infty$.

The unique equilibrium has both players investing only if they observe a signal greater than $\frac{1}{2}$. In the underlying symmetric payoff complete information game, investing is a risk dominant action (Harsanyi and Selten, 1988), exactly if $\theta \geq \frac{1}{2}$; not investing is a risk dominant action exactly if $\theta \leq \frac{1}{2}$. The striking feature of this result is that no matter how small σ is, players' behavior is influenced by the existence of the ex ante possibility that their opponent has a dominant strategy to choose each action.[2] The probability that either individual invests is

$$\Phi\left(\frac{\frac{1}{2} - \theta}{\sigma}\right);$$

Conditional on θ, their investment decisions are independent.

The previous example and analysis are due to Carlsson and van Damme (1993a). There is a many-players analog of this game, whose solution is no more difficult to arrive at. A continuum of players are deciding whether to invest. The payoff to not investing is 0. The payoff to investing is $\theta - 1 + l$, where l is the proportion of other players choosing to invest. The information structure is as before, with each player i observing a private signal $x_i = \theta + \varepsilon_i$, where the ε_i are normally distributed in the population with mean 0 and standard deviation σ. Also in this case, the unique strategy surviving iterated deletion of strictly dominated strategies has each player investing if they observe a signal above $\frac{1}{2}$ and not investing if they observe a signal below $\frac{1}{2}$. We will briefly sketch why this is the case.

Consider a player who has observed signal x and thinks that *all his opponents* are following the "switching" strategy with cutoff point k. As before, his expectation of θ will be x. As before, he will assign probability $\Phi((k - x)/\sqrt{2}\sigma))$ to

[2] Thus, a "grain of doubt" concerning the opponent's behavior has large consequences. This element has been linked by van Damme (1997) to the classic analysis of surprise attacks of Schelling (1960), Chapter 9.

any given opponent observing a signal less than k. But, because the realization of the signals are independent conditional on θ, his expectation of the proportion of players who observe a signal less than k will be exactly equal to the probability he assigns to any one opponent observing a signal less than k. Thus, his expected payoff to investing will be $x - \Phi((k-x)/\sqrt{2}\sigma)$, as before, and all the previous arguments go through.

This argument shows the importance of keeping track of the layers of beliefs across players, and as such may seem rather daunting from the point of view of an individual player. However, the equilibrium outcome is also consistent with a procedure that places far less demands on the capacity of the players, and that seems to be far removed from equilibrium of any kind. This procedure has the following three steps.

- Estimate θ from the signal x.
- Postulate that l is distributed uniformly on the unit interval [0, 1].
- Take the optimal action.

Because the expectation of θ conditional on x is simply x itself, the expected payoff to investing if l is uniformly distributed is $x - \frac{1}{2}$, whereas the expected payoff to not investing is zero. Thus, a player following this procedure will choose to invest or not depending on whether x is greater or smaller than $\frac{1}{2}$, which is identical to the unique equilibrium strategy previously outlined. The belief summarized in the second bullet point is *Laplacian* in the sense introduced in the introductory section. It represents a "diffuse" or "agnostic" view on the actions of other players in the game. We see that an apparently naive and simplistic strategy coincides with the equilibrium strategy. This is not an accident. There are good reasons why the Laplacian action is the correct one in this game, and why it turns out to be an approximately optimal action in many binary action global games. The key to understanding this feature is to consider the following question asked by a player in this game.

> "My signal has realization x. What is the probability that proportion less than z of my opponents have a signal higher than mine?"

The answer to this question would be especially important if everyone is using the switching strategy around x, since the proportion of players who invest is equal to the proportion whose signal is above x. If the true state is θ, the proportion of players who receive a signal higher than x is given by $1 - \Phi((\psi - \theta)/\sigma)$. So, this proportion is less than z if the state θ is such that $1 - \Phi((\psi - \theta)/\sigma) \leq z$. That is, when

$$\theta \leq x - \sigma\Phi^{-1}(1-z). \tag{2.4}$$

The probability of this event conditional on x is

$$\Phi\left(\frac{x - \sigma\Phi^{-1}(1-z) - x}{\sigma}\right) = z.$$

In other words, the cumulative distribution function of z is the identity function, implying that the density of z is uniform over the unit interval. If x is to serve as the switching point of an equilibrium switching strategy, a player must be indifferent between choosing to invest and not to invest given that the proportion who invest is uniformly distributed on [0, 1].

More importantly, even away from the switching point, the optimal action motivated by this belief coincides with the equilibrium action, even though the (Laplacian) belief may not be correct. Away from the switching point, the density of the random variable representing the proportion of players who invest will not be uniform. However, as long as the payoff advantage to investing is increasing in θ, the Laplacian action coincides with the equilibrium action. Thus, the apparently naive procedure outlined by the three bulleted points gives the correct prediction as to what the equilibrium action will be. In the next section, we will show that the lessons drawn from this simple example extend to cover a wide class of binary action global games.

We will focus on the continuum player case in most of this paper. However, as suggested by this example, the qualitative analysis is very similar irrespective of the number of players. In particular, the analysis of the continuum player game with linear payoffs applies equally well to any finite number of players (where each player observes a signal with an independent normal noise term). Independent of the number of players, the cutoff signal in the unique equilibrium is $\frac{1}{2}$. However, a distinctive implication of the infinite player case is that the outcome is a deterministic function of the realized state. In particular, once we know the realization of θ, we can calculate exactly the proportion of players who will invest. It is

$$\widehat{\xi}(\theta) = 1 - \Phi\left(\frac{\frac{1}{2} - \theta}{\sigma}\right).$$

With a finite number of players (I), we write $\xi_{\lambda,I}(\theta)$ for the probability that at least proportion λ out of the I players invest when the realized state is θ:

$$\xi_{\lambda,I}(\theta) = \sum_{n \geq \lambda I} \binom{I}{n} \left[\Phi\left(\frac{\frac{1}{2} - \theta}{\sigma}\right)\right]^{I-n} \left[1 - \Phi\left(\frac{\frac{1}{2} - \theta}{\sigma}\right)\right]^{n}.$$

Observe, however, that the many finite player case converges naturally to the continuum model: by the law of large numbers, as $I \to \infty$,

$$\xi_{\lambda,I}(\theta) \to 1 \quad \text{if} \quad \lambda < \widehat{\xi}(\theta)$$

and

$$\xi_{\lambda,I}(\theta) \to 0 \quad \text{if} \quad \lambda > \widehat{\xi}(\theta).$$

2.2. Symmetric Binary Action Global Games: A General Approach

Let us now take one step in making the argument more general. We deal first with the case where there is a uniform prior on the initial state, and each player's signal is a sufficient statistic for how much they care about the state (we call this the private values case). In this case, the analysis is especially clean, and it is possible to prove a uniqueness result and characterize the unique equilibrium independent of both the structure and size of the noise in players' signals. We then show that the analysis can be extended to deal with general priors and payoffs that depend on the realized state.

2.2.1. Continuum Players: Uniform Prior and Private Values

There is a continuum of players. Each player has to choose an action $a \in \{0, 1\}$. All players have the same payoff function, $u : \{0, 1\} \times [0, 1] \times \mathbb{R} \to \mathbb{R}$, where $u(a, l, x)$ is a player's payoff if he chooses action a, proportion l of his opponents choose action 1, and his "private signal" is x. Thus, we assume that his payoff is independent of which of his opponents choose action 1. To analyze best responses, it is enough to know the payoff gain from choosing one action rather than the other. Thus, the utility function is parameterized by a function $\pi : [0, 1] \times \mathbb{R} \to \mathbb{R}$ with

$$\pi(l, x) \equiv u(1, l, x) - u(0, l, x).$$

Formally, we say that an action is the *Laplacian* action if it is a best response to a uniform prior over the opponents' choice of action. Thus, action 1 is the *Laplacian* action at x if

$$\int_{l=0}^{1} u(1, l, x)dl > \int_{l=0}^{1} u(0, l, x)dl,$$

or, equivalently,

$$\int_{l=0}^{1} \pi(l, x)dl > 0;$$

action 0 is the *Laplacian* action at x if

$$\int_{l=0}^{1} \pi(l, x)dl < 0.$$

Generically, a continuum player, symmetric payoff, two-action game will have exactly one Laplacian action.

A state $\theta \in \mathbb{R}$ is drawn according to the (improper) uniform density on the real line. Player i observes a private signal $x_i = \theta + \sigma\varepsilon_i$, where $\sigma > 0$. The noise terms ε_i are distributed in the population with continuous density $f(\cdot)$, with support on the real line.[3] We note that this density need not be symmetric around the mean, nor even have zero mean. The uniform prior on the real line is "improper" (i.e., has infinite probability mass), but the conditional probabilities are well defined: a player observing signal x_i puts density $(1/\sigma)f((x_i - \theta)/\sigma)$ on state θ (see Hartigan 1983). The example of the previous section fits this setting, where $f(\cdot)$ is the standard normal distribution and $\pi(l, x) = x + l - 1$.

We will initially impose five properties on the payoffs:

A1: **Action Monotonicity**: $\pi(l, \theta)$ is nondecreasing in l.
A2: **State Monotonicity**: $\pi(l, \theta)$ is nondecreasing in θ.
A3: **Strict Laplacian State Monotonicity**: There exists a unique θ^* solving $\int_{l=0}^{1} \pi(l, \theta^*)dl = 0$.
A4: **Limit Dominance**: There exist $\underline{\theta} \in \mathbb{R}$ and $\overline{\theta} \in \mathbb{R}$, such that [1] $\pi(l, x) < 0$ for all $l \in [0, 1]$ and $x \leq \underline{\theta}$; and [2] $\pi(l, x) > 0$ for all $l \in [0, 1]$ and $x \geq \overline{\theta}$.
A5: **Continuity**: $\int_{l=0}^{1} g(l)\pi(l, x)dl$ is continuous with respect to signal x and density g.

Condition A1 states that the incentive to choose action 1 is increasing in the proportion of other players' actions who use action 1; thus there are *strategic complementarities* between players' actions (Bulow, Geanakoplos, and Klemperer, 1985). Condition A2 states that the incentive to choose action 1 is increasing in the state; thus a player's optimal action will be increasing in the state, given the opponents' actions. Condition A3 introduces a further strengthening of A2 to ensure that there is at most one crossing for a player with Laplacian beliefs. Condition A4 requires that action 0 is a dominant strategy for sufficiently low signals, and action 1 is a dominant strategy for sufficiently high signals. Condition A5 is a weak continuity property, where continuity in g is with respect to the weak topology. Note that this condition allows for some discontinuities in payoffs. For example,

$$\pi(l, x) = \begin{cases} 0, & \text{if} \quad l \leq x \\ 1, & \text{if} \quad l > x \end{cases}$$

satisfies A5 as for any given x, it is discontinuous at only one value of l.

We denote by $G^*(\sigma)$ this incomplete information game – with the uniform prior and satisfying A1 through A5. A strategy for a player in the incomplete information game is a function $s : \mathbb{R} \to \{0, 1\}$, where $s(x)$ is the action chosen if a player observes signal x. We will be interested in strategy profiles, $\mathbf{s} = (s_i)_{i \in [0,1]}$, that form a Bayesian Nash equilibrium of $G^*(\sigma)$. We will show not merely that there is a unique Bayesian Nash equilibrium of the game, but that a unique strategy profile survives iterated deletion of strictly (interim) dominated strategies.

[3] With small changes in terminology, the argument will extend to the case where $f(\cdot)$ has support on some bounded interval of the real line.

Proposition 2.1. *Let θ^* be defined as in (A3). The essentially unique strategy surviving iterated deletion of strictly dominated strategies in $G^*(\sigma)$ satisfies $s(x) = 0$ for all $x < \theta^*$ and $s(x) = 1$ for all $x > \theta^*$.*

The "essential" qualification arises because either action may be played if the private signal is exactly equal to θ^*. The key idea of the proof is that, with a uniform prior on θ, observing x_i gives no information to a player on his ranking within the population of signals. Thus, he will have a uniform prior belief over the proportion of players who will observe higher signals.

Proof. Write $\pi_\sigma^*(x, k)$ for the expected payoff gain to choosing action 1 for a player who has observed a signal x and knows that all other players will choose action 0 if they observe signals less than k:

$$\pi_\sigma^*(x, k) \equiv \int_{\theta=-\infty}^{\infty} \frac{1}{\sigma} f\left(\frac{x-\theta}{\sigma}\right) \pi\left(1 - F\left(\frac{k-\theta}{\sigma}\right), x\right) d\theta.$$

First, observe that $\pi_\sigma^*(x, k)$ is continuous in x and k, increasing in x, and decreasing in k, $\pi_\sigma^*(x, k) < 0$ if $x \leq \underline{\theta}$ and $\pi_\sigma^*(x, k) > 0$ if $x \geq \overline{\theta}$. We will argue by induction that a strategy survives n rounds of iterated deletion of strictly interim dominated strategies if and only if

$$s(x) = \begin{cases} 0, & \text{if} \quad x < \underline{\xi}_n \\ 1, & \text{if} \quad x > \overline{\xi}_n, \end{cases}$$

where $\underline{\xi}_0 = -\infty$ and $\overline{\xi}_0 = +\infty$, and $\underline{\xi}_n$ and $\overline{\xi}_n$ are defined inductively by

$$\underline{\xi}_{n+1} = \min\{x : \pi_\sigma^*(x, \underline{\xi}_n) = 0\}$$

and

$$\overline{\xi}_{n+1} = \max\{x : \pi_\sigma^*(x, \overline{\xi}_n) = 0\}.$$

Suppose the claim was true for n. By strategic complementarities, if action 1 were ever to be a best response to a strategy surviving n rounds, it must be a best response to the switching strategy with cutoff $\underline{\xi}_n$; $\underline{\xi}_{n+1}$ is defined to be the lowest signal where this occurs. Similarly, if action 0 were ever to be a best response to a strategy surviving n rounds, it must be a best response to the switching strategy with cutoff $\overline{\xi}_n$; $\overline{\xi}_{n+1}$ is defined to be the highest signal where this occurs.

Now note that $\underline{\xi}_n$ and $\overline{\xi}_n$ are increasing and decreasing sequences, respectively, because $\underline{\xi}_0 = -\infty < \underline{\theta} < \underline{\xi}_1$, $\overline{\xi}_0 = \infty > \overline{\theta} > \overline{\xi}_1$, and $\pi_\sigma^*(x, k)$ is increasing in x and decreasing in k. Thus, $\underline{\xi}_n \to \underline{\xi}$ and $\overline{\xi}_n \to \overline{\xi}$ as $n \to \infty$. The continuity of π_σ^* and the construction of $\underline{\xi}$ and $\overline{\xi}$ imply that we must have $\pi_\sigma^*(\underline{\xi}, \underline{\xi}) = 0$ and $\pi_\sigma^*(\overline{\xi}, \overline{\xi}) = 0$. Thus, the second step of our proof is to show that θ^* is the unique solution to the equation $\pi_\sigma^*(x, x) = 0$.

To see this second step, write $\Psi_\sigma^*(l; x, k)$ for the probability that a player assigns to proportion less than l of the other players observing a signal greater than k, if he has observed signal x. Observe that if the true state is θ, the proportion of players observing a signal greater than k is $1 - F((k - \theta)/\sigma)$. This proportion is less than l if $\theta \leq k - \sigma F^{-1}(1 - l)$. So,

$$\begin{aligned}\Psi_\sigma^*(l; x, k) &= \int_{\theta=-\infty}^{k-\sigma F^{-1}(1-l)} \frac{1}{\sigma} f\left(\frac{x-\theta}{\sigma}\right) d\theta \\ &= \int_{z=\frac{x-k}{\sigma}+F^{-1}(1-l)}^{\infty} f(z)\,dz, \quad \text{changing variables to} \quad z = \frac{x-\theta}{\sigma} \\ &= 1 - F\left(\frac{x-k}{\sigma} + F^{-1}(1-l)\right). \end{aligned} \tag{2.6}$$

Also observe that if $x = k$, then $\Psi_\sigma^*(\cdot; x, k)$ is the identity function [i.e., $\Psi_\sigma^*(l; x, k) = l$], so it is the cumulative distribution function of the uniform density. Thus,

$$\pi_\sigma^*(x, x) = \int_{l=0}^{1} \pi(l, x) dl.$$

Now by A3, $\pi_\sigma^*(x, x) = 0$ implies $x = \theta^*$. ■

2.2.2. Continuum Players: General Prior and Common Values

Now suppose instead that θ is drawn from a continuously differentiable strictly positive density $p(\cdot)$ on the real line and that a player's utility depends on the realized state θ, not his signal of θ. Thus, $u(a, l, \theta)$ is his payoff if he chooses action a, proportion l of his opponents choose action 1, *and the state is* θ, and as before, $\pi(l, \theta) \equiv u(1, l, \theta) - u(0, l, \theta)$. We must also impose two extra technical assumptions.

A4*: **Uniform Limit Dominance**: There exist $\underline{\theta} \in \mathbb{R}$, $\overline{\theta} \in \mathbb{R}$, and $\varepsilon \in \mathbb{R}_{++}$, such that [1] $\pi(l, \theta) \leq -\varepsilon$ for all $l \in [0, 1]$ and $\theta \leq \underline{\theta}$; and [2] there exists $\overline{\theta}$ such that $\pi(l, \theta) > \varepsilon$ for all $l \in [0, 1]$ and $\theta \geq \overline{\theta}$.

Property A4* strengthens property A4 by requiring that the payoff gain to choosing action 0 is *uniformly* positive for sufficiently low values of θ, and the payoff gain to choosing action 1 is *uniformly* positive for sufficiently high values of θ.

A6: **Finite Expectations of Signals**: $\int_{z=-\infty}^{\infty} z f(z) dz$ is well defined.

Property A6 requires that the distribution of noise is integrable.

We will denote by $G(\sigma)$ this incomplete information game, with prior $p(\cdot)$ and satisfying A1, A2, A3, A4*, A5, and A6.

Proposition 2.2. *Let θ^* be defined as in A3. For any $\delta > 0$, there exists $\overline{\sigma} > 0$ such that for all $\sigma \leq \overline{\sigma}$, if strategy s survives iterated deletion of strictly dominated strategies in the game $G(\sigma)$, then $s(x) = 0$ for all $x \leq \theta^* - \delta$, and $s(x) = 1$ for all $x \geq \theta^* + \delta$.*

We will sketch here why this general prior, common values, game $G(\sigma)$ becomes like the uniform prior, private values, game $G^*(\sigma)$ as σ becomes small. A more formal proof is relegated to Appendix A. Consider $\Psi_\sigma(l; x, k)$, the probability that a player assigns to proportion less than or equal to l of the other players observing a signal greater than or equal to k, if he has observed signal x:

$$\begin{aligned}\Psi_\sigma(l; x, k) &= \frac{\int_{\theta=-\infty}^{k-\sigma F^{-1}(1-l)} p(\theta) f(\frac{x-\theta}{\sigma})\, d\theta}{\int_{\theta=-\infty}^{\infty} p(\theta) f\left(\frac{x-\theta}{\sigma}\right) d\theta} \\ &= \frac{\int_{z=\frac{x-k}{\sigma}+F^{-1}(1-l)}^{\infty} p\,(x-\sigma z)\, f\,(z)\, dz}{\int_{z=-\infty}^{\infty} p\,(x-\sigma z)\, f(z)\, dz}, \\ &\qquad \text{changing variables to} \quad z = \frac{x-\theta}{\sigma}.\end{aligned}$$

For small σ, the shape of the prior will not matter and the posterior beliefs over l will depend only on $(x-k)/\sigma$, the normalized difference between the x and k. Formally, setting $\kappa = (x-k)/\sigma$, we have

$$\Psi_\sigma^*(l; x, x-\sigma\kappa) = \frac{\int_{z=\kappa+F^{-1}(1-l)}^{\infty} p(x-\sigma z) f(z)\, dz}{\int_{z=-\infty}^{\infty} p(x-\sigma z) f(z)\, dz},$$

so that as $\sigma \to 0$,

$$\begin{aligned}\Psi_\sigma^*(l; x, x-\sigma\kappa) &\to \int_{z=\kappa+F^{-1}(1-l)}^{\infty} f(z)\, dz \\ &= 1 - F(\kappa + F^{-1}(1-l)). \qquad (2.7)\end{aligned}$$

In other words, for small σ, posterior beliefs concerning the proportion of opponents choosing each action are almost the same as under a uniform prior. The formal proof of proposition 2.2 presented in Appendix A consists of showing, first, that convergence of posterior beliefs described previously is uniform; and, second, that the small amount of uncertainty about payoffs in the common value case does not affect the analysis sufficiently to matter.

2.2.3. Discussion

The proofs of propositions 2.1 and 2.2 follow the logic of Carlsson and van Damme (1993) and generalize arguments presented in Morris and Shin (1998). The technique of analyzing the uniform prior private values game, and then showing continuity with respect to the general prior, common values game, follows Frankel, Morris, and Pauzner (2000). (This paper is discussed further in Section 4.1.) Carlsson and van Damme (1993b) showed a version of the uniform prior result (proposition 2.1) in the finite player case (see also Kim, 1996). We briefly discuss the relation to the finite player case in Appendix B.

How do these propositions make use of the underlying assumptions? First, note that assumptions A1 and A2 represent very strong monotonicity assumptions: A1 requires that each player's utility function is supermodular in the action profile, whereas A2 requires that each player's utility function is supermodular in his own action and the state. Vives (1990) showed that the supermodularity property A2 of complete information game payoffs is inherited by the incomplete information game. Thus, the existence of a largest and smallest strategy profile surviving iterated deletion of dominated strategies when payoffs are supermodular, noted by Milgrom and Roberts (1990), can be applied also to the incomplete information game. The first step in the proof of proposition 2.1 is a special case of this reasoning, with the state monotonicity assumption A2 implying, in addition, that the largest and smallest equilibria consist of strategies that are monotonic with respect to type (i.e., switching strategies). Once we know that we are interested in monotonic strategies, the very weak assumption A3 is sufficient to ensure the equivalence of the largest and smallest equilibria and thus the uniqueness of equilibrium.

Can one dispense with the full force of the supermodular payoffs assumption A1? Unfortunately, as long as A1 is not satisfied at the cutoff point θ^* [i.e., $\pi(l, \theta^*)$ is decreasing in l over some range], then one can find a problematic noise distribution $f(\cdot)$ such that the symmetric switching strategy profile with cutoff point θ^* is *not* an equilibrium, and thus there is no switching strategy equilibrium. To obtain positive results, one must either impose additional restrictions on the noise distribution or relax A1 only away from the cutoff point. We discuss both approaches in turn.

Athey (2002) provides a general description of how monotone comparative static results can be preserved in stochastic optimization problems, when supermodular payoff conditions are weakened to single crossing properties, but signals are assumed to be sufficiently well behaved (i.e., satisfy a monotone likelihood ratio property). Athey (2001) has used such techniques to prove existence of monotonic pure strategy equilibria in a general class of incomplete information games, using weaker properties on payoffs, but substituting stronger restrictions on signal distribution. We can apply her results to our setting as follows. Consider the following two new assumptions.

A1*: **Action Single Crossing**: For each $\theta \in \mathbb{R}$, there exists $l^* \in \mathbb{R} \cup \{-\infty, \infty\}$ such that $\pi(l, \theta) < 0$ if $l < l^*$ and $\pi(l, \theta) > 0$ if $l > l^*$.

A7: **Monotone Likelihood Ratio Property**: If $\overline{x} > \underline{x}$, then $f(\overline{x} - \theta)/f(\underline{x} - \theta)$ is increasing in θ.

Assumption A1* is a significant weakening of assumption A1 to a single crossing property. Assumption A7 is a new restriction on the distribution of the noise. Recall that we earlier made no assumptions on the distribution of the noise. Denote by $\widetilde{G}(\sigma)$ the incomplete information game with a uniform prior satisfying A1*, A2, A3, A4, A5, and A7.

Lemma 2.3. *Let θ^* be defined as in A3. The game $\widetilde{G}(\sigma)$ has a unique (symmetric) switching strategy equilibrium, with $s(x) = 0$ for all $x < \theta^*$ and $s(x) = 1$ for all $x > \theta^*$.*

The proof is in Appendix C. An analog of proposition 2.2 could be similarly constructed. Notice that this result does not show the nonexistence of other, nonmonotonic, equilibria. Additional arguments are required to rule out nonmonotonic equilibria. For example, in Goldstein and Pauzner (2000a) – an application to bank runs discussed in the next section – noise is uniformly distributed (and thus satisfies A7) and payoffs satisfy assumption A1*. They show that (1) there is a unique symmetric switching strategy equilibrium and that (2) there is no other equilibrium. Lemma 2.3 could be used to extend the former result to all noise distributions satisfying the MLRP (assumption A7), but we do not know if the latter result extends beyond the uniform noise distribution.

Proposition 2.1 can also be weakened by allowing assumption A1 to fail away from θ^*. We will report one weakening that is sufficient. Let $g(\cdot)$ and $h(\cdot)$ be densities on the interval [0, 1]; *g stochastically dominates h* ($g \succeq h$) if $\int_{z=0}^{l} g(z)\,dz \leq \int_{z=0}^{l} h(z)\,dz$ for all $l \in [0, 1]$. We write $\overline{g}(\cdot)$ for the uniform density on [0, 1], i.e., $\overline{g}(l) = 1$ for all $l \in [0, 1]$. Now consider

> A8: There exists θ^* which solves $\int_{l=0}^{1} \pi(l, \theta^*)dl = 0$ such that [1] $\int_{l=0}^{1} g(l)\,\pi(l, x)dl \geq 0$ for all $x \geq \theta^*$ and $g \succeq \overline{g}$, with strict inequality if $x > \theta^*$; and [2] $\int_{l=0}^{1} g(l)\pi(l, x)dl \leq 0$ for all $x \leq \theta^*$ and $g \preceq \overline{g}$, with strict inequality if $x < \theta^*$.

We can replace A1–A3 with A8 in propositions 2.1 and 2.2, and all the arguments and results go through. Observe that A1–A3 straightforwardly imply A8. Also, observe that A8 implies that $\pi(l, \theta^*)$ be nondecreasing in l [suppose that $l > l'$ and $\pi(l, \theta^*) < \pi(l', \theta^*)$; now start with the uniform distribution $\overline{g}$ and shift mass from l' to l]. But, A8 allows some failure of A1 away from θ^*.

Propositions 2.1 and 2.2 deliver strong negative conclusions about the efficiency of noncooperative outcomes in global games. In the limit, all players will be choosing action 1 when the state is θ if $\int_{l=0}^{1} \pi(l, \theta)dl > 0$. However, it is *efficient* to choose action 1 at state θ if $u(1, 1, \theta) > u(0, 0, \theta)$. These conditions will not coincide in general. For example, in the investment example, we had $u(1, l, \theta) = \theta + l - 1$, $u(0, l, \theta) = 0$ and thus $\pi(l, \theta) = \theta + l - 1$. So in the limiting equilibrium, both players will be investing if the state θ is at least $\frac{1}{2}$, although it is efficient for them to be investing if the state is at least 0.

The analysis of the unique noncooperative equilibrium serves as a benchmark describing what will happen in the absence of other considerations. In practice, repeated play or other institutions will often allow players to do better. We will briefly consider what happens in the game if players were allowed to make cheap talk statements about the signals that they have observed in the investment example (for this exercise, it is most natural to consider a finite player case; we consider the two-player case). The arguments here follow Baliga and Morris (2000). The investment example as formulated has a nongeneric feature, which is that if a player plans not to invest, he is exactly indifferent about which action his opponent will take. To make the problem more interesting, let us perturb the payoffs to remove this tie:

Table 3.2. *Payoffs for cheap talk example*

	Invest	NotInvest
Invest	$\theta + \delta, \theta + \delta$	$\theta - 1, \delta$
NotInvest	$\delta, \theta - 1$	$0, 0$

Thus, each player receives a small payoff δ (which may be positive or negative) if the other player invests, independent of his own action. This change does not influence each player's best responses, and the analysis of this game in the absence of cheap talk is unchanged by the payoff change. But, observe that if $\delta \leq 0$, there is an equilibrium of the game with cheap talk, where each player truthfully announces his signal, and invests if the (common) expectation of θ conditional on both announcements is greater than $-\delta$ (this gives the efficient outcome). On the other hand, if $\delta > 0$, then each player would like to convince the other to invest even if he does not plan to do so. In this case, there cannot be a truth-telling equilibrium where the efficient equilibrium is achieved, although there may be equilibria with some partially revealing cheap talk that improves on the no cheap talk outcome.

2.3. Applications

We now turn to applications of these results and describe models of pricing debt (Morris and Shin, 1999b), currency crises (Morris and Shin, 1998), and bank runs (Goldstein and Pauzner, 2000a).[4] Each of these papers makes specific assumptions about the distribution of payoffs and signals. But, if one is interested only in analyzing the limiting behavior as noise about θ becomes

[4] See Fukao (1994) for an early argument in favor of using global game reasoning in applied settings. Other applications include Karp's (2000) noisy version of Krugman's (1991) multiple equilibrium model of sectoral shifts; Scaramozzino and Vulkan's (1999) noisy model of Shleifer's (1986) multiple equilibrium model of implementation cycles; and Dönges and Heinemann's (2000) model of competition between dealer markets and crossing networks in financial markets.

small, the results of the previous section imply that we can identify the limiting behavior independently of the prior beliefs and the shape of the noise.[5] In each example, we describe one comparative static exercise changing the payoffs of the game, illustrating how changing payoffs has a direct effect on outcomes and an indirect, strategic effect via the impact on the cutoff point of the unique equilibrium. We emphasize that it is also interesting in the applications to study behavior away from the limit; indeed, the focus of the analysis in Morris and Shin (1999b) is on comparative statics away from the limit. More assumptions on the shape of the prior and noise are required in this case. We study behavior away from the limit in Section 3.

2.3.1. Pricing Debt

In Morris and Shin (1999b), we consider a simple model of debt pricing. In period 1, a continuum of investors hold collateralized debt that will pay 1 in period 2 if it is rolled over and if an underlying investment project is successful; the debt will pay 0 in period 2 if the project is not successful. If an investor does not roll over his debt, he receives the value of the collateral, $\kappa \in (0, 1)$. The success of the project depends on the proportion of investors who do not roll over and the state of the economy, θ. Specifically, the project is successful if the proportion of investors not rolling over is less than θ/z. Writing 1 for the action "roll over" and 0 for the action "do not roll over," payoffs can be described as follows:

$$u(1, l, \theta) = \begin{cases} 1, & \text{if } \quad z(1-l) \leq \theta \\ 0, & \text{if } \quad z(1-l) > \theta, \end{cases}$$
$$u(0, l, \theta) = \kappa.$$

So

$$\pi(l, \theta) \equiv u(1, l, \theta) - u(0, l, \theta) = \begin{cases} 1-\kappa, & \text{if } \quad z(1-l) \leq \theta \\ -\kappa, & \text{if } \quad z(1-l) > \theta. \end{cases}$$

Now

$$\int_{l=0}^{1} \pi(l, \theta)\, dl = \begin{cases} -\kappa, & \text{if } \quad \theta \leq 0 \\ \frac{\theta}{z} - \kappa, & \text{if } \quad 0 \leq \theta \leq z \\ 1-\kappa, & \text{if } \quad z \leq \theta. \end{cases}$$

[5] The model in Goldstein and Pauzner (2000a) fails the action monotonicity property (A1) of the previous section, but they are nonetheless able to prove the uniqueness of a symmetric switching equilibrium, exploiting their assumption that noise terms are distributed uniformly. However, their game satisfies assumptions A1* and A2, and therefore whenever there is a unique equilibrium, it must satisfy the Laplacian characterization with the cutoff point θ^* defined as in A3.

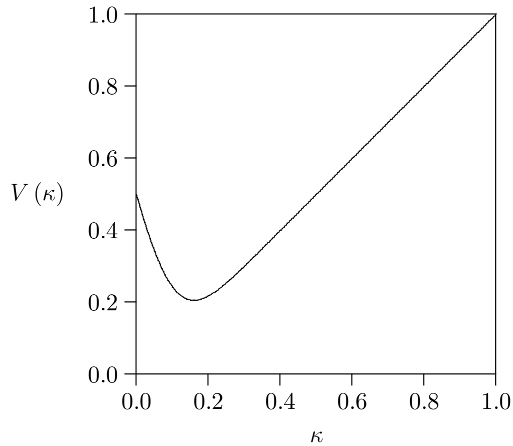

Figure 3.2. Function $V(\kappa)$.

Thus, $\theta^* = z\kappa$. In other words, if private information about θ among the investors is sufficiently accurate, the project will collapse exactly if $\theta \leq z\kappa$. We can now ask how debt would be priced ex ante in this model (before anyone observed private signals about θ). Recalling that $p(\cdot)$ is the density of the prior on θ, and writing $P(\cdot)$ for the corresponding cdf, the value of the collateralized debt will be

$$\begin{aligned} V(\kappa) &\equiv \kappa P(z\kappa) + 1 - P(z\kappa) \\ &= 1 - (1 - \kappa)P(z\kappa), \end{aligned}$$

and

$$\frac{dV}{d\kappa} = P(z\kappa) - z(1 - \kappa)p(z\kappa).$$

Thus, increasing the value of collateral has two effects: first, it increases the value of debt in the event of default (the direct effect). But, second, it increases the range of θ at which default occurs (the strategic effect). For small κ, the strategic effect outweighs the direct effect, whereas for large κ, the direct effect outweighs the strategic effect. Figure 3.2 plots $V(\cdot)$ for the case where $z = 10$ and $p(\cdot)$ is the standard normal density.

Morris and Shin (1999b) study the model away from the limit and argue that taking the strategic, or liquidity, effect into account in debt pricing can help explain anomalies in empirical implementation of the standard debt pricing theory of Merton (1974). Brunner and Krahnen (2000) present evidence of the importance of debtor coordination in distressed lending relationships in Germany [see also Chui, Gai, and Haldane (2000) and Hubert and Schäfer (2000)].

2.3.2. Currency Crises

In Morris and Shin (1998), a continuum of speculators must decide whether to attack a fixed–exchange rate regime by selling the currency short. Each speculator may short only a unit amount. The current value of the currency is e^*; if the monetary authority does not defend the currency, the currency will float to the shadow rate $\zeta(\theta)$, where θ is the state of fundamentals. There is a fixed transaction cost t of attacking. This can be interpreted as an actual transaction cost or as the interest rate differential between currencies. The monetary authority defends the currency if the cost of doing so is not too large. Assuming that the costs of defending the currency are increasing in the proportion of speculators who attack and decreasing in the state of fundamentals, there will be some critical proportion of speculators, $a(\theta)$, increasing in θ, who must attack in order for a devaluation to occur. Thus, writing 1 for the action "not attack" and 0 for the action "attack," payoffs can be described as follows:

$$u(1, l, \theta) = 0,$$
$$u(0, l, \theta) = \begin{cases} e^* - \zeta(\theta) - t, & \text{if} \quad l \leq 1 - a(\theta) \\ -t, & \text{if} \quad l > 1 - a(\theta), \end{cases}$$

where $\zeta(\cdot)$ and $a(\cdot)$ are increasing functions, with $\zeta(\theta) \leq e^* - t$ for all θ. Now

$$\pi(l, \theta) = \begin{cases} \zeta(\theta) + t - e^*, & \text{if} \quad l \leq 1 - a(\theta) \\ t, & \text{if} \quad l > 1 - a(\theta). \end{cases}$$

If θ were common knowledge, there would be three ranges of parameters. If $\theta < a^{-1}(0)$, each player has a dominant strategy to attack. If $a^{-1}(0) \leq \theta \leq a^{-1}(1)$, then there is an equilibrium where all speculators attack and another equilibrium where all speculators do not attack. If $\theta > a^{-1}(1)$, each player has a dominant strategy to attack. This tripartite division of fundamentals arises in a range of models in the literature on currency crises (see Obstfeld, 1996).

However, if θ is observed with noise, we can apply the results of the previous section, because $\pi(l, \theta)$ is weakly increasing in l, and weakly increasing in θ:

$$\begin{aligned} \int_{l=0}^{1} \pi(l, \theta)\, dl &= (1 - a(\theta))(\zeta(\theta) + t - e^*) + a(\theta)t \\ &= t - (1 - a(\theta))(e^* - \zeta(\theta)). \end{aligned}$$

Thus, θ^* is implicitly defined by

$$(1 - a(\theta))(e^* - \zeta(\theta)) = t.$$

Theorem 2 in Morris and Shin (1998) gave an incorrect statement of this condition. We are grateful to Heinemann (2000) for pointing out the error and giving a correct characterization.

Again, we will describe one simple comparative statics exercise. Consider a costly ex ante action R for the monetary authority that lowered their costs of defending the currency. For example, R might represent the value of foreign currency reserves or (as in the recent case of Argentina) a line of credit with foreign banks to provide credit in the event of a crisis. Thus, the critical proportion of speculators for which an attack occurs becomes $a(\theta, R)$, where $a(\cdot)$ is increasing in R. Now, write $\theta^*(R)$ for the unique value of θ solving

$$(1 - a(\theta, R))(e^* - \zeta(\theta)) = t.$$

The ex ante probability that the currency will collapse is

$$P(\theta^*(R)).$$

So, the reduction in the probability of collapse resulting from a marginal increase in R is

$$-p(\theta^*(R))\frac{d\theta^*}{dR} = p(\theta^*(R))\frac{\frac{\partial a}{\partial R}}{\frac{\partial a}{\partial \theta} + \frac{1-a(\theta,R)}{e^*-\zeta(\theta)}\frac{d\zeta}{d\theta}}.$$

This comparative static refers to the limit (as noise becomes very small), and the effect is entirely strategic [i.e., the increased value of R reduces the probability of attack *only* because it influences speculators' equilibrium strategies ("builds confidence") and not because the increase in R actually prevents an attack in any relevant contingency].

In Section 4.1, we very briefly discuss Corsetti, Dasgupta, Morris, and Shin (2000), an extension of this model of currency attacks where a large speculator is added to the continuum of small traders [see also Chan and Chiu (2000), Goldstein and Pauzner (2000b), Heinemann and Illing (2000), Hellwig (2000), Marx (2000), Metz (2000), and Morris and Shin (1999a)].

2.3.3. Bank Runs

We describe a model of Goldstein and Pauzner (2000a), who add noise to the classic bank runs model of Diamond and Dybvig (1983). A continuum of depositors (with total deposits normalized to 1) must decide whether to withdraw their money from a bank or not. If the depositors withdraw their money in period 1, they will receive $r > 1$ (if there are not enough resources to fund all those who try to withdraw, then the remaining cash is divided equally among early withdrawers). Any remaining money earns a total return $R(\theta) > 0$ in period 2 and is divided equally among those who chose to wait until period 2 to withdraw their money. Proportion λ of depositors will have consumption needs only in period 1 and will thus have a dominant strategy to withdraw. We will be concerned with the game among the proportion $1 - \lambda$ of depositors who have consumption needs in period 2. Consumers have utility $U(y)$ from consumption y, where the relative risk aversion coefficient of U is strictly greater than 1. They note that if $R(\theta)$ was greater than 1 and θ were common knowledge, the ex ante optimal choice of r maximizing

$$\lambda U(r) + (1 - \lambda)U\left(\frac{1 - \lambda r}{1 - \lambda}R(\theta)\right)$$

would be strictly greater than 1. But, if θ is not common knowledge, we have a global game. Writing 1 for the action "withdraw in period 2" and 0 for the action "withdraw in period 1," and l for the proportion of late consumers who do not withdraw early, the money payoffs in this game can be summarized in Table 3.3:

Table 3.3. *Payoffs in bank run game*

		$l \le \frac{r-1}{r(1-\lambda)}$	$l \ge \frac{r-1}{r(1-\lambda)}$
Early Withdrawal	0	$\frac{1-\lambda r}{(1-\lambda)(1-l)r}$	r
Late Withdrawal	1	0	$\left(r - \frac{r-1}{l(1-\lambda)}\right) R(\theta)$

Observe that, if θ is sufficiently small [and so $R(\theta)$ is sufficiently small], all players have a dominant strategy to withdraw early. Goldstein and Pauzner assume that, if θ is sufficiently large, all players have a dominant strategy to withdraw late (a number of natural economic stories could justify this variation in the payoffs).

Thus, the payoffs in the game among late consumers are

$$u(1, l, \theta) = \begin{cases} U(0), & \text{if} \quad l \le \frac{r-1}{r(1-\lambda)} \\ U\left(\left(r - \frac{r-1}{l(1-\lambda)}\right) R(\theta)\right), & \text{if} \quad l \ge \frac{r-1}{r(1-\lambda)}, \end{cases}$$

$$u(0, l, \theta) = \begin{cases} U\left(\frac{1}{1-l(1-\lambda)}\right), & \text{if} \quad l \le \frac{r-1}{r(1-\lambda)} \\ U(r), & \text{if} \quad l \ge \frac{r-1}{r(1-\lambda)} \end{cases}$$

so that

$$\pi(l, \theta) = \begin{cases} U(0) - U\left(\frac{1}{1-l(1-\lambda)}\right), & \text{if} \quad l \le \frac{r-1}{r(1-\lambda)} \\ U\left(\left(r - \frac{r-1}{l(1-\lambda)}\right) R(\theta)\right) - U(r), & \text{if} \quad l \ge \frac{r-1}{r(1-\lambda)}. \end{cases}$$

The threshold state θ^* is implicitly defined by

$$\int_{l=0}^{\frac{r-1}{r(1-\lambda)}} U(0) - U\left(\frac{1}{1 - l(1-\lambda)}\right) dl + \int_{l=\frac{r-1}{r(1-\lambda)}}^{1} U\left(\left(r - \frac{r-1}{l(1-\lambda)}\right) R(\theta)\right) - U(r)\, dl = 0.$$

The ex ante welfare of consumers as a function of r (as noise goes to zero) is

$$W(r) = P(\theta^*(r))U(1) + \int_{\theta=\theta^*(r)}^{\infty} p(\theta)\left(\lambda U(r) + (1-\lambda)U\left(\frac{1-\lambda r}{1-\lambda} R(\theta)\right)\right).$$

There are two effects of increasing r: the direct effect on welfare is the increased value of insurance in the case where there is not a bank run. But, there is also the strategic effect that an increase in r will lower $\theta^*(r)$.

Morris and Shin (2000) examine a stripped down version of this model, where alternative assumptions on the investment technology and utility functions imply that payoffs reduce to those of the linear example in Section 2.1 [see also Boonprakaikawe and Ghosal (2000), Dasgupta (2000b), Goldstein (2000), and Rochet and Vives (2000)].

3. PUBLIC VERSUS PRIVATE INFORMATION

The analysis so far has all been concerned with behavior when either there is a uniform prior or the noise is very small. In this section, we look at the behavior of the model with large noise and nonuniform priors. There are three reasons for doing this. First, we want to understand how extreme the assumptions required for uniqueness are. We will provide sufficient conditions for uniqueness depending on the relative accuracy of private and public (or prior) signals. Second, away from the limit, prior beliefs play an important role in determining outcomes. In particular, we will see how even with a continuum of players and a unique equilibrium, public information contained in the prior beliefs plays a significant role in determining outcomes, *even controlling for beliefs concerning the fundamentals.* Finally, by seeing how and when the model jumps from having one equilibrium to multiple equilibria, it is possible to develop a better intuition for what is driving results.

We return to the linear example of Section 2.1: there is a continuum of players, the payoff to not investing is 0, and the payoff to investing is $\theta + l - 1$, where θ is the state and l is the proportion of the population investing. It may help in following in the analysis to recall that, with linear payoffs, the exact number of players is irrelevant in identifying symmetric equilibrium strategies (and we will see that symmetric equilibrium strategies will naturally arise). Thus, the analysis applies equally to a two-player game.

Now assume that θ is normally distributed with mean y and standard deviation τ. The mean y is publicly observed. As before, each player observes a private signal $x_i = \theta + \varepsilon_i$, where the ε_i are distributed normally in the population with mean 0 and standard deviation σ. Thus, each player i observes a public signal $y \in \mathbb{R}$ and a private signal $x_i \in \mathbb{R}$. To analyze the equilibria of this game, first fix the public signal y. Suppose that a player observed private signal x. His expectation of θ is

$$\overline{\theta} = \frac{\sigma^2 y + \tau^2 x}{\sigma^2 + \tau^2}.$$

It is useful to conduct analysis in terms of these posterior expectations of θ. In particular, we may consider a switching strategy of the following form:

$$s(\overline{\theta}) = \begin{cases} \text{Invest}, & \text{if} \quad \overline{\theta} > \kappa \\ \text{NotInvest}, & \text{if} \quad \overline{\theta} \leq \kappa. \end{cases}$$

If the standard deviation of players' private signals is sufficiently small relative to the standard deviation of the public signal in the prior, then there is a strategy surviving iterated deletion of strictly dominated strategies. Specifically, let

$$\gamma \equiv \widetilde{\gamma}(\sigma, \tau) \equiv \frac{\sigma^2}{\tau^4}\left(\frac{\sigma^2+\tau^2}{\sigma^2+2\tau^2}\right).$$

Now we have

Proposition 3.1. *The game has a symmetric switching strategy equilibrium with cutoff κ if κ solves the equation*

$$\kappa = \Phi(\sqrt{\gamma}(\kappa - y)); \tag{3.1}$$

if $\widetilde{\gamma}(\sigma, \tau) \leq 2\pi$, then there is a unique value of κ solving (3.1) and the strategy with that trigger is the essentially unique strategy surviving iterated deletion of strictly dominated strategies; if $\widetilde{\gamma}(\sigma, \tau) > 2\pi$, then (for some values of y) there are multiple values of κ solving (3.1) and multiple symmetric switching strategy equilibria.

Figure 3.3 plots the regions in $\sigma^2 - \tau^2$ space, where uniqueness holds.

In Morris and Shin (2000), we gave a detailed version of the uniqueness part of this result in Appendix A. Here, we sketch the idea. Consider a player who has observed private signal x. By standard properties of the normal distribution

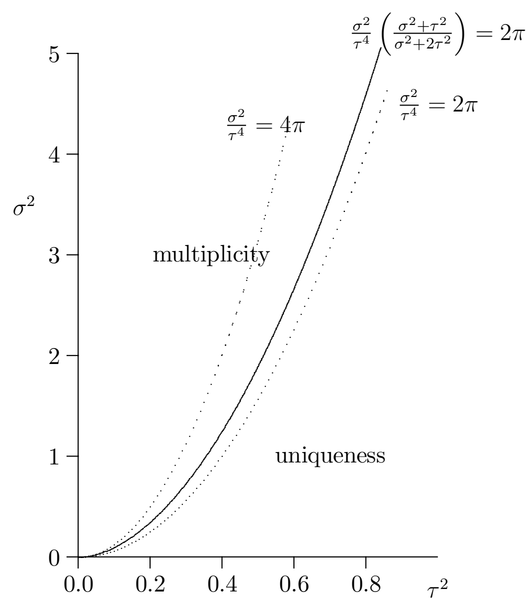

Figure 3.3. Parameter range for unique equilibrium.

(see DeGroot, 1970), his posterior beliefs about θ would be normal with mean

$$\overline{\theta} = \frac{\sigma^2 y + \tau^2 x}{\sigma^2 + \tau^2}$$

and standard deviation

$$\sqrt{\frac{\sigma^2 \tau^2}{\sigma^2 + \tau^2}}.$$

He knows that any other player's signal, x', is equal to θ plus a noise term with mean 0 and standard deviation σ. Thus, he believes that x' is distributed normally with mean $\overline{\theta}$ and standard deviation

$$\sqrt{\frac{2\sigma^2 \tau^2 + \sigma^4}{\sigma^2 + \tau^2}}.$$

Now suppose he believed that all other players will invest exactly if their expectation of θ is at least κ [i.e., if their private signals x' satisfy $(\sigma^2 y + \tau^2 x')/(\sigma^2 + \tau^2) \geq \kappa$, or $x' \geq \kappa + (\sigma^2/\tau^2)(\kappa - y)$]. Thus, he assigns probability

$$1 - \Phi\left(\frac{\kappa - \overline{\theta} + \frac{\sigma^2}{\tau^2}(\kappa - y)}{\sqrt{\frac{2\sigma^2\tau^2+\sigma^4}{\sigma^2+\tau^2}}}\right) \tag{3.2}$$

to any particular opponent investing. But his expectation of the proportion of his opponents investing must be equal to the probability he assigns to any one opponent investing. Thus, (3.2) is also equal to his expectation of the proportion of his opponents investing. Because his payoff to investing is $\theta + l - 1$, his expected payoff to investing is $\overline{\theta}$ plus expression (3.2) minus one, i.e.,

$$v(\overline{\theta}, \kappa) \equiv \overline{\theta} - \Phi\left(\frac{\kappa - \overline{\theta} + \frac{\sigma^2}{\tau^2}(\kappa - y)}{\sqrt{\frac{2\sigma^2\tau^2+\sigma^4}{\sigma^2+\tau^2}}}\right).$$

His payoff to not investing is 0. Because $v(\overline{\theta}, \kappa)$ is increasing in $\overline{\theta}$, we have that there is a symmetric equilibrium with switching point κ exactly if $v^*(\kappa) \equiv v(\kappa, \kappa) = 0$. But

$$\begin{aligned} v^*(\kappa) &\equiv v(\kappa, \kappa) \\ &= \kappa - \Phi\left(\frac{\sigma^2(\kappa - y)}{\tau^2\sqrt{\frac{2\sigma^2\tau^2+\sigma^4}{\sigma^2+\tau^2}}}\right) \\ &= \kappa - \Phi(\sqrt{\gamma}\,(\kappa - y)). \end{aligned}$$

Figure 3.4 plots the function $v^*(\kappa)$ for $y = \frac{1}{2}$ and $\gamma = 1{,}000$, 10, 5, and 0.1, respectively.

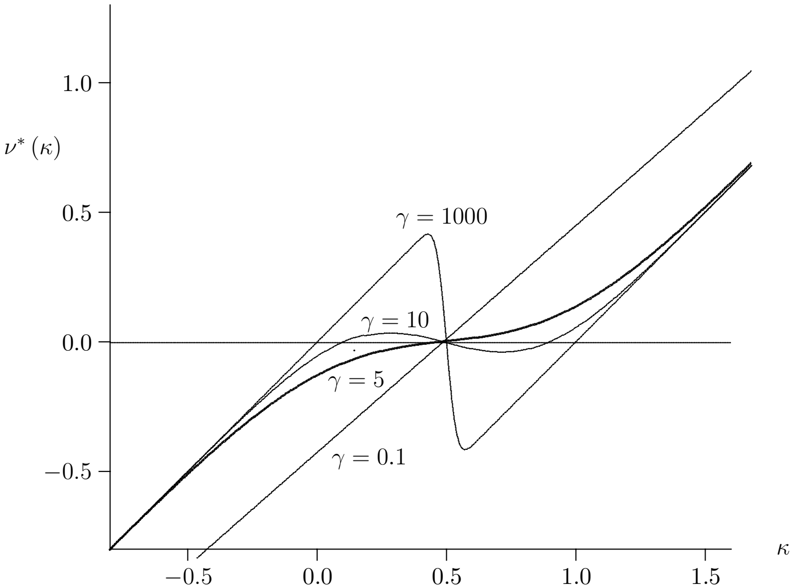

Figure 3.4. Function $\nu^*(\kappa)$.

The intuition for these graphs is the following. If public information is relatively large (i.e., $\sigma \gg \tau$ and thus γ is large), then players with posterior expectation κ less than $y = \frac{1}{2}$ confidently expect that their opponent will have observed a higher signal, and therefore will be investing. Thus, his expected utility is (about) κ. But, as κ moves above $y = \frac{1}{2}$, he rapidly becomes confident that his opponent has observed a lower signal and will not be investing. Thus, his expected utility drops rapidly, around y, to (about) $\kappa - 1$. But, if public information is relatively small (i.e., $\sigma \ll \tau$ and γ is small), then players with κ not too far above or below $y = \frac{1}{2}$ attach probability (about) $\frac{1}{2}$ to their opponent observing a higher signal. Thus, his expected utility is (about) $\kappa - \frac{1}{2}$.

We can identify analytically when there is a unique solution: Observe that

$$\frac{dv^*}{d\kappa} = 1 - \sqrt{\gamma}\phi(\sqrt{\gamma}\,(\kappa - y)).$$

Recall that $\phi(x)$, the density of the standard normal, attains its maximum of $1/\sqrt{2\pi}$ at $x = 0$. Thus, if $\gamma \leq 2\pi$, $dv^*/d\kappa$ is greater than or equal to zero always, and strictly greater than zero, except when $\kappa = y$. So, (3.1) has a unique solution. But, if $\gamma > 2\pi$ and $y = \frac{1}{2}$, then setting $\kappa = \frac{1}{2}$ solves (3.1), but $dv^*/d\kappa|_{\kappa=\frac{1}{2}} < 0$, so (3.1) has two other solutions.

Throughout the remainder of this section, we assume that there is a unique equilibrium [i.e., that $\tilde{\gamma}(\alpha, \beta) \leq 2\pi$]. Under this assumption, we can invert the equilibrium condition (3.1) to show in $(\bar{\theta}, y)$ space what the unique equilibrium

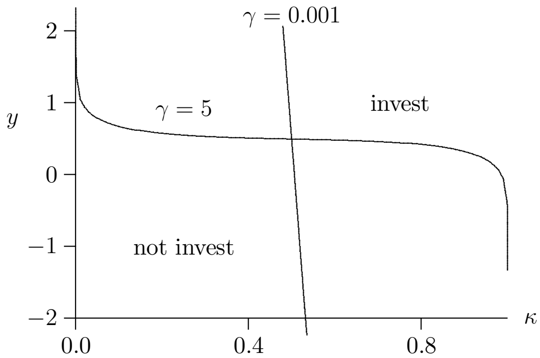

Figure 3.5. Investment takes place above and to the right of the line.

looks like:

$$y = h_\gamma(\bar{\theta}) = \bar{\theta} - \frac{1}{\sqrt{\gamma}}\Phi^{-1}(\bar{\theta}). \tag{3.3}$$

Figure 3.5 plots this for $\gamma = 5$ and $\gamma = 1/1{,}000$.

The picture has an elementary intuition. If $\bar{\theta} < 0$, it is optimal to not invest (independent of the public signal). If $\bar{\theta} > 1$, it is optimal to invest (independent of the public signal). But, if $0 < \bar{\theta} < 1$, there is a trade-off. The higher y is (for a given $\bar{\theta}$), the more likely it is that the other player will invest. Thus, if $0 < \bar{\theta} < 1$, the player will always invest for sufficiently high y, and not invest for sufficiently low y. This implies in particular that changing y has a larger impact on a player's action than changing his private signal (controlling for the informativeness of the signals). We next turn to examining this "publicity" effect.

3.1. The Publicity Multiplier

To explore the strategic impact of public information, we examine how much a player's private signal must adjust to compensate for a given change in the public signal. Equation (3.1) can be written as

$$\frac{\sigma^2 y + \tau^2 x}{\sigma^2 + \tau^2} - \Phi\left(\sqrt{\gamma}\left(\frac{\sigma^2 y + \tau^2 x}{\sigma^2 + \tau^2} - y\right)\right) = 0.$$

Totally differentiating with respect to y gives

$$\frac{dx}{dy} = -\frac{\frac{\sigma^2}{\tau^2} + \sqrt{\gamma}\phi(\cdot)}{1 - \sqrt{\gamma}\phi(\cdot)}.$$

This measures how much the private signal would have to change to compensate for a change in the public signal (and still leave the player indifferent between investing or not investing). We can similarly see how much the private signal