

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: NBER Macroeconomics Annual 2000, Volume 15

Volume Author/Editor: Ben S. Bernanke and Kenneth Rogoff, editors

Volume Publisher: MIT Press

Volume ISBN: 0-262-02503-5

Volume URL: <http://www.nber.org/books/bern01-1>

Publication Date: January 2001

Chapter Title: Rethinking Multiple Equilibria in Macroeconomic Modeling

Chapter Author: Stephen Morris, Hyun Song Shin

Chapter URL: <http://www.nber.org/chapters/c11056>

Chapter pages in book: (p. 139 - 182)

Stephen Morris and Hyun Song Shin

YALE UNIVERSITY AND OXFORD UNIVERSITY

Rethinking Multiple Equilibria in Macroeconomic Modeling

1. Introduction

It is a commonplace that actions are motivated by beliefs, and so economic outcomes are influenced by the beliefs of individuals in the economy. In many examples in economics, there seems to be an apparent indeterminacy in beliefs in the sense that one set of beliefs motivate actions which bring about the state of affairs envisaged in the beliefs, while another set of self-fulfilling beliefs bring about quite different outcomes. In both cases, the beliefs are logically coherent, consistent with the known features of the economy, and borne out by subsequent events. However, they are not fully determined by the underlying description of the economy, leaving a role for *sunspots*.

Models that utilize such apparent indeterminacy of beliefs have considerable intuitive appeal, since they provide a convenient and economical prop in a narrative of unfolding events. However, they are vulnerable to a number of criticisms. For a start, the shift in beliefs which underpins the switch from one equilibrium to another is left unexplained. This runs counter to our theoretical scruples against indeterminacy. More importantly, it runs counter to our intuition that bad fundamentals are somehow “more likely” to trigger a financial crisis, or to tip the economy into recession. In other words, sunspot explanations do not provide a basis for exploring the *correlation* between the underlying fundamentals and the resultant economic outcomes. Finally, comparative-statics analyses and the policy implications that flow from them are only as secure as the equilibrium chosen for this exercise.

We are grateful to the editors for their guidance during the preparation of this paper, and to our two discussants Andy Atkeson and H el ene Rey for their perceptive comments.

The literature on multiple equilibria is large and diverse. The recent book by Cooper (1999) provides a taxonomy for a selection of examples from macroeconomics. Technological complementarities (as in Bryant, 1983), demand spillovers (as in the “big push” model of Murphy, Shleifer, and Vishny, 1989), and thick-market externalities [as in Diamond’s (1982) search model] are some of the examples. Models of financial crises, encompassing both banking crises and attacks on currency pegs, have a similarly large and active research following. Obstfeld and Rogoff (1997) and Freixas and Rochet (1997) are good stepping-off points for this literature.

Our objective in this paper is to encourage a re-examination of the theoretical basis for multiple equilibria. We doubt that economic agents’ beliefs are as indeterminate as implied by the multiple-equilibrium models. Instead, the apparent indeterminacy of beliefs 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; and second, economic agents are assumed to be certain about each other’s 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. We will describe an approach where agents have a small amount of idiosyncratic uncertainty about economic fundamentals. Even if this uncertainty is small, agents will be uncertain about each other’s behavior in equilibrium. This uncertainty allows us as modelers to pin down which set of self-fulfilling beliefs will prevail in equilibrium.

To elaborate on this point, it is instructive to contrast a single-person decision problem with a game. In a single-person decision problem, payoffs are determined by one’s action and the state of the world. When a decision maker receives a message which rules out some states of the world, this information can be utilized directly by disregarding those states in one’s deliberations. However, the same is not true in an environment where payoffs depend on the actions of other individuals as well as on the state of the world. Since my payoff depends on your actions and your actions are motivated by your beliefs, I care about the range of possible beliefs you may hold. So, when I receive a message which rules out some states of the world, it may not be possible to disregard those states in my deliberations, since most of them may carry information concerning your beliefs. Even for small disparities in the information of the market participants, uncertainty about others’ beliefs may dictate a particular course of action as being the uniquely optimal one. In this way, it may prove possible to track the shifts in beliefs as we track the

shifts in the economic fundamentals. There is no longer a choice of what beliefs to hold. One's beliefs are dictated by the knowledge of the fundamentals and the knowledge that other agents are rational.

In this paper, we provide an elementary demonstration of why adding noise to a game with multiple equilibria removes the multiplicity. The analysis builds on the game-theoretic analysis of Carlsson and van Damme (1993) for two-player games and on the continuum-player application to currency attacks of Morris and Shin (1998). We develop a very simple continuum-player example to illustrate the argument, and show by example why this is a flexible modeling approach that can be applied to many of the macroeconomic models with multiplicity discussed above. In doing so, we hope to show that the indeterminacy of beliefs in multiple-equilibrium models is an artifact of simplifying assumptions that deliver more than they are intended to deliver, and that the approach described here is not merely a technical curiosity, but represents a better way of understanding the role of self-fulfilling beliefs in macroeconomics.

We also outline the principal benefits of the approach. One is in generating comparative statics, which in turn aids policy analysis. The other is in suggesting observational implications. Here we summarize those benefits in a general way; below, we will discuss them in the context of particular applications.

Multiple-equilibrium models in macroeconomics are often used as a starting point for policy analysis, despite the obvious difficulties of any comparative-statics analysis with indeterminate outcomes. The unique equilibrium in the approach described here is characterized by a marginal decision maker who, given his uncertainty about others' actions, is indifferent between two actions. Changing parameters in the model then delivers intuitive comparative-statics predictions and implications for optimal policy. In general, we show that inefficiencies are unavoidable in equilibrium. The question is how large such inefficiencies are. The answer turns on the underlying fundamentals of the economy as well as on the information structure of the economic agents. Thus, the notion of a *solvent but illiquid borrower* can be given a rigorous treatment, and the extent of the welfare losses associated with such illiquidity can be calculated.

The theory offers a different perspective on existing empirical work. One traditional approach in the literature is to attempt to distinguish empirically between multiple-equilibrium models and fundamentals-driven models. These ultimately reduce to tests of whether observed fundamentals are sufficient to explain outcomes or whether there is a significant unexplained component that must be attributed to self-

fulfilling beliefs. We argue that correlation between fundamentals and outcomes is exactly what one should expect even when self-fulfilling beliefs are playing an important role in determining the outcome. One will be pessimistic about others' beliefs exactly when fundamentals are weak. The standard sunspot approach, by contrast, offers no theoretical rationale as to why good outcomes should be correlated with good fundamentals (although admittedly this is consistent with the theory and often assumed).

We also suggest one distinctive observational implication. Consider an environment where agents' actions are driven by their beliefs about fundamentals and others' actions. Suppose agents are slightly uncertain about some fundamental variable when they make their decisions, but that *ex post* the econometrician is able to observe the actual realization of that fundamental variable as well as some public signal concerning it that was available to agents at the time. The theory suggests the prediction that the public signals will have an apparently disproportionate impact on outcomes, even controlling for the realization of fundamentals, precisely because it signals information to agents about other agents' equilibrium beliefs.

We start in the next section by analyzing a simple model of bank runs, in the spirit of Diamond and Dybvig (1983), to illustrate the approach in the context of a particular application. Goldstein and Puzner (1999) have developed a richer model; we abstract from a number of complications in order to bring out our methodological message. In Section 3, we show how the insights are more general and can be applied in a variety of contexts. In particular, we discuss models of currency crises and pricing debt in the presence of liquidity risk.

2. *Bank Runs*

There are three dates, $\{0, 1, 2\}$, and a continuum of consumers, each endowed with 1 unit of the consumption good. Consumption takes place at either date 1 or date 2. There is a measure λ of *impatient* consumers who derive utility only from consumption at date 1, and a measure 1 of *patient* consumers for whom consumption at date 1 and at date 2 are perfect substitutes. The consumers learn of their types at date 1. At date 0, the probability of being patient or impatient is proportional to the incidence of the types. Thus, there is probability

$$\frac{\lambda}{1 + \lambda}$$

of being an impatient consumer, and complementary probability of being the patient consumer. All consumers have the log utility function, and the utility of the impatient type is

$$u(c_1) = \log c_1,$$

where c_1 is consumption at date 1, while the utility of the patient type is

$$u(c_1 + c_2) = \log(c_1 + c_2)$$

where c_2 is consumption at date 2.

The consumers can either store the consumption good for consumption at a later date, or deposit it in the bank. Those consumers who have invested their wealth in the bank have a decision at date 1, after learning of their type. They can either leave their money deposited in the bank, or withdraw the sum permitted in the deposit contract (to be discussed below). The bank can either hold the deposits in cash (with rate of return 1) or invest the money in an illiquid project, with gross rate of return $R > 1$ obtainable at date 2. We assume that this technology is only available to the bank. If proportion ℓ of the resources invested in the illiquid investment are withdrawn at date 1, then the rate of return is reduced to $R \cdot e^{-\ell}$, reflecting the costs of premature liquidation. Writing $r \equiv \log R$, this rate of return can be written as $e^{r-\ell}$. We assume that $0 < r < 1$.

2.1 OPTIMAL CONTRACT

We proceed to solve for the optimal contract in this context. The aim is to maximize the ex ante expected utility

$$\frac{\lambda}{1 + \lambda} u(c_1) + \frac{1}{1 + \lambda} u(c_2) \tag{2.1}$$

by choosing the amount c_1 that can be withdrawn on demand at date 1. We assume that the bank is required to keep sufficient cash to fund first-period consumption under the optimal contract. Thus, the first constraint is

$$\lambda c_1 + \frac{c_2}{R} \leq 1 + \lambda, \tag{2.2}$$

which states that the amount held in cash (λc_1) plus the amount invested in the project (c_2/R) cannot exceed the total resources. The second is the incentive compatibility constraint

$$u(c_1) \leq u(c_2), \quad (2.3)$$

which states that patient consumers will, indeed, choose to leave their money in the bank. Ignoring the incentive compatibility constraint, we obtain $c_1 = 1$ and $c_2 = R$. Then,

$$u(c_1) = 0 < r = u(c_2),$$

so that the incentive compatibility constraint is satisfied strictly. Thus, the optimal deposit contract stipulates that any depositor can withdraw the whole of their 1 unit deposit at date 1. Because the investment is assumed to be available only to the bank, such a contract can only be implemented through the bank. Under such a contract, it is a weakly dominant action for every consumer at date 0 to deposit their wealth in the bank. At worst, they will get their money back at date 1, and possibly do better if the consumer turns out to be a patient type. Thus, at date 0, all consumers deposit their money in the bank.

2.2 THE COORDINATION GAME BETWEEN PATIENT CONSUMERS

Diamond and Dybvig (1983) observed that, unfortunately, the optimal contract gives rise to multiple equilibria at date 1. At date 1, the impatient consumers will clearly have a dominant strategy to withdraw. Given this behavior, the patient consumers are playing a coordination game. If a patient consumer withdraws, he gets a cash payoff of 1, giving utility of $0 = u(1)$. This payoff is independent of the number of patient consumers who withdraw. If a patient consumer does not withdraw, then the payoff depends on the proportion of patient consumers who withdraw. If a proportion ℓ withdraw, his cash payoff to leaving money in the bank is $e^{r-\ell}$, which gives utility $r - \ell$. Thus, utility is linearly decreasing in the proportion of patient consumers who withdraw. If a patient consumer expects all other consumers *not* to withdraw (i.e., $\ell = 0$), then his utility from not withdrawing is $r > 0$. Thus there is an equilibrium where all patient consumers conform to the optimal deposit contract and leave their money in the bank. But if a patient consumer expects all other patient consumers to withdraw (i.e., $\ell = 1$), then his utility from not withdrawing is $r - 1 < 0$. Thus there is also an equilibrium where all patient consumers withdraw.

2.3 UNCERTAIN RETURN AND UNIQUE EQUILIBRIUM

Postlewaite and Vives (1987) and Chari and Jagannathan (1988) both examine how bank runs become a unique equilibrium when asymmetric information is added to the model. We follow Goldstein and Pauzner (1999) in introducing a small amount of uncertainty concerning the log return r , holding fixed the deposit contract described above. It should be noted that as soon as we depart from the benchmark case, there is no guarantee that the existing deposit contract is optimal. Neither the portfolio choice of the bank nor the amount that can be withdrawn at date 1 need be optimal in the new context. The objective here is to examine the equilibrium outcome and the welfare losses that result when the benchmark contract is imposed on an environment with noisy signals.

Suppose that r is a normal random variable, and that r has mean \bar{r} and precision α (i.e., variance $1/\alpha$). We carry forward the assumption that the return is neither too small nor too large—we assume that \bar{r} lies in the range:

$$0 < \bar{r} < 1.$$

The depositors have access to very precise information about r before they make their withdrawal decisions, but the information is not perfect. Depositor i observes the realization of the signal

$$x_i = r + \epsilon_i, \tag{2.4}$$

where ϵ_i is normally distributed with mean 0 and precision β , and independent across depositors.

With the introduction of uncertainty, we need to be explicit about what is meant by equilibrium in the bank-run game. At date 1, depositor i not only observes his type, but also observes his signal x_i , and forms the updated belief concerning the return r and the possible signals obtained by other depositors. Based on this information, depositor i decides whether to withdraw or not. A *strategy* for a depositor is a rule of action which prescribes an action for each realization of the signal. A profile of strategies (one for each depositor) is an equilibrium if, conditional on the information available to depositor i and given the strategies followed by other depositors, the action prescribed by i 's strategy maximizes his conditional expected utility. Treating such realization of i 's signal as a possible "type" of this depositor, we are solving for the Bayes Nash equilibria of the imperfect-information game. To economize on the statement of the results, we assume that if withdrawal yields the same ex-

pected utility as leaving money in the bank, then the depositor prefers to leave money in the bank. This assumption plays no substantial role in what follows.

Since both r and x are normally distributed, a depositor's updated belief of r upon observing signal x is

$$\rho = \frac{\alpha \bar{r} + \beta x}{\alpha + \beta} \tag{2.5}$$

In contrast to the benchmark case in which there is no uncertainty, the introduction of uncertainty eliminates multiplicity of equilibrium if private signals are sufficiently accurate. The result depends on the prior and posterior precision of r . Specifically, let

$$\gamma \equiv \frac{\alpha^2 (\alpha + \beta)}{\beta (\alpha + 2\beta)}, \tag{2.6}$$

and write $\Phi(\cdot)$ for the standard normal distribution function. Our main result states that there is a unique equilibrium in this context, provided that γ is small enough.

THEOREM. *Provided that $\gamma \leq 2\pi$, there is a unique equilibrium. In this equilibrium, every patient consumer withdraws if and only if $\rho < \rho^*$, where ρ^* is the unique solution to*


$$\rho^* = \Phi\left(\sqrt{\gamma}(\rho^* - \bar{r})\right).$$

In the limit as γ tends to zero, ρ^ tends to $\frac{1}{2}$.*

Provided that the depositors' signals are precise enough (β is high relative to α), every depositor follows the switching strategy around the critical value ρ^* . This critical value is obtained as the intersection of a cumulative normal distribution function with the 45° line, as depicted in Figure 1. In the limiting case when the noise becomes negligible, the curve flattens out and the critical value ρ^* tends to 0.5. The critical value ρ^* then divides the previously indeterminate region $[0, 1]$ around its midpoint.

Let us sketch the argument behind this result. For ρ^* to be an equilibrium switching point, a depositor whose updated belief is exactly ρ^* ought to be indifferent between leaving his money deposited in the bank and withdrawing it. The utility of withdrawing is zero, and is non-random. The utility of leaving money in the bank is

Figure 1 SWITCHING POINT ρ^*



$$r - \ell \tag{2.7}$$

which is random and depends on ℓ , the proportion of the patient depositors that withdraw. At the switching point ρ^* , the expectation of $r - \ell$ conditional on ρ^* must therefore be zero. The expectation of r conditional on ρ^* is simply ρ^* itself. Thus, consider the expectation of ℓ conditional on ρ^* . Since noise is independent of the true return r , the expected proportion of patient depositors who withdraw is equal to the probability that any particular depositor withdraws. And since the hypothesis is that every depositor follows the switching strategy around ρ^* , the probability that any particular depositor withdraws is given by the probability that this depositor's updated belief falls below ρ^* .


When patient depositor i has posterior belief ρ_i , what is the probability that i attaches to some other depositor j having posterior belief lower than himself? Figure 2 illustrates the reasoning.

Conditional on ρ_i , return r is normal with mean ρ_i and precision $\alpha + \beta$. Since $x_j = r + \epsilon_j$, the distribution of x_j conditional on ρ_i is normal with mean ρ_i and precision

$$\frac{1}{\frac{1}{\alpha + \beta} + \frac{1}{\beta}} = \frac{\beta(\alpha + \beta)}{\alpha + 2\beta} \tag{2.8}$$

But $\rho_j = (\alpha \bar{r} + \beta x_j) / (\alpha + \beta)$, so that the distribution of $\rho_j | \rho_i$ is as depicted in Figure 2, and the probability that ρ_j is less than ρ_i conditional on ρ_i is given by the shaded area. Moreover,

Figure 2 BELIEFS CONDITIONAL ON ρ_i



$$\rho_j < \rho_i \Leftrightarrow \frac{\alpha \bar{r} + \beta x_j}{\alpha + \beta} < \rho_i \Leftrightarrow x_j < \rho_i + \frac{\alpha}{\beta} (\rho_i - \bar{r}), \tag{2.9}$$

so the question of whether ρ_j is smaller than ρ_i can be reduced to the question of whether x_j is smaller than $\rho_i + (\alpha/\beta)(\rho_i - \bar{r})$. Hence,

$$\begin{aligned} \text{Prob}(\rho_j < \rho_i | \rho_i) &= \text{Prob}\left(x_i < \rho_i + \frac{\alpha}{\beta}(\rho_i - \bar{r}) \mid \rho_i\right) \\ &= \Phi\left(\sqrt{\frac{\beta(\alpha + \beta)}{\alpha + 2\beta}}\left(\rho_i + \frac{\alpha}{\beta}(\rho_i - \bar{r}) - \rho_i\right)\right) \\ &= \Phi(\sqrt{\gamma}(\rho_i - \bar{r})). \end{aligned} \tag{2.10}$$


So the shaded area in Figure 2 can be represented in terms of the area under a normal density which is centered on the ex ante mean \bar{r} . Figure 3 illustrates.

If ρ^* is an equilibrium switching point, the expectation of $r - \ell$ conditional on ρ^* must be zero. Since

$$E(r - \ell | \rho^*) = \rho^* - \Phi(\sqrt{\gamma}(\rho^* - \bar{r})), \tag{2.11}$$

ρ^* must be the point at which $\Phi(\sqrt{\gamma}(\rho - \bar{r}))$ intersects the 45° line, exactly as depicted in Figure 1. Provided that γ is small enough, the slope of $\Phi(\sqrt{\gamma}(\rho - \bar{r}))$ is less than one, so that there can be at most one point of intersection. Since the slope of the cumulative normal is given

Figure 3 DENSITY $\phi(\sqrt{\gamma}(\rho_i - \bar{r}))$



by the corresponding density function (which has the maximum value of $\sqrt{\gamma/2\pi}$), we can guarantee that there is a unique intersection point provided that γ is less than 2π . All that remains is to show that if there is a unique symmetric equilibrium in switching strategies, there can be no other equilibrium. Appendix A completes the argument.

2.4 COMPARATIVE STATICS AND POLICY ANALYSIS

The uniqueness of equilibrium makes it possible to perform secure comparative-statics analysis. We will illustrate this with a simple exercise in our example, where an early-withdrawal penalty t is imposed on consumers who withdraw in period 1.

In order to set a benchmark to measure our results against, consider the case with no uncertainty. The log return r is commonly known, and there is multiplicity of equilibria. The introduction of the early-withdrawal penalty has little effect in this case. The only effect is to shift the range of returns where multiple equilibria exist from $[0, 1]$ to $[\log(1 - t), \log(1 - t) + 1]$. Without a theory guiding us as to which outcome results in the game, it is hard to evaluate the welfare consequences of this policy. The most we can say is that when r is close to 1 [i.e., in the marginal interval $(\log(1 - t) + 1, 1]$], the tax will remove the multiplicity of equilibrium, and the efficient outcome that consumers do not withdraw will occur for sure. When r is slightly less than 0 [i.e., in the marginal interval $(\log(1 - t), 0]$], the tax will allow multiple equilibria.

In contrast to the lack of meaningful comparative statics when r is common knowledge, we can say much more when r is observed with noise. In particular, contrast the case with no uncertainty with the case in which noise is negligible (i.e. the limiting case where $\gamma \rightarrow 0$). The theorem tells us that patient consumers will withdraw if and only if $\rho < \log(1 - t) + \frac{1}{2}$. This allows us to calculate the incidence of withdrawals at any realized value of r . Policy affects outcomes for interior values of the

parameters, by shifting the boundary of the two populations, not merely at extremal parameter values.

We can also use this unique equilibrium to examine policy trade-offs. Recall that the efficient outcome at date 1 is for withdrawal by patient consumers to take place only if $r < 0$. If noise concerning r is very small, we achieve this outcome with very high probability by setting $t = 1 - e^{-1/2}$ [so that $\log(1 - t) + \frac{1}{2} = 0$]. But of course achieving efficiency in the withdrawal decision comes at the cost of reducing the value of the contract to consumers. The explicit form for the equilibrium allows us to calculate the ex ante expected utility of consumers. For any given t , it is $1/(1 + \lambda)$ times

$$[\lambda + \Phi(\sqrt{\alpha}(\log(1 - t) + \frac{1}{2} - \bar{r}))] \log(1 - t) + \int_{\log(1-t)+1/2}^{\infty} r \phi(\sqrt{\alpha}(r - \bar{r})) \sqrt{\alpha} dr,$$

while the revenue from the penalty is

$$[\lambda + \Phi(\sqrt{\alpha}(\log(1 - t) + \frac{1}{2} - \bar{r}))] t.$$


An increase in the penalty can be welfare-enhancing for consumers (even if they derive no benefit from the tax revenue). Goldstein and Pauzner (1999) examine contracts where early-withdrawal penalties are received by consumers who leave their money until date 2. This further enhances the desirability of early-withdrawal penalties from the consumers' point of view.

2.5 OBSERVABLE IMPLICATIONS

We have presented a highly simplified model of bank runs. Even in this model, though, we can start thinking about observable implications of this theory. The main prediction is that despite the self-fulfilling aspect of the bank run, each depositor will withdraw his money exactly when his beliefs about the riskiness of bank deposits crosses some threshold, implying that the size of equilibrium bank runs will be negatively correlated with returns. Consider the incidence of deposit withdrawals as given by the equilibrium value of ℓ . This incidence is a random variable that depends on the realized return r . A depositor withdraws whenever his posterior belief falls below the critical value ρ^* , which happens whenever

$$\frac{\alpha \bar{r} + \beta x_i}{\alpha + \beta} < \rho^*.$$

Figure 4 PROPORTION $\ell(r)$ OF WITHDRAWALS



In other words, a depositor withdraws whenever the realization of his signal x_i falls below the critical value

$$x^*(\rho^*, \bar{r}) \equiv \frac{\alpha + \beta}{\beta} \rho^* - \frac{\alpha}{\beta} \bar{r}. \tag{2.12}$$

Since $x_i = r + \epsilon_i$, the incidence of withdrawal is a function of the realized return r , and is given by

$$\ell(r) = \Phi(\sqrt{\beta} (x^*(\rho^*, \bar{r}) - r)). \tag{2.13}$$

Figure 4 illustrates.

Clearly, the incidence of withdrawal is high when the return is low. Fundamentals plays a key explanatory role. Gorton (1988) studies bank panics in the U.S. national-banking era (1863–1914). He interprets the data in the light of the traditional dichotomy between fundamentals and sunspots as a cause of panics:


A common view of panics is that they are random events, perhaps self-confirming equilibria in settings with multiple equilibria, caused by shifts in the beliefs of agents which are unrelated to the real economy. An alternative view makes panics less mysterious. Agents cannot discriminate between the riskiness of various banks because they lack bank-specific information. Aggregate information may then be used to assess risk, in which case it can occur that all banks may be perceived to be riskier. Consumers then withdraw enough to cause a panic. . . . [This latter] hypothesis links panics to occurrences of threshold value of some variable depicting the riskiness of bank deposits.

He concludes that the latter theory performs well. The highly simplified model of bank runs presented here suggests a reinterpretation of the evidence. The theory suggests that depositors will indeed withdraw their money when the perceived riskiness of deposits crosses a threshold value. But nonetheless, the banking panic is self-fulfilling in the sense that individual investors only withdraw because they expect others to do so. The theory suggests both that banking panics are correlated with poor fundamentals and that inefficient self-fulfilling panics occur. Of course, it is possible to make assumptions about sunspots that mimic these predictions; but the theory presented here places tighter restrictions on outcomes than sunspot theory.

One would like to come up with distinctive implications that are harder to mimic with judiciously chosen sunspots. We will suggest one example in this bank-deposit context. Suppose that we were able to observe both the prior mean of the log return \bar{r} and the realized log return r ; the prior mean \bar{r} is a public signal that is observable by all depositors when they make their withdrawal decisions. Our theory predicts that for any given level of fundamentals r , the proportion of consumers running would be decreasing in \bar{r} . This is apparent from our theorem, since a fall in the ex ante mean \bar{r} shifts the curve $\Phi(\sqrt{\gamma}(\rho - \bar{r}))$ to the left, so that its intersection with the 45° degree line is shifted to the right. Figure 5 illustrates this shift. Thus, when the fundamentals are commonly known to be weak (i.e. \bar{r} is low), the equilibrium strategy dictates much more aggressive withdrawals, even controlling for one's posterior belief about r .

A prediction of the model, then, would be that if we could divide the fundamentals variables in Gorton's analysis into those that were most

Figure 5 SHIFT IN ρ^*



readily available to depositors contemporaneously and those that were not, we should expect the most readily available variables to have the biggest effect. We will come across another instance of the impact of public information below.

3. *Complementarities and Macroeconomics*

The example above was constructed around a simple coordination game played by a continuum of players. Much of the macroeconomics literature on complementarities, multiple equilibria, and sunspots similarly reduces in the end to coordination games played by large populations. In this section, we illustrate how other issues can be addressed using similar methods.

Consider the following class of problems. A continuum of individuals must choose between a safe action and a risky action. If an individual chooses the safe action, his payoff is a constant. If he chooses the risky action, his payoff is an increasing function of the “state of fundamentals” r but a decreasing function of the proportion of the population who choose the safe action, ℓ . In the bank-run example above, the payoff was linear in both r and ℓ . This linearity allowed us to give simple characterizations of the equilibrium. But as long as the payoff to the risky action is increasing in r and decreasing in ℓ , there will be a unique equilibrium of the type described above when information is sufficiently accurate. We will give an informal description of two applications that fit this general setup that we have analyzed elsewhere.

3.1 CURRENCY CRISES

A continuum of speculators must decide whether to attack a fixed exchange rate. The cost to the monetary authority of defending the peg depends on the fundamentals of the economy and the proportion of speculators who attack the currency. If the monetary authority has some fixed benefit of maintaining the peg, then for each realization of fundamentals, there will be some critical mass of speculators sufficient to induce abandonment of the currency. If the peg is abandoned, the exchange rate will float to some level that depends on the fundamentals. A speculator may choose to attack by selling a fixed amount of the currency short. If he attacks, he must pay a transaction cost but receives the difference between the peg and the floating rate if the attack is successful and there is a devaluation.

This stylized model is in the spirit of the self-fulfilling-attacks literature (see, for example, Obstfeld, 1996). If the state of fundamentals is common knowledge, there are three ranges of fundamentals to consider. If funda-

mentals are sufficiently low, devaluation is guaranteed. If fundamentals are sufficiently high, there will be no devaluation. But for some intermediate range of fundamentals, there are multiple equilibria. Morris and Shin (1998) show how if there is a small amount of noise concerning fundamentals, there is a unique equilibrium.

Now consider a policy that makes it harder for an attack to be successful. For example, the monetary authority might accumulate reserves. A naive calculation of the value of those reserves might involve calculating the likelihood of contingencies in which those extra reserves would make the difference in the authority's ability to defend against an attack. This is analogous to seeing when a tax on early withdrawals would remove the existence of a withdrawal equilibrium in the bank-run model. But taking into account the strategic analysis, we see that the true benefit of accumulating reserves is as a confidence-building measure. If the accumulation of reserves is publicly observed, speculators will anticipate that other speculators will be less aggressive in attacking the currency. So in regions of fundamentals where a self-fulfilling attack is in fact feasible, it will not occur.

The theory also generates intuitive predictions about which events lead to currency attacks. Deteriorating fundamentals, even if observed by most participants, will have less effect if the fact that fundamentals are deteriorating is not common knowledge. Very public signals that fundamentals have deteriorated only a small amount may have a large impact. This is because a speculator observing a bad signal not only anticipates that the monetary authority will have a harder time defending against an attack, but also anticipates that other speculators will be attacking. This explanation is quite commonplace. But the theoretical model that we have described captures this argument exactly.

3.2 PRICING DEBT

Our methods may also help us to understand some of the anomalies noted in the empirical literature on the pricing of defaultable debt. One influential approach has been to note that a lender's payoff is analogous to the payoff that arises from holding a short position in a put option on the borrower's assets. Hence, option-pricing techniques can be employed to price debt, as shown in the classic paper by Merton (1974). Nevertheless, the empirical success of this approach has been mixed, with the usual discrepancy appearing in the form of the overpricing (by the theory) of the debt, and especially of the lower-quality, riskier debt. The anomaly would be explained if it can be shown that the default trigger for asset values actually *shifts* as the underlying asset changes in value, and shifts in such a way that disadvantages lower-quality debt.

The incidence of inefficient liquidation seen in our bank-run example suggests that similar inefficiencies might arise in the coordination problem between creditors facing a distressed borrower. This would give us a theory of *solvent* but *illiquid* borrowers, enabling us to address the empirical anomalies. This is attempted in Morris and Shin (1999).

When the fundamentals are bad, coordination to keep a solvent borrower afloat is more difficult to achieve, and the probability of inefficient liquidation is large. This is another manifestation of the importance of public information in achieving coordination alluded to in the previous section. The disproportionate impact of public information can be illustrated in the following example of a borrower in distress.

Consider a group of lenders who are funding a project. Time is discrete, and advances by increments of $\Delta > 0$. The fundamentals of the project at date t are captured by the random variable r_t . Conditional on its current realization, the next realization of r_t is i.i.d., normally distributed around its current realization, with variance Δ . In other words, $\{r_t\}$ is a sequence of snapshots of a simple Brownian motion at time intervals of Δ . To economize on notation, we denote by r the current value of the fundamentals, and by r_+ its value in the next period. At each date, every lender chooses whether or not to continue funding the project. The project fails if and only if

$$\ell > r,$$

where ℓ is the proportion of creditors who pull out of the project. Hence, when $r > 1$ the project is viable irrespective of the actions of the creditors. If $r < 0$, the project fails irrespective of the actions of the creditors. However, when r lies between 0 and 1, the fate of the project depends on how severe the creditor run is. At each date, a lender receives a payment of 1 if the project has survived. When the project fails, a lender receives zero. By pulling out, a lender receives an intermediate payoff λ , where $0 < \lambda < 1$. We also suppose that a creditor who withdraws when the project is still viable rejoins the project in the next period (having missed a single payment of 1). This assumption ensures that the creditors face a sequence of one-shot games.

None of the creditors observe the current fundamentals perfectly. Each has signal

$$x_i = r + \epsilon_i,$$

where ϵ_i and ϵ_j are independent for $i \neq j$, and ϵ_i is normal with mean 0 and variance Δ^2 . The noise in the signal x is thus small compared to the

underlying uncertainty. The lenders, however, observe the previous realization of r perfectly. This will serve as the public information on which much of the analysis will hinge. As the time interval Δ becomes small, the noise disappears at a faster rate than the overall uncertainty governing r . Each lender chooses an action based on the realized signal x and the (commonly known) previous realization of r .

This game has a unique equilibrium (the proof is sketched in Appendix B) in which there is a critical value of fundamentals r_+^* for which the project fails next period whenever $r_+ < r_+^*$. We call r_+^* the *collapse point* for the project. It is given by the (unique) solution to

$$r_+^* = \Phi \left(r_+^* - r + \Phi^{-1}(\lambda)\sqrt{1+\Delta} \right). \quad (3.1)$$

The collapse point is obtained as the intersection between the 45° line and the distribution function for a normal with unit variance centered on $r - \Phi^{-1}(\lambda)\sqrt{1+\Delta}$. The following points are worthy of note.

1. r_+^* is a function of the current realization r . Hence, public information plays a crucial role in determining the trigger point for collapse.
2. The continuous time limit as $\Delta \rightarrow 0$ is well defined.
3. r_+^* is a *decreasing* function of r . So, when fundamentals deteriorate, the probability of collapse increases not only because the fundamentals are worse, but also because the trigger point has moved unfavorably.

This last feature is possibly quite significant. For an asset whose fundamentals are bad (i.e., r is low), the probability of default is higher than would be the case in the absence of coordination problems among creditors. Such a pattern would explain why one would misprice such an asset in a model that assumes a fixed default point. The mispricing takes the form of *overpricing* the riskier bonds—exactly the empirical anomaly discussed in the literature.

There is a more general lesson. The onset of financial crises can be very rapid, and many commentators note how the severity of a crisis is disproportionate to the deteriorating fundamentals. In our account, such apparently disproportionate reactions arise as an essential feature of the model. When fundamentals deteriorate, coordination is less easy to achieve. We can explore this effect further by examining the comparative statics of the probability of collapse. The probability of collapse next period conditional on the current fundamentals r is

$$\Phi \left(\frac{r_+^* - r}{\sqrt{\Delta}} \right).$$

As r falls, the probability of collapse increases at the rate

$$\frac{\phi}{\sqrt{\Delta}(1 - \phi)},$$

where ϕ is the standard normal density at $(r_+^* - r)/\sqrt{\Delta}$. The increase in the probability of collapse can be quite large when r hovers close to the collapse point, and the onset of failure can thus be quite rapid. As compared to the naive model which does not take into account the dependence of the collapse point on the current fundamentals, this is larger by a factor of $1/(1 - \phi)$. When r is close to the collapse point r_+^* , this is roughly $\sqrt{2\pi}/(\sqrt{2\pi} - 1) \approx 1.66$.

The inverse relationship between the current value of fundamentals and the collapse point is suggestive of the precipitous falls in the price of defaultable securities during financial crises.

The continuous time limit of the model makes possible further simplifications in the analysis. Taking the limit as $\Delta \rightarrow 0$, the fundamentals r evolve as a simple Brownian motion, and the collapse point r_+^* for the next period converges to the collapse point in the current period. So (3.1) can be written

$$r^* = \Phi \left(r^* - r + \Phi^{-1}(\lambda) \right)$$

Collapse occurs when r hits r^* , i.e. at $r^* = \lambda$.

3.3 HOW SPECIAL IS THE ANALYSIS?

In this paper, we have described stylized examples with normally distributed states and signals, binary choices by a symmetric continuum of players, and payoffs linear in the state and proportion of players choosing each action. These assumptions allowed us to give simple characterizations of the unique equilibrium. However, the analysis is arguably quite general. If one is only interested in the limiting case where noise in signals is very small, the exact shape of the noise or prior beliefs about the state do not matter. Asymmetries among the players can also be incorporated. Corsetti et al. (1999) examine the role of a large trader in currency markets in an asymmetric game. The qualitative features of the analysis are very similar between continuum and finite player cases. Indeed, in the special case of the payoffs in the bank-run model, where only the proportion of other players choosing each action matters, the analysis is literally unchanged. That is, if we had a finite number of depositors, with proportion $\lambda/(1 + \lambda)$ impatient and proportion $1/(1 + \lambda)$

patient, the unique equilibrium would have patient consumers using the same cutoff point for withdrawals. Dealing with many actions is more delicate (see Frankel, Morris, and Pauzner, 2000), although the analysis extends straightforwardly in some instances. Carlsson and Ganslandt (1998) describe what happens when noise is added to Byrant's (1983) model of technological complementarities.

4. *Concluding Remarks*

We draw two conclusions from our analysis. The first is that applied theorists should be wary of selecting an arbitrary outcome for further attention when conducting comparative-statics exercises and in drawing policy implications. The mere fact that an outcome is Pareto-superior to another is no good reason for it to be selected, and we should expect to see some inefficiencies as a rule. The notion of a "solvent but illiquid bank" can be given a rigorous treatment, and we hope that our discussions can contribute to policy debates in the area.

Our second conclusion is a methodological one. Contrary to the impression given by multiple-equilibrium models of the apparent autonomy of beliefs to float freely over the fundamentals, we believe that such autonomy of beliefs is largely illusory when information is modeled in a more realistic way. No doubt some researchers may find this regrettable, since one degree of freedom is lost in the exercise of providing a narrative of unfolding events. However, there are compensations for this loss, and we hope that these benefits will be recognized by researchers. One promising line of inquiry is to explore the correlations between the underlying fundamentals and the degree of optimism of the economic agents. Empirical investigations will then have a much firmer basis.

Appendix A

When there is a unique symmetric equilibrium in switching strategies, there can be no other equilibrium. An argument is sketched here. Denote by $u(\rho, \hat{\rho})$ the expected utility from leaving one's money in the bank conditional on posterior ρ when all other patient depositors follow a switching strategy around $\hat{\rho}$. Conditional on ρ , the expected proportion of depositors who withdraw is given by the probability that any particular depositor receives a signal lower than the critical value $\hat{\rho}$. From the argument in the text, this probability is given by

$$\Phi \left(\sqrt{\frac{\beta(\alpha+\beta)}{\alpha+2\beta}} \left(\hat{\rho} + \frac{\alpha}{\beta} (\hat{\rho} - \bar{r}) - \rho \right) \right) = \Phi \left(\sqrt{\gamma} \left(\hat{\rho} - \bar{r} + \frac{\beta}{\alpha} (\hat{\rho} - \rho) \right) \right). \quad (\text{A.1})$$

Hence, $u(\rho, \hat{\rho})$ is given by

$$u(\rho, \hat{\rho}) = \rho - \Phi \left(\sqrt{\gamma} \left(\hat{\rho} - \bar{r} + \frac{\beta}{\alpha} (\hat{\rho} - \rho) \right) \right). \quad (\text{A.2})$$

If r is negative, the utility from withdrawing is higher than that from leaving money in the bank, irrespective of what the other depositors decide. So, if the posterior belief ρ is sufficiently unfavorable, withdrawing is a dominant action. Let $\underline{\rho}_1$ be the threshold value of the belief for which withdrawal is the dominant action. Any belief $\rho < \underline{\rho}_1$ will then dictate that a depositor withdraws. Both depositors realize this, and each rules out strategies of the other depositor which leave money in the bank for signals lower than $\underline{\rho}_1$. But then, leaving money in the bank cannot be optimal if one's signal is lower than $\underline{\rho}_2$, where $\underline{\rho}_2$ solves

$$u(\underline{\rho}_2, \underline{\rho}_1) = 0. \quad (\text{A.3})$$

This is so because the switching strategy around $\underline{\rho}_2$ is the best reply to the switching strategy around $\underline{\rho}_1$, and even the most optimistic depositor believes that the incidence of withdrawals is higher than that implied by the switching strategy around $\underline{\rho}_1$. Since the payoff to withdrawing is increasing in the incidence of withdrawal by the other depositors, any strategy that leaves money in the bank for signals lower than $\underline{\rho}_2$ is dominated. Thus, after *two* rounds of deletion of dominated strategies, any strategy that leaves money in the bank for signals lower than $\underline{\rho}_2$ is eliminated. Proceeding in this way, one generates the increasing sequence

$$\underline{\rho}_1 < \underline{\rho}_2 < \dots < \underline{\rho}_k < \dots, \quad (\text{A.4})$$

where any strategy that leaves money in the bank for a signal $\rho < \underline{\rho}_k$ does not survive k rounds of deletion of dominated strategies. This sequence is increasing, since $u(\cdot, \cdot)$ is increasing in its first argument and decreasing in its second. The smallest solution $\underline{\rho}$ to the equation $u(\rho, \rho) = 0$ is the least upper bound of this sequence, and hence its limit. Any strategy that leaves money in the bank for signal lower than $\underline{\rho}$ does not survive iterated dominance.

Conversely, if ρ is the largest solution to $u(\rho, \rho) = 0$, there is an exactly analogous argument from “above,” which demonstrates that a strategy that withdraws for signals larger than ρ does not survive iterated dominance. But if there is a *unique* solution to $u(\rho, \rho) = 0$, then the smallest solution just is the largest solution. There is precisely one strategy remaining after eliminating all iteratively dominated strategies. Needless to say, this also implies that this strategy is the only *equilibrium* strategy.

Appendix B

The posterior belief of the current value of r is normal with mean

$$\rho = \frac{x_i + \Delta r_-}{1 + \Delta}$$

and precision $(1 + \Delta)/\Delta^2$, where r_- denotes the previous realization of r . Denote by $U(\rho)$ the payoff to continuing with the project conditional on ρ when all creditors are following the ρ -switching strategy. It is given by

$$U(\rho) = \Phi\left(\frac{\sqrt{1+\Delta}(r^* - \rho)}{\Delta}\right), \quad (\text{B.1})$$

where r^* is the trigger value of fundamentals at which the project collapses. r^* satisfies $r^* = \ell$. But if other speculators follow the ρ -switching strategy, ℓ is the proportion of creditors whose signal is lower than the marginal value of x that implies the switching posterior ρ . This gives

$$r^* = \Phi\left(\rho - r_- + \frac{\rho - r^*}{\Delta}\right). \quad (\text{B.2})$$

From these two equations, we can show by implicit differentiation that $U'(\rho) > 0$. There is a unique solution to $U(\rho) = \lambda$, and the equilibrium is unique for the same reasons as cited for the main theorem. To solve explicitly for the collapse point r^* , we solve the pair of equations given by (B.2) and $U(\rho) = \lambda$. This gives

$$r^* = \Phi\left(r^* - r_- + \Phi^{-1}(\lambda) \sqrt{1+\Delta}\right),$$

as required.

REFERENCES

- Bryant, J. (1983). A simple rational expectations Keynes type model. *Quarterly Journal of Economics* 98:525–529.
- Carlsson, H., and M. Ganslandt. (1998). Noisy equilibrium selection in coordination games. *Economic Letters* 60:23–34.
- Carlsson, H., and E. van Damme. (1993). Global games and equilibrium selection. *Econometrica* 61:989–1018.
- Chari, V. V., and R. Jagannathan. (1988). Banking panics, information and rational expectations equilibrium. *Journal of Finance* 43:749–761.
- Cooper, R. (1999). *Coordination Games: Complementarities and Macroeconomics*. Cambridge: Cambridge University Press.
- Corsetti, G., A. Dasgupta, S. Morris, and H. S. Shin. (1999). Does one Soros make a difference? A theory of currency crises with large and small traders. Yale University.
- Diamond, D., and P. Dybvig. (1983). Bank runs, deposit insurance and liquidity. *Journal of Political Economy* 91:401–119.
- Diamond, P. (1982). Aggregate demand management in search equilibrium. *Journal of Political Economy* 90:881–894.
- Freixas, X., and J.-C. Rochet. (1997). *Microeconomics of Banking*. MIT Press.
- Frankel, D., S. Morris, and A. Pauzner. (2000). Equilibrium selection in global games with strategic complementarities. Yale University.
- Goldstein, I., and A. Pauzner. (1999). Endogenous probability of bank runs in a rational expectations model. Tel Aviv University.
- Gorton, G. (1988). Banking panics and business cycles. *Oxford Economic Papers* 40:751–781.
- Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance* 29:449–470.
- Morris, S., and H. S. Shin. (1997). Approximate common knowledge and coordination: Recent lessons from game theory. *Journal of Logic, Language and Information* 6:171–190.
- , and ———. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review* 88:587–597.
- , and ———. (1999). Coordination risk and the price of debt. Unpublished paper. <http://www.nuff.ox.ac.uk/users/Shin/working.htm>.
- Murphy, K., A. Schleifer, and R. Vishny. (1989). Industrialization and the big push. *Journal of Political Economy* 97:1003–1026.
- Obstfeld, M. (1996). Models of currency crises with self-fulfilling features. *European Economic Review* 40:1037–1047.
- , and K. Rogoff. (1997). *Foundations of International Macroeconomics*. MIT Press.
- Postlewaite, A., and X. Vives. (1987). Bank runs as an equilibrium phenomenon. *Journal of Political Economy* 95:485–491.

Comment

ANDREW ATKESON

Federal Reserve Bank of Minneapolis and University of Minnesota

1. Introduction

Macroeconomists have used coordination games with multiple equilibria to describe any number of phenomena in which we appear to see large changes in economic outcomes with little or no apparent change in the underlying economic fundamentals. Usually, in macroeconomic applications, these games are shown to have multiple equilibria and the argument is made that large changes in economic outcomes can follow from changes in agents' expectations about what other agents will do rather than from changes in economic fundamentals alone.

Morris and Shin present a simple and dramatic insight into the structure of simple coordination games. With only a few assumptions, they show that if agents see a noisy signal of the true state of the world and thus have some uncertainty about the exact structure of the coordination game that they are playing as well as some uncertainty about what every other agents knows about the coordination game that they are playing, then these games in fact have a unique equilibrium corresponding to each underlying state of the world. This result suggest that macroeconomists should reassess whether their previous findings of multiple equilibria in these coordination games are robust to small changes in the structure of information available to agents.

Morris and Shin go on to show that if the noise introduced into the coordination game is small, the selected equilibrium has the feature that there is a threshold state of the world around which the economic outcome changes very rapidly with small changes in the state, while in the other regions of the state space, the economic outcome is quite stable as the underlying state of the world varies. This second feature of the equilibrium selected by Morris and Shin's apparatus suggests that their work may be more than a criticism of the robustness of previous multiple-equilibrium literature and may, in fact, have important implications for a number of applications.

My discussion of this paper by Morris and Shin has four parts. First, I present what I think is the simplest environment in which to apply their apparatus. Second, I go over a proof of their result in this environment that is slightly different than the proof presented in the paper. I hope that it will help any reader interested in understanding the logic of Morris and Shin's results to see the argument from a different angle. Third, I describe

what I find to be the most interesting feature of the equilibrium that is selected and also go over an example of how one might use this technology to do comparative statics. Fourth, I describe what I believe is the main impediment to use of this technology for modeling macroeconomic phenomena.

To jump ahead for a moment, this fourth and final part of my discussion does not focus on the applied question of whether the models that Morris and Shin have proposed for currency crises and the pricing of corporate debt in related papers are relevant for analyzing those phenomena. Instead I focus on the broader question of whether one can introduce markets and prices, clearly essential parts of any macroeconomic application, into what, to date, has been a purely game-theoretic analysis. Morris and Shin, in their introduction, criticize previous applications of coordination games in macroeconomics for relying on assumptions that “allow agents’ actions and beliefs to be perfectly coordinated in a way that invites multiplicity of equilibria.” The noise that they introduce into coordination games has the effect of preventing coordination of agents’ actions and beliefs. In a market economy, however, prices serve precisely to coordinate actions (so that supply equals demand), and in a dynamic market economy, asset prices play an important role in coordinating agents’ beliefs, since these prices tend to aggregate information across individuals.

It is not clear to me how the argument presented by Morris and Shin would carry over to a model with markets. Their arguments require agents to have diverse beliefs about the probabilities of future outcomes in equilibrium, and this typically does not happen in models in which agents see the market signals about those probabilities embodied in asset prices. The nature of this difficulty in translating Morris and Shin’s technology to a market environment should become clearer after we review the details of how this technology works.

2. *A Simple Coordination Game*

Let us review how Morris and Shin’s technology works in the context of what seems a natural application of the game theory. Consider a crowd that faces riot police in the street. Individuals in the crowd must decide whether to riot or not. If enough people riot, the riot police are overwhelmed, and each rioter gets loot $W > 0$. If too few people riot, the riot police contain the riot, and each rioter gets arrested with payoff $L < 0$. Individuals who choose not to riot leave the crowd and get safe payoff 0. The strength of the riot police depends on the state of the world, θ , and the strictly increasing function $a(\theta)$ indexes the fraction of the crowd that

must riot to overwhelm the police. Let $\underline{\theta}$ denote the point at which $a(\theta)$ crosses 0, and $\bar{\theta}$ the point at which $a(\theta)$ crosses 1.

The equilibria of this game when the state θ is common knowledge are as follows. If $\theta \leq \underline{\theta}$, then it is a dominant strategy for each individual to riot, since the riot police in this case are so weak that they cannot stop even a single rioter [$a(\theta) \leq 0$]. Thus, if θ is in this region of the state space, everyone riots and gets payoff W for sure. If $\theta > \bar{\theta}$, then it is a dominant strategy for each individual not to riot, since the police can contain the crowd even if everyone riots [$a(\theta) > 1$]. Thus, if θ is in this region of the state space, no one riots and everyone gets payoff 0 for sure. In the middle of the state space, with $\underline{\theta} < \theta \leq \bar{\theta}$, there are two possible equilibria corresponding to each value of the state θ . In the first of these equilibria, everybody riots. In this case, the fraction of the crowd that riots is $1 \geq a(\theta)$, so the police are overwhelmed and everybody gets payoff $W > 0$. In the second of these equilibria, nobody riots. In this case the fraction of the crowd that riots is $0 < a(\theta)$, so the police contain the crowd and any individual who riots is arrested. Hence, nobody riots, and everybody in the crowd gets $0 > L$. This game clearly has multiple equilibria in a region of the state space, and when the state variable is in this region, the economic outcome depends on agents' expectations of what other agents will do and not on the underlying economic fundamental $a(\theta)$.

3. *An Alternative Presentation of the Proof of Their Result*

Morris and Shin introduce the following changes into this coordination game. They assume that individuals do not know the state of the world, θ . Instead, each individual starts with a common prior that θ is normally distributed with some mean m_θ and variance $1/\alpha$ (precision α). (I think of the randomness in θ as arising from the problem that the precise strength of the squad of riot police available to any particular crowd in any particular street at any particular time depends somewhat on chance.) Each individual in the crowd then receives an idiosyncratic signal $x_i = \theta + \epsilon_i$ of the state θ , where ϵ_i is normally distributed with mean 0 and variance $1/\beta$ (precision β) and is i.i.d. across individuals. Given these assumptions, we have two distributions that play a key role in the analysis. First is the distribution of signals x_i across agents conditional on the realization of the state θ . With the assumptions above, this is a normal distribution, but we can write it more generally as a c.d.f. $\text{Prob}(x \leq x^*|\theta)$, which we will assume to be a strictly positive, continuous, decreasing function of θ for any value of x^* . Second is the posterior distribution over θ for an agent who has seen signal x . This is obtained

from Bayes's rule and, under the assumptions above, is a normal distribution; but it can also be written more generally as a c.d.f. $\text{Prob}(\theta \leq \theta^* | x)$. We also assume that this is a continuous and decreasing function of x for any value of θ^* .

Morris and Shin's result in the context of this simple game can then be stated as follows. Assume that there is a unique solution x^* , θ^* to the following two equations:

$$\text{Prob}(x \leq x^* | \theta) = a(\theta^*), \tag{1}$$

$$\text{Prob}(\theta \leq \theta^* | x^*)W + [1 - \text{Prob}(\theta \leq \theta^* | x^*)]L = 0. \tag{2}$$

Then there is unique equilibrium described by x^* and θ^* . The signal x^* is a threshold signal such that all individuals who get signals $x \leq x^*$ riot, and those who get signal $x > x^*$ do not riot. The state θ^* is a threshold state such that the crowd overwhelms the police, so that rioters get payoff W if $\theta \leq \theta^*$ and the police contain the crowd, and rioters are arrested and get payoff L if $\theta > \theta^*$.

In the paper, Morris and Shin make assumptions on the precision of the signal relative to the precision of the prior in stating the result. In proving their proposition they show that there is a unique solution to the analogues to equations (1) and (2) if we assume that the precision of the signal, denoted β , is sufficiently high relative to the precision of the prior, denoted α , and the slope of the function $a(\theta)$. The necessary and sufficient condition for their result, however, appears to be that these two equations have a unique solution.

One way to prove this proposition is by iterated deletion of dominated strategies. I find this proof the easiest to understand. It goes as follows.

First observe that individuals who get sufficiently low and high signals, which I denote x_0 for the low signal and x^0 for the high signal, are so confident of their posterior beliefs that $\theta \leq \underline{\theta}$ or $\theta > \bar{\theta}$, that they find it a dominant strategy to riot or not riot, respectively, regardless of what everyone else does. The low signal x_0 is the highest value of x such that

$$\text{Prob}(\theta \leq \underline{\theta} | x)W + \text{Prob}(\theta > \underline{\theta} | x)L \geq 0.$$

The interpretation here is that, even if one believed that everyone else in the crowd was not going to riot, and thus any individual rioter would be arrested in the event that $\theta > \underline{\theta}$, the posterior probability that $\theta \leq \underline{\theta}$ for someone who saw $x \leq x_0$ is high enough to make it worthwhile to run the risk of rioting.

Analogous reasoning defines x^0 . Even with the belief that everyone

else always riots and thus that rioters will get W if $\theta \leq \bar{\theta}$, someone who saw signal $x > x^0$, where x^0 is the smallest x such that

$$\text{Prob}(\theta \leq \bar{\theta}|x)W + \text{Prob}(\theta > \bar{\theta}|x)L \leq 0,$$

would not find the potential reward of rioting likely enough to justify the risk. These two observations give us the first round of deletion of dominated strategies: any equilibrium strategy must have all agents with signals $x \leq x_0$ rioting and those with signals $x > x^0$ not rioting, because, for agents with such signals, rioting and not rioting are optimal strategies regardless of what everyone else does.

In the subsequent rounds of our iterated deletion of dominated strategies, we take as given the restriction on dominated strategies obtained from the previous round. That is, any individual contemplating the actions of others must believe that everyone who has signals $x \leq x_0$ will riot and no one who has signals $x > x^0$ will riot. If everyone who has signals $x \leq x_0$ riots, then the fraction of the crowd that riots in state θ must be at least $\text{Prob}(x \leq x_0|\theta)$. Given our assumptions on this c.d.f., this fraction of rioters is always positive and is a continuous and declining function of θ . Thus, there is a maximum value of the state, which I denote $\theta_0 > \underline{\theta}$, such that

$$\text{Prob}(x \leq x_0|\theta) \geq a(\theta).$$

Accordingly, a rational individual must realize that at least in all states of nature $\theta \leq \theta_0$, enough of the crowd will riot to overwhelm the police, and such an individual thus finds it a dominant strategy to riot as long as his signal $x \leq x_1$, where x_1 is the largest signal x such that

$$\text{Prob}(\theta \leq \theta_0|x)W + \text{Prob}(\theta > \theta_0|x)L \geq 0.$$

Likewise, each agent realizes that at least a fraction $\text{Prob}(x > x^0|\theta)$ of the crowd will not riot in state θ , and thus the rioters must lose and be arrested in all states greater than or equal to θ^0 , where $\theta^0 < \bar{\theta}$ is the maximum value of θ such that

$$\text{Prob}(x \leq x^0|\theta) \geq a(\theta).$$

Accordingly, it is a dominant strategy for a rational agent not to riot when his signal exceeds x^1 , where x^1 is the smallest x such that

$$\text{Prob}(\theta \leq \theta^0|x)W + \text{Prob}(\theta > \theta^0|x)L \leq 0.$$

With these observations we iteratively delete dominated strategies: given x_0 and x^0 as threshold signals below which everyone riots and above which no one riots, we have shown that any equilibrium strategy must have the crowd winning at least in states $\theta \leq \theta_0$ and losing at least in states $\theta > \theta^0$, and thus rational agents should riot when their signals $x \leq x_1$ and not riot when their signals $x > x^1$. These new threshold signals x_1 and x^1 then take the place of x_0 and x^0 as restrictions on the behavior of every other agent, and we go through these calculations again, deriving new restrictions on the equilibrium strategies.

This iterative procedure of restricting the equilibrium strategies defines increasing sequences $\{x_n, \theta_n\}_{n=0}^\infty$ and decreasing sequences $\{x^n, \theta^n\}_{n=0}^\infty$ that progressively put tighter and tighter bounds on the equilibrium strategies. To finish the proof of Morris and Shin's proposition, we need only show that these sequences converge to common limit points, which I will denote x^* and θ^* . Showing this proves the proposition because it forces the conclusion that all agents with signals $x \leq x^*$ riot, while no agents with signals $x > x^*$ riot, and that the crowd wins the riot in all states $\theta \leq \theta^*$, and loses in all states $\theta > \theta^*$.

To show that the sequences above have common limit points, we observe that any limit points x^* and θ^* of either of these two sequences must be a solution to the two equations (1) and (2). But, if these two equations have a unique solution, then we are done, since that forces the conclusion that these two sequences have a common limit point.

The algebra behind Morris and Shin's result that equations (1) and (2) have a unique solution when the signals x are precise relative to the prior is straightforward. To do the algebra under the assumption of normality, observe that the term

$$\text{Prob}(x \leq x^* | \theta^*) = \Phi\left(\sqrt{\beta}(x^* - \theta^*)\right),$$

where Φ is a standard normal c.d.f., and use the fact that an agent who sees signal x has a posterior over θ that is normal with mean $(\alpha m_\theta + \beta x) / (\alpha + \beta)$ and precision $\alpha + \beta$, to get that

$$\text{Prob}(\theta \leq \theta^* | x^*) = \Phi\left(\sqrt{\alpha + \beta}\left(\theta^* - \frac{\alpha m_\theta + \beta x^*}{\alpha + \beta}\right)\right).$$

Use equation (2) to get

$$x^* = \frac{\alpha + \beta}{\beta} \theta^* - \frac{\alpha}{\beta} m_\theta - \frac{\sqrt{\alpha + \beta}}{\beta} \Phi^{-1}\left(\frac{-L}{W - L}\right), \tag{3}$$

and plug this into (1) to get one equation in the threshold state θ^* :

$$\Phi\left(\frac{\alpha}{\sqrt{\beta}}(\theta^* - m_\theta) - \frac{\sqrt{\alpha+\beta}}{\sqrt{\beta}}\Phi^{-1}\left(\frac{-L}{W-L}\right)\right) = a(\theta^*). \quad (4)$$

Equation (3) gives us the threshold signal x^* at which an agent is indifferent between rioting and not rioting given threshold state θ^* , and the left-hand side of equation (4) gives us the fraction of the crowd who receives signals less than or equal to x^* . Any solutions to equation (4) must lie in the interval $[\underline{\theta}, \bar{\theta}]$. Both sides of this equation are increasing functions: the left-hand side looks like a normal c.d.f. with steepness determined by $\alpha/\sqrt{\beta}$, and the right-hand side has whatever slope is assumed to reflect how the strength of the police varies with the state. The fact that there is at most one solution when β is large relative to α follows from the fact that the left-hand side becomes flat in θ over the interval $[\underline{\theta}, \bar{\theta}]$ in the limit as $\alpha/\sqrt{\beta}$ goes to zero. Note that if $\alpha/\sqrt{\beta}$ is large, then this equation typically has three solutions in the interval $[\underline{\theta}, \bar{\theta}]$ (since the c.d.f. looks more like an S over this interval), and thus the iterated deletion of dominated strategies does not pin down a unique equilibrium.

4. *The Selected Equilibrium and Comparative Statics*

Consider now what the unique equilibrium outcome looks like as a function of the state of nature θ . Note first that, whatever θ is, some portion of the crowd will riot and some portion of the crowd will not. All that varies with the state θ is the size of the fraction of the crowd that riots and whether the rioters overwhelm the police or are arrested.

The fraction of the crowd that riots in state θ is $\text{Prob}(x \leq x^* | \theta) = \Phi(\sqrt{\beta}(x^* - \theta))$. This fraction, as a function of the state θ , is one minus a normal c.d.f. and thus looks like a reverse S-curve. If the noise ϵ has a small variance, then this fraction begins to look like a step function: close to 1 for $\theta < x^*$ and close to 0 for $\theta > x^*$, with a steep transition from high to low as θ crosses the threshold signal x^* . Thus, the equilibrium relationship between the actions of the crowd and the strength of the police is highly nonlinear. For large ranges of values of the state θ , we have changes in the strength of the police, $a(\theta)$, but little or no change in the fraction of the crowd that riots. On the other hand, for values of θ close to the threshold signal x^* , we have a large and sudden change in the number of people rioting. This is a very interesting result, since it suggests that sudden shifts in agents' expectations with small changes in the state may play an important role in determining equilibrium outcomes despite the fact that the equilibrium is unique.

Now let us go through an example of how to use this technology to do comparative statics. The natural exercise in this example is to ask what effect changes in the average strength of the police (parametrized by m_θ) have on the equilibrium incidence of riots, computed as the ex ante probability that the state θ is below the threshold θ^* , or $\Phi(\sqrt{\alpha}(\theta^* - m_\theta))$. Differentiating equation (4) gives us the result that as long as the left-hand side of (4) is flatter in θ than the right-hand side (the same condition that ensures uniqueness of the solution), then $d\theta^*/dm_\theta < 0$. What this implies, of course, is that strengthening the police has two beneficial effects: first, it lowers the probability that the crowd will win the riot, holding fixed the threshold state θ^* , and second, it leads to a reduction in the threshold state θ^* , further reducing the probability that the crowd will overwhelm the police. Morris and Shin play up this shift in the threshold state in their application of this technology to the pricing of corporate debt. Note, of course, that this second effect, this shifting of the threshold state, is smaller, the larger is β relative to α . In the limit as $\alpha/\sqrt{\beta}$ goes to zero, this second effect disappears.

5. *The Problem with Introducing Prices into the Model*

So far in our analysis, individuals in this crowd have no information other than their own signal to consider when they decide whether to riot or not. This would be different, of course, if we introduced markets and prices into the model. Imagine, for the sake of this discussion, that individuals also could see asset prices, and assume specifically that there is a traded asset with payout contingent on the claims that the insurance company that covers the property threatened by the rioters must pay. For simplicity, assume that the claims that the insurance company would have to pay following a riot take on only two values: a large value in the event that the crowd overwhelms the police, and a small value in the event that the police keep the crowd under control. Imagine, as well, that assets trade continuously, so that individuals in the crowd can see asset prices after θ is realized but before they need to decide whether to riot.

On the one hand, if this asset ends up being priced in equilibrium in a way that accurately reflects its subsequent payout, it will have one price in all states $\theta \leq \theta^*$ (reflecting that the insurance claims will be large) and another price in all states $\theta > \theta^*$ (reflecting that the insurance claims will be 0). This, of course, will be a problem for our previous analysis. Every individual should be able to look at this asset price and know whether the crowd is going to overwhelm the police or not. Depending on the price, then, either every individual should strictly prefer to riot, or to not

riot. Agents' actions and beliefs would be coordinated, since there would be no reason for any individual to act differently on the basis of his own signal. The logic of Morris and Shin's argument goes out the window.

On the other hand, if this asset does not get priced in equilibrium in a way that allows agents to infer whether the crowd will overwhelm the police or not, we must ask why it is not priced that way. How do we set up the model so that the asset price does not aggregate the information that all of the individuals in the economy have and thus reveal the true state?

The idea that individuals in a crowd considering whether to riot or not would consult asset prices via the newspaper or their handy wireless Internet connections seems farfetched. That, in part, was my motivation for picking this example for my discussion. The analysis of Morris and Shin seems as if it might work pretty well for this example. In the macroeconomic examples that Morris and Shin point to in their paper, however, asset prices are clearly a necessary part of the picture, and it is not at all clear how their arguments apply.

In Morris and Shin's example regarding speculative attacks on currencies, one would think that forward exchange rates (interest-rate differentials) and options on exchange rates are readily observed by all market participants when they consider whether to attack or not. Their example in their earlier paper (Morris and Shin, 1988), like the riot example above, has agents holding diverse beliefs about the probability that the currency will be devalued and deciding whether or not to attack on the basis of those beliefs. But, if, given the fundamentals today, the equilibrium uniquely pins down whether the currency will soon be devalued or not, then it seems that those interest-rate differentials and exchange-rate option prices should reflect today which of the two outcomes will occur. If those prices do accurately reflect which outcome will occur, agents should coordinate their decision to attack or not according to them: everyone should attack if the asset prices indicate a devaluation will occur, and no one should attack if they indicate that a devaluation will not occur. It does not make sense in this application to assume that agents will take different actions (attacking or not) on the basis of their private signals if publicly observed asset prices accurately reveal which outcome will actually occur. It thus does not make sense to apply the argument proposed by Morris and Shin to the analysis of currency attacks unless we can tell some story as to why interest-rate differentials and exchange-rate options do not reveal an imminent devaluation of the currency even if that devaluation must occur in equilibrium with probability one.

In Morris and Shin's example regarding corporate debt, discussed in detail in a cited working paper (Morris and Shin, 1999), the price of the firm's equity and the secondary market price of the firm's debt will clearly reflect some market assessment of the likelihood that the firm will be liquidated in equilibrium. If the outcome, liquidation or not, is uniquely pinned down by the fundamentals, then these prices should reveal that, and agents should be able to coordinate their actions accordingly.

Finally, in the bank-run example presented in this paper, the price of the bank's equity should reveal whether there will be a run or not, since this outcome is pinned down in equilibrium. Accordingly, agents should look at this price in deciding whether to run or not, and it seems natural to suspect that their actions and beliefs might be coordinated upon the observation of this price.

The question then stands, how do we integrate prices into the analysis and yet preserve the diversity of posterior beliefs across agents that is key to pinning down a unique equilibrium? Perhaps the answer to this question will depend on the specific application: it seems plausible that rioters are not integrating asset prices into their analysis of whether to riot or not; it seems less plausible to assume that currency traders are ignoring interest-rate differentials and option prices in deciding whether to attack a currency or not. Finding an answer to this question seems to me to be the obvious next step in refining this potentially useful technology for analyzing macroeconomic coordination games.

*Comment*¹

HÉLÈNE REY

London School of Economics and CEPR

1. Introduction

It is a real pleasure to comment on a paper which is of great interest, addresses a fundamental issue in macroeconomics, and is also very elegant.

In a series of articles, Steve Morris and Hyun Song Shin have developed a fruitful line of research that extends and applies sophisticated game-theoretic concepts to traditional macroeconomic problems. In this paper, which may be seen to some degree as a synthesis of their approach, they use a simple bank-run model as a framework to ask a very

1. These comments benefited from discussions with Harald Hau, Thomas Philippon, Richard Portes, David Romer, and Mike Woodford.

important question: Are multiple equilibria in economics the unintended consequence of too simplistic assumptions?

The answer provided by the paper is unambiguously yes. The authors write, for example: "We doubt that economic agents' beliefs are as indeterminate as implied by the multiple-equilibrium models. Instead, the apparent indeterminacy of beliefs 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; and second, economic agents are assumed to be certain about each other's behavior in equilibrium." The paper then claims that introducing a small amount of idiosyncratic uncertainty is enough to destroy the perfect coordination of agents' actions and beliefs and therefore to eliminate the possibility of multiple equilibria. Since our world seems indeed to be one of imperfect and asymmetric information, this realistic generalization of our traditional macroeconomic models appears to banish multiple equilibria once and for all. They become an "artifact of simplifying assumptions that deliver more than they are intended to deliver," as the authors put it.

In my discussion, I will emphasize that Morris and Shin's paper does not in fact eliminate the possibility of multiple equilibria. I will also discuss the robustness of their results more precisely and perform some comparative-statics exercises. Finally, I will comment on the empirical applicability of their model and its relations to the literature on multiple equilibria.

2. *Unique Equilibrium?*

Morris and Shin set up a Diamond–Dybvig bank-run model with a slightly more sophisticated information structure than usual. Returns follow a normal distribution with a given precision α ; this is *public information*. On the other hand, each agent gets a signal with precision β regarding the realization of the return; this is *private information*. When the fundamentals are common knowledge, it is well known that the Diamond–Dybvig model gives rise to multiple equilibria. By introducing a little bit of noise (a very small degree of asymmetric information), the authors show that the equilibrium is unique. So a very minor modification to an otherwise standard model is able to eliminate the multiplicity of equilibria.

This is a very strong result. I will argue, however, that the minor deviation from the benchmark model chosen by the authors brings with it a lot of interesting and sometimes puzzling results, some of them not emphasized in the paper. In particular, if one does not look exclusively at

the limiting case on which the authors are focusing but at the general case of their own model, the possibility for multiple equilibria reappears very naturally.

In Morris and Shin’s paper, the condition characterizing the equilibrium is

$$\rho^* = \Phi(\sqrt{\gamma}(\rho^* - \bar{r})),$$

where ρ^* is the cutoff point below which patient consumers withdraw their money from the bank, \bar{r} is the mean of the returns, $\Phi(\cdot)$ is the cumulative normal distribution, and γ is a constant given below. Graphically, this equilibrium is illustrated in Figure 1 of the paper. It is immediately apparent that the 45° line and the cumulative normal distribution will intersect only once if the slope of the cumulative normal is “not too steep.” Formally, a sufficient condition for this to happen is

$$\gamma = \frac{\alpha^2(\alpha + \beta)}{\beta(\alpha + 2\beta)} \leq 2\pi.$$

When the precision of the private information, β , is very high (β goes to infinity for a given α , meaning that γ becomes very small), the authors interpret their model as being a very small deviation from the standard Diamond–Dybvig model with common knowledge. In that case the Morris–Shin model gives the discontinuity result emphasized in the paper: If private information is very precise, then the two curves intersect only once and we have a unique equilibrium. If, on the other hand, private information is infinitely precise, then we are in the standard Diamond–Dybvig case and there are multiple equilibria. This is an interesting and surprising result, and the authors present it very well in the paper.

But this is not the end of the story. Note that there are two different ways to approach the common knowledge case from within the Morris–Shin framework (Figure 1). We can approach common knowledge either by letting the precision of the *private* signal go to infinity, as in the paper, or by letting the precision of the *public* information go to infinity. In that latter case, α would be going to infinity for a given β and the slope of the cumulative normal distribution would become very steep as in Figure 2. In that case, there can be multiple equilibria. More generally, it is obvious that as long as the precision of the public information is high compared to that of the private signal, multiple equilibria will still exist. This result is intuitive: the more precise public information is, the closer we

Figure 1 COMMON KNOWLEDGE AS A LIMIT

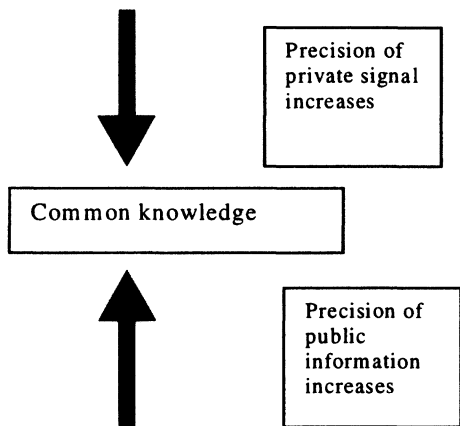



Figure 2 MULTIPLE EQUILIBRIA IN THE MORRIS–SHIN FRAMEWORK



are to the standard case of common knowledge among economic agents, which is known to generate multiple equilibria. At another level, it is however somehow paradoxical to think that the economies that are generating the more accurate publicly available information are also the ones that are the more prone to multiple equilibria. And, conversely, it is also

puzzling that for a given degree of precision of the private signals, economies with very diffuse public information will converge to a unique equilibrium.

To summarize, the central claim of the paper (the discontinuity result), that a very minor deviation from the standard models with common knowledge is enough to eliminate the possibility of multiple equilibria, is not the whole story: if one does accept that the Morris–Shin framework is a better representation of reality, one has to recognize that this model also delivers multiple equilibria for some parameter domains. Furthermore, common knowledge can be seen as a limiting case in two different ways: in one case, one converges towards common knowledge with a unique equilibrium, in the other case, one converges towards common knowledge with multiple equilibria.

3. *Comparative-Statics Results and Dynamics*

If we limit ourselves to the parameter region where uniqueness of equilibrium prevails, we can perform comparative-statics exercises, which pave the way towards policy recommendations. A first thing to look at is the effect of the precision of private and public information on the cutoff signal x^* below which patient depositors will withdraw their money from the bank. It turns out that an increase in the precision of either type of information may either lower or raise the value of the critical signal for a given return. This result is puzzling.

Another interesting exercise is to look at the impact of a change in public information versus the impact of a change in private information. One can even characterize by how much a private signal should change to balance the impact of a change in public information so that the strategies of the agents are kept unchanged. Since one of the key aspects of public information in Morris and Shin's paper is that it coordinates the expectations of agents, one would expect that public news would have a bigger relative effect than private news. This intuition is correct provided one is able to control for the relative precision of the private and public informations, which requires knowing the magnitudes of α and β .

The model presented is a one-shot game (a repetition of one-shot games in the second part of the paper). It would obviously be very nice to do a dynamic extension of the framework. Careful thought should then be given to the process of information revelation. Let us imagine that economic agents play the game presented in the paper at date t . At date $t + 1$, they will have observed the number of people having run on the bank at date t , which is given by

$$\ell(r) = \Phi \left(\sqrt{\beta} \left((\rho^* - r + \frac{\alpha}{\beta} (\rho^* - \bar{r})) \right) \right).$$

As soon as the proportion of people withdrawing money is observed, the realized return becomes common knowledge, since all the other parameters are known. If there is some persistence in the return variable, then the precision of the public information is increasing over time (α is increasing for a given β), and we may exit the unique equilibrium region. Extending the model dynamically therefore requires keeping enough “fuzziness” in the public information.

4. Empirical Applicability

An interesting feature of Morris and Shin’s approach is the ability to address policy issues, thanks to the comparative-statics exercises performed around the unique equilibrium. For practical purposes, it is therefore very important to know whether the economy is in a unique- or a multiple-equilibrium region, which depends on the value of the parameter γ . Since γ is not homogeneous in α and β , figuring out the relative precision of the two types of information is not enough. It matters whether α is 17 rather than 13 or whether β is 9 rather than 10. Moreover, just like the number of equilibria themselves, we have seen that the comparative-statics results depend on the absolute magnitude of the precision of the public and the private information. Unfortunately, it seems extremely difficult to get an idea of what these numbers are in reality. They partly depend on the interpretation one has of the model itself. Should we think of the private-information element of the model as differences in psychology across individuals, so that traders reading the same economic news may form different views on the economy depending on their temperament? Or should we think of it as the degree of precision of “inside information”?

This aspect put aside, we should ask ourselves whether Morris and Shin’s approach has empirical implications which can clearly be distinguished from those of the models exhibiting multiple equilibria. The authors argue that their model provides testable implications so that it suggests a correlation between fundamentals and outcomes, unlike multiple-equilibrium models, where the shift from one equilibrium to the other may be due to pure sunspots. This point is interesting. Note, however, that multiple-equilibrium models also provide some correlations between fundamentals and outcomes. In a self-fulfilling specula-

tive-attack model, for example, the parameter space is divided into three regions: one where the fundamentals are so good that there can be no attack, one where the fundamentals are so bad that there is an attack for sure, and an intermediate region where there are multiple equilibria.

Therefore I would argue that the key empirical implication of the Morris–Shin model is not that fundamentals are correlated with outcomes, nor that multiple equilibria do not exist—as discussed above—but rather that the degree of information aggregation matters. Having recognized this fact, there are nice natural experiments which could be used to test the model. One could for example look at the role of polls or surveys in a situation with strategic complementarities (like foreign-exchange traders). One could also study the impact of the introduction of a futures market on the evolution of spot prices, the idea being that prices of futures would aggregate the information of market participants. The difficulty of putting numbers on the precisions of public and private information and therefore of pinning down the exact implications of the model—which vary across parameter regions—will, however, remain.

5. *Interpretation of Multiple Equilibria*

The main message of multiple-equilibrium models may be that even when the fundamentals of the economy are almost the same, outcomes can be very different. The sense of this basic message seems empirically quite relevant. The ERM crisis of 1992, for example, has often been given as an example of self-fulfilling speculative attack. By fundamentals we usually mean all the variables describing the economy (GDP, prices, exchange rates, etc.) except the information structure. A great virtue of the Morris–Shin model is that it introduces the information structure into the set of the fundamental variables. The question is then whether the model can deliver the flavor of the multiple-equilibrium models while keeping the uniqueness of the equilibrium.

The paper shows that for some parameter values, introducing some noise makes the equilibrium unique. In that uniqueness region, small changes in the information structure do change the threshold value below which an attack occurs, but not dramatically so (in general). In Figure 1 of the paper, for example, one can see that a small change in the information structure (or in the mean return) will change the slope of the normal distribution (or shift it, respectively). But this will not result in a big variation in ρ^* , the posterior belief below which the patient consumers withdraw their money, unless the slope of the normal distribution is

quite steep, which is exactly the case when one is close to the region where multiple equilibria exist. In other words, the Morris–Shin model can have the flavor of multiple-equilibrium models, but that is only provided one is in a parameter region away from the limit case considered by the authors and close to the multiplicity domain.

In the absence of even more sophisticated ways to model information aggregation and the endogeneity of the information structure, we are still left with a multiple-equilibrium region where we cannot say much about equilibrium selection. Perhaps a phenomenon like the 1992 ERM crisis could be modeled as unique equilibrium if dramatic shifts in information aggregation were incorporated explicitly. One way forward could be to think harder about the information aggregation process: here private information is costless to acquire and is automatically given to all agents. Costly and voluntary information acquisition should ideally be related to the other fundamentals of the economy.

6. *Conclusion*

The paper makes a very important contribution to the literature on strategic complementarities. First, Morris and Shin's approach can be applied to a wide spectrum of issues. We have many macroeconomic models which exhibit multiple equilibria, whether they are used to discuss bank runs, speculative attacks, industrialization, inflation, poverty traps, or thick-market externalities. As the authors point out, this multiplicity is a problem if one wants to perform comparative-static exercises. What is the impact of increasing a tax rate, for example, when the system can switch from one equilibrium to another in a random fashion? Determining the equilibrium to which an economic system will converge is a key issue for policy makers, and this is where the Morris and Shin's approach is so valuable. But as I have pointed out, the Morris–Shin model is not as opposed to the multiple-equilibrium literature as the authors claim. This is not a criticism, and it underlines that the model has many interesting and rich features, which can be exploited further. The model is not very operational yet as far as empirical tests are concerned, mainly because it is hard to pin down the magnitude of the key parameters which determine the domain of existence of equilibria as well as the comparative-static results. It also lacks true dynamics. The Morris–Shin framework has however already been (rightly) very influential in the way we think about coordination and information aggregation in macroeconomic models and will certainly generate a lot of interesting new results in very diverse areas.

Discussion

In responding to the discussants' comments, Stephen Morris agreed that there are two ways of proving the existence of a unique equilibrium in their framework. The first, employed in their paper, is the direct approach of showing that no equilibrium other than the one they find can exist. The other approach is discussant Andrew Atkeson's iterated elimination of dominated strategies. Morris noted that Atkeson's approach makes the equilibrium a little less mysterious but may give the impression that agents in the model need to do sophisticated reasoning, which is not the case, as no other threshold will work as an equilibrium. Morris agreed with the discussant H el ene Rey that the implications of alternative assumptions on the information structure, especially the distinction between public and private information, are the most important area of future research. On the importance of public information, he attributed to Robert Shiller the claim that "bubbles started when newspapers became widely available." The importance of financial news networks of all kinds is not only that an individual receives information, but that he knows that others are also receiving that information.

Daron Acemoglu began the general discussion by suggesting that one way of interpreting multiple-equilibrium results is that they arise from sparse formal models with limited fundamentals. The fact that models with few fundamentals imply multiple equilibria does not imply necessarily that multiple equilibria are a feature of the real world. The virtue of this paper is that it suggests that minimal increases in the complexity of our models may reduce or eliminate multiple equilibria. Acemoglu also offered an intuition about why the model works: In discrete choice models there are multiple equilibria because when everybody else does something the return is sufficiently high, and when everybody else doesn't do something the return is sufficiently low, so that "following the crowd" is a winning strategy for everyone. Noise creates a thick tail of people who will always do or not do something independently of others' actions, leading ultimately to a unique equilibrium. This reasoning suggests why we want the parameter β to be large but not α , because when α is large, tails are not thick enough relative to common information. The authors agreed with this intuition.

Mike Woodford argued against making the inference that any finding of multiple equilibria depends on extreme assumptions. He said the paper does not show that multiple-equilibrium models are not robust; it only shows that one can construct examples with private signals where there is a unique equilibrium. Woodford suggested that minor perturba-

tions of private-information models will yield multiple equilibria that are qualitatively similar to those in common-knowledge models; hence Morris and Shin's analysis does not demonstrate that the conclusions of models without private signals are not robust. Woodford also thought the authors overstressed the finding that, in their model, a small change in the public signal can have a large effect on the equilibrium outcome, even though the equilibrium is unique. He observed that this result requires the parameter γ to be large, which means that the model is very close to one with multiple equilibria, implying in turn that the unique equilibrium is very fragile. Fundamentally, then, the mechanisms that generate big swings are similar in this model and in models exhibiting multiple equilibria.

The authors agreed with Woodford that there are many ways to perturb the common-knowledge model that preserve multiple equilibria. However, Shin argued that their model is special (and not a small perturbation) in its result that people remain uncertain in equilibrium about the actions of others, so that common knowledge is destroyed. Shin suggested that their model should be viewed as a cousin of the "second-generation" models, which preserves the flavor of those models but includes a rigorous equilibrium argument that pins down the model's prediction. Morris added that finding out under what information structures equilibrium is unique is a worthwhile project in itself. He noted that models with strategic complementarities will already have a strong multiplier effect, so that interpreting a crisis as a switch from one equilibrium to another amounts to throwing in an extra, and perhaps unnecessary, strategic complementarity. Their approach eliminates this extra degree of freedom.

Pierre Gourinchas suggested that in multiperiod models there might be a feedback process: agents may observe fundamentals and get very precise information. Asset prices that efficiently aggregate information may have the same effect. If precise information on fundamentals becomes publicly available, the results of the paper are undermined. Morris agreed there would be progressive information revelation through prices in a dynamic model; he thought a useful direction would be to write down a more complicated model in which this happens, but in which people also receive private information over time.

Paolo Pesenti was enthusiastic about the approach, arguing that it represents the best chance so far to build the foundations of a theory of confidence crises. Such a theory would have far-reaching implications about how we think about financial stability and other important issues. He hoped that this approach would eliminate the current dichotomy in the

literature between fundamentals-based and non-fundamentals-based models of crises.

Ken Rogoff recalled that, in his Ely lecture, Larry Summers complained about models of crises that have the implication that, in certain regions, economics can say nothing. Summers called for a model with the feature that the worse the fundamentals are, the more likely a crisis. Rogoff noted that, at a formal level, the Morris–Shin model has this desirable property. But, he asked, suppose we establish at a theoretical level that equilibrium is unique but can't connect it to any fundamental that we can plausibly observe. Then would our empirical approach be in any way different, relative to the case in which we are guided by models with multiple equilibria?

