Predicting Self-Fulfilling Financial Crises*

Christopher Gandrud† Thomas B. Pepinsky‡

FIRST DRAFT: July 3, 2015
THIS DRAFT: October 23, 2016

Abstract

This paper studies how self-fulfilling dynamics affect the predictability of financial crises. We build a model in which market participants play an investment coordination game with common economic shocks and private information about their own willingness to cooperate. An observer attempts to predict the occurrence of a financial crisis in the future based on the players' actions in the present. Under favorable conditions, good market conditions prevent an observer from learning about players' types by observing their actions. This induces a negative correlation between economic conditions in the present and the variance of players' types in the following period, which defines the predictability of financial crises. Under more extreme positive market conditions, the observer cannot use present observations of the players to learn about the probability of a financial crisis in the future. We test the implications of the model using a new continuous measure of financial market stress–"FinStress"–developed by Gandrud and Hallerberg (2015). We find support for a key implication of our theory: the *variance* of financial market stress is larger following periods of good economic conditions than following poor economic conditions. These conclusions have implications for both empirical analyses of the predictors of financial crises and for policymakers seeking to prevent crises. The findings in our paper suggest that regulators should focus on actors' types (e.g. with stress tests) more than macro-economic conditions. More broadly, our analysis provides a new explanation for why social scientists and area specialists are generally poor at predicting events that require mass coordination, such as financial crises, coups, and revolutions.

*Thanks to Andrew Little for patient guidance and enormously helpful discussions, and to David Singer and James Morrison for critical feedback on earlier drafts. In addition, we thank seminar participants at City, University of London, the University of Mannheim, and the 2015 American Political Science Association Annual Convention. Thanks also to the Deutsche Forschungsgemeinschaft for financially supporting Christopher's research (No. HA5996/2-1). The formal model was developed by Thomas Pepinsky. The empirical tests of the model were conducted by Christopher Gandrud. We are jointly responsible for all errors.

†Lecturer, Department of International Politics, City, University of London and Research Fellow, Hertie School of Governance, `christopher.gandrud@city.ac.uk`

‡Associate Professor, Department of Government, Cornell University, Ithaca, NY 14853, `pepinsky@cornell.edu`.

1 Introduction

Understanding financial crises lies at the core of modern economic policymaking. The recent financial crisis and the ensuing policy response has demonstrated that crisis politics is one area in which academic social science has extorted a major influence in policymaking (see e.g. Chinn and Frieden 2011; Drezner 2014; Mian and Sufi 2014; Reinhart and Rogoff 2011). Indeed, a robust empirical literature has emerged in recent decades that predicts financial crises (Berg and Patillo 1999; Danielsson, Valenzuela, and Zer 2015; Eichengreen, Rose, and Wyplosz 1996; Frankel and Saravelos 2012; Frankel and Rose 1996; Leblang and Satyanath 2006; Minsky 1982). Yet a predominant strand of theoretical literature on financial crises holds that they are often self-fulfilling: a financial crisis occurs when market actors believe that a financial crisis will occur (for the classic statement, see Obstfeld 1986). Despite the widespread acceptance of the existence of self-fulfilling financial crises, there is little analysis of the implications of self-fulfilling dynamics for the problem of prediction.

In this paper we study how self-fulfilling dynamics affect our ability to predict financial crises. We build a model of financial policymaking in which market participants play a coordination game with private information about their willingness to hold an asset, and an external analyst or regulator attempts to predict the occurrence of a financial crisis in future from the observed behavior of the players today. Each market participant generally wishes to hold an asset so long as her counterpart holds the asset as well, and to sell the asset if her counterpart sells. The occurrence of a financial crisis—here, an outcome where both participants sell—depends on both private information held by each participant, participants' beliefs about the types of other participants, and common economic conditions.

We show that under favorable conditions, good market conditions reduce an observer's ability to learn about participants' types by observing that they hold the asset. This induces a negative correlation between economic conditions in the present and the variance of participants' types in the following period. Because crises occur in the model as a function of both market conditions and participant strategies, a higher variance of participants' types reduces the predictability of financial crises. Under more extreme market conditions, equilibrium play is uninformative: an observer cannot use participants' strategies in the present to learn about the probability of a financial crisis in the future.

The intuition behind this result is that favorable economic conditions induce cooperation even among "weak" participants by lowering the costs of an opponent's defection. As a consequence, an analyst or regulator who observes good financial market conditions and no financial crisis in the present cannot infer if participants are strong or weak. By contrast, when that analyst observes cooperation during unfavorable economic conditions, he can infer that participants' types are sufficiently strong to induce cooperation. We term this the *screening effect*. Market participants too, moreover, will update their beliefs about their opponents' strength in an analogous fashion, with the consequence that they will be more likely to cooperate than they previously were. We term this the *confidence effect*. The net result of the screening and confidence effects is that in all non-crisis periods, weaker economic conditions in the present produce higher levels of cooperation in the future, while also revealing more information about participants' types to the observer.

We test this novel theoretical prediction by exploiting a unique new dataset on financial market stress–"FinStress" from Gandrud and Hallerberg 2015–that allows us to calculate the variance of financial market conditions in any given year. We show that, consistent with our theory, we observe a positive correlation between market conditions in period t and the variance of financial market conditions in the following period $t+1$. This new empirical result illustrates the power of our theoretical approach for understanding financial market volatility and the relationship between financial market conditions and future crises.

Our theoretical approach builds on the recent literature on global games (Carlsson and Damme 1993; Morris and Shin 2003) and related models of cooperation with private uncertainty (Baliga and Sjöström 2011, 2012). The most attractive feature of this approach is that cooperation games such as the one that we study have a unique equilibrium under mild conditions. This is critical for our argument because it rules out—by assumption—equilibrium multiplicity as a fundamental determinant of the predictability of crises. When we relax a key assumption required for equilibrium uniqueness, allowing for the more extreme market conditions mentioned above, we obtain the starker result that equilibrium play is uninformative. We argue that the distinction between the predictability of crises based on inferences about the distribution of player types, and epistemic uncertainty associated with multiple equilibria in coordination games, can be linked to Knight's (1921) distinction between risk and uncertainty. A key benefit of our modeling approach is that both emerge from a single framework.

These theoretical findings have clear implications for both policy and practice. For empirical predictions of financial crises, the main result is that "good" observable economic conditions should not predict a *lower probability* of a crisis in the future, but instead should predict a *higher variance* in the probability of a crisis in the future. This observation may help to explain why empirical findings about leading indicators of financial crises are often inconsistent, for in a regression of fundamentals on the probability of a future criss, even the sign is indeterminate.

Our analysis also provides a distinct counterpoint to the widespread condemnation of social scientists—primarily economists, but also those working in allied disciplines—for failing to predict the Global Economic Crisis. It is clear that cognitive shortcuts, social conventions, disciplinary failures, and many other pathologies of human decisionmaking led both academics and financial market actors to ignore (or at least underestimate) the vulnerability of the global financial system to financial crisis in the early 2000s (see among many others Colander et al. 2009; Helleiner 2011; Nelson and Katzenstein 2014). Our approach, however, uncovers a more discouraging feature of financial markets. Even a clear-headed observer subject to none of the biases attributed by critics to the economics profession, and completely disinterested in market outcomes, should be unable to predict financial crises during the very periods in which such a prediction would be most useful. This conclusion parallels Kuran's (1989; 1991) analysis of the failure of social scientists to predict the events of 1989 in Eastern Europe, although it rests on very different theoretical mechanisms. Our analysis indicates that improving predictions requires better knowledge of either the path of future common economic shocks, or greater knowledge about individual market actors.

2 A Simple Model of Coordination and Crisis

To model coordination dynamics in financial markets, we rely on the stag hunt as a heuristic. The stag hunt is a coordination game in which players prefer to cooperate on one choice (described as hunting together for a stag) over another choice (hunting together for a hare), but strictly prefer both to not cooperating at all. It is a game of "strategic complements" (Baliga and Sjöström 2012), meaning that each participant wants to choose the same action that her counterpart chooses, even though they mutually prefer hunting a stag to hunting a hare. It has recently garnered attention as a particularly useful way to conceptualize bank runs (McAdams 2009): both participants wish to keep their deposits in a bank, but if either believes that the other will withdraw, she wishes to withdraw as well. Many types of financial crises share similar features. In the exposition below we will ignore the actual value of the asset (a deposit, a currency, or any type of security) besides to assume that it is valuable enough that both players would prefer to hold it in the event that they were certain that the other player would hold it. This focuses the analysis exclusively on self-fulfilling dynamics, holding aside changes in players' behavior that emerge from the objective loss of a financial asset's value.

The stag hunt heuristic is also useful because it draws on the same modeling framework used in the literature on self-fulfilling currency crises (see Obstfeld 1996). It captures the basic insight that if a currency is believed to be solid because a government will defend it, traders would prefer to hold it. However, if traders know that others will sell, then they prefer to do the same. We do not explicitly model the decision of a policymaker to defend or abandon a currency in this paper, but do incorporate fluctuations in market conditions that may be interpreted as either objective measures of market conditions or common subjective beliefs about value of cooperation.

The model that we outline below relies on a number of simplifying assumptions that stack the deck in favor of crises being predictable. Among other assumptions, all information about the game except for players' individual payoffs is publicly known to all players and to the observer, there is no communication between players, the game has only two stages, and the model produces a unique Bayesian Nash equilibrium. As a result, we will not address some of the main concerns found in the theoretical literature on coordination games, including communication and reputation, and we will only address multiple equilibria as an extension of the main results.

2.1 A Baseline Coordination Game

We begin with a full information game in which two market participants or "Players," A and B, simultaneously decide whether to hold or sell an asset. The players' payoffs are given in Table 1, with the payoff to each player from both holding normalized to zero as a reference point.

Table 1: A Coordination Game

		Player B	
		Hold	*Sell*
Player A	*Hold*	0,0	-3,-1
	Sell	-1,-3	-2,-2

In this game, the two pure strategy Nash equilibria are $\{Hold, Hold\}$ and $\{Sell, Sell\}$.[1] The payoff-dominant equilibrium of $\{Hold, Hold\}$ represents a "normal" equilibrium, which

[1] There is also a mixed strategy equilibrium, but this is "unstable" (see Echinque and Edlin 2004).

the equilibrium $\{Sell, Sell\}$ is the "crisis equilibrium."

2.2 Coordination with Private Information

In the baseline case, all payoffs are public information. To analyze how an analyst or regulator would learn about the likelihood of a future crisis, we begin first by introducing private information into the model. Specifically, we introduce individual-specific private information σ_i about each individual's payoffs for selling the asset. The private information represents any increase or decrease in the present value to player i of selling the asset. Intuitively, these may be interpreted as players' need for liquidity (a large liquidity preference increases the value of selling the asset), or more generally, a measure of a player's preference for cooperation. Each player i observes her own type σ_i, but not σ_j. We also introduce a common shock κ. The common shock represents any increase or decrease in the present value of selling the asset that is shared among players. Intuitively, κ may be interpreted as "financial market conditions," which affect the liquidity demands of both players at the same time. The value of the common shock κ is known to both players.

The payoffs from this game appear in Table 2. To clarify the logic of the game, player i must decide whether to hold or to sell the asset based on her beliefs about the likelihood that player j will hold the asset. She knows that this depends on σ_j, and that player j makes that decision based on her beliefs about σ_i. Both players also observe the value of κ, and they also know that $\sigma_{i,j}$ are drawn from the same common, single-peaked distribution F_σ, with support on the interval of $[\underline{\sigma}, \overline{\sigma}]$. Because this is a coordination game in which each

Table 2: The Private Information Game

		Player B *Hold*	Player B *Sell*
Player A	*Hold*	κ, κ	$-\alpha$, $\beta - \sigma_B$
	Sell	$\beta - \sigma_A$, $-\alpha$	$-\sigma_A$,$-\sigma_B$

player wishes to Hold (Sell) when the other Holds (Sells), $\alpha, \beta > 0$.

The game in Table 2 is an extension of the coordination game in Baliga and Sjöström

(2012), and analysis will follow theirs very closely. We begin with two assumptions about the structure of payoffs and the shape of F_σ.

Assumption 1 (Strategic Complements) *The parameters of the Private Information Game satisfy the following constraints:*

a. $\beta < \alpha$

b. $\underline{\sigma} < \beta - \kappa < \alpha < \overline{\sigma}$

Assumption 2 (Sparsity of Types) *For all* $\sigma \in [\underline{\sigma}, \overline{\sigma}]$, $f(\sigma) < \frac{1}{\alpha+\kappa-\beta}$

Assumption 1 ensures that the payoff for i choosing Hold is strictly higher when player j chooses Hold (condition *a*), and that there is some positive probability that player i will Hold in all cases ($\sigma_i < \beta - \kappa$) or Sell in all cases ($\sigma_i > \alpha$) (condition *b*). Assumption 2 guarantees that the type space is sparse enough that the probability of encountering any type is sufficiently high. This assumption together with the previous one rules out the possibility of multiple equilibria (see Bueno de Mesquita 2014), a topic to which we will return below.

The game departs from Baliga and Sjöström (2012) in the "common shock" parameter κ. Notice that if $\kappa > \beta - \sigma_i$ and $\alpha > \sigma_i$, then player i has a dominant strategy of Hold regardless of j's action. Likewise, if $\kappa < \beta - \sigma_i$ and $\alpha < \sigma_i$, then player i has a dominant strategy of Sell. However, the restriction on $\beta - \kappa$ from Assumption 1(b) ensures that there is no situation in i prefers to sell *regardless of* i*'s type*.

A strategy s for player i maps the distribution of types to the *actions* H̲old and S̲ell: $s_i : [\underline{\sigma}, \overline{\sigma}] \to \{H, S\}$. i's expected utility if j plays H with probability p_j is

$$\begin{aligned} U_i(H) =& \kappa \cdot p_j + -\alpha \cdot (1 - p_j) \\ =& \kappa p_j - \alpha + \alpha p_j \end{aligned} \tag{1}$$

$$
\begin{aligned}
U_i(S) =&(\beta - \sigma_i) \cdot p_j + (-\sigma_i) \cdot (1 - p_j) \\
=&\beta p_j - \sigma_i
\end{aligned}
\tag{2}
$$

Player i chooses H over S iff (1) $>$ (2), or

$$
-\alpha + \sigma_i + (\alpha + \kappa - \beta)p_j > 0 \tag{3}
$$

We first consider all Bayesian Nash equilibria in which player i chooses H if $\sigma_i \geq \sigma^*$, where σ^* represents the threshold in $[\underline{\sigma}, \overline{\sigma}]$ that makes i indifferent between holding and selling if the other player uses the same strategy. Suppose that j plays H with probability $p_j = 1 - F(\sigma_j)$. Then i's best response function $\Gamma_i(\sigma_j)$ is

$$
\begin{aligned}
\Gamma_i(\sigma_j) =& -\alpha + \sigma + (\alpha + \kappa - \beta) \cdot (1 - F(\sigma_j)) \\
=&\beta - \kappa + (\alpha + \kappa - \beta)F(\sigma_j)
\end{aligned}
\tag{4}
$$

Because $\sigma_{i,j}$ are drawn from a common distribution, and all other parameters of the model are common knowledge, the best responses are symmetrical. Via Theorem 1 in Baliga and Sjöström (2012), the strategy $s_i(\sigma_i) = H$ if $\sigma_i \geq \sigma^*$ is the unique Bayesian Nash equilibrium.

Theorem 1 (Baliga and Sjöström 2012) *If Assumption 1 and Assumption 2 hold, there is a unique Bayesian Nash equilibrium to the Private Information Game in which* $s_i(\sigma_i) \to \{H\}$ *if* $\sigma_i \geq \sigma^*$*, characterized by* $\sigma^* = \beta - \kappa + (\alpha + \kappa - \beta)F(\sigma^*)$*, otherwise* $s_i(\sigma_i) \to \{S\}$.

Figure 1 plots the best response functions for players A and B. The value of σ^* will depend on the parameters of the model, but can be calculated by finding their point of intersection. In this particular example, $\sigma^* = 0.652$

2.3 Common Shocks and Strategy

The common shock parameter κ has a natural interpretation as measuring the overall state of the financial sector or the economy more generally. The intuition behind this pa-

Figure 1: Best Responses with $\kappa = 0$

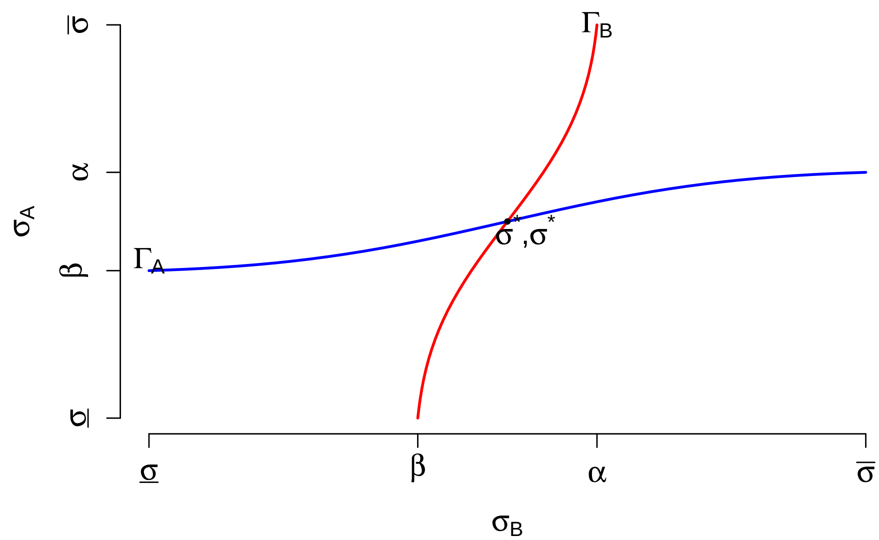

Notes: The unique Bayesian Nash equilibrium lies at the intersection of Γ_A and Γ_B. With parameters $\alpha = 1.5, \beta = .5, \kappa = 0, \underline{\sigma} = -1, \overline{\sigma} = 3, F \sim N(1.5, 1)$, this strategy $s_i(\sigma_i) \to \{H\}$ if $\sigma_i \geq 1$.

rameter is that when an economy is doing better—so κ is high—this decreases an actor's incentive to sell the asset by increasing the relative benefit for holding. Substantively, this may reflect a subjective sense of optimism shared by both player players about the likelihood of a crisis, or an objective increase in the asset's returns. It is straightforward to verify that changing the value of κ changes the best response functions for the two players by examining the relationship between κ and σ^*. Implicitly differentiating Equation 4,

$$\frac{\partial \sigma^*}{\partial \kappa} = \frac{F(\sigma) + 1}{(\alpha - \kappa + \beta) f(\sigma) - 1} \tag{5}$$

Equation (5) is less than 0 if Assumption 2 holds, which guarantees that the denominator is negative. This indicates that σ^* is decreasing in κ; substantively, this means that better the state of the economy, or the more optimistic both players are about the state of the economy, the more negative player i's type must be to entice her to sell the asset. And because this applies to both players, it makes j more likely to hold as well, and shifts equilibrium behavior as shown in Figure 2.

Figure 2: Best Responses with κ Shocks

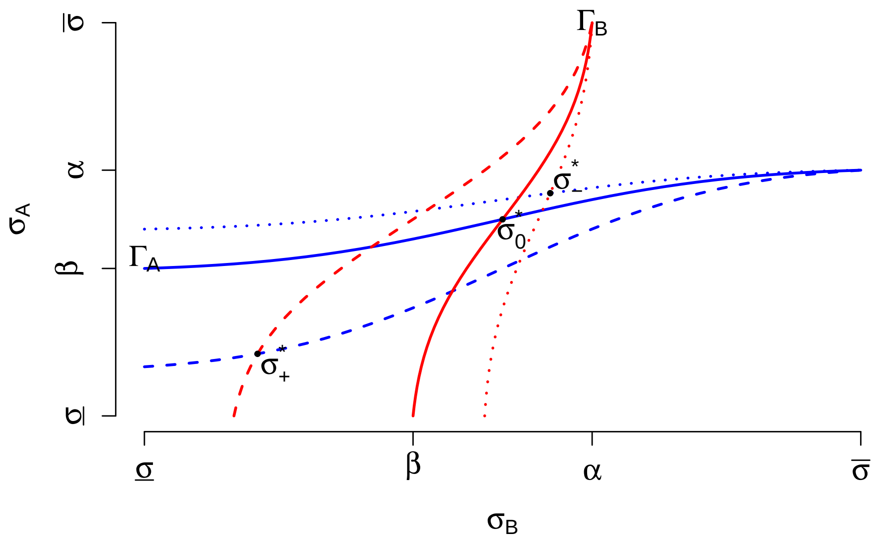

Notes: As κ increases, σ^* decreases. The solid lines are the best response functions from Figure 1. The dashed lines are the best response functions for a positive κ shock. The dotted lines are the best response functions for a negative κ shock.

3 Predicting Future Crises

We now turn to our main theoretical results. Our goal is to characterize the likelihood of a financial crisis in the future—which in the model depends on players' types $\sigma_{i,j}$ and the value of a future shock κ_2—based on information available to him or her in the present. This is a natural analogue for the actual tasks that social scientists, financial regulators, and others undertake when attempting to predict financial crises: they use publicly available information in the present to make predictions about the likelihood of financial crisis, in the future.

We model this task by adding a third player, the Observer, who knows the structure of the game as well as the value of the common shock κ, just as players A and B do. However, the Observer does not know the value of either σ_A or σ_B. He observes only the equilibrium outcome of the game in period 1, which is one of the following action pairs: $\{H, H\}$, $\{H, S\}$, $\{S, H\}$, $\{S, S\}$. The structure of the interaction between Observer and the players unfolds as follows.

1. Nature draws σ_A, σ_B and reveals them to players A and B.

2. Nature draws κ_1 from distribution F_κ and reveals it to A, B, and the Observer.

3. A and B play the game in Table 2.

4. The Observer and players A and B all updates their beliefs about the values of σ_A, σ_B conditional on what they observe.

5. The Observer predicts whether a financial crisis will occur given possible future realizations of the common shock κ_2, also drawn independently from F_κ.

To reiterate, the task of the Observer is to predict the likelihood of a crisis in period 2, knowing that the common shock changes each period, that the players types are constant across periods, and that each player i will also update her beliefs about player j's type.

The interesting case is when both players choose Hold, $\{H, H\}$ in period 1, for the occurrence of a crisis obviates the need to predict it.[2] Conditional on observing $\{H, H\}$ in period 1, the Observer can infer that the values for of σ_A and σ_B both lie in the range $[\sigma^*, \overline{\sigma}]$. Figure 3 illustrates how this works. The area under the curve F_σ, bounded by $\underline{\sigma}, \overline{\sigma}$, represents the distribution of possible types in period 1. Having observed $\{H, H\}$ the Observer knows that the lowest possible type for each player is now $\sigma^*(\kappa_1)$, which truncates the distribution of possible types from below, as shown. Let $r(\kappa_1) = \overline{\sigma} - \sigma^*(\kappa_1)$, the range of values greater than $\sigma^*(\kappa_1)$.

[2]We leave aside here the cases where the Observer views outcomes of $\{H, S\}$ or $\{S, H\}$.

Figure 3: Types Conditional on κ_1

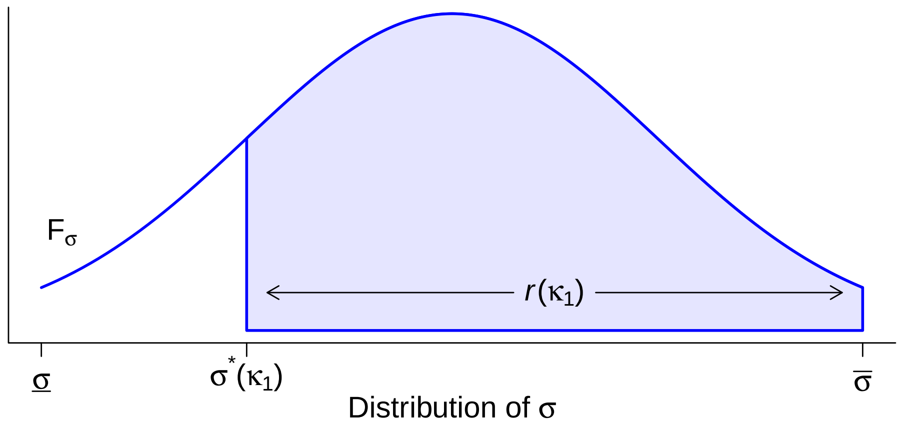

Notes: $r(\kappa_1)$ describes the range of types conditional on a realization of κ_1 and a $\{H, H\}$ outcome in period 1.

3.1 Screening and Confidence

To distinguish between the two separate mechanism through which shocks affect the predictability of crises, assume for the time being that Players A and B do not make similar inferences, and therefore continue to believe that the opponent's type ranges from $[\underline{\sigma}, \overline{\sigma}]$. This means that the players have the same strategies in period 2 as in 1. Then, if $\kappa_2 > \kappa_1$, then the probability of a $\{S, S\}$ outcome in period 2 is zero, otherwise it is positive. Denote $\sigma^*(\kappa_1)$ as the cutoff value given κ_1. If $\kappa_2 < \kappa_1$ then the value $\sigma^*(\kappa_2)$ lies in the interval $[\sigma^*(\kappa_1), \overline{\sigma}]$, which confirms that if the common shock in 1 has resulted in no crisis, then the common shock in 2 must be lower—"more negative"—in period 2 to generate a crisis.

Via Equation (5), we know that $\sigma^*(\kappa_1)$ is decreasing in κ_1, meaning that $r(\kappa_1)$ is increasing in κ_1. Recall that κ_1 and κ_2 are drawn independently from the same distribution. As a result, the range of values of κ_2 that can generate a crisis is strictly greater as κ_1 grows larger. A greater range implies a lower degree of truncation in F_σ, and hence a higher variance in types, just so long as the function F_σ is log-concave (An 1998).[3]

[3]This includes a very broad range of common distributions, including the normal and uniform. Furthermore, many distributions—such as the Student's t distribution—that are not globally log-concave are log-concave when truncated on some interval (Bagnoli and Bergstrom 2005)

Proposition 1 (Common Shocks Screen Types) *If Assumption 1 and Assumption 2 hold in periods 1 and 2, then larger common shocks κ_1 result in larger ranges $r(\kappa_1)$, and hence a higher variance in player types in period 2 for any for any log-concave distribution F_σ.*

$r(\kappa_1)$ thus measures the "predictability" of a crisis. As this value grows larger, corresponding to less truncation of F_σ, the Observer can be less certain about whether players will cooperate in period 2.

Now suppose that each player *does* learn about her opponent's type. This means that given shock κ_1, player i knows (just as the Observer does) that player j's type lies in the interval $[\sigma^*, \overline{\sigma}]$. This changes her best response function because it truncates the distribution of types F_σ. Denote the new distribution $F_{\sigma,2} = F_\sigma \in [\sigma^*(\kappa_1), \overline{\sigma}]$. To illustrate how

Figure 4: Best Responses in Period 2 Conditional on κ_1

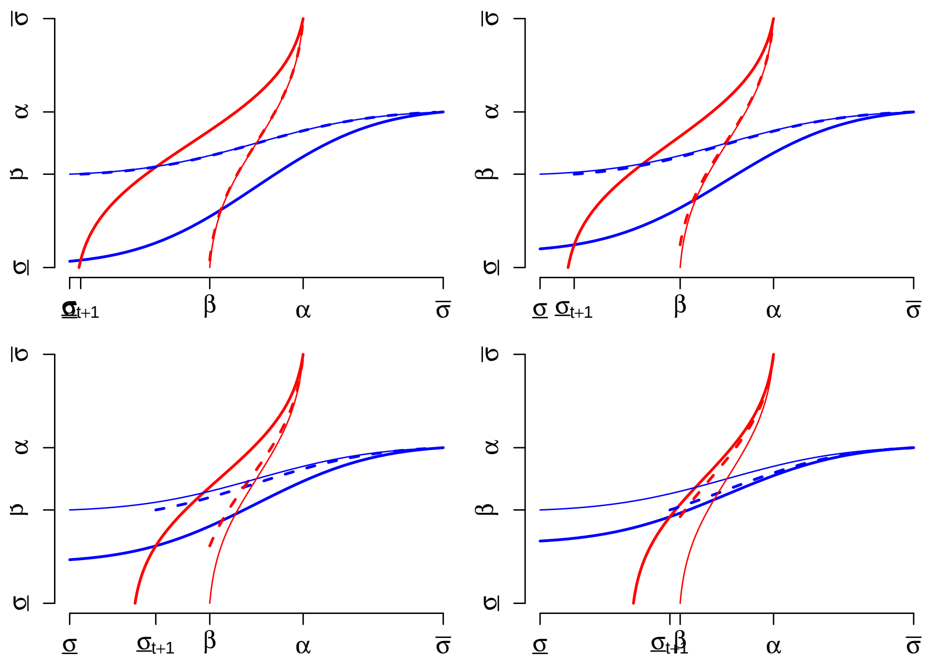

Notes: As κ_1 grows larger, σ_2^* grows smaller. The thicker solid lines are the best response functions in period 1 given four values of κ_1. The dashed lines are the best response functions in period 2, with $\kappa_2 = 0$. The thin solid lines are the best response functions for $\kappa_1 = 0$, which illustrate the difference between σ_1^* and σ_2^*.

players' best response functions change as $F_{\sigma,2}$ changes, Figure 4 plots the players' best response functions in period 2 given shock κ_1. The less positive the shock in period 1, the greater the truncation in the distribution $F_{\sigma,2}$, and the more certain player i is that j is of a stronger type. This increases the likelihood that both players will Hold. For any given κ_2, then, smaller shocks in κ_1 increase the range of cooperation in 2.

There are, in sum, two separate effects of common shocks in period 1. Smaller shocks *screen types* by helping to place bounds on the range of $\sigma_{i,j}$. This effect does not depend on coordination or self-fulfilling dynamics. However, by screening types, in turn, smaller shocks also *increase confidence* among both players that their opponent is of a stronger type, a result that stems directly from coordination dynamics. From the Observer's perspective, the larger the shock κ_1, the less he can infer about how strong or weak the players are, and the less confident he will be that for any value of κ_2, each player will Hold. Both of these effects decrease his ability to predict a crisis in the period 2.

The main implication of this analysis is that when attempting to predict financial crises with self-fulfilling features, the probability of a crisis depends on both the future state of the economy and the strength or resilience of market participants in the context of future economic conditions. If participants' types were known, then it it would be possible to predict whether they would sell the asset in any future state of the economy. But since types are not known, market conditions in the present help both players and the observe to make inferences about players' types in the future. The more favorable current conditions are, the more difficult it is to infer just how strong or resilient market participants are, a consequence of both the screening effect (if conditions had been poor, then only the strong would have survived) and the confidence effect (if conditions had been poor, then participants themselves would adopt more cooperative strategies). The result is that favorable market conditions predict *greater uncertainty* about the likelihood of future crises.

3.2 Extensions

In the Appendix we outline several theoretical extensions that relax some of the assumptions made so far. Here, we address two that are particularly useful: allowing for larger shocks that create epistemic uncertainty, and embedding financial "euphoria" into our analysis.

3.2.1 Boom Times and Epistemic Uncertainty

Assumption $1b$ ensures that there is some probability that player i both sells and holds by restricting the value of κ to satisfy the constraint that $\underline{\sigma} < \beta - \kappa < \alpha < \overline{\sigma}$. This assumption guarantees the uniqueness of the Bayesian Nash equilibrium, but in doing so it rules out some of the more interesting features of coordination games with multiple equilibria (see Bueno de Mesquita 2014 for a discussion). In the present application, it rules out the possibility that a common shock is so positive that players would have a dominant strategy to hold the asset regardless of their types, which is true if $\kappa > \beta - \underline{\sigma}$.

Relaxing that assumption generates a richer set of strategies, and accordingly complicates predictions. To see this, replace Assumption 1 with Assumption 1′ as follows:

Assumption 1′ *The parameters of the Private Information Observer Game satisfy the following constraints:*

a. $\beta < \alpha$

b. $\underline{\sigma} < \beta < \alpha < \overline{\sigma}$

The difference between Assumption $1b$ and Assumption $1'b$ is simply that the latter makes no references to the values of κ. Assumption 1′ still ensures that the game is one of strategic complements, but allows for the common shock to be positive (negative) enough that players will Hold (Sell) the asset regardless of their types $\sigma_{i,j}$.

As a result, there are now three possible scenarios in period 1.

Scenario 1: $\kappa_1 < \beta - \underline{\sigma}$ and $\kappa_1 > \beta - \alpha$. Under this "Normal times" scenario, the game proceeds as above.

Scenario 2: $\kappa_1 < \beta - \alpha$. In this "Hard times" scenario, an equilibrium exists where player i sells regardless of her type.

Scenario 3: $\kappa_1 > \beta - \underline{\sigma}$. In this "Boom times" scenario, an equilibrium exists where player i holds regardless of her type.

The proliferation of equilibria in Scenarios 2 and 3 has substantial implications for the Observer's task of inferring players' types based on their observed actions. Concentrate on Scenario 3 ("Boom times"). If $\kappa_1 > \beta - \underline{\sigma}$ and the Observer observes $\{H, H\}$, what can he conclude? Only that each player is playing a strategy of always Hold, or that she is playing a cutoff strategy. If the former, he can make no new inferences about players' types; if the latter, he can. But it is impossible for the Observer to assign probability distributions over these equilibria, and hence to make inferences about player types, even though (by assumption) the distribution F_σ and the current realization of κ are known.

This is a fundamental source of the "epistemic uncertainty" discussed by Nelson and Katzenstein (2014).[4] Unlike most treatments of uncertainty in financial decision making, however, this source of uncertainty is not a constant problem for the Observer that stems from some fundamental inability to assign probabilities to future states of the world. Rather, it only emerges during Boom times, when a large positive value for holding the asset prevents there from being a unique Bayesian Nash equilibrium to the incomplete information coordination game. Of course, if it were the case that the Observer did not know the distribution F_σ then the Observer would indeed be in a world of Knightian uncertainty.

[4]Nelson and Katzenstein's (2014) use of the term "epistemic uncertainty" differs from standard usage in statistics and decision science. See O'Hagan (2004) on the distinction between "epistemic" and "aleatory" uncertainty: the former corresponds to uncertainty that stems from features of the world that could be known but are not, whereas aleatory uncertainty refers to features of the world that are unknowable. This use of epistemic uncertainty can be made consistent with the Observer's inference problem by stipulating that the Observer could in principle know players' types and strategies—they *are* known to the players, after all—but does not because players either cannot or will not share that information.

The main theoretical point this extension raises is that players' strategies themselves are a conceptually distinct source of uncertainty that regulators and analysts may face. Sufficiently positive common shocks—sufficiently large booms—prevent any learning about players' types, and the probability of a crisis in the future, from players' observed actions.

3.2.2 Financial "Euphoria"

In the baseline model, players' types are constant. They do not change across periods, either as a result of market events, or due to some other source of variation. This simplification facilitates a direct exposition of how the Observer and the players make inferences about types given players' actions. But it does come at the expense of a simplified account in which there is no relationship at all between firm strength and market conditions.

Scholars such as Kindleberger (2000) and Minsky (1982) have observed that long spells of good economic conditions can lead to a kind of financial "euphoria." To relate this to our analysis, we conceptualize financial euphoria as an erroneous belief about players' risk tolerance. These beliefs may be about their own type, or about other players' type. In our theoretical model, as in Kindleberger and Minsky, such beliefs are *irrational* in the sense that they do not follow from the logical inferences that players and the observer may draw about players' types given observed play. However, if we add such a dynamic to our model, it follows that players become more willing to choose $Hold$ following positive economic shocks than they would otherwise be. If the Observer is aware that financial euphoria is possible—we presume that the players themselves are unaware of their own euphoria—then he knows that he has even less information about players' types. Such a dynamic makes financial crises even less predictable from the perspective of the Observer.

This example shows how it is possible to embed additional insights into our modeling framework, while preserving our key predictions. In the case of Kindleberger-Minsky financial euphoria, we have assumed that players have erroneous beliefs about types. In Appendix 1.1 we outline three other ways to conceptualize how players' types may evolve over time and the consequences for our core theoretical predictions.

4 Empirical Evidence

The central result of our model is that crises become more unpredictable, i.e. there is greater variance in the future states of the world that will experience financial crises, when economic conditions are more favorable in the present. In this section we empirically test this novel theoretical prediction with a new continuous monthly indicator of financial market stress originally developed by Gandrud and Hallerberg (2015).

4.1 Measuring Variance in Financial States

In order to empirically test the model's primary implication we need to measure both economic conditions at one point in time and the variance of financial conditions subsequently. Measuring economic conditions is relatively straightforward using a number of different widely available approaches. We first focus on gross domestic product growth rates (GDP). Higher growth rates clearly reflect better economic conditions, and GDP growth rates have the benefit of summarizing in aggregate the behavioral responses of market actors to current market conditions. This data is available at both annual and quarterly intervals from the World Bank[5] and the OECD,[6] respectively. We also examine stock price volatility, where lower volatility would indicate "better" economic conditions. This data is taken from the Bloomberg via the World Bank's Global Financial Development Database (World Bank 2015). We currently only have yearly stock price volatility data.

Because we are trying to model how an observer—e.g. a regulator—would be able to predict future crises, we also examine a key indicator that regulators themselves observe when attempting to measure financial market conditions. We focus here on a key

[5] See `http://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG`. Accessed November 2015.
[6] See `https://stats.oecd.org`. Accessed November 2015.

"CAMELS" indicator[7]: mean country-year impaired assets to gross loans ratios. Financial systems with higher impaired assets ratios–e.g. loans that have delayed repayment–have higher probabilities of bank insolvencies. Examining impaired asset ratios helps ensure that our analysis is a realistic approximation of the actual macro-prudential regulatory task of predicting financial crises.

Measuring the variability of future states is less straightforward. Most indicators of financial crises are at the country-year level. These include binary measures of banking crisis by Reinhart and Rogoff (2009; 2011), Laeven and Valencia (2013), and their predecessors (see Gandrud and Hallerberg 2015 for a detailed discussion); continuous measures of the probability that banking systems will become insolvent based on Z-scores (see Andrianova et al. 2015 and Lepetit and Strobel 2013); and less widely used continuous variables developed by others, including Rosas (2009). However, yearly data makes it difficult to construct a valid indicator of the variance of future states as it would require finding the variance across a number of years, and this is not the relevant window that observers are considering when they attempt to predict a crisis. Even worse, short windows that more accurately approximate the time horizon across which market actors predict crises are especially problematic when calculating variances with yearly data. For example, if we use a two year window with a binary crisis indicator, there are only two possible values that the variance may take: 0 and 0.5.[8]

Ideally, we would have an indicator of financial market conditions that is continuous and recorded at more frequent intervals.

There are many individual pricing measures–such as bank credit default swap prices, bank equity prices, or stock market returns more generally (e.g. Danielsson, Valenzuela,

[7]CAMELS indicators measures banks' capital adequacy, asset quality, management capacity, earnings, the liquidity of its assets, and its sensitivity to market risks (hence CAMELS). Andrianova et al. 2015 recently aggregated bank-year level CAMELS data from the Bankscope database into country-year averages. See `https://bankscope.bvdinfo.com`. Accessed December 2015.

[8]That is, if there is either no crisis in the two years $[0,\ 0]$ or a crisis in both of the years $[1,\ 1]$, then the variance is 0. If there is a crisis in one of the years and not one in the other–$[1,\ 0]$ or $[0,\ 1]$–then the variance is 0.5.

and Zer 2015)–that vary at sub-yearly intervals. However, making meaningful cross-country macro-comparisons of the variance in these prices can be difficult. Given that different countries have different banking systems that rely to greater and lesser extents on these types of financing, the importance of variance in each individual instrument for indicating market stability varies across countries.

Gandrud and Hallerberg's (2015) "FinStress" measure of financial market stress meets both the continuous and frequent periodicity criteria, while also allowing us to make meaningful cross-country comparisons. They use a text analysis tool called kernel principal component analysis to analyze monthly Economist Intelligence Unit country banking and financial system reports. From this analysis they created a continuous measure of financial market stress that ranges between zero and one, with higher numbers indicating more stress. As the measure summarizes qualitative assessments of country-level conditions, it allows us to make more meaningful comparisons across countries than individual price measures would. Because the indicator is monthly, we can easily calculate the variance of FinStress over reasonably short windows.[9] When using yearly data on the right-hand side of our regressions, we calculate the variance of FinStress across the future period; the twelve months of $year + 1$. When using quarterly data on the right-hand side of the estimation model we similarly find FinStress' variance across $quarter + 1$. We denote this dependent variable as $Var(FinStress_{t+1})$ where t is either a year or a quarter. FinStress is currently available from 2003 through 2011 for over 180 countries, though quarterly economic data was only accessible for OECD countries.

[9]Note that we also ran the models using FinStress standard deviations rather than variances. The results from these models (not shown) were substantively similar.

Table 3: Regression result from predicting FinStress Variance using annual explanatory variable data

	Dependent variable:							
	Var(FinStress)$_{year+1}$							
	Full Sample	Full Sample	Full Sample	Full Sample	OECD	OECD	OECD	OECD
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Var(FinStress)$_{year+0}$	0.014	−0.011	−0.083*	0.076**	0.036	−0.027	−0.048	0.052
	(0.029)	(0.030)	(0.047)	(0.038)	(0.068)	(0.070)	(0.072)	(0.077)
GDP Growth (%)	0.046**	0.028			0.336***	0.207**		
	(0.022)	(0.023)			(0.083)	(0.090)		
FinStress Mean$_{year}$		−5.347***	−6.263***			−10.146***	−6.916*	
		(1.299)	(2.404)			(3.150)	(3.531)	
Stock Price Volatility			−0.061***				−0.144***	
			(0.022)				(0.039)	
Impaired Loans (log)				−0.277				−1.636***
				(0.196)				(0.530)
Fixed Effects	y	y	y	y	y	y	y	y
Observations	1,349	1,349	599	833	248	248	231	205
R^2	0.004	0.018	0.044	0.009	0.077	0.119	0.153	0.061
Adjusted R^2	0.003	0.016	0.038	0.008	0.067	0.103	0.132	0.052

Note: *p<0.1; **p<0.05; ***p<0.01

Table 4: Regression result from predicting FinStress Variance using quarterly explanatory variable data (OECD only)

	Dependent variable:	
	$\text{Var(FinStress)}_{quarter+1}$	
	(1)	(2)
$\text{Var(FinStress)}_{quarter+0}$	0.090^{***}	0.066^{**}
	(0.027)	(0.027)
GDP Growth (%)	0.102^{***}	0.043
	(0.027)	(0.029)
$\text{FinStress Mean}_{quarter+0}$		-5.569^{***}
		(1.065)
Fixed Effects	y	y
Observations	1,237	1,237
R^2	0.023	0.045
Adjusted R^2	0.022	0.043
Note:	*p<0.1; **p<0.05; ***p<0.01	

4.2 Predicting Financial Market Stress: A Regression-based Approach

We first use partial adjustment linear regressions to examine if present economic conditions were associated with future financial market stress variance. The results of these regressions are shown in tables 3 and 4. We multiply $Var(FinStress)$ by 1,000 to make the estimated coefficients easier to interpret.[10] All models include FinStress variance in the present period, as well as country fixed-effects.[11]

Across the regressions we find evidence for the same conclusion: better conditions today are associated with higher variance in future financial market stress. When statistically

[10] Due to the fact that FinStress can only vary between zero and one, in the full sample the annual variance variable originally ranged from 0 to 0.02 with a median at 0.003.

[11] Including a lagged dependent variable on the right-hand in a panel regression with few time periods per country could create Nickell Bias (Nickell 1981), though Beck, Katz, and Mignozzetti 2014 argue this is rarely a substantively meaningful problem in comparative research. Nonetheless, we ran the models without the lagged dependent variable. Results from these models (not shown) are substantively similar. We also ran models with year fixed-effects in addition to country fixed effects. The results of the mean FinStress level remained substantively the same.

significant, GDP growth is positively associated with $Var(FinStress)$ in the following period; that is, higher GDP growth in the present period is associated with more variance in financial market stress in the following period. We also looked at the mean of FinStress in the present year/quarter as an alternative measure of current financial market conditions. Lower values indicate less stress and hence better financial market conditions, and so should be negatively correlated with FinStress variance in the following period. This is precisely what we find. Stock market volatility is strongly negatively associated with $Var(FinStress)$ when there is higher stock price volatility, i.e. there are worse equity market conditions, then there is lower $Var(FinStress)$.

Moving on to key CAMELS regulatory risk variables, we find impaired loans,[12] e.g. non-performing and restructured loans, to be negatively associated with $Var(FinStress)$ and this effect to be strongly statistically significant in the OECD sub-sample. When there are more impaired loans—an important threat to bank solvency—in the banking system, there is less financial market stress variance in the following period. This appears to be particularly likely in more developed countries. In the full sample we also find some weak evidence that liquid asset—e.g. cash—ratios are positively associated with $Var(FinStress)$. This result may seem counter-intuitive; if banks had more liquid assets, then they would be less likely to fail. However, Andrianova et al. (2014) find that banks in less developed countries tend to have high liquid asset ratios because their banks are very wary of making new loans due to high credit risks. So, very high liquid asset ratios in our model reflect worse economic conditions.

4.3 Predicting Financial Market Stress: A Drift-Diffusion-Jump Approach

As a robustness check we use an approach from the literature on time-series forecasting called nonparametric drift-diffusion-jump models (DDJ, Carpenter and Brock 2011; Dakos et al. 2012). DDJ models allow us to approximate processes of change in a time series

[12]The variable was logged due to its highly positively skewed distribution.

without needing to make explicit assumptions about the underlying process that creates these changes. The findings from this approach corroborate those from our regression analysis: there is more variation in financial market stress during good times compared to bad.

Drift is a measure of the local rate of change. Diffusion codes small changes that happen at each time increment. Jumps are larger shocks that occur intermittently and are uncorrelated in time. The approach we take to estimating the DDJ model is from Carpenter and Brock 2011.[13]

Based on our theory, we expect that jumps will be more common in countries' FinStress scores during non-crisis periods and more diffusion during crisis periods. To test this we first graphically compared the distributions of jump and diffusion parameters estimated for each country's FinStress time series during what Laeven and Valencia[14] classify as crisis and non-crisis periods. Figure 5 shows these densities. We have also included a measure of total variance, which is a summary of both jump and diffusion parameters.

We can see that the distribution of estimated jump parameters in 'non-crisis' periods is shifted upward from the distribution of jump parameters in 'crisis' periods. Conversely, the distribution of diffusion parameters in crisis periods is shifted upward from non-crisis periods. Finally, the distribution of total variance in crisis periods is lower than non-crisis periods. We found these distributions to be statistically significantly different in the described direction at all conventional levels using one-sided Kolmogorov–Smirnov tests.[15]

To further understand this finding, it is useful to refer to Figure 6 which shows the estimated jump parameters for a wide selection of countries that recently had banking crises

[13]The model approximates the unknown process generating FinStress scores: $dx_t = f(x_t,\ \theta_t)dt + g(x_t,\ \theta_t)dw + dJ_t$. dx_t is the change in the FinStress score x for a country at time t. θ_t is a critical transition parameter. The drift function is given by $f(x_t\theta_t)dt$. The diffusion function is given by $g(x_t\theta_t)dw$. J is a jump process. Please see Dakos et al. (2012, 7) for further details. We estimated the model using the `ddjnonparam_ews` function from the **earlywarnings** R package (Dakos and Lahti 2013). Note that we estimated the parameters for each country's time series separately.

[14]Despite the previously discussed shortcomings, they are the most recently updated and comprehensive binary measure of crises.

[15]Again, we ran the tests using the `ks.test` function from base R.

Figure 5: Diffusion, Jump, and Total Variance Estimate Distributions Across Crisis and Non-Crisis Periods from Laeven and Valencia 2013

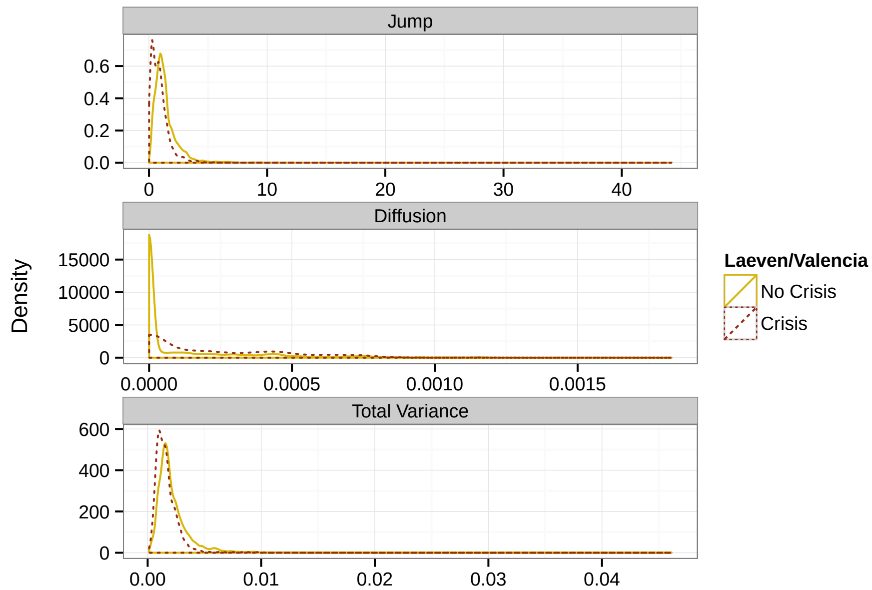

as classified by Reinhart and Rogoff 2011 and, alternatively, Laeven and Valencia 2013. Notice that many of the periods that are classified by these authors as being banking crises do indeed begin with relatively high jumps. What happens after the crisis "settles in" is interesting. Periods before crises are often very "jumpy". In crisis periods, conversely, estimated jumps tend to be close to 0. Germany, Greece, Ireland, Luxembourg, Spain, and the United Kingdom are stark examples of this pattern. Countries where there are relatively many jumps during what Laeven and Valencia 2013 and/or Reinhart and Rogoff 2011 classify as banking crises, such as Switzerland, actually have very short periods of of elevated FinStress. So, hidden behind the blunt binary crisis measures are events that follow the general pattern of there being less frequent FinStress jumps during periods of elevated stress and vice versa. These findings are strong support for the implications of our formal model. During "good" times actors' behaviour is less informative about whether they will hold or sell their assets given a shock and thus we observe greater swings in financial market stress.

Figure 6: Estimated Jumps in Perceptions of Financial Market Conditions During Crisis vs. Non-Crisis Periods

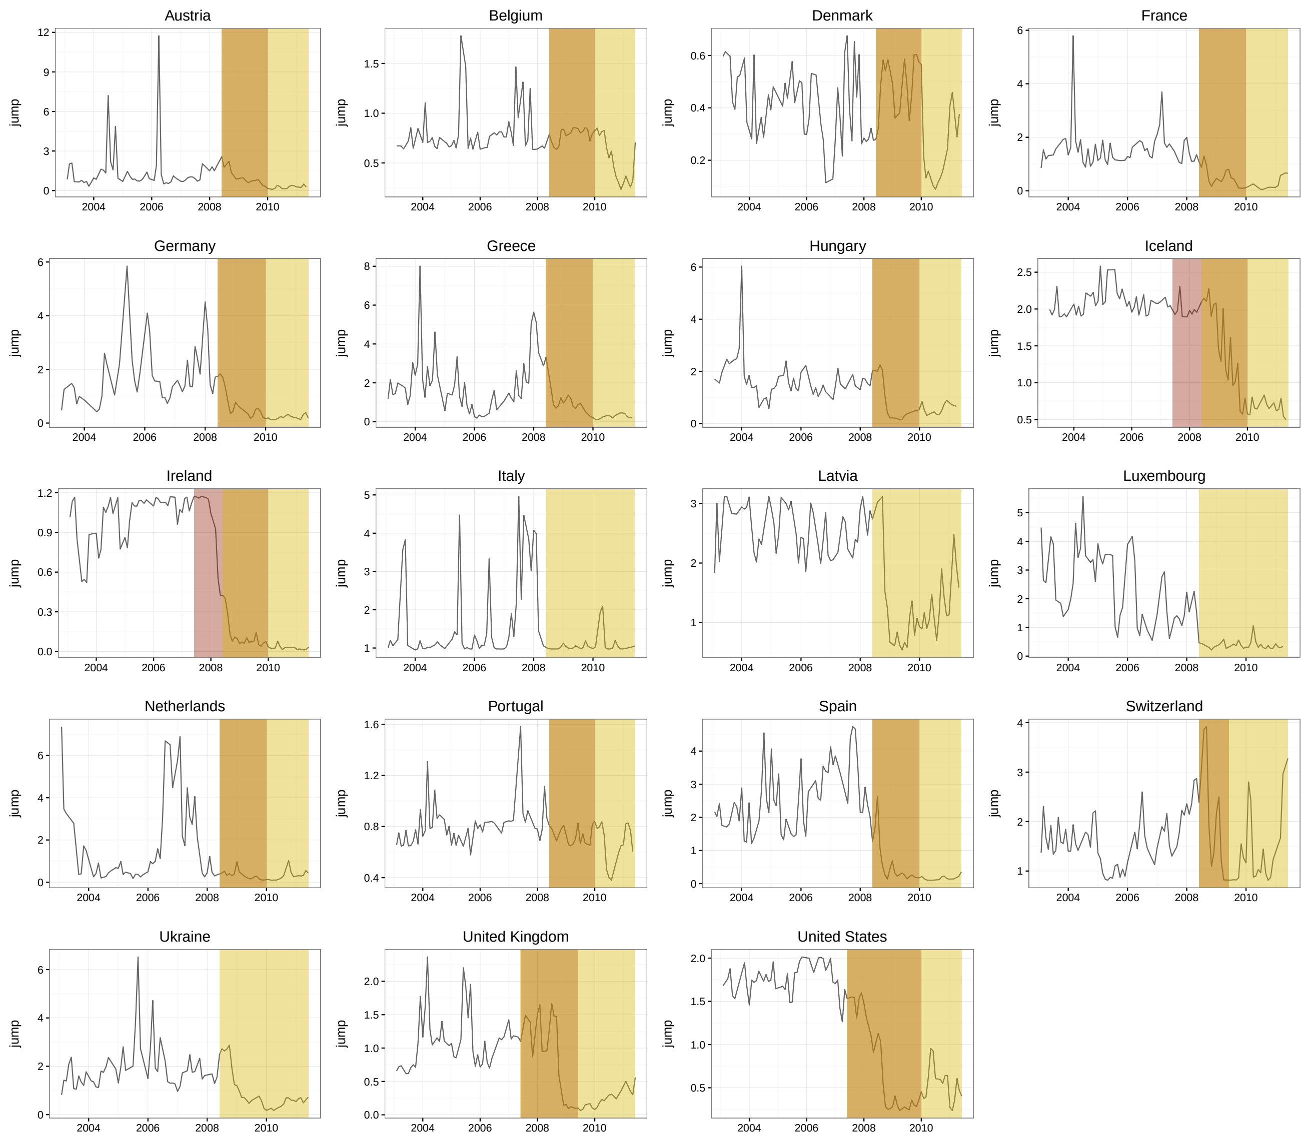

Solid lines show the estimated jump parameter.

Yellow shaded areas indicate periods that Laeven and Valencia 2013 classify as systemic banking crises. Note that crises are automatically terminated at the end of 2011 due to the series not extending beyond this point, not necessarily because the crisis finished.

Red shaded areas indicate periods that Reinhart and Rogoff 2011 classify as banking crises. Note that crises are automatically terminated at the end of 2009 due to the series not extending beyond this point, not necessarily because the crisis finished.

Orange areas indicate periods where a crisis is recorded for both measures.

5 Conclusion

This paper has studied the problem of predicting self-fulfilling financial crises, using coordination games as a simple heuristic device for capturing how financial market actors interact with one another. We follow the recent literature that models coordination games using a global games approach by incorporating private information about players' types into the coordination game. Our conceptual innovation is to model the prediction problem explicitly by including an Observer whose task is to predict financial crises in the future based on observed actions in the present, and to characterize what he can learn about the future given common information about fundamentals and observed play in the present. Defining the predictability of a crisis as the range of possible future states of the world that would generate a crisis, our results, both theoretic and empirical, show that future crises become more unpredictable as current economic conditions become more favorable. We theoretically decompose the sources of uncertainty to two effects: a *screening effect* in which weak players cooperate when economic conditions are strong, and a *confidence effect* in which each player becomes more confident about his opponent's strength conditional on her having held the asset in the present. A more general model that allows for multiple equilibria in the event of sufficiently large common shocks, in turn, allows for epistemic uncertainty as a distinct source of uncertainty in predicting future crises.

There are many coordination dilemmas of interest to social scientists beyond financial crises, and the failure of social scientists to predict discrete changes in mass behavior has long been a prominent critique of contemporary social science. Prominent examples beyond the Global Economic Crisis include the Arab Spring revolts, the Asian Financial Crisis of the late 1990s, and the revolutions in Eastern Europe in 1989, among others. Yet there are numerous models of coordination and market behavior in which the unpredictability of coordination failures is a feature of the model rather than a problem to be explained. Cass and Shell (1983) formalize the effects of "sunspots" on macroeconomic fluctuations, and link this concept to Keynes's (1936) notion of "animal spirits." In these cases, beliefs or sentiments vary for reasons that are irreducible to economic or individual preferences, so fluctuations occur for reasons that are not *and cannot be* captured in any model. Adopting a different line of argument, Timur Kuran (1989; 1991) has argued that individuals may falsify their preferences for supporting an incumbent political regime based on the belief that other citizens support the regime. In this case, the incentive to hide private information prevents any observer from predicting future changes in mass sentiment towards the regime unless the observer has a mechanism for enticing individuals to reveal their private beliefs to the observer and for predicting when an information cascade will begin. Goodwin (2011) applies Kuran's insight to political scientists' failure to anticipate the Arab Spring, and Bellin (2012) gives a more critical perspective still in which the lesson to be drawn from the Arab Spring is simply that individual and collective agency renders prediction nearly impossible.

Our analysis offers a complementary, but distinct, perspective on the limits of prediction in repeated coordination games. Like sunspots or animal spirits, a key driver of uncertainty is fluctuations in common conditions that are external to the model itself. Like Kuran, common conditions interact with private information about players' types, but unlike Kuran the mechanism that ultimately causes unpredictability is not their public falsification of their beliefs. Instead, the core mechanism that we highlight is the simple fact that players' strategies depend on the future values of the common shock parameter κ, which in the financial crisis example represents future economic conditions but in other examples may represent collective beliefs, mass sentiments, public actions by a regime, or other common changes to the economic and political environment that shape the value of cooperation. This provides us with insights that existing approaches do not yet capture: crises or other changes in equilibrium behavior are not fundamentally predictable or not, they are *more or less predictable* as a function of past conditions.

We use these insights to derive a novel prediction about the relationship between market conditions and financial crises, and found strong evidence in favor of this account using a new measure of financial market stress and a wide range of indicators of financial market conditions. These results help to explain the absence of a clear link between economic conditions and crises. Policy variables, on the other hand, should be better predictors of financial crises insofar as they follow a "first generation" logic (Eichengreen 2003; see Chari and Kehoe 2003 for theoretical model in which both fundamentals and herding generate crises). By restricting attention to coordination dynamics in this paper, we isolate how they affect the predictability of financial crises without denying the importance of other factors.

Our empirical results should be understood as the first step towards demonstrating the possibilities and limits of predicting financial crises. The recent literature on global games and related applications has been able to draw out other implications for the empirical testing of such theoretical models. Notably, Angeletos and Pavan (2013) show how to make "predictions about the joint distribution of a limited set of variables: fundamentals, policy choices, and regime outcomes" (886) in coordination games with multiple equilibria. While there are many differences between their model and ours, one particularly important one is that they do not characterize the problem of an observer making a prediction about future crises, but rather of understanding the *contemporary* link between fundamentals and regime outcomes, which in our model corresponds to economic conditions and financial crises. Our predictions do coincide, however, for the contemporaneous relationship between market conditions and equilibrium outcomes. Embedding the prediction problem into their general model represents one promising direction for future research.

References

An, Mark Yuying (1998). "Logconcavity versus Logconvexity: A Complete Characterization". *Journal of Economic Theory* 80.2, pp. 350 –369.

Andrianova, Svetlana et al. (2014). "Why Do African Banks Lend So Little?" *Oxford Bulletin of Economics and Statistics* 77.3, pp. 339–359.

Andrianova, Svetlana et al. (2015). "A New International Database on Financial Fragility". *University of Leicester Department of Economics Working Paper*, pp. 1–54.

Angeletos, George-Marios and Alessandro Pavan (2013). "Selection-free predictions in global games with endogenous information and multiple equilibria". *Theoretical Economics* 8.3, pp. 883–938.

Bagnoli, Mark and Ted Bergstrom (2005). "Log-Concave Probability and Its Applications". *Economic Theory* 26.2, pp. 445–469.

Baliga, Sandeep and Tomas Sjöström (2011). *Conflict Games with Payoff Uncertainty*. Working paper, Department of Economics, Rutgers University.

— (2012). "The Strategy of Manipulating Conflict". *American Economic Review* 102.6, pp. 2897–2922.

Beck, Nathaniel L, Jonathan N Katz, and Umberto G Mignozzetti (2014). "Of Nickell Bias and its Cures: Comment on Gaibulloev, Sandler, and Sul". *Political Analysis* 22.2, pp. 274–278.

Bellin, Eva (2012). "Reconsidering the Robustness of Authoritarianism in the Middle East: Lessons from the Arab Spring". *Comparative Politics* 44.2, pp. 127–149.

Berg, Andrew and Catherine Patillo (1999). "Predicting currency crises: The indicators approach and an alternative". *Journal of International Money and Finance* 18.4, pp. 561–586.

Carlsson, Hans and Eric van Damme (1993). "Global Games and Equilibrium Selection". *Econometrica* 61.5, pp. 989–1018.

Carpenter, SR and WA Brock (2011). "Early warnings of unknown nonlinear shifts: a nonparametric approach". *Ecology* 92, pp. 2196–2201.

Cass, David and Karl Shell (1983). "Do Sunspots Matter?" *Journal of Political Economy* 91.2, pp. 989–1018.

Chari, V. V. and Patrick J. Kehoe (2003). “Hot Money”. *Journal of Political Economy* 111.6, pp. 1262–1292.

Chinn, Menzie D and Jeffry A Frieden (2011). *Lost Decades: The Making of America's Debt Crisis and the Long Recovery*. New York: WW Norton & Company.

Colander, David et al. (2009). “The Financial Crisis and the Systemic Failure of the Economics Profession”. *Critical Review* 21.2-3, pp. 249–267.

Dakos, Vasilis and Leo Lahti (2013). “R Early Warning Signals Toolbox”. *The R Project for Statistical Computing*. http://cran.r-project.org/web/packages/earlywarnings/index.html.

Dakos, Vasilis et al. (2012). “Methods for Detecting Early Warnings of Critical Transitions in Time Series Illustrated Using Simulated Ecological Data”. *PLoS ONE* 7.7, e41010–20.

Danielsson, Jon, Marcela Valenzuela, and Ilknur Zer (2015). “Learning from History: Volatility and Financial Crises”. Available from: `http://www.riskresearch.org/files/DanielssonValenzuelaZer2015.pdf`. Accessed December 2015.

Drezner, Daniel W (2014). *The System Worked: How the World Stopped Another Great Depression*. New York: Oxford University Press.

Echinque, Federico and Aaron Edlin (2004). “Mixed equilibria are unstable in games of strategic complements”. *Journal of Economic Theory* 118.1, pp. 61–79.

Eichengreen, Barry (2003). “Three generations of crises, three generations of crisis models”. *Journal of International Money and Finance* 22.7, pp. 1089–1094.

Eichengreen, Barry, Andrew Rose, and Charles Wyplosz (1996). “Contagious Currency Crises: First Tests”. *Scandinavian Journal of Economics* 98.4, pp. 463–484.

Frankel, Jeffrey and George Saravelos (2012). “Can leading indicators assess country vulnerability? Evidence from the 2008–09 global financial crisis”. *Journal of International Economics* 87.2, pp. 216 –231.

Frankel, Jeffrey A. and Andrew K. Rose (1996). “Currency crashes in emerging markets: An empirical treatment”. *Journal of International Economics* 41.3-4, pp. 351–366.

Gandrud, Christopher and Mark Hallerberg (2015). "What is a Financial Crisis? Efficiently Measuring Real-Time Perceptions of Financial Market Stress with an Application to Financial Crisis Budget Cycles". *CESIfo Working Paper* 5632. https://www.cesifo-group.de/DocDL/cesifo1_wp5632.pdf.

Goodwin, Jeff (2011). "Why We Were Surprised (Again) by the Arab Spring". *Swiss Political Science Review* 17.4, pp. 452–456.

Helleiner, Eric (2011). "Understanding the 2007–2008 Global Financial Crisis: Lessons for Scholars of International Political Economy". *Annual Review of Political Science* 14.67-87.

Keynes, John M. (1936). *The General Theory of Employment, Interest and Money*. London: Macmillan.

Kindleberger, Charles P. (2000). *Manias, Panics, and Crashes: A History of Financial Crises*. New York: Wiley.

Knight, Frank H. (1921). *Risk, Uncertainty, and Profit*. New York: Houghton Mifflin.

Kuran, Timur (1989). "Sparks and prairie fires: A theory of unanticipated political revolution". *Public Choice* 61.1, pp. 41–74.

— (1991). "Now out of never: The element of surprise in the East European revolution of 1989". *World Politics* 44.1, pp. 7–48.

Laeven, Luc and Fabián Valencia (2013). "Systemic Banking Crisis Database". *IMF Economic Review* 61.2, pp. 225–270.

Leblang, David and Shanker Satyanath (2006). "Institutions, Expectations, and Currency Crises". *International Organization* 60.1, pp. 245–262.

Lepetit, Laetitia and Frank Strobel (2013). "Bank insolvency risk and time-varying Z-score measures". *Journal of International Financial Markets, Institutions and Money* 25, pp. 73–87.

McAdams, Richard H. (2009). "Beyond the Prisoners' Dilemma: Coordination, Game Theory, and Law". *Southern California Law Review* 82, pp. 209–259.

Mesquita, Ethan Bueno de (2014). *Regime Change and Equilibrium Multiplicity*. Working paper, Harris School, University of Chicago.

Mian, Atif and Amir Sufi (2014). *House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent It from Happening Again*. Chicago: The University of Chicago Press.

Minsky, Hyman P (1982). *Can "It" Happen Again?: essays on instability and finance*. New York: Routledge.

Morris, Stephen and Hyun Song Shin (2003). "Global Games: Theory and Applications". In: *Advanced in Economics and Econometrics*. Ed. by M. Dewatripont, L. Hansen, and S. Turnovsky. Vol. 1. Cambridge University Press, pp. 57–114.

Nelson, Stephen C. and Peter J. Katzenstein (2014). "Uncertainty, Risk, and the Financial Crisis of 2008". *International Organization* 68.2, pp. 361–392.

Nickell, Stephen (1981). "Biases in dynamic models with fixed effects". *Econometrica* 49, 1417–1426.

Obstfeld, Maurice (1986). "Rational and Self-fulfilling Balance-of-Payments Crises". *American Economic Review* 76.1, pp. 72–81.

— (1996). "Models of currency crises with self-fulfilling features". *European Economic Review* 40.3-5, pp. 1037–1047.

O'Hagan, Tony (2004). "Dicing with the unknown". *Significance* 1.3, pp. 132–133.

Reinhart, Carmen and Kenneth Rogoff (2009). *This Time is Different: Eight Centuries of Financial Folly*. Princeton: Princeton University Press.

— (2011). *This Time is Different: Eight Centuries of Financial Folly*. Princeton: Princeton University Press.

Rosas, Guillermo (2009). "Dynamic Latent Trait Models: An application to Latin American Banking Crises". *Electoral Studies* 28, pp. 375–387.

Taleb, Nassim Nicholas (2004). *Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets*. New York: Random House.

World Bank (2015). *The Global Financial Development Database*. http://data.worldbank.org/data-catalog/global-financial-development. Accessed December 2015.

Appendix 1 Theoretical Extensions

This appendix discusses several extensions to the main model discussed in the text by allowing for (1) changes in firm strength and (2) the possibility that common shocks are not independent. We also briefly address further extensions that would allow for additional periods, additional players, and correlated player types.

Appendix 1.1 Hard Times and Endogenous Types

There are several ways to model how players' types might evolve over time. One is to allow random variations in player types, such that in period 2, $\sigma_{i,2} = \sigma_{i,1} + \epsilon_i$, where ϵ_i is drawn from some distribution F_ϵ that ensures that $\sigma_{i,2} > \underline{\sigma}$. In this case, assuming that F_ϵ is known to all the players but that the value of ϵ_i is only known to player i in period 2, the main effect is to widen $F_{\sigma,2}$ by extending the range of possible types for $\sigma_{i,2}$. This result is intuitive: the more uncertainty over types in period 2, the less information is conveyed by players' actions in period 1. In the limit, if $\sigma_{i,1}$ is completely uninformative about $\sigma_{i,2}$, then neither the players nor the Observer can make inferences about the probability of a future crisis given present actions.

An alternative is to assume that players' types change over time in response to market conditions, which would imply that future types are endogenous to present common shocks. Put otherwise, what if hard times hurt firm strength? We may model this by assuming that $\sigma_{i,2} = \sigma_{i,1} + \eta$, where η is a "common type shock" that is positively correlated with κ and, as above, is drawn from some distribution F_η that ensures that $\sigma_{i,2} > \underline{\sigma}$. There are several consequences to allowing types to be endogenous to market conditions in this way. First, notice that the posterior distribution of player types $F_{\sigma,2}$ now depends on how much $\sigma^*(\kappa_1)$ shifts relative to η. Depending on the precise relationship between κ and η, this may partially counteract the screening effects identified above, or even completely overwhelm them for the largest common shocks.[16] Second, because negative common shocks will decrease the level of truncation in the distribution $F_{\sigma,2}$ relative to the constant types case, they may, again, partially or even completely overwhelm the confidence effects identified above.

A related way to model changes in players' types is to assume that common type shocks are asymmetrical: good market conditions do not strengthen players, but bad market conditions will weaken them, meaning that $\eta < 0$ iff $\kappa < 0$, otherwise $\eta = 0$. The implication of this approach are that positive shocks will increase uncertainty about the probability of a future crisis, while the effects of the negative shocks depend on $\sigma^*(\kappa_1)$ shifts relative to η. The summary point is that endogenous types can—depending on how they are modeled—counteract some of the dynamics identified above. However, the modeling assumptions matter, and a wide range of approaches exist in which the core predictions of the model survive even when player types are endogenous to market conditions.

Appendix 1.2 Correlated Common Shocks

Another important feature of the baseline model is that common shocks κ are unrelated to one another, drawn independently from the same distribution F_κ. On one hand, this assumption may be seen as too generous by describing future shocks as coming from a known distribution, when that distribution is unknown. On the other hand, it might be too restrictive. It is useful to conceptualize a model of the evolution of market conditions in which good market conditions in the present imply a higher likelihood of good market conditions in the immediate future, even if this remains a probabilistic prediction. In such a model, common shocks are correlated over time.

There are two ways to approach this issue. One is to stipulate that in the model, common shocks are correlated. The other is to stipulate that market actors and regulators *behave as if they believe* that common shocks are correlated, even if it is not true in the model.

[16]Screening effects will be overwhelmed by endogenous type changes when $\eta|(\kappa_1 = \kappa) > \sigma^*(\kappa)|(\kappa_1 = \kappa)$.

This latter view accords with Taleb's (2004) arguments about market actors and regulators alike discounting the likelihood of extreme negative events. In either case, this implies adjusting the baseline model by stipulating that both players and the Observer believe that $\kappa_{t+1} = g(\kappa_t)$, where $g(\cdot)$ is some function that maps realizations of κ in the current period to realizations in the next period.

The consequences of correlated common shocks are straightforward. The greater the correlation between κ_1 and κ_2 produced by $g(\cdot)$, the more certain the Observer is that a positive shock in the current period will be positive in the future. Now, the Observer has two pieces of information through which to make inferences: his updated range of possible player types that he derived from observing players' actions with on κ_1, and also an updated belief about the future values of κ. Depending on both the shape of the distribution F_κ and the function $g(\cdot)$, this can counteract the screening effect by increasing the probability that the next shock will be sufficiently positive enough to induce cooperation relative to the baseline case where common shocks are independently and identically distributed. Consider an example where F_κ is a truncated standard normal distribution, and $g()$ maps this to a normal distribution with the same truncation points, but with mean κ_1 and variance $\left(1 - (abs(\kappa_1)\right)^2$. Figure 7 compares F_κ with two distributions representing two positive shocks κ_1. Both shocks here decrease the likelihood of extreme common shocks in the future, but each also shifts the mass of the probability distribution F_{κ_2} in a positive direction. Whether this will moderate the screening effect by decreasing the variance of $F_{\sigma,2}$, in turn, depends on the shape of F_σ. And of course, this is only one possible way for $g(\cdot)$ to produce correlated common shocks. But it is useful to observe that in the limiting case, where $g(\kappa_1) = \kappa_1$, then the Observer knows with certainty that players will cooperate in next period, and therefore, that the probability of a crisis in that period is zero.

Figure 7: F_{κ_2} with Correlated Common Shocks

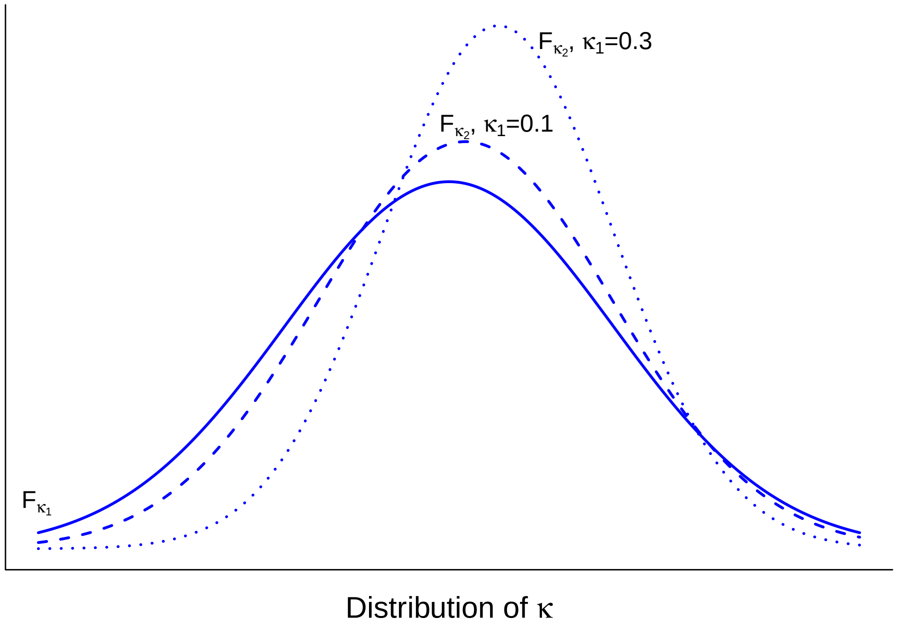

Notes: This figure illustrates one way to incorporate correlated common shocks into the baseline model by showing how the distribution of F_κ changes in response to any realization of κ_1. Here, $F_{\kappa_1} \sim TruncN(0,1)$, and $F_{\kappa_2} \sim TruncN\Big(\kappa_1, \big(1 - abs(\kappa_1)\big)^2\Big)$.

Appendix 1.3 Other Extensions: Repeated Games, N Players, and Correlated Types

Beyond these three extension to the baseline model, there are others that might yield additional results into the ways in which self-fulfilling dynamics affect the predictability of crises. Three in particular are worthy of comment: additional periods in the game, additional players in the game, and correlated player types. I do not analyze these extensions here, but I do suggest here how these extensions may affect the basic results of this model.

1. *N Players* The baseline model includes only two market participants, which is obviously a large departure from a model of a modern economy with multiple anonymous financial sector actors. An extension of the model to an arbitrarily large number of players is possible by defining a crisis not as an outcome where A and B both

sell, but rather as an outcome in which, for a population of size N, proportion π of all players sell. This captures the idea that some number of market participants choosing to see is not a crisis, and i may continue to prefer to hold the asset under these conditions. However, past some threshold (perhaps defined probabilistically) a sufficient number of players selling makes player i prefer to sell. In such a model, rather than confronting a single player j and anticipating whether she Holds or Sells, player i now plays all players $\tilde{}i$, with payoffs defined according to whether or not π players choose Sell.

2. *Repeated Games* The other obvious simplification is that the baseline model contains only two periods. Extending the model to additional time periods is possible just so long as each time period, Assumptions 1 and 2 hold. One interesting consequence is that the history of shocks determines the predictability of a crisis in each period: if a neutral common shock of $\kappa = 0$ in period t_0 were followed by a series of positive shocks $\kappa_{t>0} > 0$ in periods $t_{1 \ldots \tau}$, neither the players nor the Observer would learn anything further in periods $1 \ldots \tau$ about the probability of a crisis in each following period, because the range of possible player types remains $r(\kappa = 0)$.

   Note also that so long as market participants are not engaged in a signaling game, there is no distinction between a finite period game and an infinite period game. Introducing signaling dynamics—by allowing patient players to "bluff" their opponents by holding the asset even when the common shock is negative enough that they would sell—would introduce further complications and additional dynamics. These would warrant separate treatment.

3. *Correlated Types* A final extension is to allow players' types to be correlated. The purpose of allowing correlated types is to entertain the possibility that players assign a higher probability to other market participants being strong types when they themselves are strong types. If player i believes that player j is a stronger type, and j does

the same, then this increases the likelihood that strong types will cooperate.