

Private Sunspots in Games of Coordinated Attack*

Yuliyana Mitkov
University of Bonn
ymitkov@uni-bonn.de

September 9, 2024

Abstract

I endogenize the probability of self-fulfilling outcomes in a game where the only uncertainty comes from extrinsic sunspots. There is a group of players wishing to coordinate on the same action and another player, the regime defender, whose action affects the payoff from coordination. The coordinating players' actions can be based on a sunspot state, which, unlike in the classic sunspot approach, is observed with a small, idiosyncratic noise (a *private sunspot*). I show how private sunspots, combined with the action of the regime defender, can be used to derive a unique coordination probability in any equilibrium where sunspots influence actions. I show how this approach can be used to determine the probability of a sunspot-driven bank run.

Keywords: Coordination problems, sunspots, strategic uncertainty

JEL Codes: D70, D84, G01

*I thank the Co-Editor Pierre-Olivier Weill and two anonymous referees for their helpful comments. I am also thankful to Huberto Ennis, Todd Keister, Jing Zeng, and the participants at the Alghero Workshop on Institutions, Individual Behavior, and Economic Outcomes. Funding by the German Research Foundation (DFG) through CRC TR 224 (Project C03) and EXC 2126/1-390838866 under Germany's Excellence Strategy is gratefully acknowledged.

1 Introduction

Several economic phenomena, such as currency attacks, bank runs, sovereign defaults, and technology adoption, can be understood as collective action games where the players can coordinate on one of two equilibria with very different welfare consequences (Diamond and Dybvig, 1983; Obstfeld, 1996; Calvo, 1988; Katz and Shapiro, 1986; Cole and Kehoe, 2000). Multiple equilibria emerge in those settings because of strategic *complementarities*: the benefit of an action for a player increases with the number of players choosing the same action (Bulow et al. (1985)). Equilibrium multiplicity poses issues since the effect of a given policy on equilibrium outcomes may be indeterminate. One approach to studying games with multiple equilibria is introducing extrinsic uncertainty to the model - a sunspot reflecting agents' sentiment. In this approach, each agent is influenced by the sunspot only because he expects the others to be.¹ However, a shortcoming of the sunspot-based approach is that the probability of coordinating on a specific action is exogenous: it depends on the probability of the associated sunspot state. Stated differently, the standard sunspot approach offers no theoretical rationale for why good outcomes should be correlated with good fundamentals and vice versa (Morris and Shin, 2000).

I propose a systematic way to endogenize the probability of a coordination event within the sunspot-based approach. Specifically, I perturb the original public sunspot game by assuming that each coordinating player receives a signal of the realization of the sunspot state, which is arbitrarily close to the true realization - a *private sunspot*.

I show how and when this private sunspot approach generates a unique probability of

¹Sunspot equilibria originated with Cass and Shell (1983) and has been applied to macroeconomics (Azariadis, 1981; Woodford, 1986), monetary economics Smith (1988), learning Woodford (1990), business cycles Benhabib and Farmer (1994), and bank runs (Peck and Shell, 2003), among others. For an overview, see Shell (1989).

coordinating on a given action when sunspots matter (namely, when players' actions are contingent on sunspots). The general approach, presented in Sections 2 and 3, assumes a continuum of homogeneous coordinating players (the *agents*) taking a binary action (to attack or not) and another player, the regime *defender*, taking a continuous action that affects each agent's net payoff from an attack. The setup is similar to Morris and Shin (2003) with the following important differences. (i) There is no intrinsic (i.e., fundamental) uncertainty. (ii) The defender makes a strategic choice that affects the benefit of the attack for the agents and rules out an attack as the strictly dominant action. (iii) The defender's action is unobservable to the agents when deciding whether to attack. (iv) The defender cannot credibly commit to a specific action but will best respond to the agents' strategies.

As is well-known, the strategies of the coordinating players could be conditioned on a *sunspot state* - an extrinsic random variable. Even though the state is payoff-irrelevant, each coordinating player would base his actions on the realized sunspot if he expects others to do the same. With the usual sunspot-based approach, however, the equilibrium probability of an attack is generally indeterminate and can be any number within an interval. Such a prediction is unsatisfactory and reveals a weakness in that approach. Notice that if the sunspot state is perfectly observed, there is no *strategic uncertainty*: given his signal, each coordinating player can perfectly predict the actions of the others. This is an unappealing assumption since some strategic uncertainty (however small) is likely to persist. The private sunspot approach introduces a small degree of strategic uncertainty: each coordinating player is never sure of the exact private sunspot received by each of the other coordinating players. As a result, the strategies of the coordinating players must satisfy an additional condition, which, together with the defender's actions, pins down the probability of an attack.

The private sunspot approach allows the coordinating players to hold idiosyncratic

sentiments about the prospect of coordinating on a given outcome. One interpretation of this approach is as a modeling device whose goal is to sharpen the predictions of sunspot equilibria. Another is that the correlating device (the sunspot structure) is noisy and unreliable because its realizations cannot be measured precisely or are open to interpretation. This point has been made by Angeletos (2008), who analyses a model with imperfectly observed sunspots (see also Angeletos and La'O (2013)). My approach differs in several ways. First, no regime defender exists in Angeletos (2008), whereas this player is instrumental here. Second, Angeletos (2008) is interested in how private sunspots induce variation in the equilibrium actions, even if all players share the same information about the fundamentals. In contrast, the analysis here perturbs the original game by adding small noise in the player's coordination device. Ex-ante, the coordinating players' probability of choosing different actions is arbitrarily close to zero.

One popular approach to resolving the multiplicity of equilibrium when strategic complementarities are present is the global games approach (Rubinstein, 1989; Carlsson and Van Damme, 1993; Morris and Shin, 1998; Goldstein and Pauzner, 2005). That approach requires the underlying model to have a particular structure, which is not satisfied for the environments I study here. In particular, the fundamental value determining the payoff from attacking is a *strategic* choice of the defender in my model rather than an exogenous random variable. The approach proposed here is appropriate for situations where the defender is *not* a Stackelberg leader, either because he lacks commitment or his actions are unobservable to the coordinating players. This setup is similar to Jann and Schottmüller (2021), where an active defender invests in costly, unobservable defenses. They show that the game has a unique Nash equilibrium where no attack occurs if the number of potential attackers is large enough. My approach differs in several ways. First, they explicitly rule out correlating devices

(such as sunspots), crucial to my analysis. Second, my results do not depend on the number of attackers. Third, the private sunspot approach generates an attack with positive probability as part of equilibrium.

The private sunspot approach has another feature that distinguishes it from global games. Conditional on the fundamentals, a global game selects a particular equilibrium; either an attack happens, or it does not (in the limiting case where the noise goes to zero). The implication is that a *small* change in the fundamental can lead to a large, *discontinuous* change in the outcome. The private sunspots approach, in contrast, assigns a non-trivial probability to each outcome, which captures the idea that the coordination process is somewhat random. Moreover, that probability responds *continuously* to a change in the fundamentals. The approach here can be viewed as endogenously generating the equilibrium selection mechanism advocated by Ennis and Keister (2005a) for analyzing government policy in models with complementarities and multiple equilibria.²

Finally, to illustrate the advantages of the private sunspots approach, I apply the method to a version of the canonical bank runs model of Diamond and Dybvig (1983). Specifically, I use a version of the model in which a policymaker without commitment can intervene and resolve the bank when it faces a run, as in Ennis and Keister (2009, 2010). The policymaker in this framework plays the role of the defender in my framework as it reschedules payments to depositors to achieve ex-post efficiency. If depositors run on the bank in some states, the private sunspots approach delivers a unique equilibrium probability of a run. I show how this probability has natural comparative statics, as changes in economic fundamentals lead to continuous changes in the probability of a run. I conclude by discussing the advantages of my

²In related work, Ennis and Keister (2005b) show how an equilibrium selection mechanism with these general properties can result from an adaptive learning process with boundedly rational agents. In contrast, the approach here is fully consistent with rationality.

approach relative to the global games approach to bank runs pioneered by Goldstein and Pauzner (2005).

2 The Setup

There is a continuum of coordinating players called the *agents*. Each agent chooses an action $a_i \in \{0, 1\}$ where $a_i = 1$ is an attack on the regime.³ There is another player called the regime *defender* choosing an action $\theta \in [\underline{\theta}, \bar{\theta}]$. All agents have the same payoff function $u : \{0, 1\} \times [\underline{\theta}, \bar{\theta}] \times [0, 1] \rightarrow \mathbb{R}$ where $u(a, \theta, \alpha)$ is an agent's payoff if he chooses action a , the defender chooses an action θ , and proportion $\alpha = \int_0^1 a_i di$ of the other agents choose action 1. The defender's payoff is $W : [\underline{\theta}, \bar{\theta}] \times [0, 1] \rightarrow \mathbb{R}$ where $W(\theta, \alpha)$ is his payoff if he chooses an action θ and proportion α of the agents choose action 1. Define $\Delta u : [\underline{\theta}, \bar{\theta}] \times [0, 1] \rightarrow \mathbb{R}$ as the payoff for an agent from choosing action 1 minus the payoff from choosing action 0. That is, $\Delta u(\theta, \alpha) \equiv u(1, \theta, \alpha) - u(0, \theta, \alpha)$. I impose that $\Delta u(\theta, \alpha)$ is bounded, and in addition:

A1: *Complementarities.* $\Delta u(\theta, \alpha)$ is weakly increasing in α for all $\theta \in [\underline{\theta}, \bar{\theta}]$.

A2: *Action monotonicity.* There is $\hat{\alpha} \in (0, 1)$ such that $\Delta u(\theta, \alpha)$ is (a) weakly increasing in θ for $\alpha < \hat{\alpha}$ and (b) strictly increasing in θ for $\alpha > \hat{\alpha}$.

A3: *Continuity.* $\int_0^1 \Delta u(\theta, \alpha) d\alpha$ is continuous in θ .

According to A1, there are strategic complementarities: the incentive for an agent to choose action 1 increases in the proportion of agents choosing that action, namely α .

According to A2, the incentive for an agent to choose action 1 weakly increases in the

³The continuum of agents is for simplicity. The results hold for a finite number of agents, as shown in Section 4.

defender's action θ and strictly increases when α is sufficiently close to one. Thus, higher θ makes an attack more appealing.⁴ Finally, $\int_0^1 \Delta u(\theta, \alpha) d\alpha$ is the expected net benefit from an attack for an agent assigning a uniform probability over the proportion of agents choosing to attack. According to A3, this net benefit varies continuously with the defender's action θ .

Next, recall that the defender's choice of θ is unobservable by the agents, implying that their actions cannot be contingent on the defender's action. They would, of course, infer θ in equilibrium. Thus, a *no-attack* Nash equilibrium exists whenever there is some $\theta_{NA} \in [\underline{\theta}, \bar{\theta}]$ such that (i) $\theta_{NA} \in \operatorname{argmax}_{\theta} W(\theta, 0)$ and (ii) $\Delta u(\theta_{NA}, 0) \leq 0$. According to (i), the defender best responds with θ_{NA} when he expects all agents to choose action 0. According to (ii), each agent best responds with action 0 when the defender picks θ_{NA} and no other agent attacks. Similarly, a *sure-attack* Nash equilibrium exists whenever there is some $\theta_A \in [\underline{\theta}, \bar{\theta}]$ such that (i) $\theta_A \in \operatorname{argmax}_{\theta} W(\theta, 1)$ and (ii) $\Delta u(\theta_A, 1) \geq 0$. According to (i), the defender best responds with θ_A when all agents choose action 1. According to (ii), each agent best responds with action 1 when the defender selects θ_A and all other agents attack. I assume there is at least one Nash equilibrium in pure strategies.

3 Sunspot equilibria

Sunspot equilibria are introduced as follows: nature first draws a payoff-irrelevant random variable s that no one observes. The variable s , with support $S \subseteq \mathbb{R}$ and CDF $F(\cdot)$ is the underlying *sunspot state*. Each agent then privately observes a payoff-irrelevant random variable \hat{s}_i , which, conditional on s , is i.i.d. with support $\hat{S} \subseteq \mathbb{R}$ and CDF $\hat{F}(\cdot|s, \epsilon)$. Henceforth, \hat{s}_i is agent i 's *private sunspot*. The parameter

⁴For example, $\Delta u(\theta, \alpha) = -c$ for $\alpha < \hat{\alpha}$ and $\Delta u(\theta, \alpha) = \theta - c$ for $\alpha \geq \hat{\alpha}$ where $\hat{\alpha} \in (0, 1)$.

$\epsilon \geq 0$ captures the precision of the private sunspots and is normalized so that lower values of ϵ correspond to greater precision. Public sunspots correspond to $\epsilon = 0$. The defender *does not* observe the sunspot state or receive an informative private sunspot. The sunspot structure $\Phi^\epsilon = (S, F, \hat{S}, \hat{F})$ is common knowledge.⁵

Definition 1. *An equilibrium with private sunspots consists of a sunspot structure $\Phi^\epsilon = (S, F, \hat{S}, \hat{F})$, a strategy for each agent $\hat{a}^* : \hat{S} \rightarrow \{0, 1\}$, and an action for the defender θ^* such that*

$$\hat{a}^*(\hat{s}) \in \operatorname{argmax}_{\hat{a} \in \{0, 1\}} \lim_{\epsilon \rightarrow 0} \int_S u(\hat{a}, \theta^*, \alpha^*(s, \epsilon)) dP(s|\hat{s}, \epsilon) \quad (1)$$

and

$$\theta^* \in \operatorname{argmax}_{\theta \in [\underline{\theta}, \bar{\theta}]} \lim_{\epsilon \rightarrow 0} \int_S W(\theta, \alpha^*(s, \epsilon)) dF(s) \quad (2)$$

where $\alpha^*(s, \epsilon) = \int_{\hat{S}} \hat{a}^*(\hat{s}) d\hat{F}(\hat{s}|s, \epsilon)$ is the proportion of agents choosing to attack for given underlying sunspot state s and noise precision ϵ and $P(s|\hat{s}, \epsilon)$ is the CDF of s conditional on private sunspot \hat{s} and noise precision ϵ as implied by Bayes rule.

I focus on the vanishing noise case where the precision of the agents' private sunspots ϵ is arbitrarily close to zero. Specifically, I require each agent's equilibrium strategy to be an optimal choice in the limit as $\epsilon \rightarrow 0$, holding the strategies of all other agents fixed. I further impose the following assumptions on the sunspot structure.

- A4: Sunspot structure.** (i) the sunspot state s is a continuous random variable with density $f(\cdot)$, (ii) agent i 's private sunspot is $\hat{s}_i = s + \epsilon\eta_i$ where η_i are i.i.d. continuous random variables that are independent of

⁵Note that Φ^ϵ is not part of the model primitives but a modeling device. As in Angeletos (2008), one can think of private sunspots as idiosyncratic sentiments due to disagreements, different interpretations, or measurement errors.

the sunspot state. The distribution of the noise terms η_i is $H(\cdot)$ with support $[-b, b] \subseteq \mathbb{R}$ and density $h(\cdot)$.⁶

Finally, sunspots will be said to *matter* if agents' actions are contingent on their private sunspots. Henceforth, I only focus on equilibria where sunspots matter.

3.1 Strategies

Since the sunspot structure is continuous and the noise is arbitrarily small, we can assume without loss of generality that the strategy of each agent will be contingent on M threshold points $\chi_1 \dots \chi_M$ such that each threshold is a switch from one action to another (the action at a threshold can be either of the two). For example, the following is a single-threshold strategy such that an agent attacks if and only if his private sunspot is greater than or equal to χ .

$$\hat{a}(\hat{s}) = \begin{cases} 0 & \text{if } \hat{s} < \chi \\ 1 & \text{if } \hat{s} \geq \chi \end{cases} \quad (3)$$

As the noise in the private sunspots vanishes $\epsilon \rightarrow 0$ either all agents attack with probability $q = 1 - F(\chi)$, or there is no attack with the complement probability, where $F(\cdot)$ is the CDF of the sunspot state. In general, denote by $\mathcal{A} \equiv \{\hat{s} \in \hat{S} : \hat{a}(\hat{s}) = 1\}$ the set of private sunspots leading to an attack. Then, as $\epsilon \rightarrow 0$, either all agents attack with probability $q = P(s \in \mathcal{A})$ or there is no attack with probability $1 - q$ where $P(s \in \mathcal{A})$ is the probability that the sunspot state belongs to \mathcal{A} . The defender's

⁶The support of the sunspot state s and the noise terms η_i can be a bounded interval or the entire real line. In addition, η_i need not be mean zero or even symmetrically distributed.

best response in (2) is then given by

$$\hat{\theta}(q) \in \operatorname{argmax}_{\theta \in [\underline{\theta}, \bar{\theta}]} (1 - q)W(\theta, 0) + qW(\theta, 1). \quad (4)$$

A5: *Defender's action.* $\hat{\theta}(q)$ is unique for each $q \in [0, 1]$. The function $\hat{\theta}(q)$ is continuous and strictly decreasing in q .

A6: *Laplacian action monotonicity.* $\int_0^1 \Delta u(\hat{\theta}(0), \alpha) d\alpha > 0 > \int_0^1 \Delta u(\hat{\theta}(1), \alpha) d\alpha$.

According to A5, $\hat{\theta}(q)$ decreases in q , which, in turn, reduces the agents' incentive to attack since $\Delta u(\theta, \alpha)$ increases in θ . Next, recall that an agent's beliefs are said to be *Laplacian* when he assigns a uniform probability distribution over the proportion of agents that attack, i.e., $\alpha \sim U[0, 1]$ (see, e.g., Morris and Shin (2003)). According to A6, an agent with the Laplacian beliefs would attack when $q = 0$ since, in that case, the defender's best response $\hat{\theta}(0)$ is relatively high, making attacking the preferable action. On the other hand, an agent with these same beliefs will not attack if $q = 1$ since the defender's best response, in that case, $\hat{\theta}(1)$ is relatively low, making no attack the preferable action.

3.2 Public sunspots

I first examine the more familiar type of sunspot equilibria where the sunspot state s is observed perfectly by the agents, corresponding to setting $\epsilon = 0$ in A4. Consequently, there is no strategic uncertainty in that case: each agent (after observing the sunspot) knows what the other agents will do.

Proposition 1. *Public sunspots.* *Suppose A1 - A3 and A5 - A6 are satisfied. Suppose each agent perfectly observes the sunspot state. Then there are q_1 and q_2 ($q_1 < q_2$) such that and for any $q \in [q_1, q_2] \in [0, 1]$ there exists a sunspot equilibrium with an*

attack probability equal to q .

This proposition illustrates a critical shortcoming of the public sunspot approach - the probability of an attack can be any value in the interval $[q_1, q_2]$ and, in particular, there is a continuum of equilibria where the attack probability is strictly between zero and one.⁷ The public sunspot approach does not provide a way of selecting among these equilibria or linking the probability of an attack to the model's primitives.

3.3 Private sunspots

I now show that the model yields much sharper predictions if the agents observe the sunspot state with a small noise. In that case, as long as agents' actions are contingent on sunspots, the equilibrium attack probability is unique and can be directly related to the model's primitives.

Proposition 2. *Private sunspots. Suppose A1 - A6 are satisfied. Suppose the agents observe the sunspot state with a vanishing noise, and their actions are contingent on their private sunspots. Then, the equilibrium attack probability q^* is unique and given as the solution of*

$$\int_0^1 \Delta u(\hat{\theta}(q^*), \alpha) d\alpha = 0. \quad (5)$$

Proof. Assume each agent follows the single-threshold strategy in (3) and attacks if and only if his private sunspot is greater than or equal to χ^* , where χ^* will be determined in equilibrium. Recall from Definition 1 that I require each agent's strategy to be an optimal choice in the limit as the noise in the private sunspots vanishes, holding the strategies of the other agents fixed. Hence, I need to establish two things. First, as $\epsilon \rightarrow 0$, each agent best responds by an attack (no attack) if his private sunspot

⁷The lower bound q_1 is zero whenever the no-attack Nash equilibrium exists, and the upper bound q_2 is one whenever the sure-attack Nash equilibrium exists.

is greater (smaller) than χ^* (thus χ^* is an equilibrium). Second, the corresponding equilibrium attack probability q^* must satisfy condition (5) (hence χ^* is unique).

In particular, as $\epsilon \rightarrow 0$, the CDF of the proportion of agents that attack converges to a distribution with mass q^* on full attack and mass $1 - q^*$ on no attack where q^* is the probability that the sunspot is greater than or equal to χ^* . That is, $P(s \geq \chi^*)$, where since $F(\cdot)$ is the CDF of the sunspot state, we have

$$q^* \equiv P(s \geq \chi^*) = 1 - F(\chi^*).$$

I will show that q^* must satisfy the condition in (5). From A4 agent i 's private sunspot is $\hat{s}_i = s + \epsilon\eta_i$. Then, if the sunspot state is s , the proportion of agents that observe a private sunspot greater than or equal to χ^* is

$$P(s + \epsilon\eta_i \geq \chi^* | s, \epsilon) = 1 - H\left(\frac{\chi^* - s}{\epsilon}\right),$$

where $H(\cdot)$ is the CDF of the noise term η_i . We then get for $\alpha \in [0, 1]$:

$$1 - H\left(\frac{\chi^* - s}{\epsilon}\right) \leq \alpha \quad \Leftrightarrow \quad s \leq \chi^* - \epsilon H^{-1}(1 - \alpha). \quad (6)$$

Denote by $G(\alpha | \hat{s}, \epsilon)$ the probability that an agent with a private sunspot \hat{s} assigns to the event that the proportion of agents that attack is at most α . From (6) we get for $\alpha \in [0, 1]$:

$$G(\alpha | \hat{s}, \epsilon) = P(s \leq \chi^* - \epsilon H^{-1}(1 - \alpha) | \hat{s}, \epsilon). \quad (7)$$

Denote by $\Pi(\hat{s}, \epsilon)$ the net expected payoff of an agent with a private sunspot equal

to \hat{s} . We have

$$\Pi(\hat{s}, \epsilon) = \int_S \Delta u(\hat{\theta}(q^*), \alpha(s, \epsilon)) dP(s|\hat{s}, \epsilon) = \int_0^1 \Delta u(\hat{\theta}(q^*), \alpha) dG(\alpha|\hat{s}, \epsilon).$$

The threshold strategy in (3) is consistent equilibrium if and only if $\lim_{\epsilon \rightarrow 0} \Pi(\hat{s}, \epsilon) \leq 0$ for $\hat{s} < \chi^*$ and $\lim_{\epsilon \rightarrow 0} \Pi(\hat{s}, \epsilon) \geq 0$ for $\hat{s} > \chi^*$. That is, as the noise vanishes, each agent best responds by an attack for $\hat{s} < \chi^*$ and by no attack for $\hat{s} > \chi^*$, implying that an agent with $\hat{s} = \chi^*$ will be indifferent between the two actions

$$\Pi(\chi^*, \epsilon) = \lim_{\epsilon \rightarrow 0} \int_0^1 \Delta u(\hat{\theta}(q^*), \alpha) dG(\alpha|\chi^*, \epsilon) = 0.$$

where $G(\alpha|\chi^*, \epsilon)$ is the probability that an agent with a private sunspot equal to the threshold χ^* assigns to the event that the proportion of agents that attack is at most α . Thus, q^* will be determined by the equation in (5) if the CDF $G(\cdot|\chi^*, \epsilon)$ converges to $U[0, 1]$ as the noise in the private sunspots vanishes $\epsilon \rightarrow 0$. That is, we must show

$$\lim_{\epsilon \rightarrow 0} G(\alpha|\chi^*, \epsilon) = \alpha \quad \text{for all } \alpha \in [0, 1].$$

First, recall that the density of the sunspot state s is $f(\cdot)$, and the density of the noise term η_i is $h(\cdot)$. We can then express $G(\alpha|\hat{s}, \epsilon)$ as follows:

$$G(\alpha|\hat{s}, \epsilon) = \int_{-\infty}^{\chi^* - \epsilon H^{-1}(1-\alpha)} p(s|\hat{s}, \epsilon) ds = \frac{\int_{-\infty}^{\chi^* - \epsilon H^{-1}(1-\alpha)} f(s) h\left(\frac{\hat{s}-s}{\epsilon}\right) ds}{\int_{-\infty}^{\infty} f(s) h\left(\frac{\hat{s}-s}{\epsilon}\right) ds}, \quad (8)$$

where $p(s|\hat{s}, \epsilon)$ is the posterior of the sunspot state s as implied by Bayes rule

$$p(s|\hat{s}, \epsilon) = \frac{f(s) h\left(\frac{\hat{s}-s}{\epsilon}\right)}{\int_{-\infty}^{\infty} f(s) h\left(\frac{\hat{s}-s}{\epsilon}\right) ds}.$$

I now proceed along the lines of Morris and Shin (2003). In particular, implementing a change of variable $z = \frac{\hat{s}-s}{\epsilon}$ into (8) yields

$$G(\alpha|\hat{s}, \epsilon) = \frac{\int_{(\hat{s}-\chi^*)/\epsilon+H^{-1}(1-\alpha)}^{\infty} f(\hat{s}-\epsilon z)h(z)dz}{\int_{-\infty}^{\infty} f(\hat{s}-\epsilon z)h(z)dz}. \quad (9)$$

To characterize the limit of the above CDF as ϵ goes to zero, we need to consider three cases.

Case 1: $\hat{s} > \chi^*$. For $\alpha \in (0, 1)$ the limit of integration in (9) goes to infinity as ϵ goes to zero. Thus, for all $\alpha \in (0, 1)$:

$$\lim_{\epsilon \rightarrow 0} G(\alpha|\hat{s}, \epsilon) = 0.$$

Also, since $G(\cdot|\hat{s}, \epsilon)$ is a CDF $\lim_{\epsilon \rightarrow 0} G(0|\hat{s}, \epsilon) = 0$ and $\lim_{\epsilon \rightarrow 0} G(1|\hat{s}, \epsilon) = 0$, implying that the limiting distribution in that case assigns all probability mass on $\alpha = 1$.

Case 2: $\hat{s} = \chi^*$. Here ϵ does not enter the limit of integration in (9). Then, as $\epsilon \rightarrow 0$, $f(\hat{s}-\epsilon z)$ converges to $f(\hat{s})$ and, as a result, $f(\cdot)$ drops out from the expression in (9) and we get

$$\lim_{\epsilon \rightarrow 0} G(\alpha|\chi^*, \epsilon) = \frac{\int_{H^{-1}(1-\alpha)}^{\infty} h(z)dz}{\int_{-\infty}^{\infty} h(z)dz} = 1 - H(H^{-1}(1-\alpha)) = \alpha. \quad (10)$$

As the noise vanishes, the CDF $G(\cdot|\chi^*, \epsilon)$ converges to the uniform distribution $U[0, 1]$. Stated differently, an agent with a private sunspot equal to the threshold assigns a uniform probability over the proportion of agents that attack.

Case 3: $\hat{s} < \chi^*$. For $\alpha \in (0, 1)$ the limit of integration in (9) goes to minus infinity as ϵ goes to zero. We then get $\lim_{\epsilon \rightarrow 0} G(\alpha|\hat{s}, \epsilon) = 1$ for all $\alpha \in [0, 1]$ implying that the limiting distribution in that case assigns all probability mass on $\alpha = 0$.

By combining the above three cases we get $\lim_{\epsilon \rightarrow 0} G(\alpha | \hat{s}, \epsilon) \preceq \lim_{\epsilon \rightarrow 0} G(\alpha | \chi^*, \epsilon)$ for all $\hat{s} < \chi^*$ and $\lim_{\epsilon \rightarrow 0} G(\alpha | \hat{s}, \epsilon) \succeq \lim_{\epsilon \rightarrow 0} G(\alpha | \chi^*, \epsilon)$ for all $\hat{s} > \chi^*$ where \preceq denotes *first-order stochastic dominance*. Then, since $\Delta u(\hat{\theta}(q^*), \alpha)$ is weakly increasing in α (from A1), it follows that $\lim_{\epsilon \rightarrow 0} \Pi(\hat{s}, \epsilon) \leq \lim_{\epsilon \rightarrow 0} \Pi(\chi^*, \epsilon) = 0$ for $\hat{s} < \chi^*$ and $\lim_{\epsilon \rightarrow 0} \Pi(\hat{s}, \epsilon) \geq \lim_{\epsilon \rightarrow 0} \Pi(\chi^*, \epsilon) = 0$ for $\hat{s} > \chi^*$. Consequently, the strategy in (3) with a threshold point $\chi^* = F^{-1}(1 - q^*)$ where q^* solves the equation in (5) is consistent with equilibrium.

It remains to show that q^* exists and is unique. It will be useful to define the function:

$$\varphi(q) \equiv \int_0^1 \Delta u(\hat{\theta}(q), \alpha) d\alpha.$$

From A3 $\int_0^1 \Delta u(\theta, \alpha) d\alpha$ is continuous in θ whereas, from A5, $\hat{\theta}(q)$ is continuous q . Thus, $\varphi(q)$ will be continuous in q . Also, let q_1 and q_2 be such that $q_1 < q_2$. From A5 $\hat{\theta}(q)$ is strictly decreasing in q hence $\hat{\theta}(q_1) > \hat{\theta}(q_2)$ whereas from A2 $\Delta u(\hat{\theta}, \alpha)$ is weakly increasing in θ for $\alpha < \hat{\alpha}$ and strictly increasing for $\alpha > \hat{\alpha}$ where $\hat{\alpha} \in (0, 1)$. Hence, $\Delta u(\hat{\theta}(q_1), \alpha) \geq \Delta u(\hat{\theta}(q_2), \alpha)$ with strict inequality whenever $\alpha > \hat{\alpha}$. We then have

$$\begin{aligned} \varphi(q_1) - \varphi(q_2) &= \int_0^1 \left(\Delta u(\hat{\theta}(q_1), \alpha) - \Delta u(\hat{\theta}(q_2), \alpha) \right) d\alpha \\ &\geq \int_{\hat{\alpha}}^1 \left(\Delta u(\hat{\theta}(q_1), \alpha) - \Delta u(\hat{\theta}(q_2), \alpha) \right) d\alpha > 0. \end{aligned}$$

The function $\varphi(q)$ is thus continuous and strictly decreasing in q . Finally, from A6: $\varphi(0) > 0 > \varphi(1)$. Then, by the intermediate value theorem, there exists a unique $q^* \in (0, 1)$ such that $\varphi(q^*) = 0$.

The reverse case, where each agent attacks if and only if his private sunspot is less than or equal to a given threshold, is treated analogously and is omitted. Finally, the multiple-threshold case requires some additional technical details and is relegated to the Appendix.

□

4 Discussion

The private sunspot approach implies that the equilibrium in which the agents' actions are contingent on sunspots is unique and generates a probability of an attack that responds continuously to a change in the parameters. It thus captures the intuitive idea that the coordination process is somewhat random and cannot be perfectly predicted given the fundamentals (see, e.g., Ennis and Keister (2005a,b)).

4.1 Defender's role

The private sunspot approach is not applicable if θ is not determined as part of the equilibrium. Specifically, the equilibrium value of the attack probability q^* leads the defender to choose $\theta^* = \hat{\theta}(q^*)$ such that an agent with the Laplacian beliefs about the proportion of other agents that attack (i.e., $\alpha \sim U[0, 1]$) is indifferent between an attack and no attack $\int_0^1 \Delta u(\theta^*, \alpha) d\alpha = 0$. Many applications naturally feature a player whose action affects the benefit of an attack for the coordinating players. For example, the central bank will choose the level of reserves in a currency attack model, the bank will choose its early payment in a bank run model, and the incumbent will choose the amount spent on defense in a regime change model.

4.2 The discrete case

It is not hard to adapt the analysis to a discrete number of agents. To illustrate, assume there are two agents and a defender. An agent has a payoff of zero if he does not attack, a payoff of $-c$ if he is the only one attacking, and a payoff of $\psi(\theta)$ if both agents attack where $\psi(\cdot)$ is some strictly increasing function of θ . For simplicity, the sunspot structure is taken to be uniform: $s \sim U[0, 1]$ and $\hat{s}_i = s + \epsilon\eta_i$ where

$\eta_i \sim U[-1, 1]$.⁸ Assume each agent attacks if and only if his private sunspot is greater than or equal to χ and, as before, focus on vanishing noise $\epsilon \rightarrow 0$. Thus, if agent i 's private sunspot is $\hat{s}_i = \chi$, the probability he assigns to the other agent's private sunspot \hat{s}_j being greater than or equal to χ will be $\frac{1}{2}$. Then, the equilibrium probability of an attack q^* is determined as the solution of the following equation

$$\frac{1}{2}\psi(\hat{\theta}(q^*)) = c \quad \text{or} \quad \hat{\theta}(q^*) = \psi^{-1}(2c), \quad (11)$$

where $\hat{\theta}(q)$ is the defender's best response given an attack with probability q . The solution q^* exists, is unique, and decreasing in c , whenever $\hat{\theta}(0) > \psi^{-1}(2c) > \hat{\theta}(1)$ and $\hat{\theta}(q)$ is continuous and strictly decreasing in q . In contrast, all values of q such that $\hat{\theta}(q) \geq \psi^{-1}(c)$ will be consistent with a public-sunspot equilibrium.

4.3 Higher-order beliefs

The approach to establishing the uniqueness of the private sunspot equilibrium is reminiscent of the global game's literature (Carlsson and Van Damme, 1993). It is known that the global games selection rule is not robust to perturbations of higher-order beliefs (Weinstein and Yildiz, 2007). Here, I show that similar issues plague the private sunspot approach.⁹ Take the two-agent example from the previous subsection. Assume the sunspot state is uniform $s \sim U[0, 1]$ as before, but now for given s with probability λ , each agent's private sunspot is s and with probability $1 - \lambda$ each agent's private sunspot is drawn independently, from $U[s - \epsilon, s + \epsilon]$. Then, for any $\epsilon > 0$, agent i 's assigns probability $\lambda + \frac{1}{2}(1 - \lambda)$ that agent j 's private sunspot is weakly higher $P(\hat{s}_j \geq \hat{s}_i) = \frac{1}{2}$. The equilibrium attack probability q^* is then determined as

⁸Then, conditional on \hat{s}_i , agent's i 's posterior about the sunspot state is $s|\hat{s}_i \sim U[\hat{s}_i - \epsilon, \hat{s}_i + \epsilon]$ and his posterior about the private sunspot of the other agent is $\hat{s}_j|\hat{s}_i \sim U[\hat{s}_i - 2\epsilon, \hat{s}_i + 2\epsilon]$.

⁹I thank an anonymous referee for raising this point.

the solution of

$$(\lambda + \frac{1}{2}(1 - \lambda))\psi(\hat{\theta}(q^*)) = c.$$

Notice that q^* now depends on the sunspot structure through the parameter $\lambda \in [0, 1]$. The only case such that q^* is independent of the sunspot structure is $\lambda = 0$ corresponding to the private sunspot selection rule. The private sunspot method is thus appropriate for applications where sunspots matter (i.e., agents' actions are contingent on them), but one wishes to model the probability of an attack as independent of the sunspot structure and, instead, dependent only on the fundamentals.

5 An Application

This section demonstrates the versatility of the private sunspots approach by applying it to the limited commitment version of the Diamond and Dybvig (1983) model in Ennis and Keister (2009, 2010). Most Diamond-Dybvig models assume full commitment by the bank and policymakers to a course of action, even in a run. However, it is well-known that a bank with commitment can eliminate runs by promising to suspend payments as soon as a run is detected (see, e.g., Diamond and Dybvig (1983)).¹⁰ At the same time, the actions taken during financial crises are often characterized by delayed response and only partial suspensions. Ennis and Keister (2009, 2010) show how such delays arise naturally under limited commitment and how the anticipation of delays can cause runs.¹¹ However, runs in their setup are either unanticipated

¹⁰The literature generates runs in environments with commitment by assuming banks must give a pre-specified payment until they run out of funds Postlewaite and Vives (1987); Cooper and Ross (1998); Allen and Gale (2004); Goldstein and Puzner (2005). In practice, however, bank liabilities are frequently altered in a crisis (see Ennis and Keister (2009)).

¹¹The limited commitment approach to the Diamond-Dybvig framework has been used to study a range of topics, including how financial fragility is affected by interest rates (Li, 2017), inequality (Mitkov, 2020), asset opacity (Izumi, 2021), competition (Gao and Reed, 2021), and banking regulation (Keister and Mitkov, 2023).

or their probability is exogenously based on sunspots. I will show how the private sunspots approach is natural in this setup and delivers a unique run probability that depends on the parameters in an intuitive way. I then relate my analysis to the global games approach to Diamond-Dybvig that builds on Goldstein and Pauzner (2005).¹²

5.1 The environment

There are three time periods $t = 0, 1, 2$. There is a continuum of agents, called the *depositors*, and indexed by $i \in [0, 1]$. Each depositor has preferences given by

$$u(c_1 + \omega_i c_2) = \frac{(c_1 + \omega_i c_2)^{1-\gamma}}{1-\gamma}, \quad (12)$$

where c_t is consumption in period t and ω_i is a binomial random variable with support $\Omega = \{0, 1\}$. As is standard, the coefficient of relative risk aversion is greater than one $\gamma > 1$. If $\omega_i = 0$, the depositor is *impatient* and values consumption only in period 1, whereas if $\omega_i = 1$, he is *patient* and values consumption in periods $t = 1, 2$. Each depositor learns his type privately in period 1. Each depositor is impatient with probability π , and the fraction of impatient depositors is also π .

Technology. Each depositor is endowed with one unit of the good in period 0 and there is a constant-returns-to-scale technology for transforming goods in period 0 into goods in periods 1 and 2. One unit of the good placed in this technology in period 0 yields $R > 1$ units in period 2 but only 1 unit in period 1.

Sequential service. There is banking technology that allows depositors to pool their endowment to insure against idiosyncratic liquidity risks. As in Wallace (1988), the

¹²Gu (2011) considers asymmetric observations of the sunspot state in a setup very different from mine. In particular, the bank in Gu (2011) has full commitment, and there are multiple sunspot-based equilibria.

depositors are isolated from each other, and no trade can occur among them. Each depositor can visit the banking technology to receive a payment from the pooled resources (i.e., to withdraw). Those choosing to withdraw in period 1 arrive one at a time and must consume immediately upon arrival. This sequential service constraint implies the payment to a depositor can only depend on the information available to the banking technology when this payment is being determined.¹³

Banking authority. A benevolent *banking authority* (BA) operates the banking technology. The BA anticipates that a fraction π of the depositors will be impatient and withdraw in period 1 and always act to maximize the depositors' expected utilities. Importantly, as in Ennis and Keister (2009, 2010), the BA cannot commit to actions that are not ex-post optimal and instead chooses payments as a best response given its information and taking as given the profile of withdrawal strategies of the depositors. As will become clear, the BA is the defender.

Timing. In period 0, all endowments are deposited. At the start of period 1, the depositors are isolated from each other. After observing his type ω_i and private sunspot \hat{s}_i , depositor i chooses to contact the banking technology in period 1 or 2. Those depositors choosing to withdraw in period 1 arrive at the bank in the order given by their index i . Thus, depositor $i = 0$ knows that he has the opportunity to be the first to withdraw in period 1 whereas depositor $i = 1$ knows that his opportunity comes last. The depositor's position in this order is private information.¹⁴ The BA

¹³If there is no sequential service, the bank would first collect all withdrawal requests and then assign payments. Then, if payments can be made contingent on that information, runs will not occur as part of equilibrium since if all depositors request early payment, the bank's best response is to give 1 to each in period 1. But then, each patient depositor prefers to leave his share in the bank to get a larger payment in period 2.

¹⁴The approach here follows Green and Lin (2003) and Ennis and Keister (2010) among others since it simplifies the analysis while capturing in a tractable way the notion that the depositors may have some information about their position in the withdrawal order. The results will be very similar if the depositors first choose when to withdraw and are then randomly assigned positions in the withdrawal order.

determines the payment to each depositor as they arrive and as a best response to the situation. In particular, the BA detects a run as soon as period 1 withdrawals exceed π and, as explained below, would reschedule payments for the remaining depositors to reflect this new information.

The efficient allocation. To derive a benchmark allocation, suppose a benevolent planner observes all depositors' types and controls their withdrawal actions. The planner gives c_1^* in period 1 to each impatient depositor and c_2^* in period 2 to each patient depositor. These payments will be chosen to solve

$$\max_{\{c_1, c_2\}} \pi u(c_1) + (1 - \pi)u(c_2), \quad (13)$$

subject to $(1 - \pi)c_2 = R(1 - \pi c_1)$. The planner's solution satisfies $1 < c_1^* < c_2^* < R$. As is well-known, this solution can be implemented as an equilibrium by a bank that does not observe depositors' types. In particular, the bank pays c_1^* to each of the first π depositors in period 1. In period 2, the bank's remaining resources mature and are evenly divided among the remaining depositors.

5.2 Sunspot equilibria

Consider now the decentralized economy, where each depositor chooses his withdrawal strategy as part of a non-cooperative game, and the BA chooses payments as a best response to the strategy profile for the depositors.

Sunspot structure. Fix a sunspot structure $\Phi^\epsilon = (S, F, \hat{S}, \hat{F})$ and suppose the depositors observe their private sunspots before withdrawals begin. A strategy for depositor i is a mapping from his realized type ω_i and his private sunspot \hat{s}_i to a decision of whether to withdraw in period 1 or period 2. The analysis is simplified by

the following observations. First, each impatient depositor has a strictly dominant strategy to withdraw in period 1, implying that the measure of withdrawals in period 1 will be at least π . Second, the BA detects a run as soon as the measure of withdrawals exceeds π - but not before that since at least π withdrawals always happen. Finally, one can show that any run in this setup is necessarily partial and restricted to those patient depositors with an opportunity to withdraw before the BA detects a run, namely those with an index $i \leq \pi$. So, consider the following strategy profile for the depositors. (i) Each impatient depositor withdraws early. (ii) Each patient depositor with $i \leq \pi$ withdraws early (i.e., runs on the bank) if and only if his private sunspot falls in the attack set $A \subseteq \hat{S}$. (iii) Each patient depositor with an index $i > \pi$ withdraws late.

Next, I will derive the BA's best response to this strategy profile and then apply the private sunspot approach to derive the equilibrium run probability.

Remaining payments. Denote by $\alpha \in [0, 1]$ the fraction of patient depositors among the first π to contact the bank, implying that the measure of depositors that run on the bank is $\alpha\pi$. Also, let $\hat{\pi}_\alpha$ denote the fraction of the remaining $1 - \pi$ depositors that are impatient. We have $\frac{d\hat{\pi}_\alpha}{d\alpha} > 0$, $\hat{\pi}_0 = 0$, and $\hat{\pi}_1 = \pi$.¹⁵ The BA does not observe α but instead makes inferences based on the flow of withdrawals. If withdrawals stop at π , the BA infers there is no run ($\alpha = 0$), and all remaining depositors are patient. The BA would then give $\hat{c}_{2NR} = R(\frac{1-\pi c_1}{1-\pi})$ in period 2 to each patient depositor implying that the sum of expected utilities of the remaining depositors is

$$V_{NR} \left(\frac{1-\pi c_1}{1-\pi} \right) = (1 - \pi)u \left[R \left(\frac{1-\pi c_1}{1-\pi} \right) \right]. \quad (14)$$

¹⁵The exact expression for $\hat{\pi}_\alpha$ is derived in the proof of Proposition 3

On the other hand, if withdrawals continue after the first π , the BA infers a run is happening (i.e., $\alpha > 0$), and, therefore, not all impatient depositors have been served. Payments for the remaining depositors will then be rescheduled to (c_{1R}, c_{2R}) depending on the BA's updated belief about the distribution of α . Specifically, since the BA's prior on the event $\{\alpha \in (0, 1)\}$ goes to zero as the noise in the depositor's private sunspots vanishes, Bayes rule implies that the BA's posterior on the event $\{\alpha \in (0, 1)\}$ also goes to zero. Then, the rescheduled payments in a run (c_{1R}, c_{2R}) converge to the solution of the program

$$V_R\left(\frac{1-\pi c_1}{1-\pi}\right) = \max_{\{\hat{c}_{1R}, \hat{c}_{2R}\}} \pi u(\hat{c}_{1R}) + (1-\pi)u(\hat{c}_{2R}) \quad (15)$$

$$\text{s.t. } \pi \hat{c}_{1R} + (1-\pi)\hat{c}_{2R}/R = \frac{1-\pi c_1}{1-\pi}. \quad (16)$$

After a run is detected, the BA would set \hat{c}_{1R} and \hat{c}_{2R} to maximize the sum of expected utilities of the remaining depositors given that $\alpha = 1$. Then, for given $\alpha \in (0, 1]$, the measure of remaining impatient depositors is $(1-\pi)\hat{\pi}_\alpha$ and each of those is paid \hat{c}_{1R} whereas the measure of the remaining patient depositors is $(1-\pi)(1-\hat{\pi}_\alpha)$ and each is paid $\hat{c}_{2\alpha}$ where

$$\hat{c}_{2\alpha} = \frac{R(1-\pi c_1 - (1-\pi)\hat{\pi}_\alpha \hat{c}_{1R})}{(1-\pi)(1-\hat{\pi}_\alpha)}. \quad (17)$$

Thus, given c_1 and α , the late payment for each remaining patient depositor, after incorporating the response of the BA, will be determined as in the above. The function $\hat{c}_{2\alpha}$ is equal to \hat{c}_{2NR} for $\alpha = 0$, strictly decreasing in α , and equal to \hat{c}_{2R} for $\alpha = 1$. Then, since the solution of the program in (15) - (16) satisfies $\hat{c}_{2R} > \hat{c}_{1R}$ we get $\hat{c}_{2\alpha} > \hat{c}_{1R}$ for all $\alpha \in [0, 1]$. Next, each of the remaining patient depositors would best respond by withdrawing late - as specified in the depositors' strategy profile. Finally, I turn to the incentive of the depositors, who have a chance to withdraw before the

BA detects a run.

Equilibrium run probability. Recall that all impatient depositors would withdraw early, whereas all patients with a chance to withdraw after the BA has detected a run (i.e., $i > \pi$) would withdraw late. So, consider a patient depositor with an opportunity to withdraw before the BA has detected a run (i.e., $i \leq \pi$). For given c_1 and α this depositor gets a payoff of $u(c_1)$ from running on the bank and a payoff of $u(\hat{c}_{2\alpha})$ from waiting where $\hat{c}_{2\alpha}$ is given in (17). Thus, the net payoff from a run for such a patient depositor for given c_1 and α is

$$\Delta u(c_1, \alpha) = u(c_1) - u(\hat{c}_{2\alpha}). \quad (18)$$

The BA does not learn new information during the first π withdrawals and, since depositors are risk-averse, would give a common amount c_1 to each of the first π depositors. Then, once the measure of withdrawals reaches π , one of two things can happen. First, withdrawals stop, and the bank's remaining matured resources are spread equally among the patient depositors. Second, withdrawals continue, and the BA immediately reschedules payments to \hat{c}_{1R} and \hat{c}_{2R} to solve the program in (15). Notice that all remaining payments are uniquely pinned down once c_1 and α have been determined. As the noise in the private sunspots goes to zero, we have $\alpha = 1$ with probability q or $\alpha = 0$ with probability $1 - q$. Given any $q \in [0, 1]$, the BA selects c_1 to maximize the sum of expected utilities of all depositors, anticipating how it reacts after a run is detected. That is,

$$c_1(q) \in \operatorname{argmax}_{c_1 \in [0, 1/\pi]} \pi u(c_1) + (1 - \pi) \left[(1 - q) V_{NR} \left(\frac{1 - \pi c_1}{1 - \pi} \right) + q V_R \left(\frac{1 - \pi c_1}{1 - \pi} \right) \right]. \quad (19)$$

where $V_{NR} \left(\frac{1 - \pi c_1}{1 - \pi} \right)$ and $V_R \left(\frac{1 - \pi c_1}{1 - \pi} \right)$ are defined in (14) and (15) respectively.

Proposition 3. *Suppose the depositors observe the sunspot state with a vanishing noise. Then the equilibrium bank run probability q^* (when sunspots matter i.e., $q^* \in (0, 1)$) is determined as the solution of the following equation*

$$\int_0^1 \Delta u(c_1(q^*), \alpha) d\alpha = 0,$$

where $\Delta u(c_1, \alpha) d\alpha$ is defined in (18) and $c_1(q)$ in (19).


Figure 1 displays q^* as a function of the coefficient of relative risk aversion γ (left panel), the return of the investment technology R (middle panel), and the aggregate proportion of impatient depositors π (right panel).¹⁶ There are several notable implications. First, a higher coefficient of relative risk aversion γ leads to a higher q^* since it pushes the bank to engage in more maturity transformation by setting a higher early payment. In particular, a higher early payment implies that the bank will be in worse financial shape when a run is discovered (after π withdrawals), which implies that the payments for depositors who wait will be lower. This fact creates a stronger incentive to run, thus pushing up the equilibrium run probability. Second, R has a non-monotone effect on q^* . As R increases, the bank will give a larger early payment. However, there are two competing effects on the expected late payment and, thus, on the return to staying invested. First, if there is no run, the late payment will increase relative to the early payment. Second, when a run occurs, more investment will be liquidated (due to the larger early payments), reducing the late relative to the early payment. As the middle panel in Figure 1 shows, the second effect dominates for relatively small values of R , whereas the first effect dominates for relatively large.¹⁷

¹⁶I set $\pi = 0.5$ in the left and the middle panel, $R = 2$ in the left and the right panel, and $\gamma = 5$ in the middle and the right panel.

¹⁷The non-monotonicity in R reflects forces similar to those identified by Li (2017) by arbitrarily focusing on the largest possible sunspot-based run probability.

Finally, π also has a non-monotone effect on q^* . Recall that π in this setup also measures how many depositors will get to withdraw before the bank discovers and responds to a run - a large π implies that the response to a run comes later when the bank is in worse financial shape. This force also appears in models with public sunspots (see, e.g., Ennis and Keister (2010) and Keister (2016)) where fragility always increases in π . The private sunspots approach introduces strategic uncertainty, bringing a second competing effect. A depositor who receives the threshold signal will be relatively more optimistic about the value of waiting because he anticipates that some agents may not be participating in the run. The magnitude of this effect is increasing in π because the bank reacts more strongly to a run when π is large. This second effect, therefore, decreases the incentive to run as π becomes larger. The right panel in Figure 1 shows that, for this example, the first effect dominates when π is small, but the second effect dominates as π becomes larger, eventually pushing the equilibrium run probability to zero.

Figure 1: Bank run probability



5.3 Relation to Goldstein and Pauzner (2005)

It is informative to conclude this section by comparing the approach here to the large literature that follows Goldstein and Pauzner (2005) in changing the Diamond-Dybvig model to fit into the global games framework of Carlsson and Van Damme (1993).

Some of these changes are questionable on theoretical grounds.

First, Goldstein and Pauzner (2005) assumes the bank continues to pay withdrawing depositors at face value until it is out of resources. They need to take this approach because having a fixed contract allows them to use global game techniques. This, however, is not how a financial crisis usually unfolds in practice, as discussed at the start of this section. I show how a more realistic version of the Diamond and Dybvig (1983) model, one where the bank can adjust payments as the crisis unfolds, fits naturally into the private sunspots framework.

Second, Goldstein and Pauzner (2005) assumes that the bank is fully liquid in some states of the world, meaning it could repay all of its depositors immediately, at face value, and without defaulting. This is another strong assumption required for the global game techniques to work but is not required by the private sunspots method.

Third, the private sunspots approach delivers more intuitively appealing comparative statics. Conditional on fundamentals, the framework based on Goldstein and Pauzner (2005) predicts the probability of a run is either 0 or 1 (as the noise vanishes). In contrast, in my case, the probability of a run is either zero or strictly between 0 and 1, and it varies smoothly with the parameters. For example, Cipriani et al. (2024) shows that in the 2023 regional banking crisis in the U.S., some banks experienced runs while others with similar fundamentals did not. This observation is consistent with the private sunspots model with vanishing noise but not with vanishing noise in global games.

Finally, global games tie fundamental *and* strategic uncertainty. But suppose (as in the model from this section) that the bank can react to an incipient run by changing the payment schedule. As the noise in the depositors' private signals vanishes, a banker who is well-informed about the fundamentals can also accurately

predict whether the depositors would run and, if necessary, cut payments to preserve resources. Such pre-emptive actions generally imply that the equilibrium run probability collapses to zero. To avoid this conclusion, one must assume that either (i) the bank must rely on less precise information about its fundamentals than the investors or (ii) the bank is not allowed to act on interim information. These are unappealing assumptions.¹⁸

6 Conclusion

I propose a method to endogenize the probability of self-fulfilling outcomes based on sunspots, which I call the private sunspots approach. The framework is a non-cooperative game where each agent chooses between two actions: to attack or not. A regime defender also participates, taking action in response to the agents' strategies, which in turn affects the agents' payoffs. Although the equilibrium attack probability is self-fulfilling, the private sunspot approach pins down this probability as a function of the model's parameters. This is possible because the approach introduces strategic uncertainty: each agent observes the sunspot state with vanishing noise. Consequently, the agents' strategies must satisfy an additional equilibrium condition, which, combined with the defender's actions, uniquely determines the equilibrium attack probability. This approach is particularly well-suited for applications where the defender makes strategic choices without commitment, as exemplified by the bank run scenario discussed in the previous section.

¹⁸The private sunspots approach assumes the bank is unaware of investor sentiment and must make inferences based on withdrawal demand (as in [Ennis and Keister \(2009, 2010\)](#)). This is perhaps more plausible than the bank having less precise information about its fundamentals than the depositors.

Appendix

Proof of Proposition 1.

Since the sunspot state is perfectly observed, the agents' strategies can be directly contingent on it. So, assume each agent attacks if and only if $s \geq \chi$. The probability of an attack is then $q = 1 - F(\chi)$ where $F(\cdot)$ is the CDF of the sunspot state. Notice that focusing on a single-threshold strategy is without loss of generality since one can generate any $q \in (0, 1)$ through $\chi = F^{-1}(1 - q)$. An attack with probability q will be consistent with a public sunspot equilibria whenever the following is satisfied

$$\Delta u(\hat{\theta}(q), 0) \leq 0 \leq \Delta u(\hat{\theta}(q), 1), \quad (20)$$

where $\hat{\theta}(q)$ is the defender's best response to an attack probability of q . The first inequality in (20) ensures that each agent best responds by not attacking when all other agents do not attack $\alpha = 0$, whereas the second that each agent best responds with an attack when all other agents attack $\alpha = 1$. Next, define

$$\varphi(q) \equiv \int_0^1 \Delta u(\hat{\theta}(q), \alpha) d\alpha. \quad (21)$$

From A2 - A3 $\int_0^1 \Delta u(\theta, \alpha) d\alpha$ is continuous and strictly increasing in θ whereas, from A5, $\hat{\theta}(q)$ is continuous and strictly decreasing in q . Hence, $\varphi(q)$ is continuous and strictly decreasing in q . Further, from A6 we have $\varphi(0) > 0 > \varphi(1)$. Then, by the intermediate value theorem, there exists (in this case a unique) $q^* \in (0, 1)$ such that $\varphi(q^*) = 0$. That is,

$$\varphi(q^*) = \int_0^1 \Delta u(\hat{\theta}(q^*), \alpha) d\alpha = 0.$$

Next, since $\Delta u(\hat{\theta}(q^*), \alpha)$ is non-decreasing in α we get from the above:

$$\Delta u(\hat{\theta}(q^*), 0) \leq 0 \leq \Delta u(\hat{\theta}(q^*), 1).$$

Hence, q^* will be consistent with a public sunspot equilibrium. Then, since $\Delta u(\theta, 0)$ is non-decreasing in θ whereas $\hat{\theta}(q^*)$ is non-increasing in q , the following set of conditions will be satisfied. First, if the probability of an attack is greater than q^* and no other agent attacks $\alpha = 0$, then each agent best responds by not attacking.

$$\Delta u(\hat{\theta}(q), 0) \leq \Delta u(\hat{\theta}(q^*), 0) \leq 0 \quad \text{for each } q \in (q^*, 1] \quad (22)$$

Second, if the probability of an attack is less than q^* and all other agents attack $\alpha = 1$, then each agent best responds by attacking.

$$\Delta u(\hat{\theta}(q), 1) \geq \Delta u(\hat{\theta}(q^*), 1) \geq 0 \quad \text{for each } q \in [0, q^*) \quad (23)$$

Letting q_1 (q_2) denote the smallest (largest) value of q consistent with public-sunspot equilibrium, we have $q_1 \leq q^* \leq q_2$ and, to establish that there is a continuum of public-sunspot equilibria, we must show that either $q_1 < q^*$ or $q_2 > q^*$. Recall that either the no-attack or the sure-attack Nash equilibrium is assumed to exist. If the no-attack Nash equilibrium exist we have $\Delta u(\hat{\theta}(0), 0) \leq 0$. The previous implies $\Delta u(\hat{\theta}(q), 0) \leq 0$ for each $q \in [0, 1]$ which, combined with (22), implies that each $q \in [0, q^*]$ is a public-sunspot equilibrium (i.e., $q_1 = 0$). On the other hand, if the sure-attack Nash equilibrium exists, we have $\Delta u(\hat{\theta}(1), 1) \geq 0$. Then, $\Delta u(\hat{\theta}(q), 1) \geq 0$ for each $q \in [0, 1]$ which, combined with (23), implies that each $q \in [q^*, 1]$ is a public-sunspot equilibrium (i.e., $q_2 = 1$).

Proof of Proposition 2.

The main text deals with the single-threshold case, and it remains to establish the multi-threshold case. So, consider a strategy containing $M > 1$ thresholds $\chi_1^* < \chi_2^* < \dots < \chi_M^*$. I consider the equilibrium as ϵ goes to zero while the number of thresholds M remains fixed. Also, suppose that each agent attacks whenever his private sunspot is less than or equal to the lowest of those thresholds. That is,

$$\hat{a}(\hat{s}) = 1 \quad \text{for all } \hat{s} \leq \chi_1^*. \quad (24)$$

Hence, $\hat{a}(\hat{s}) = 0$ for $\hat{s} \in (\chi_1^*, \chi_2^*)$, $\hat{a}(\hat{s}) = 1$ for $\hat{s} \in [\chi_2^*, \chi_3^*]$, and so on. The flip case $\hat{a}(\hat{s}) = 0$ for $\hat{s} \leq \chi_1^*$ is very similar and omitted. I will show that q^* will still be determined by the equation in (5) as in the single-threshold case. Consider any agent i with a private sunspot \hat{s} and define the random variables $Y(\hat{s})$, $X(\hat{s}, y)$, and $X(\hat{s})$ as follows:

- $Y(\hat{s}) = y$ is the realized proportion of agents with a sunspot less than or equal to agent i 's private sunspot \hat{s} . Let $\Psi(y|\hat{s})$ denote the CDF of $Y(\hat{s})$.
- $X(\hat{s}, y) = x$ is the realized proportion of agents that attack given that agent i 's private sunspot is \hat{s} and y is the fraction of agents with a private sunspot less than or equal to \hat{s} . Let $Q_y(x|\hat{s})$ denote the CDF of $X(\hat{s}, y)$.
- $X(\hat{s}) = x$ is the realized proportion of agents that attack given that agent i 's private sunspot is \hat{s} . Let $Q(x|\hat{s})$ denote the CDF of $X(\hat{s})$.

Notice that the CDF $\Psi(\cdot|\hat{s})$ depends on the sunspot structure but not on the agents' strategies, whereas the CDFs $Q_y(x|\hat{s})$ and $Q(x|\hat{s})$ depend on the sunspot structure and the agents' strategies. In particular, $Q(x|\hat{s})$ is a mixture distribution where $Q_y(x|\hat{s})$

are the component CDF and $Y(\hat{s})$ is the mixing variable with CDF $\Psi(y|\hat{s})$. That is,

$$Q(x|\hat{s}) = \int_0^1 Q_y(x|\hat{s}) d\Psi(y|\hat{s}) \quad \text{for } x \in [0, 1]. \quad (25)$$

Also, notice that (25) is an equivalent way of defining the CDF $G(\alpha|\hat{s})$ in (8). The net expected payoff of an agent whose private sunspot is equal to \hat{s} is then

$$\Pi(\hat{s}) = \int_0^1 \Delta u(\hat{\theta}(q^*), x) dQ(x|\hat{s}) \quad \text{for } \hat{s} \in \hat{S}.$$

If the strategy in (24) is consistent with equilibrium, an agent with a private sunspot equal to one of those thresholds will be indifferent between an attack and no attack. That is, $\Pi(\chi_k^*) = 0$ for all $k \in \{1, \dots, M\}$. I will show that as $\epsilon \rightarrow 0$, $\Psi(\cdot|\hat{s})$ converges to $U[0, 1]$ for all $\hat{s} \in \hat{S}$. That is,

$$\lim_{\epsilon \rightarrow 0} \Psi(y|\hat{s}) = y \quad \text{for all } y \in [0, 1] \text{ and all } \hat{s} \in \hat{S}.$$

Indeed, if the underlying sunspot is s , the proportion of agents who observe a private sunspot less than or equal to \hat{s} is

$$P(s + \epsilon\eta_i \leq \hat{s}) = H\left(\frac{\hat{s}-s}{\epsilon}\right),$$

where $H(\cdot)$ is the CDF of the noise term η_i . This proportion will be less than or equal to $y \in [0, 1]$ if and only if $s \geq \hat{s} - \epsilon H^{-1}(y)$. In other words, agent i with private sunspot \hat{s} assigns probability

$$P(s \geq \hat{s} - \epsilon H^{-1}(y) | \hat{s}) = \Psi(y|\hat{s})$$

to the event that the proportion of agents with a private sunspot less than or equal to his is less no greater than y . Then

$$\Psi(y | \hat{s}) = \int_{\hat{s} - \epsilon H^{-1}(y)}^{\infty} p(s | \hat{s}) ds = \frac{\int_{\hat{s} - \epsilon H^{-1}(y)}^{\infty} f(s) h(\frac{\hat{s} - s}{\epsilon}) ds}{\int_{-\infty}^{\infty} f(s) h(\frac{\hat{s} - s}{\epsilon}) ds},$$

where $h(\cdot)$ is the PDF of the noise terms η_i . Changing the variable of integration in the above to $z = \frac{\hat{s} - s}{\epsilon}$ yields:

$$\Psi(y | \hat{s}) = \frac{\int_{-\infty}^{H^{-1}(y)} f(\hat{s} - \epsilon z) h(z) dz}{\int_{-\infty}^{\infty} f(\hat{s} - \epsilon z) h(z) dz}.$$

As $\epsilon \rightarrow 0$, $f(\hat{s} - \epsilon z)$ converges to $f(\hat{s})$ and the density $f(\cdot)$ drops out from $\Psi(y | \hat{s})$. Then, for each $y \in [0, 1]$ we get as $\epsilon \rightarrow 0$

$$\Psi(y | \hat{s}) = \frac{\int_{-\infty}^{H^{-1}(y)} h(z) dz}{\int_{-\infty}^{\infty} h(z) dz} = H(H^{-1}(y)) = y.$$

Hence, $\lim_{\epsilon \rightarrow 0} \Psi(y | \hat{s}) = y$ for $y \in [0, 1]$ and $\hat{s} \in \hat{S}$ implying that as the noise vanishes the CDF $\Psi(\cdot | \hat{s})$ converges to $U[0, 1]$. Then, since $Y(\hat{s}) \sim \Psi(\cdot | \hat{s})$ converges in distribution to $U[0, 1]$ for all \hat{s} the mixture CDF in (25) becomes

$$\lim_{\epsilon \rightarrow 0} Q(x | \hat{s}) = \int_0^1 Q_y(x | \hat{s}) dy \quad \text{for all } x \in [0, 1].$$

Next, fix the lowest threshold point χ_1^* and recall that $\hat{a}(\hat{s}) = 1$ for all $\hat{s} \leq \chi_1^*$. The strategy profile in (24) thus implies $P(X(\chi_1^*, y) \geq y)$ for all $y \in [0, 1]$. In other words, $X(\chi_1^*, y)$ is obtained from $Y(\chi_1^*)$ by adding a non-negative random variable to each realization of $Y(\chi_1^*)$. Hence, $P(X(\chi_1^*) \geq Y(\chi_1^*)) = 1$ which implies $X(\chi_1^*) \succeq Y(\chi_1^*)$ where \succeq means first-order stochastic dominance (see, e.g., Section 6 in Mas-Colell

et al. (1995)). Then since $\Delta u(\hat{\theta}(q^*), y)$ is non-decreasing in y :

$$\int_0^1 \Delta u(\hat{\theta}(q^*), x) dQ(x|\chi_1^*) \geq \int_0^1 \Delta u(\hat{\theta}(q^*), y) dy. \quad (26)$$

where the above uses $Y(\hat{s}) \sim U[0, 1]$ for all \hat{s} . Also, let χ_k^* for $k \in \{2, \dots, M\}$ be any threshold point greater than χ_1^* . The strategy profile in (24) then implies $P(y \geq X(\chi_k^*, y)) = 1$ for all $y \in [0, 1]$. Hence, $P(Y(\chi_k^*) \geq X(\chi_k^*)) = 1$. Hence, $Y(\chi_k^*) \succeq X(\chi_k^*)$ which, in turn, implies

$$\int_0^1 \Delta u(\hat{\theta}(q^*), y) dy \geq \int_0^1 \Delta u(\hat{\theta}(q^*), x) dQ(x|\chi_k^*). \quad (27)$$

Combining (26) - (27) with $\Pi(\chi_k^*) = 0$ for all k yields:

$$0 = \int_0^1 \Delta u(\hat{\theta}(q^*), x) dQ(x|\chi_1^*) \geq \int_0^1 \Delta u(\hat{\theta}(q^*), y) dy \geq \int_0^1 \Delta u(\hat{\theta}(q^*), x) dQ(x|\chi_k^*) = 0. \quad (28)$$

The first line follows from $\Pi(\chi_1^*) = 0$, the second line from (26), the third line from (27), and the final line from $\Pi(\chi_k^*) = 0$. We then get from (28) that the following must be satisfied in equilibrium:

$$\int_0^1 \Delta u(\hat{\theta}(q^*), y) dy = 0.$$

That is, the equilibrium attack probability is still determined by (5) as in the single-threshold case. In other words, the number of thresholds in the agents' strategy does not affect the equilibrium attack probability since the equilibrium location of those thresholds will be such that the equilibrium attack probability corresponds to its value in the single-threshold case.

Proof of Proposition 3.

I first show that equilibrium runs will be restricted to patient depositors with an opportunity to withdraw before the BA detects a run. To see why, suppose withdrawals continue beyond the first π , leading the BA to infer that a run is happening. Then, if the BA anticipates some fraction ρ of the remaining depositors to withdraw in period 1, it reschedules the payment plan for the remaining depositors $(\hat{c}_{1\rho}, \hat{c}_{2\rho})$ maximize the sum of their expected utilities

$$\rho u(\hat{c}_{1\rho}) + (1 - \rho)v(\hat{c}_{2\rho})$$

subject to the budget constraint

$$\rho \hat{c}_{1\rho} + (1 - \rho)\hat{c}_{2\rho}/R = \frac{1 - \pi c_1}{1 - \pi}.$$

The solution of this program will be characterized by the above budget constraint and the following first-order condition

$$u'(\hat{c}_{1\rho}) = Ru'(\hat{c}_{2\rho}).$$

Since $R > 1$, we have $\hat{c}_{1\rho} < \hat{c}_{2\rho}$ for all $\rho \in [0, 1]$ implying that a patient depositor with an opportunity to withdraw after the BA detects a run strictly prefers to wait. Next, the late payment to each patient depositors $\hat{c}_{2\alpha}$, defined in (17), depends on c_1 , \hat{c}_{1R} and $\hat{\pi}_\alpha$ where $\hat{\pi}_\alpha$ is obtained from the depositors' strategy profile and given by

$$\hat{\pi}_\alpha = \frac{\pi}{1 - \pi} \left(1 - \frac{\pi}{\pi + \alpha(1 - \pi)} \right).$$

Thus $\hat{\pi}_0 = 0$ for $\alpha = 0$, $\hat{\pi}_\alpha$ strictly increases in α , and $\hat{\pi}_1 = \pi$. Next, since the depositors' utility function u is of the constant relative risk aversion form, the solution of the program in (15) - (16) is given by

$$\hat{c}_{1R} = \frac{1-\pi c_1}{(1-\pi)(\pi+(1-\pi)R^{(1-\gamma)/\gamma})} \quad \text{and} \quad \hat{c}_{2R} = R^{1/\gamma} \frac{1-\pi c_1}{(1-\pi)(\pi+(1-\pi)R^{(1-\gamma)/\gamma})}.$$

We have $\hat{c}_{1R} < \hat{c}_{2R}$ and, in addition, \hat{c}_{1R} and \hat{c}_{2R} are decreasing functions of c_1 . Thus, $\hat{c}_{2\alpha}$ defined in (17), equals \hat{c}_{2NR} for $\alpha = 0$, is strictly decreasing in α , and equals \hat{c}_{2R} for $\alpha = 1$. Also, $\hat{c}_{2\alpha}$ is strictly decreasing in c_1 . Hence, A1 and A2 hold since $\Delta u(c_1, \alpha) = u(c_1) - u(\hat{c}_{2\alpha})$ is strictly increasing in α and c_1 . Next, the optimal choice of c_1 in (19) is characterized by the following first-order condition

$$u'(c_1) = (1 - q)Ru'(\hat{c}_{2NR}) + qRu'(\hat{c}_{2R}). \quad (29)$$

Given that u is constant relative risk aversion we can solve for c_1 in terms of q to get

$$c_1(q) = \frac{1}{\pi + A(q)^{1/\gamma}} \quad \text{where} \quad A(q) = (1 - q)\lambda_0^{-\gamma} + q\lambda_1^{-\gamma}$$

where λ_0 and λ_1 are such that $\lambda_0 > \lambda_1$ and given by

$$\lambda_0 = \frac{1}{(1-\pi)R^{(1-\gamma)/\gamma}} \quad \text{and} \quad \lambda_1 = \frac{1}{(1-\pi)(\pi+(1-\pi)R^{(1-\gamma)/\gamma})}.$$

Observe that A5 holds since $c_1(q)$ is unique for each $q \in [0, 1]$ and decreasing in q . Finally, as $q \rightarrow 1$, we get from (29) that $u'(c_1(q))$ converges to $Ru'(\hat{c}_{2R})$ implying $c_1 < \hat{c}_{2R}$. Then, $\hat{c}_{2R} \leq \hat{c}_{2\alpha}$ for all $\alpha \in [0, 1]$ implies $c_1(q) < \hat{c}_{2\alpha}$ for all $\alpha \in [0, 1]$. As a result, $\int_0^1 [u(c_1(q)) - u(\hat{c}_{2\alpha})] d\alpha < 0$ for all q sufficiently close to one implying that the equilibrium run probability will be strictly less than one $q^* < 1$ (that is, q^* is then

either 0 or strictly between 0 and 1).

References

- Allen, Franklin and Douglas Gale (2004) “Competition and financial stability,” *Journal of Money, Credit and Banking*, 453–480.
- Angeletos, George-Marios (2008) “Private sunspots and idiosyncratic investor sentiment,” Technical report, National Bureau of Economic Research.
- Angeletos, George-Marios and Jennifer La’O (2013) “Sentiments,” *Econometrica*, 81 (2), 739–779.
- Azariadis, Costas (1981) “Self-fulfilling prophecies,” *Journal of Economic Theory*, 25 (3), 380–396.
- Benhabib, Jess and Roger EA Farmer (1994) “Indeterminacy and increasing returns,” *Journal of Economic Theory*, 63 (1), 19–41.
- Bulow, Jeremy I, John D Geanakoplos, and Paul D Klemperer (1985) “Multimarket oligopoly: Strategic substitutes and complements,” *Journal of Political Economy*, 93 (3), 488–511.
- Calvo, Guillermo A (1988) “Servicing the public debt: The role of expectations,” *The American Economic Review*, 647–661.
- Carlsson, Hans and Eric Van Damme (1993) “Global games and equilibrium selection,” *Econometrica*, 989–1018.
- Cass, David and Karl Shell (1983) “Do sunspots matter?” *Journal of Political Economy*, 91 (2), 193–227.
- Cipriani, Marco, Thomas M Eisenbach, and Anna Kovner (2024) “Tracing Bank Runs in Real Time,” *FRB of New York Staff Report* (1104).
- Cole, Harold L and Timothy J Kehoe (2000) “Self-fulfilling debt crises,” *The Review of Economic Studies*, 67 (1), 91–116.
- Cooper, Russell and Thomas W Ross (1998) “Bank runs: Liquidity costs and investment distortions,” *Journal of Monetary Economics*, 41 (1), 27–38.
- Diamond, Douglas W and Philip H Dybvig (1983) “Bank runs, deposit insurance, and liquidity,” *Journal of Political Economy*, 91 (3), 401–419.
- Ennis, Huberto M and Todd Keister (2005a) “Government policy and the probability of coordination failures,” *European Economic Review*, 49 (4), 939–973.

- (2005b) “Optimal fiscal policy under multiple equilibria,” *Journal of Monetary Economics*, 52 (8), 1359–1377.
- (2009) “Bank runs and institutions: The perils of intervention,” *American Economic Review*, 99 (4), 1588–1607.
- (2010) “Banking panics and policy responses,” *Journal of Monetary Economics*, 57 (4), 404–419.
- Gao, Jiahong and Robert R Reed (2021) “Sunspot bank runs and fragility: The role of financial sector competition,” *European Economic Review*, 139, 103877.
- Goldstein, Itay and Ady Pauzner (2005) “Demand–deposit contracts and the probability of bank runs,” *Journal of Finance*, 60 (3), 1293–1327.
- Green, Edward J and Ping Lin (2003) “Implementing efficient allocations in a model of financial intermediation,” *Journal of Economic Theory*, 109 (1), 1–23.
- Gu, Chao (2011) “Noisy sunspots and bank runs,” *Macroeconomic Dynamics*, 15 (3), 398–418.
- Izumi, Ryuichiro (2021) “Opacity: Insurance and fragility,” *Review of Economic Dynamics*, 40, 146–169.
- Jann, Ole and Christoph Schottmüller (2021) “Regime change games with an active defender,” *Games and Economic Behavior*, 129, 96–113.
- Katz, Michael L and Carl Shapiro (1986) “Technology adoption in the presence of network externalities,” *Journal of Political Economy*, 94 (4), 822–841.
- Keister, Todd (2016) “Bailouts and financial fragility,” *The Review of Economic Studies*, 83 (2), 704–736.
- Keister, Todd and Yuliyana Mitkov (2023) “Allocating losses: Bail-ins, bailouts and bank regulation,” *Journal of Economic Theory*, 210, 105672.
- Li, Yang (2017) “Interest rates and financial fragility,” *Journal of Economic Dynamics and Control*, 82, 195–205.
- Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green et al. (1995) *Microeconomic Theory*, 1: Oxford university press New York.
- Mitkov, Yuliyana (2020) “Inequality and financial fragility,” *Journal of Monetary Economics*, 115, 233–248.
- Morris, Stephen and Hyun Song Shin (1998) “Unique equilibrium in a model of self-fulfilling currency attacks,” *American Economic Review*, 587–597.

- (2000) “Rethinking multiple equilibria in macroeconomic modeling,” *NBER Macroeconomics Annual*, 15, 139–161.
- (2003) “Global games: Theory and applications.”
- Obstfeld, Maurice (1996) “Models of currency crises with self-fulfilling features,” *European Economic Review*, 40 (3-5), 1037–1047.
- Peck, James and Karl Shell (2003) “Equilibrium bank runs,” *Journal of Political Economy*, 111 (1), 103–123.
- Postlewaite, Andrew and Xavier Vives (1987) “Bank runs as an equilibrium phenomenon,” *Journal of Political Economy*, 95 (3), 485–491.
- Rubinstein, Ariel (1989) “The Electronic Mail Game: Strategic Behavior Under” Almost Common Knowledge”,” *The American Economic Review*, 385–391.
- Shell, Karl (1989) “Sunspot equilibrium,” *The New Palgrave: General Equilibrium*, New York, London, WW Norton, 274–280.
- Smith, Bruce D (1988) “Legal restrictions,” sunspots,” and Peel’s bank act: the real bills doctrine versus the quantity theory reconsidered,” *Journal of Political Economy*, 96 (1), 3–19.
- Wallace, Neil (1988) “Another attempt to explain an illiquid banking system: With sequential service taken seriously,” *Quarterly Review*.
- Weinstein, Jonathan and Muhamet Yildiz (2007) “A structure theorem for rationalizability with application to robust predictions of refinements,” *Econometrica*, 75 (2), 365–400.
- Woodford, Michael (1986) “Stationary sunspot equilibria in a finance constrained economy,” *Journal of Economic Theory*, 40 (1), 128–137.
- (1990) “Learning to believe in sunspots,” *Econometrica*, 277–307.