Falling Dominoes:
A Theory of Rare Events and Crisis Contagion*

HENG CHEN
University of Hong Kong
WING SUEN
University of Hong Kong

December 15, 2014

Abstract. Crises, such as revolutions and currency attacks, rarely occur; but when they do they typically arrive in waves. The rarity of crises itself is an important contagion mechanism in a multiple-country dynamic global game model. When players are uncertain about the true model of the world, observing a rare success elsewhere can substantially change their expectations concerning the payoffs from attacking or defending the regime. Such dramatic revisions in beliefs, amplified by strategic complementarity in actions, may lead to a series of attacks in other countries. The crisis period can be long-lasting, but will eventually come to an end.

Keywords. spread of revolutions, crisis contagion, model uncertainty, confidence collapse

JEL Classification. D74, D83, D84, G01

*Chen: School of Economics and Finance, University of Hong Kong, Pokfulam, Hong Kong (email: hengchen@hku.hk); Suen: School of Economics and Finance, University of Hong Kong, Pokfulam, Hong Kong (email: wsuen@econ.hku.hk.). We thank Tiancheng Chen, Shangguan Ruo, and Fanqi Shi for their excellent research assistance.

1. Introduction

Crises are typically infrequent events but they tend to be contagious, sometimes even among countries or regions that are seemingly unrelated. These two salient features are particularly evident for currency attacks, mass revolutions, and bank runs. In this paper we offer a theory to explain the link between the rarity of crises and the contagion phenomenon.

Eichengreen, Rose, and Wyplosz (1996a) offer systematic empirical evidence which corroborates these features for currency crises. With a panel of quarterly data covering twenty industrial economies from 1959 to 1993, they can only identify a small sample of 77 crises out of 2800 total observations.[1] They estimate a binary probit model and find robust evidence for contagion: controlling for a wide range of macroeconomic variables, the existence of a currency attack elsewhere raises the probability of an attack on the domestic currency.[2] Figure 1(a) shows the number of currency attacks in each quarter during the operation of European Monetary System (EMS).[3] Three distinct clusters of attacks are evident in this figure.

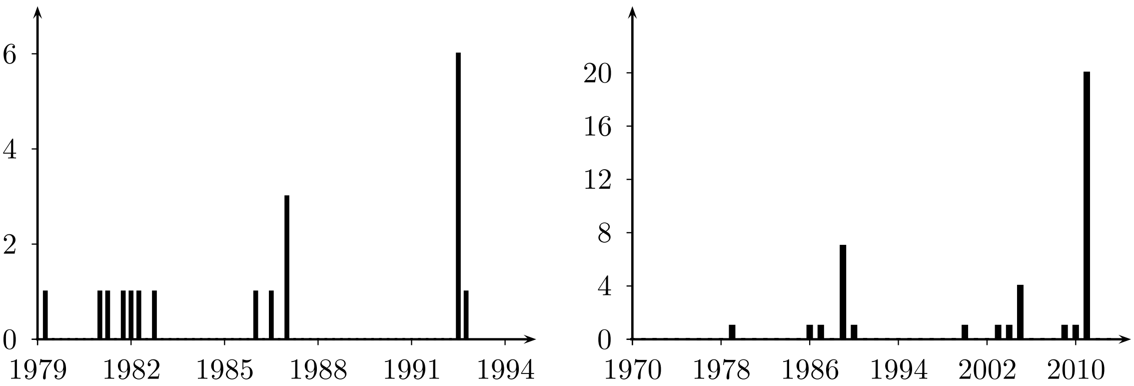

(a) Currency crises in EMS

(b) Mass revolutions

Figure 1. *Rarity and Contagion.*

Revolutions as a mass movement are also a relatively rare occurrence (Walt 1992), and they typically arrive in waves (Katz 1997).[4] An investigation of revolutions after 1970 reveals that most of them clustered in three waves: the collapse of communism,

[1] They compile a speculative pressure index, which is a weighted average of percentage changes in exchange rate, foreign reserve and interest rate. Currency attacks are identified as quarters in which the index is at least one and a half standard deviations above the sample mean.

[2] Other empirical evidence and anecdotes for contagion of speculative attacks abound. See, for example, Bordo and Murshid (2000).

[3] EMS (Stage I) is chosen because it has well defined operating time and member states. Crisis data are obtained from Eichengreen, Rose, and Wyplosz (1996b).

[4] Following Kuran (1989), political revolution is defined as political regime change brought about in a short period of time through a massive action against the status quo with popular participation.

the "color revolutions," and the Arab Spring. See Figure 1(b).[5] A similar pattern is also evident for bank runs.[6]

Another interesting observation about contagion is that the spread of crisis is not necessarily confined to countries with close economic linkages or political ties. For instance, it is difficult to justify the transmission of crisis from Russia to Brazil or from Mexico to Argentina with any real linkage between these countries (Bordo and Murshid 2000, Krugman 1999). Similarly, "revolutionary spirit" often ripples across countries that seem to be totally unconnected. For example, the chain of revolutions in Eastern Europe in 1990 spurred mass movements and led to political changes in Africa.[7] Strikingly, the call for "Jasmine Revolution" that started in Tunisia and inspired Arab countries even echoed in China, where it stirred protests in a handful of major cities, including Beijing and Shanghai.[8]

In this paper, we develop a theory to explain why rare crises tend to cluster and why crises can spread among countries that are seemingly unconnected either economically, financially, or politically. Toward this end, we construct a multiple-country dynamic global game model, which is general enough to address the similar pattern for speculative attacks, political revolutions, and bank runs.[9]

The key element of our theory is that people are uncertain about how the world works, but they can learn it from observing equilibrium outcomes in other countries. This feature of "model uncertainty" is captured by assuming that players maintain two hypothetical descriptions of the world, but cannot be completely sure about which one is true. One of the worlds is tranquil: winning a battle against autocratic regimes or central banks is particularly difficult. The other world is frantic: overthrowing the regime is relatively easy and successes are commonplace.

Suppose that the true world is indeed tranquil and players believe so with a probability close to one. In this situation, success takes place with a very low probability. When a rare success is actually observed elsewhere, the little doubt harbored by these

[5] We perform the "runs test" to test the randomness or serial dependence of elements in this sequence. We find the null hypothesis that these observations are independent can be rejected at the five percent significance level. To perform the test conservatively, we set all the positive values in the sequence to 1. If we start the period of investigation from 1980 or exclude the period of the Arab Spring, we can still reject the null hypothesis. We thank Nandiyang Zhang for making this data set available to us.

[6] In the United States, the median time interval separating nationwide banking panics was 16 years in the pre-Federal Reserve era, and even fewer panics existed in the post-Federal Reserve period (Jalil 2013). It was rather unusual to observe a standalone bank run that did not affect other banks. In fact, the commonly used working definition of banking panics is a cluster of bank runs and suspensions.

[7] See "Wind of change, but a different one," *The Economist*, 14 July 1990.

[8] See "Calls for a 'Jasmine Revolution' in China persist," *The New York Times*, 23 February 2011.

[9] Morris and Shin (1998) develop the global game framework to address issues related to currency crises. The framework has been adapted to model political regime changes (Edmond 2013; Chen, Lu, and Suen 2014) and bank runs (Goldstein and Pauzner 2005).

Bayesian players that the true world might be frantic grows dramatically and their confidence in the world being tranquil collapses. As a consequence, citizens in the second country will be more aggressive in attacking the regime—not only because they feel that success is more likely, but also because they think their fellow citizens will reason in the same way. Anticipating a massive attack, the regime believes it is less worthwhile to defend itself. Another collapse in the second country would deepen this crisis further: the regime in the third country collapses with an even higher probability. However, given the true world is tranquil in nature, the law of large numbers eventually ends the "frantic period," and public confidence recovers or revolutionary euphoria fades as the number of failed attacks increases over time.

We should emphasize that our theory spotlights the role of rarity: only rare crises can shake the public confidence dramatically. If successful attacks are common, this learning mechanism does not produce any quantitatively sizable contagion effect. Furthermore, the contagion mechanism we describe here is effective even for countries with uncorrelated fundamentals, as long as people believe that the same model of the world applies. Citizens believe that a rare crisis in one country may signal systematic weaknesses in other countries that bear institutional similarities to the originating one, because these countries are supposed to operate in a similar fashion with the same rules of the game.[10]

While much has been written about the contagion of currency crisis, there is little research on why revolutions often occur in waves. Political scientists have noticed the need for a contagion model to explain the spread of revolution (Francisco 1993). Our contagion model can be applied to both situations. Kuran (1997) argues that revolution contagion is caused by the "domino effect," as citizens in one country are inspired by revolutions in other countries. However, a concrete link is yet to be specified between actions taken in one country and those in others. Our paper characterizes exactly how such inspiration can be modeled as Bayesian updating following rare events, and how strategic complementarities reinforce the effects of belief revisions.

In this model, one successful revolution or speculative attack may lead to rampant collapse of confidence or outbreak of euphoria and, as a result, a series of attacks in other countries. Such a prediction is supported by surveys conducted among market participants. Eichengreen and Wyplosz (1993) report that more than 90 percent of the respondents agreed that the Exchange Rate Mechanism crisis of 1992–93 was contagious, and 77 percent of them attributed the contagion phenomenon to the reason that "the markets had 'tasted blood' [realized that there were profits to be made]." That

[10]This implication is consistent with empirical findings in studies on the contagion of currency crisis. For example, Dasgupta, Leon-Gonzalez, and Shortland (2011) show that crises tend to spread among economies which are institutionally similar.

is, after observing an encouraging success, market participants believed that it would be very easy to conquer another currency and make profit. Similarly, revolutionary euphoria is often stimulated after a success is observed in neighboring countries. For example, the Tunisia revolution was often considered as one such catalyst which inspired people in other countries and ignited the Arab spring.[11]

Economists, commentators, and policy makers have long maintained that psychological factors may be critical in the international transmission of crises. Kindleberger and Aliber (2011) argue that "there are psychological connections, as when investor euphoria or pessimism in one country affects investors in others." They regard specifically the spread of crisis from Russia to Brazil as largely psychological. Among many others, Reinhart and Rogoff (2009) also blame the "fickle nature of confidence" for contagion of currency crises. We adopt the position that the so-called "psychological" factors can be rationalized. In our model, we formalize the effect of "euphoria or pessimism" spillover, with all the players being Bayesian learners.

In economics, multiple equilibria models with self-fulfilling feature, in the spirit of Diamond and Dybvig (1983) or Obstfeld (1986), are typically considered convenient to shed light on the contagion issue (Masson 1999): the success of currency attack in a neighboring country triggers the domestic economy to "jump" to a bad equilibrium where everybody frantically attacks. However, the missing part in this argument is how equilibrium is selected. By adopting a global game framework, our model does not exhibit multiple equilibria. Nevertheless the evolution of beliefs over different models of the world provides an explicit link to determine the switch between periods of tranquility and periods of frenzy.

This model differs from Drazen's (2000) theory of political contagion, which holds that the outbreak of crisis in a member country of a government "club" lowers the value of its membership and therefore undermines other member governments' commitment to the fixed rate arrangement. Our model features the decision of governments to defend the regime. Upon observing a crisis elsewhere, a government is more likely to give up defending as it expects larger scale of attack, which in turn invites even more attackers.

Our model also differs from contagion models that feature herding or information cascades (e.g., Calvo and Mendoza 2000). We allow players to observe equilibrium outcomes in precedent countries, but not actions of other individuals or aggregate action. More importantly, our model does not rely on having players discard their private information to follow their predecessors' actions. The private beliefs of all players

[11] Anecdotal accounts and media coverages which document this phenomenon abound. For instance, see "Thousands Protest across Egypt, Inspired by Tunisia," *NPR*, 25 January, 2011; and "Arab Revolutions: from Tunisia to Egypt, is this the Beginning of a Trend?" *The Huffington Post*, 1 February, 2011.

are incorporated into the equilibrium outcome, but when an unexpected outcome occurs the beliefs of subsequent sets of players are altered dramatically.

This paper is also different from the dynamic coordination games of Angeletos, Hellwig, and Pavan (2007) and Chamley (1999), in which the same set of agents extract information about a latent variable (i.e., the fundamental) from history and take action repeatedly in each period. In our model, players in each country act only once, and they are uncertain not only about a latent variable but also a latent structure (i.e., the way the world operates). The observed history is only useful for updating the belief about the latent structure. Therefore, the dynamics are completely captured by the evolution of the belief over the model space. In one extension of this model, we also show that our mechanism still works even when we allow for heterogeneity in the latent structure across countries; see Section 4.2.

2. The Model with One Country

In this section we lay out the benchmark model with a single country and expound its qualitative properties. Section 3 characterizes the linkage between countries and details how the rarity of crises causes contagion.

2.1. Players and Payoffs

Consider an economy populated by a unit mass of ex ante identical agents, indexed by $i \in [0,1]$, who play against another player, the regime. Agent i chooses one of two actions: attack ($a_i = 1$) or not attack ($a_i = 0$). The aggregate mass of attackers is denoted A. Simultaneously, the regime can also choose to defend ($y = 1$) or not defend ($y = 0$).[12] If the regime defends, the attack can be either a success ($S = 1$), where the regime is forced to surrender or to devalue the domestic currency, or a failure ($S = 0$), where the status quo is maintained. If the regime chooses to give up fighting, the attack is always successful ($S = 1$), regardless of the size of the attack A.

An agent's payoff depends both on whether the attack is successful and on whether the agent chooses to attack. A positive cost, $c \in (0,1)$, has to be paid if she attacks. Upon a success, agents who attack receive a benefit, $b = 1$, and those who do not participate receive no benefit.[13] If the regime defends, it has to pay a country-specific

[12]As in much of the literature, we assume that both players move at the same time. The alternative scenarios where the regime takes action before or after agents attack, have been dealt with in Angeletos, Hellwig, and Pavan (2006) and Goldstein, Ozdenoren, and Yuan (2011), respectively.

[13]In the context of revolution, this payoff structure abstracts from free-riding issues, which is common in the literature that models revolution as coordination games, e.g., Bernhardt and Shadmehr (2011) and Edmond (2013). A more general payoff structure can incorporate the fact that the fruits of a revolution are a public good, with a sufficient condition that will ensure that citizens still have incentives to act against the regime despite the free-riding problem. Once such a condition is satisfied, the concern of free-riding only increases the opportunity cost of participation but does not change the qualitative

cost d.[14] The payoff to the regime is V if it survives the attack; and 0 otherwise. If the regime does not defend, it collapses and the payoff to the regime is normalized to 0. The net utility u for agent i and the net utility v for the regime are:

$$u(a_i, S) = \begin{cases} 1 - c & \text{if } a_i = 1 \text{ and } S = 1, \\ -c & \text{if } a_i = 1 \text{ and } S = 0, \\ 0 & \text{if } a_i = 0; \end{cases} \qquad v(y, S) = \begin{cases} -d & \text{if } y = 1 \text{ and } S = 1, \\ V - d & \text{if } y = 1 \text{ and } S = 0, \\ 0 & \text{if } y = 0. \end{cases}$$

2.2. Information Structure

Let θ represent the strength of the regime. Citizens are ex ante identical and become ex post heterogeneous after each of them observes a noisy private signal about the strength of the regime,

$$x_i = \theta + \varepsilon_i,$$

where the strength θ is drawn by nature from a normal distribution, $\mathcal{N}(\mu, \sigma^2)$, and the idiosyncratic noise $\varepsilon_i \sim \mathcal{N}(0, \sigma_x^2)$ is independent of θ, and is independently and identically distributed across i. In contrast, the regime can observe perfectly the fundamental θ, which is selected by nature.[15] Moreover, the regime knows the cost of defending d, while agents cannot observe it but know that d is drawn from a distribution function H on the support $[0, \overline{d}]$ where $\overline{d} \leq V$.

The main departure from a standard global game is that, besides the uncertainty about regime strength θ, the rule of the game is also not common knowledge. In other words, we allow for model uncertainty—neither the regime nor the agents are certain about how the world operates. Specifically, we assume that players maintain two alternative hypotheses: the world can be either a *tranquility world*, denoted T, or a frenzy world, denoted F. They also know that one of the two is the true model for how the world operates. In Section 4.3, we show that our results continue to hold in the general case where there are a wide range of hypothetical descriptions about the world, with one of them being the true description.

Tranquility and frenzy worlds differ only in the success determination condition

features of the model.

[14]To keep a currency peg typically requires sacrificing domestic goals, which are in conflict with a fixed exchange rate. Such a cost can vary across countries. In a context of revolution, costs have to be incurred to maintain a repressive regime if the regime chooses to defend. Such costs include, but are not confined to, the cost of maintaining troops, as repressive policies may have negative implications for economic performance. In Section 4.3, we discuss a case where the defense cost are the same across country but unknown to agents. We show that the contagion mechanism that we propose will be strengthened.

[15]This assumption has been widely adopted in the literature, e.g., Hellwig, Mukherji, and Tsyvinski (2006). Our results do not rely on this specific informational asymmetry between the regime and agents. See discussion on this in Section 4.3.

when the regime chooses to defend and confront the attackers. Specifically, it is easier to overcome the regime's defense in F-world than in T-world: the outcome of the confrontation is determined by

$$S_j(A,\theta) = \mathbb{1}\left(A(\theta) \geq D_j(\theta)\right),$$

where $D_j(\theta)$ is an increasing function of θ, and $D_T(\theta) \geq D_F(\theta)$ for any θ.

One may interpret the possibility of these two worlds as uncertainty about the defense technology. In the context of mass revolution, θ may stand for the size of troops defending it. Even holding it constant, there is some uncertainty about how big a revolt is needed to topple the regime: in a frenzy world a small mass of attackers is sufficient to produce a revolt that would overcome the troops, while in a tranquility world a much larger mass of attackers is required. An alternative interpretation is that there is uncertainty about the consequences of attack. In the context of currency crisis, θ may stand for the amount of foreign reserves. An attack on currency may produce unpredictable economic pains on the real sector in addition to the depletion of foreign reserves.[16] In a frenzy world these economic pains may be so large that the central bank devalues the currency well before its reserves are exhausted; in a tranquility world the central bank can hold on to the fixed exchange rate for much longer.

Both the regime and the agents share a common prior belief π that they live in tranquility world. This section simply treats π as exogenous and studies its role in equilibrium via comparative statics. In the next section we characterize the evolution of this belief, which is the key to the contagion mechanism in our model.

2.3. Equilibrium

Definition 1. *An equilibrium consists of a set of posterior beliefs which are derived using Bayes' rule, an attack decision $a(x_i,\pi)$, a defense decision $y(\theta,\pi,d)$, and a mass of attackers $A(\theta,\pi)$, such that*

$$\begin{aligned}
y(\theta,\pi,d) &= \underset{y\in\{0,1\}}{\operatorname{argmax}} \left\{[\pi(1-S_T(A,\theta)) + (1-\pi)(1-S_F(A,\theta))]V - d\right\} y; \\
a(x_i,\pi) &= \underset{a_i\in\{0,1\}}{\operatorname{argmax}} \left[\left(\int_{-\infty}^{\infty} \pi u(a_i, S_T(A,\theta)) + (1-\pi)u(a_i,S_F(A,\theta))\right) f(\theta|x_i)\,\mathrm{d}\theta\right]; \\
A(\theta,\pi) &= \int_{-\infty}^{\infty} a(x_i,\pi)\frac{1}{\sigma_x}\phi\left(\frac{x_i-\theta}{\sigma_x}\right)\mathrm{d}x_i,
\end{aligned}$$

where $f(\theta|x_i)$ is the posterior density of θ.

[16] For example, the increased cost of servicing the public debt is one reason why a government accedes to devaluation pressures. Obstfeld (1994) offers a description of different types of economic pains that the real sector in a country under attack may suffer.

We focus on monotone equilibrium throughout (which is shown to be the only equilibrium). In such an equilibrium, there exists a cutoff x^* such that agent i attacks the regime if and only if $x_i \leq x^*$. An agent who observes a low value of x_i believes that the strength of the regime is low, which raises the expected payoff from attacking the regime. In a monotone equilibrium, the size of the attack $A(\theta, \pi)$ decreases in θ. When the regime chooses to defend, there exist thresholds θ_T^* and θ_F^* such that the attack would be successful if and only if $\theta \leq \theta_j^*$ in world j.

Given our assumption about the attack technology, the regime is more likely to be toppled when confronting attackers in a frenzy world than in a tranquility world. Thus, it always holds that $\theta_T^* \leq \theta_F^*$. When θ is above θ_F^*, the attack would fail as long as the regime chooses to defend. Since $V > d$, the regime always chooses to defend in this case. When θ is below θ_T^*, defending the regime is futile because the attackers would succeed despite the effort of the regime to maintain the status quo. So the regime always chooses not to defend in this case. When θ is between θ_T^* and θ_F^*, defense is effective only if it is a tranquility world. Consequently, the regime adopts the following monotonic decision rule:

$$y(\theta, \pi, d) = \begin{cases} 1 & \text{if } \theta \geq \theta_F^*, \\ & \text{or } \theta \in [\theta_T^*, \theta_F^*) \text{ and } \pi V \geq d, \\ 0 & \text{otherwise.} \end{cases}$$

Let P_j be the probability of success from the perspective of an agent with private information x if she believes that the true world is j. Then,

$$\begin{aligned} P_T(\theta_T^*, \theta_F^*, x, \pi) &= \Phi\left(\frac{\theta_T^* - X}{\sqrt{\beta}\sigma}\right) + \int_{\theta_T^*}^{\theta_F^*} (1 - H(\pi V)) \frac{1}{\sqrt{\beta}\sigma} \phi\left(\frac{t - X}{\sqrt{\beta}\sigma}\right) \mathrm{d}t, \\ P_F(\theta_F^*, x, \pi) &= \Phi\left(\frac{\theta_F^* - X}{\sqrt{\beta}\sigma}\right); \end{aligned}$$

where $X = \beta x + (1-\beta)\,\mu$ is the posterior mean of θ, $\beta = \sigma^2/(\sigma_x^2 + \sigma^2)$ is the weight that the agent attaches to her private information, and Φ is the standard normal distribution function. The first term in P_T represents the probability that θ is less than θ_T^* (in which case the regime never defends and the attack is always successful). The second term is the probability that the regime gives up (the attack would have failed if the regime chose to defend). The probability that the regime collapses in the frenzy world, P_F, is simply given by the probability that θ is less than θ_F^* (because in this case the attack is successful regardless of whether the regime defends or not).

A marginal agent with private information x^* is indifferent between attacking or not attacking. For such an agent, the probability of success is a weighted average of P_T and P_F and the weight attached to P_T is π. In other words, the indifference condition can be written as:

$$c = \pi P_T(\theta_T^*, \theta_F^*, x^*, \pi) + (1-\pi) P_F(\theta_F^*, x^*, \pi). \tag{1}$$

The equilibrium regime survival thresholds must satisfy the following critical mass conditions:

$$\Phi\left(\frac{x^* - \theta_T^*}{\sigma_x}\right) = D_T(\theta_T^*), \tag{2}$$

$$\Phi\left(\frac{x^* - \theta_F^*}{\sigma_x}\right) = D_F(\theta_F^*). \tag{3}$$

A monotone equilibrium can be characterized by the triple $(x^*, \theta_T^*, \theta_F^*)$ that solves the system of equations (1), (2) and (3).

It will also be useful to obtain the "objective" probability (i.e., the probability for one with no private information about θ or d) that the regime collapses. Let p_T and p_F represent this probability in the tranquility world and in the frenzy world, respectively. We have

$$p_T = \Phi\left(\frac{\theta_T^* - \mu}{\sigma}\right) + \int_{\theta_T^*}^{\theta_F^*} (1 - H(\pi V)) \frac{1}{\sigma} \phi\left(\frac{t-\mu}{\sigma}\right) \mathrm{d}t; \tag{4}$$

$$p_F = \Phi\left(\frac{\theta_F^* - \mu}{\sigma}\right). \tag{5}$$

Proposition 1. *If $D_j'(\theta)\phi(0)\sigma_x < \sigma^2$ for $j = T, F$, a monotone equilibrium exists and is the only equilibrium. Further, x^*, θ_T^*, and θ_F^*, as well as p_T and p_F, decrease in π.*

Proof. See the Appendix. ∎

From equations (2) and (3), it is easy to see that we must have $\theta_T^* \leq \theta_F^*$ in a monotone equilibrium, because $D_F(\theta) \leq D_T(\theta)$ for all θ and both functions are increasing. Since $\theta_T^* \leq \theta_F^*$, equations (4) and (5) also imply that we must have $p_T \leq p_F$. Likewise, for any agent with private information x_i, we have $P_T(\theta_T^*, \theta_F^*, x_i, \pi) \leq P_F(\theta_F^*, x_i, \pi)$. In other words, the probability of a successful attack, either subjective or objective, is lower in a tranquility world than in a frenzy world.

A decline in public belief π has two effects. First, since $P_T \leq P_F$, equation (1) implies that it raises the expected payoff from attacking. Second, when θ is between θ_T^* and θ_F^*, the expected payoff for the regime from fighting against attack, πV, also falls. As agents expect the regime to give up defending more often, this raises P_T and, by equation (1), raises the expected payoff from attacking. Consequentially, both effects cause agents to become more aggressive, i.e., x^* increases. These effects are strengthened through complementarity in action between agents and the regime: expecting an attack of larger scale, the regime gives up with higher probability; anticipating this, agents attack even more aggressively. Moreover, complementarity in action among the agents themselves as in a standard global game model further amplifies these effects: an increase in x^* raises the critical success thresholds θ_T^* and θ_F^*, which in turn increases the payoff from attacking and causes x^* to rise further. Therefore, the attack is more likely to be successful, no matter the true world is tranquility or frenzy. That is, both p_T and p_F increase.

3. Contagion across Countries

3.1. The Evolution of Public Belief

Now we turn to a dynamic multi-country version of the benchmark model. Countries are identical except that the fundamental strength of the regime θ and the cost of defending d for each country are independent draws.[17] They are lined up in a fixed order and take action sequentially.[18] We use subscript t on some variables to indicate the country concerned.

Players in a country can observe whether attacks in preceding countries are successful or not and update their beliefs accordingly. We let π_t represent the public belief in country t that the world is a tranquility world. This is the belief of a player in country t after observing the history of successes in preceding countries but before observing her private information (x_i for agents or θ for the regime).[19] By Bayes' rule, the belief updating process can be written recursively as:

$$\frac{\pi_{t+1}}{1-\pi_{t+1}} = \begin{cases} \frac{\pi_t}{1-\pi_t}\frac{p_T(\pi_t)}{p_F(\pi_t)} & \text{if } S_t = 1, \\ \frac{\pi_t}{1-\pi_t}\frac{1-p_T(\pi_t)}{1-p_F(\pi_t)} & \text{if } S_t = 0. \end{cases} \tag{6}$$

In equation (6), p_T/p_F is the likelihood ratio of observing a successful attack in country t. It should be noted that the equilibrium values of p_T and p_F depend on θ_T^* and

[17] We share the view of Drazen (2000) that contagion is "the phenomenon whereby a currency crisis *itself* in one country makes a currency crisis in another country more likely." It is conceptually distinct from common shocks that affect all countries and is also different from the so-called "spillover effect" in which the fundamental in one country affects those in others.

[18] This assumption is intended to capture the usual pattern of how crises diffuse: they do not occur at different countries simultaneously. It bears similarity to standard herding models (Banerjee 1992 and Bikhchandani, Hirshleifer, and Welch 1992), where the timing of actions is fixed. To focus on the key mechanism of contagion, we abstract from the endogenous timing of actions.

[19] Observations on the outcomes of previous countries do not help agents to update their belief on the fundamental in their own countries, because the fundamentals are uncorrelated across countries.

θ_F^*, which in turn depend on π_t. Likewise, the likelihood ratio of observing an unsuccessful attack in country t is $(1-p_T)/(1-p_F)$, which is also a function of π_t. The evolution of public belief described by equation (6) depends on the stochastic process $\{S_t\}$. Assume that the true world is a T-world. The probability law governing this stochastic process is:

$$S_t = \begin{cases} 1 & \text{with probability } p_T(\pi_t), \\ 0 & \text{with probability } 1 - p_T(\pi_t). \end{cases} \tag{7}$$

Recall that an attack is more likely to be successful in a frenzy world than in a tranquility world, i.e., $p_T/p_F < 1$. Therefore, a successful attack in country t is evidence in favor of the hypothesis that the world is frenzy. As a result, the public belief π_{t+1} becomes lower than π_t. But a lower π_{t+1} in turn causes agents to become more aggressive and the regime to defend less often in country $t+1$. It follows from Proposition 1 that the objective probability p_T of a successful attack in country $t+1$ increases following a successful attack in country t. This link provides a mechanism for crises to be contagious across countries.

3.2. Rare Event and Collapse in Confidence

Since a successful attack always causes players to lower their belief that the world is tranquil, the contagion mechanism described in the previous subsection always tends to produce positive serial correlation in the occurrence of crises. However, the *quantitative significance* of this serial correlation critically depends on the magnitude of belief updating. If the likelihood of success is similar in the two worlds (p_T is nearly as large as p_F), then even when a success is observed the amount of downward revision in belief is tiny. In contrast, if p_T is small, crises will be rare in a tranquility world. When the likelihood ratio p_T/p_F is small, the occurrence of a crisis will trigger a substantial downward revision in belief. Suppose that the difference in the defense technology $D_j(\theta; k_j)$ is characterized by the parameter k_j, where $j = T, F$. To ensure that the rareness of crises can be obtained in T-world, we let D_T goes to infinity when k_T goes to infinity. One example is a linear defense technology with $D_j(\theta) = k\theta + k_j$, where k is positive and $k_T > k_F$.[20]

Proposition 2. *As k_T increases, the regime survival threshold θ_T^* in world T falls without bound, while θ_F^* and x^* decrease but remain bounded. For k_T sufficiently large, an increase in k_T causes (a) success to become rarer in world T (i.e., p_T decreases); and (b) a greater downward revision in the belief in world T upon observing success (i.e., p_T/p_F decreases). In the limit,*

[20] Any alternative assumptions that ensure the rarity of crises can also work and would not change our results qualitatively.

the belief in country $t+1$ upon observing $S_t = 1$ is

$$\lim_{k_T \to \infty} \pi_{t+1} = \frac{\pi_t(1 - H(\pi_t V))}{\pi_t(1 - H(\pi_t V)) + 1 - \pi_t}.$$

Because the likelihood ratio p_T/p_F approaches $1 - H(\pi V)$ as k_T goes to infinity, the right tail of the distribution of d is relevant for belief updating. We say that H has a *thinner right tail* than the uniform distribution on $[0, V]$ if there exists $\hat{d}$ such that, for all $d > \hat{d}$,

$$1 - H(d) < 1 - \frac{d}{V}.$$

Ex ante, a regime is less likely to incur high defense cost when H has a thin right tail. This makes it more likely to successfully defend itself against attack in the tranquility world, which helps make both p_T and p_T/p_F small. The next result follows immediately from the formula for π_{t+1} stated in Proposition 2.

Corollary 1. *For any $\epsilon > 0$, there exists $\bar{k}_T$ such that for all $k_T > \bar{k}_T$, the updated belief π_{t+1} following $S_t = 1$ satisfies:*

1. *If H has support on $[0, \bar{d}]$ with $\bar{d} < V$, then for any $\pi_t \in (\bar{d}/V, 1)$, $\pi_{t+1} < \epsilon$.*
2. *If H has a thinner right tail than the uniform distribution on $[0, V]$, then for any $\pi_t \in (\hat{d}/V, 1)$, $\pi_{t+1} < \pi_t/(1 + \pi_t) + \epsilon$.*

Case (1) of Corollary 1 says that the public belief can drop from near 1 to near 0 following a single successful attack, provided $\bar{d} < V$. Even when this condition is not satisfied, a rare success can trigger a substantial collapse in belief under the relatively mild assumption that the distribution H has a thin right tail. For example, in case (2) of Corollary 1, if $\pi_t = 0.9 > \hat{d}/V$, then $\pi_{t+1} < 0.48$ following a successful attack in country t.

When k_T is very large, success is expected to be difficult to achieve in the tranquility world. Once players attach a high probability to the tranquility hypothesis, few agents would attack and the regime would defend more often. As a result, success indeed occurs with a very low probability. The larger is k_T, the rarer is success. But when such a rare success happens, agents drastically lower their belief that they live in the tranquility world, precisely because they do not expect such an event to occur if their maintained hypothesis is true.

The consequence of a rare success elsewhere and the associated dramatic belief revision is described in Proposition 1. Agents become much more aggressive and decrease their estimates of regime's survival substantially, and the regime is also less confident in defense. As a result, the probability of the regime collapsing increases sharply after the first domino falls. In other words, a tranquil world may suddenly become frantic just because people's beliefs change.

The transmission mechanism we describe here captures popular notions about crisis contagion often heard in policy discussions: a crisis elsewhere leads to "abrupt shifts in investor confidence" (Pesenti and Tille 2000) and weakens the domestic government's "commitment to a fixed exchange rate" (Krugman 1999). Economists tend to view such abrupt changes through the lens of multiple equilibria, but remain silent on how agents shift from one equilibrium to another. In this unique equilibrium model, as long as people are not completely sure about their model of the world, Bayesian belief updating can cause a sudden switch in behavior.

More importantly, our result emphasizes that the fragility of beliefs is closely related to the rarity of crises: it is precisely the unlikely occurrence of a crisis in a tranquility world that prompts people to fundamentally reassess whether the world is indeed tranquil. In a model with complementarity between beliefs and actions, such reassessment can become self-fulfilling, at least in the short run.

3.3. Propagation and End of Crisis Period

It is important to point out that one successful attack in a country not only affects the next country, but also triggers the plummeting of confidence for a number of subsequent countries and produces a *sustained period* of crisis.

Interestingly, although the period of turbulent time can be long-lasting, our model predicts that in the long run it must come to an end. Intuitively, since the true world is T, Bayesian learning implies that the public belief will on average recover after a collapse in confidence. To see this clearly, write the belief updating equation (6) in reciprocal form and take expectation conditional on the true world being T, we obtain:

$$\begin{aligned}\mathbb{E}\left[\frac{1-\pi_{t+1}}{\pi_{t+1}}\,\middle|\, T\right] &= p_T(\pi_t)\left(\frac{1-\pi_t}{\pi_t}\frac{p_F(\pi_t)}{p_T(\pi_t)}\right) + (1-p_T(\pi_t))\left(\frac{1-\pi_t}{\pi_t}\frac{1-p_F(\pi_t)}{1-p_T(\pi_t)}\right)\\ &= \frac{1-\pi_t}{\pi_t}.\end{aligned}$$

Thus, $\mathbb{E}[1/\pi_{t+1}|T] = 1/\pi_t$. But since $1/\pi$ is strictly convex in π, Jensen's inequality implies

$$\mathbb{E}[\pi_{t+1}|T] > \pi_t.$$

Therefore, the public belief will eventually recover to a level close to one (until another accidental success strikes) and the world becomes tranquil again. This feature of our model is absent in most other models of contagion, which focus on how a crisis starts and propagates, but pay less attention to how it ends.[21]

3.4. Numerical Examples

The evolution of beliefs and equilibrium outcomes described by equations (6) and (7) is not deterministic. In this subsection we provide numerical simulations to illustrate how the occurrence of one rare event interacts with the contagion mechanism to produce a period of crises across countries.

For the benchmark exercise, we set the cost of attacking at $c = 0.5$. The variances of private information and prior are the same, $\sigma_x^2 = \sigma^2 = 1$. The payoff for the regime is $V = 1$ if it survives. The distribution H of the cost of defending is uniform on $[0, \overline{d}]$ and $\overline{d} = V$. We characterize the defense technology in the two alternative worlds with a linear function, i.e., $D_j(\theta) = k\theta + k_j$, and let $k_T = 5$ and $k_F = -2$, while $k = 3$. The mean of the distribution of the fundamental is set to be $\mu = 0.5$ in the benchmark case. The parameters for T-world are chosen such that a successful attack is a rare event: the regime collapses with a probability of 1.5 percent if the true world is tranquil and if players fully believe so. The parameters for the F-world would imply that the regime would collapse with high probability if players fully believe in the F-world. We assume that players are very confident that the true world is tranquil, but they are not 100 percent certain. We therefore set the initial public belief at $\pi_0 = 0.99$. In other words, they entertain a small possibility that the world is frantic.

We compute the dynamics of π (the public belief) and p_T (the probability of successful attack in the true world) for thirty countries under two scenarios: there is a success in the initial country ($S_0 = 1$), or a failure in the initial country ($S_0 = 0$). To contrast the case where success is rare with the case where success is commonplace, we repeat the exercise by setting $\mu = -1.5$, so that success is common in the tranquility world as well (the probability of success would be 50 percent if players fully believe that it is a T-world). We simulate each of the four sub-cases 2,000 times. After averaging the outcomes of each country in these 2,000 histories, we obtain the average sample path for π and p_T.

Crisis, propagation, and slow recovery. The left column of Figure 2 shows the average sample paths of π and p_T in the case where crises are rare. An accidental success in country 0 shakes the public confidence and unfolds a "turbulent time": π drops from 99 percent to 78 percent, and the probability of success in country 1 shoots up from less than 1.5 percent to more than 13.5 percent. If we start the simulation with

[21] In a currency attack model with multiple equilibria, for example, one may attribute the beginning of a crisis to individuals suddenly jumping to a bad equilibrium altogether (Masson 1999). But then one must argue that they always coordinate to jump back to the good equilibrium to end the crisis.

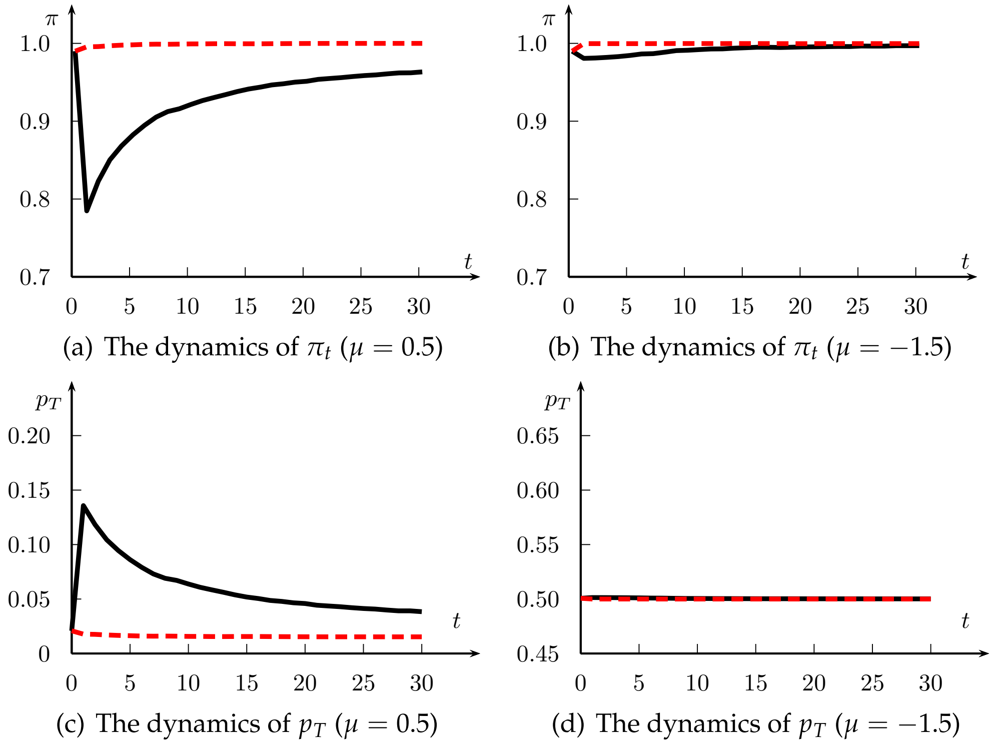

Figure 2. *Contrasting Patterns of Contagion when Success is Rare or Commonplace.*

The solid lines stand for the scenario with $S_0 = 1$, and the dashed lines for $S_0 = 0$. The two graphs in the left column display results from the case where success is rare, with $\mu = 0.5$, while the right column shows the case where success is commonplace, with $\mu = -1.5$.

lower values of π_0, or with a distribution H that has a thinner tail than the uniform distribution, then our quantitative results will be even more dramatic.[22]

The sustained crisis period following a rare success is evident in Figures 2(a) and 2(c), where we see that π and p_T do not revert to their baseline levels even after thirty countries take actions. One reason of the observed slow recovery is that the belief updating process is not symmetric. In our example, the public belief changes from $\pi_0 = 0.99$ to $\pi_1 = 0.78$ after a successful attack in country 0. But even if the attack in country 1 fails, the public belief only goes back up to $\pi_2 = 0.87$, still considerably lower than the initial level of 0.99.

This asymmetry is caused by the rarity of success. In country 0, a successful attack is fifty times more likely in an F-world than in a T-world ($p_T/p_F = 0.027$). Success in the initial country is very strong evidence against tranquility, and hence triggers a large downward revision is belief. In country 1, a failed attack is two times more likely in a T-world than in an F-world ($(1-p_T)/(1-p_F) = 1.99$). This means that failure in the subsequent country is only mild evidence against frenzy, and hence it triggers a mild rebound in confidence.

[22]For example, if $\pi_0 = 0.95$ and $S_0 = 1$, then $\pi_1 = 0.59$ and p_T for country 1 is more than 26 percent.

Another effect which contributes to the slow recovery is that, with a significantly higher p_T in country 1, the probability becomes larger that the public belief will go down even further following another success in country 1. Thus, when we focus on the average sample path, the public belief in country 2 is on average $\pi_2 = 0.82$. See Figure 2(a).

The dashed line in Figure 2(a) presents the case where the rare success does not occur in country 0. In this case the average sample path for π_t monotonically increases, since failure is evidence for tranquility. But the increase is almost not discernible, because failure is so common that the additional evidence does not have much effect.

Commonplace success. The right column of Figure 2 shows the average sample paths of π and p_T for the case where success is commonplace. In this case, one successful attack in the initial country still triggers a fall in confidence and an increase in the probability of success in subsequent countries, but the effects are very small. When successes are commonplace even in the true world, the occurrence of a successful attack in the initial country is only weak evidence in favor of the frenzy world hypothesis. In our example, the public belief only goes down from $\pi_0 = 0.99$ to $\pi_1 = 0.98$. As a result, the effects on the behavior of players in other countries are negligible. The contrast between this case and the case of rare events shows that it is the combination of rarity and the contagion mechanism that produce the clustering of crises.

Crisis deepening. As our theory predicts, a rare successful attack in country 0 increases the chance of another success in country 1 substantially. If such a success happens to materialize in country 1 again, people's confidence in a tranquility world falls further and the probability of observing subsequent successes will be even higher. Figure 3 shows how the crisis can deepen after more than one success in a row. The dashed line shows the situation for the case with $S_0 = S_1 = 1$. Following another success in country 1, the public belief is further lowered to $\pi_2 = 0.46$ and p_T shoots up to 35 percent.

The dotted line shows the case with $S_0 = S_1 = S_2 = 1$. Observed that after two successes in a row, the third success leads to a much smaller revision in public belief. That is because the probability of success has gone up to 35 percent, so that successes are no longer rare. Thus, further successes can only produce a small effect on beliefs and hence on actions.

Runs tests. One of the main predictions of our model is that successful attacks tend to occur in serial clusters: a success in one country raises the probability of success in the next country. A commonly used method to detect such serial dependence of binary events is the runs test (Wald and Wolfowitz 1940). For each of the 2,000 simulated

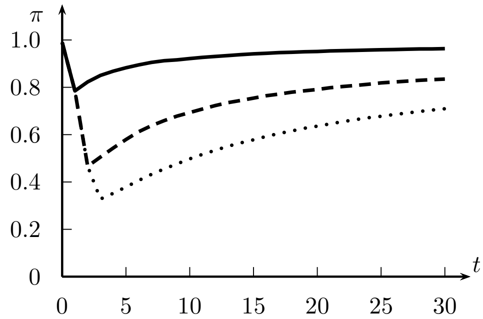

(a) The dynamics of π_t

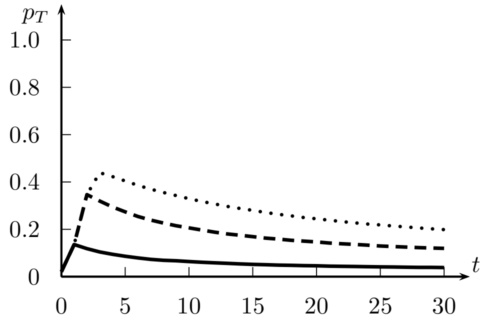

(b) The dynamics of p_T

Figure 3. *Crisis Deepening.*

The average sample paths of π and p_T after one, two or three consecutive successes. Solid line is the case for $S_0 = 1$, dashed line for $S_0 = S_1 = 1$, and dotted line for $S_0 = S_1 = S_2 = 1$.

histories of successes and failures, we group all consecutive successes or consecutive failures into a "run" and count the total number of runs in each history of thirty countries. When successes tend to be clustered, the number of runs will be relatively small compared to the expected number.[23] Thus, if the runs test statistic is significantly negative, it is evidence in favor of the hypothesis that successes are positively serially correlated. For each of the four scenarios, we compute the proportion of simulated histories that fails the one-tailed runs test at the 5 percent significance level. Table 1 summarizes the results.

Table 1. *The Runs Test.*

	$S_0 = 1$	$S_0 = 0$
Successes are rare ($\mu = 0.5$)	61.2%	3.9%
Successes are commonplace ($\mu = -1.5$)	4.6%	5.5%

Note: Histories with all failures are excluded because they pass the test automatically (the runs test has no power).

In the rarity case, a predominant proportion (61 percent) of simulations fail the runs test after observing $S_0 = 1$. This indicates that an initial success in one country tends to be followed by clusters of successes in other countries. In contrast, when $S_0 = 0$, only about 4 percent of the simulations fail the test, which is what one would expect from pure sampling variability if there is no serial correlation.

In the case where successes are commonplace, the proportion of simulations that fail the runs test after observing $S_0 = 1$ is around 5 percent. This means that, although the contagion mechanism is still present in this case, its effects are so weak that they

[23] If there are n successes and m failures in the sample, the expected number of runs is $2nm/(n+m) + 1$ under the null hypothesis that successes and failures are drawn independently.

are not statistically discernible in small samples. Note also that observing success or failure in the initial country does not matter much for the clustering of crises. In either case, a statistician would be hard-pressed to detect the presence of contagion.

4. Discussion

4.1. Model Uncertainty and Correlation in Fundamental

Our contagion mechanism is purely informational. While acknowledging their importance, we abstract from those commonly assumed correlation of economic fundamentals and trade or financial market linkages which can cause and exacerbate contagion across countries (e.g., Gerlach and Smets 1995; Allen and Gale 2000). Instead, the only connection between countries is that they are subject to the same set of rules that govern the operation of the world. We show that even such a seemingly tenuous linkage can be potentially a source of contagion, which is largely overlooked in the existing literature.

How is our model distinguishable from an alternative model with correlation in fundamentals across countries but without model uncertainty? Consider, for example, a model in which the fundamental θ_t follows an autoregressive process of order 1, and citizens in country t can observe a private signal about θ_t as well as the realized value of θ_{t-1}. Call this a *correlation model*. Obviously, citizens in country t are more likely to attack (and the regime is less likely to defend) when θ_{t-1} is low. Since θ_t is serially correlated across countries, successes will also be serially correlated in the correlation model. Moreover, provided that the long run average of θ_t is high, successes are rare ex ante.

There are, however, subtle differences between our model and the correlation model. First, our contagion mechanism relies solely on an informational linkage. If citizens in country t do not observe the outcome of success or failure in country $t-1$, then equilibrium outcomes in different countries will be totally unrelated. In the correlation model, by contrast, outcomes will be correlated even if citizens in one country do not observe what happened in another country. The correlation model imposes a mechanical link that produces contagion, which does not depend on (but can be magnified by) the endogenous response of players. Second, as we stress, the contagion mechanism would be not identifiable quantitively if crises are commonplace in our model. In the correlation model, successes and failures are always dependent across countries, no matter whether successes are rare or commonplace.[24] Third, our model

[24] We specify, solve and simulate a model with correlation in fundamental. We perform the runs test on the simulated data and conclude that there are always serial dependence in the sequence of outcomes regardless of the ex ante probability of success.

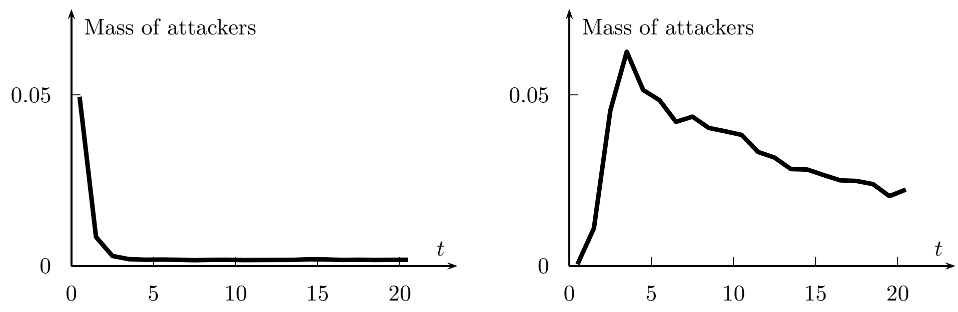

(a) The correlation model (b) The benchmark model

Figure 4. *The Size of Attack in the Correlation Model and Benchmark Model.*

The figure shows the average sample path of the size of attack based on 2,000 simulated histories. The correlation model specifies an AR1 process with $\rho = 0.5$ and $\mu = 2$. The fundamental strength in the initial country is set to 0 to set off a successful attack. The parameters used in the benchmark model are the same as those in Figure 2(a)

.

exhibits *path dependence* while the correlation model only exhibits *state dependence*, to use the distinction made by Page (2006). In the correlation model, the outcome in country t depends on other countries only through the state transition to θ_t; the actual realization of successes or failures in other countries is irrelevant. In our model, the entire history $\{S_1, \ldots, S_{t-1}\}$ matters. The history $\{1, 0, 1\}$, for example, produces a very different outcome for country 4 than does the history $\{0, 1, 1\}$.

In practice, when the underlying state is unobserved by the econometrician, it is difficult to distinguish between path dependence and state dependence, especially when the time series comprise only binary data about success or failure. The difference becomes more apparent if the econometrician also has access to data on the size of the attack in different countries. In the correlation model, the mass of attackers A_t shoots up when the realization of the fundamental θ_{t-1} is exceedingly low. Because θ_t tends to regress back to the mean, the size of attackers declines on average in the subsequent countries—regardless of whether the attack succeeds or fails in country t. In our model, the mass of attackers A_t also shoots up once a success is observed in the previous country, but what happens in subsequent countries depend on what happens in country t—A_{t+1} will be higher than A_t if $S_t = 1$ and will be lower than A_t if $S_t = 0$. More interestingly, the size of attack *on average* can go up even further in the subsequent countries, even though we show in Section 3.3 that the public belief must decline on average.[25] Figure 4 shows that our model and the correlation model have very different implications for the mass of attackers following a success in one country.

[25]This is possible because A_t depends on the public belief π_t non-linearly.

We do not claim that the mechanism we expound here is the only one that matters, but our theory can shed light on situations that are otherwise difficult to explain using direct linkages. Regarding currency attacks, the often-cited crisis transmission from Russia to Brazil falls into that category. In the context of revolution, citizens can learn about how a dictatorship state works from the crisis experience in other countries, even though the regime strength in each affected country may not necessarily be correlated. Because revolutions are seldom observed, neither citizens nor the ruling elites have a good understanding of the technology of revolutions, which justifies our assumption on model uncertainty. Our model offers an interpretation of the revolution wave of the Arab Spring: the unexpected success in Tunisia triggered a large shift in the belief on the revolution technology, which caused citizens to become more aggressive and some of the ruling elites to lose their will to suppress civilian uprisings, and therefore ushered in a sequence of revolutions.

One natural implication of this model is that the degree of "similarity" between countries (perceived by attackers) predicts the strength of our contagion mechanism, provided that crises are rare. Countries would not affect each other if they were governed by rules that are entirely different. Conversely, if the rules of the game are similar across countries, players can learn the robustness of these rules based on their observations about successful or failed attacks against other countries. This prediction is consistent with a number of empirical findings. For example, Eichengreen, Rose, and Wyplosz (1996a) show that a crisis within the EMS had a larger contagion effect on member countries than on non-member countries. Dasgupta, Leon-Gonzalez, and Shortland (2011) find that crises tend to spread among countries which bear institutional similarities.

4.2. Model Uncertainty and State Switching

In the benchmark model, we assume that the true world is indeed tranquil and agents harbor only a slight doubt about it. Implicitly, we assume that countries with the same institutional features, i.e., countries with fixed exchange rate system or countries with dictatorship, may work in a similar fashion, even though the strength of each country may not be correlated. In this extension, we assume that the state of the world can actually shift from tranquility to frenzy and the other way around. Agents still only observe the past outcomes but not the state of the world. In other words, we maintain the assumption of model uncertainty and entertain the idea that countries may not share the same set of characteristics.

Specifically, we adopt a Markov state-switching model. If the world in current period is T, the chance to stay in T-world next period is q_T; if the world is F, the

Table 2. *State Transition Matrix.*

	T in Period $t+1$	F in Period $t+1$
T in Period t	q_T	$1-q_T$
F in Period t	$1-q_F$	q_F

probability of staying in F-world next period is q_F. See Table 2.

The belief about which world agents live in is summarized by the public belief π_t. Given the outcome in period t, the public belief is updated in the following fashion:

$$\pi_{t+1} = \begin{cases} \frac{\pi_t p_T}{\pi_t p_T + (1-\pi_t) p_F} q_T + \frac{(1-\pi_t) p_F}{\pi_t p_T + (1-\pi_t) p_F} (1-q_F) & \text{if } S_t = 1, \\ \frac{\pi_t (1-p_T)}{\pi_t (1-p_T) + (1-\pi_t)(1-p_F)} q_T + \frac{(1-\pi_t)(1-p_F)}{\pi_t (1-p_T) + (1-\pi_t)(1-p_F)} (1-q_F) & \text{if } S_t = 0. \end{cases}$$

Considering the following three cases. First, suppose both q_T and q_F are also close to 1. This means that T-world rarely switches to F-world or vice versa. In this case, our mechanism of model uncertainty is still at work. Suppose the public belief in country t that the true state of the world is T is close enough to 1, a success would revise the public belief downwards substantially. However, ex ante, the public belief gradually converges to the long-run stationary distribution of the two worlds, which is

$$\pi_\infty = \frac{1-q_F}{1-q_F+1-q_T}.$$

When both q_T and q_F are close to 1, π_∞ is close to $1/2$, and the chance of success in each world, $p_T(\pi_\infty)$ and $p_F(\pi_\infty)$ will be both substantially higher than 0. In other words, in this model, crisis events are not rare in the long run.[26] To illustrate, we repeat the numerical exercise shown in Figure 2, by starting with T-world at $t=1$ and allowing the state of the world to transit from one to another, following the transition matrix specified in Table 2. Figure 5 shows the results when $q_T = q_F = 0.99$. Notice that our contagion mechanism is still at work even in this case when the initial belief π_0 is sufficiently close to 1, as evidenced by the large initial drop in public belief and the subsequent recovery. However, also note that there is a long-run tendency for the public belief to converge to its stationary state level.

Consider next the case where q_T is close to 1 but q_F is close to 0. In this case, the long-run stationary probability of T-world is high. However, even if citizens in country $t+1$ believe that the world has switched to a frenzy world in country t (following a successful attack in that country, for example), they believe that their own country will

[26]Likewise, if q_T is small but q_F is big, then successes are not rare in the long run.

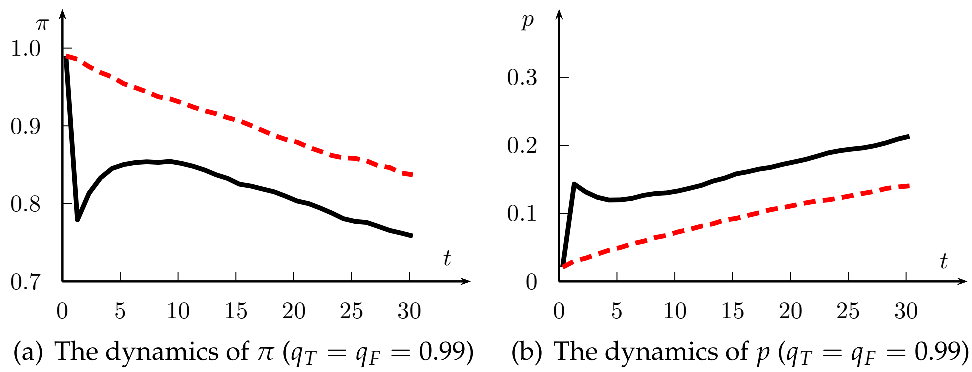

Figure 5. *State Switching with Model Uncertainty.*

Panel (a) repeats the numerical example of Figure 2(a), with the additional mechanism of Markov state-switching. Panel (c) is the counterpart to Figure 2(c). This figure shows the case where both q_T and q_F are close to 1*.*

have switched back to a tranquility world with high probability. Therefore, a success observed in country t revises the public belief of agents in country $t+1$ only a little. In this case, crises are likely to be stand-alone incidents.

In the two cases mentioned above, either the rarity of success is not ensured in the long run or contagion does not take place. Neither case is a good candidate for explaining why rare crises tend to be contagious. Both features of rarity and contagion can be present in a third case when q_T is close to 1 and q_F is not too big or too small (e.g., $q_F = 0.7$). In this case, the model is quantitatively similar to our benchmark model, where a success in the pervious country revises the public belief downwards dramatically, which gradually recovers and converges to a value below but close to 1. Therefore, we may consider our model as an approximation to this case where the operation of the world can change. Even in this case, note that the key to contagion is still the fact that agents cannot observe the true world, which is exactly the same mechanism as we describe in the benchmark model. If instead citizens in each country can observe the true world in the Markov state-switching model, then the probability of success will simply fluctuate between a low level (when the world is tranquil) and a high level (when the world is frantic). It will not exhibit the pattern of sudden fall and gradual recovery shown in Figure 2(c) of our benchmark model.

4.3. Robustness Issues

Multiple worlds. Our results do not rely on the number of "models of the world" being maintained by players. Firstly, even when we allow for more than two hypotheses being maintained, it is still the most extreme one among these hypotheses that will gain the most credibility following a rare success in another country. To see this, suppose in addition to the tranquility world and the frenzy world, there is also a *middle* *world*, M, in which it is less difficult to overthrow a regime than in tranquility world, but more difficult than in frenzy world. Equilibrium in such a model has the property that the success probabilities in these three worlds are ordered: $p_T < p_M < p_F$. Upon observing a success, Bayes' rule dictates that the belief in the T-world, π^T, must fall and π^F must rise, while π^M may increase or decrease. Importantly, since $p_M < p_F$, the posterior ratio of π^F/π^M must be higher than its prior ratio. In other words, a successful attack is evidence in favor of the frenzy world relative to the middle world.

Secondly, our results continue to hold in a more general setting where some of the possible worlds maintain that the world is even more tranquil than what the true model implies. We establish this point by considering a general model of model uncertainty in which players maintain a wide range of alternative hypotheses surrounding the true one. Specifically, suppose there are K worlds, ordered such that $p_1 < \ldots < p_K$, with k^* being the true model. After observing a success, the posterior belief in the true hypothesis, π^{k^*}, may go up or down. However, the posterior belief distribution $\boldsymbol{\pi} = (\pi^1, \ldots, \pi^K)$ must dominate the prior belief distribution according to the likelihood ratio order, because for any $k > l$,

$$\frac{\pi_{t+1}^k}{\pi_{t+1}^l} = \frac{\pi_t^k}{\pi_t^l}\frac{p_k}{p_l} > \frac{\pi_t^k}{\pi_t^l}.$$

Further, if θ is between θ_l^* and θ_{l+1}^*, then defense is worthwhile for the regime only if the state is greater than the true survival threshold. Since $\theta_1^* < \ldots < \theta_K^*$, the regime defends if and only if

$$\sum_{k=1}^{l} \pi^k V \geq d.$$

Hence, the payoff from defense falls and the regime is less likely to defend. This in turn implies that, from the perspective of agent i, the probability of successful attack, $P_k(\boldsymbol{\theta}^*, x_i, \boldsymbol{\pi})$, increases for each k. Moreover the expected payoff from attack is:

$$\sum_{k=1}^{K} \pi^k P_k(\boldsymbol{\theta}^*, x_i, \boldsymbol{\pi}).$$

Since $P_1 < \ldots < P_K$, her expected payoff increases following a successful attack. Thus, even in this more general framework where players believe that the world can be more tranquil than it really is, a successful attack in one country always causes agents to become more aggressive and hence raises the probability of success in another country.[27]

Hypothesis testing. In our model, players are fully Bayesian and they maximize

[27] However, we need at least one hypothesis in which the world is more frantic than it truly is in order to generate a quantitatively significant contagion effect. If $k^* = K$, then successes would be rare even if players fully believe in world K.

expected utility taking into account all possible models of the world. Alternatively, we may assume that players only hold one model at a time when making a decision, but they also maintain a belief over the model space. They update their belief over possible models in light of the observed evidence, and keep using the existing model as long as the belief in this model does not fall below a certain threshold. When the belief in the existing model falls below the threshold, they pick an alternative model with the highest belief. Ortoleva (2012) calls this a "hypothesis-testing representation" and provides an axiomatic foundation for such a decision rule. We argue that the contagion mechanism in this paper still operates if players adopt a decision rule that satisfies the hypothesis-testing representation, but the onset and the end of the crisis period will be even more abrupt than in our model.

To see this, consider a hypothesis-testing representation in which the tranquility hypothesis is maintained if $\pi > \hat{\pi}$. As long as $\pi > \hat{\pi}$, equilibrium is characterized by (x^*, θ_T^*) (whose values are equal their counterparts in our model with $\pi = 1$) and the associated p_T. But players can also compute what the equilibrium (x^*, θ_F^*) would be (which is equal to their counterparts in our model with $\pi = 0$) and the associated p_F in the frenzy world. We can assume that players in another country update their belief about possible worlds before they choose their actions. Proposition 2 still holds. For k_T sufficiently large, p_T/p_F is small, so that observing a rare success in another country can push players' belief in the tranquility world below $\hat{\pi}$, and they switch to adopt the frenzy model as a result. Under this decision rule, the equilibrium probability of success can only take two values, depending on whether π is above or below the threshold. The period when it takes the higher value can be considered a crisis period.

Information of the regime. In the benchmark model, we assume that the regime knows its own strength θ but does not know the true model of the world. As we have shown, in this case complementarity arises between the actions of agents and the regime. In a sense, how strong the regime is depends not only the fundamental strength θ, but also its willingness to fight the attackers. Thanks to this additional layer of complementarity, the regime can become more fragile simply because agents tend to believe it is easier to topple.

Alternative assumptions about what the regime knows deliver the same qualitative results. Suppose, for example, that the regime neither knows θ nor the true model. Then it chooses to defend if and only if

$$d \leq \left[\pi\left(1 - \Phi\left(\frac{\theta_T^* - \mu}{\sigma}\right)\right) + (1-\pi)\left(1 - \Phi\left(\frac{\theta_F^* - \mu}{\sigma}\right)\right)\right] V \equiv \delta(\theta_T^*, \theta_F^*, \pi).$$

When π declines, δ decreases. In other words, ex ante, it is more likely to observe the regime to give up when the public belief is revised downwards.

Next, suppose the regime knows both θ and the true model of the world. The regime would defend if and only if $\theta \geq \theta_j^*$, where world j is the true world. Thus, from the perspective of an agent with private information x, she believes that the probability of successful attack is $P_T = \Phi\left((\theta_T^* - X)/(\sqrt{\beta}\sigma_x)\right)$ in world T, while P_F remains the same as in the benchmark case. Under this set of assumptions, because the regime never gives up unless fighting is futile, the equilibrium is the same as in a global game model where the regime is not an active player. There would be no complementarity between the actions of the agents and the regime, but the complementarity among the actions of the agents themselves remains.

Finally, suppose that the regime does not know θ but knows the true model of the world. Then the regime chooses to defend if and only if $(1 - \Phi((\theta_j^* - \mu)/\sigma))V \geq d$. From the perspective of an agent with private information x, an attack would fail in world j if $\theta > \theta_j^*$ and the regime chooses to defend. Therefore, for $j = T, F$, we have

$$P_j = 1 - \left(1 - \Phi\left(\frac{\theta_j^* - X}{\sqrt{\beta}\sigma_x}\right)\right) H\left(\left(1 - \Phi\left(\frac{\theta_j^* - \mu}{\sigma}\right)\right) V\right).$$

Note that P_j is increasing in θ_j^*. Therefore, we still have $P_T < P_F$ (and, similarly, $p_T < p_F$). Therefore, a successful attack is evidence against a tranquility world, and encourages agents in other countries to attack more aggressively. The contagion mechanism works in a way very similar to our benchmark model.

Defense cost. In the benchmark case, the assumption that the defense cost of each country, d, is independent and drawn from a known distribution H is introduced for realism. An alternative assumption is that the defense cost may be the same across countries but agents do not know its exact value and they need to update their beliefs, based on outcomes in previous countries. We show that this additional mechanism can strengthen our results quantitatively without affecting them qualitatively.

To see this point, suppose the defense cost distribution perceived by the agents in country t is $H_t(d)$. On condition that there is a success in country t, the density function of defense cost is updated by Bayes' rule:

$$h_{t+1}(d|S_t = 1) = \begin{cases} \eta\Phi\left(\frac{\theta_F^* - \mu}{\sigma}\right) h_t(d) & \text{if } d \geq \pi_t V, \\ \eta\left[(1 - \pi_t)\Phi\left(\frac{\theta_F^* - \mu}{\sigma}\right) + \pi_t\Phi\left(\frac{\theta_T^* - \mu}{\sigma}\right)\right] h_t(d) & \text{if } d < \pi_t V; \end{cases}$$

where η is a normalizing constant. Since $\theta_F^* > \theta_T^*$, observing $S_t = 1$ induces a first-order stochastic increase in the distribution $H_{t+1}(\cdot|S_t = 1)$ relative to $H_t(\cdot)$. Intuitively, upon a success observed in country t, agents believe that the defense cost is more likely to be higher than $\pi_t V$ and less likely to be smaller than that. Such updating makes agents in country $t+1$ even more aggressive, because they believe that the probability of the regime choosing defense, $H_{t+1}(\pi_{t+1}V)$, is even smaller.

5. Concluding Remarks

It is not difficult to understand why people can drastically change their belief and hence their behavior following a rare event. A case in point is the plummeting of air travel shortly after the September 11 attack, which took years to recover to the pre-attack level.[28] Arguably, that was the safest time to travel because airport security was extremely tight. Nevertheless, our theory suggests that it was a rational response for people to avoid air travel: the totally unexpected terrorist attack had shaken their belief that the air travel system is secure. They tend to believe the system is risk-prone and need to accumulate further evidence to restore their confidence.

In the case of air travel, an increase in perception of risks was associated with increased efforts to counteract those risks. So beliefs and actions were not reinforcing each other. That is probably why we did not observe a wave of terrorist attacks of the same type. But in coordination games such as currency attacks, mass revolutions, and bank runs, an increase in the perception of fragility of the regime can lead to increased efforts to topple it. A successful attack elsewhere does raise such a perception dramatically for subsequent similar regimes. In sum, we show that the rarity of a single crisis, coupled with the complementarity between actions and beliefs, can be a powerful contagion mechanism and produce a wave of crises.

In a sense this paper attempts to revive the role of belief fragility in our thinking about currency crises or revolutions, which was central for the "second generation" of crisis models. In a standard global game framework, the fundamental strength of the regime is the key to the equilibrium outcome. In our model, the belief over possible models of world is as critical as the fundamental strength. A sudden shift in this belief can also arise, but it is provoked by rational belief updating over the model space, rather than a switch or jump between multiple equilibria. Further, since our model retains the advantage of having a unique equilibrium, we can bypass the issue of equilibrium selection and therefore offer an explicit link for the evolution of equilibrium.

[28]See "How 9/11 Changed America: in Statistics," *BBC News*, 1 September, 2006.

Appendix

Proof of Proposition 1. First, we assume a monotone equilibrium and show that it is unique when the condition holds that $D_j'(\theta)\phi(0)\sigma_x < \sigma^2$. Given x^*, let $\theta_j^*(x^*)$ satisfy the critical mass conditions (2) and (3). Differentiate these equations with respect to x^*, we obtain

$$\frac{\mathrm{d}\theta_j^*}{\mathrm{d}x^*} = \frac{\frac{1}{\sigma_x}\phi\left(\frac{x^*-\theta_j^*}{\sigma_x}\right)}{D_j'(\theta_j^*) + \frac{1}{\sigma_x}\phi\left(\frac{x^*-\theta_j^*}{\sigma_x}\right)} > 0.$$

Since $\phi(t)$ attains at maximum at $t = 0$, $\mathrm{d}\theta_j^*/\mathrm{d}x^*$ is bounded above by $\phi(0)/(D_j'(\theta_j^*)\sigma_x + \phi(0))$, which is less than $\beta = \sigma^2/(\sigma_x^2 + \sigma^2)$.

Let $g(x) = \pi P_T(\theta_T^*(x), \theta_F^*(x), x) + (1-\pi)P_F(\theta_F^*(x), x) - c$. This function is strictly decreasing because $\mathrm{d}\theta_j^*/\mathrm{d}x^* < \beta$. Moreover, $\lim_{x\to\infty} g(x) = -c$ and $\lim_{x\to-\infty} g(x) = 1 - c$. Therefore, for any $c \in (0,1)$, there exists a unique x^* that satisfies $g(x^*) = 0$. According to the critical mass conditions (2) and (3), for any x^*, there exists a unique pair (θ_T^*, θ_F^*).

Second, we show that non-monotone equilibria do not exist. Given a mass of attackers A, the attack succeeds in world j if $\theta \leq D_j^{-1}(A)$. The regime's best response to A is to choose $y(A,\theta) = 1$ if (a) $\theta > D_F^{-1}(A)$, or (b) $\theta > D_T^{-1}(A)$ and $d < \pi V$; and to choose $y(A,\theta) = 0$ otherwise. Given A and the regime's strategy $y(\cdot)$, the gain from attack for agent i is

$$G(A, x_i) = \pi \Pr[\theta \leq D_T^{-1}(A) \text{ or } y(A,\theta) = 0 \mid x_i]) + (1-\pi)\Pr[\theta \leq D_F^{-1}(A) \mid x_i] - c.$$

Since $D_j^{-1}(\cdot)$ is increasing and $y(\cdot,\theta)$ is non-increasing, $G(A, x_i)$ increases in A. Since $y(A,\theta)$ is non-decreasing in θ and the distribution of θ is stochastically increasing in x_i, $G(A, x_i)$ decreases in x_i. Let $A(\tau)$ represent the aggregate mass of attackers when agents adopt strategy τ, and consider the monotone strategy τ_κ, with $\tau_\kappa(x_i) = 1$ if $x_i < \kappa$ and $\tau_\kappa(x_i) = 0$ if $x_i > \kappa$. Then $G(A(\tau_\kappa), \kappa) = g(\kappa)$, and we have already shown that there is a unique x^* such that $g(x^*) = 0$. Therefore, the monotone strategy τ_{x^*} is an equilibrium strategy.

Now, consider any equilibrium strategy τ' of the game. Let $\underline{x}$ be the smallest x_i such that $G(A(\tau'), x_i) \leq 0$ and let $\overline{x}$ be the largest x_i such that $G(A(\tau'), x_i) \geq 0$. We have $\underline{x} \leq \overline{x}$. Since $G(A(\tau'), x_i) > 0$ for any $x_i < \underline{x}$, we have $\tau'(x_i) = 1$ for $x_i < \underline{x}$. Thus, $A(\tau') \geq A(\tau_{\underline{x}})$, which in turn implies $g(\underline{x}) = G(A(\tau_{\underline{x}}), \underline{x}) \leq G(A(\tau'), \underline{x}) \leq 0$. But since $g(\cdot)$ is decreasing, we must have $\underline{x} \geq x^*$. A parallel argument establishes that $\overline{x} \leq x^*$. Together, they imply $\underline{x} = \overline{x} = x^*$. Thus, $\tau' = \tau_{x^*}$. In other words, any equilibria except the monotone equilibrium do not exist.

Use the implicit theorem and differentiate equation $g(\cdot)$ with respect to π to obtain:

$$g'(x^*)\frac{\mathrm{d}x^*}{\mathrm{d}\pi} = (P_F - P_T)\left(1 + \frac{\pi V h(\pi V)}{H(\pi V)}\right).$$

Since $g'(\cdot) < 0$ and $P_F > P_T$, $\mathrm{d}x^*/\mathrm{d}\pi < 0$. Moreover, for $j = T, F$, $\mathrm{d}\theta_j^*/\mathrm{d}\pi = (\mathrm{d}\theta_j^*/\mathrm{d}x^*)(\mathrm{d}x^*/\mathrm{d}\pi) < 0$, which in turn implies $\mathrm{d}p_j/\mathrm{d}\pi < 0$. ■

Proof of Proposition 2. Write $g(x; k_T) = \pi P_T(\theta_T^*(x; k_T), \theta_F^*(x), x) + (1-\pi)P_F(\theta_F^*(x), x) - c$, where $\theta_j^*(\cdot)$ solves the respective critical mass condition in world j. Implicit differentiation of the indifference condition $g(x^*; k_T) = 0$ gives

$$g'(x^*)\frac{\mathrm{d}x^*}{\mathrm{d}k_T} = -\pi \frac{\partial P_T}{\partial \theta_T^*}\frac{\partial \theta_T^*}{\partial k_T}.$$

Thus $\mathrm{d}x^*/\mathrm{d}k_T < 0$, because $g' < 0$, $\partial P_T/\partial\theta_T^* < 0$, and $\partial\theta_T^*/\partial k_T < 0$. This in turn implies:

$$\begin{aligned}
\frac{\mathrm{d}\theta_T^*}{\mathrm{d}k_T} &= \frac{\mathrm{d}\theta_T^*}{\mathrm{d}x}\frac{\mathrm{d}x^*}{\mathrm{d}k_T} + \frac{\partial\theta_T^*}{\partial k_T} < 0,\\
\frac{\mathrm{d}\theta_F^*}{\mathrm{d}k_T} &= \frac{\mathrm{d}\theta_F^*}{\mathrm{d}x}\frac{\mathrm{d}x^*}{\mathrm{d}k_T} < 0.
\end{aligned}$$

Now, combine the critical mass conditions in the two worlds to obtain:

$$\Phi\left(\frac{x^* - \theta_T^*}{\sigma_x}\right) - \Phi\left(\frac{x^* - \theta_F^*}{\sigma_x}\right) = D_T(\theta_T^*; k_T) - D_F(\theta_F^*; k_F).$$

Since the left-hand-side of the above is bounded, we must have $\theta_F^* - \theta_T^*$ going to infinity as k_T goes to infinity. Let $\hat{g}(x) = (1-\pi)P_F(\theta_F^*(x), x) - c$ and let $\overline{x}$ solves $\hat{g}(x) = 0$. Since $\hat{g}(x) < g(x)$ for any x, we have $x^* > \overline{x}$ for any k_T, and thus $\theta_F^* > \theta_F^*(\overline{x})$ for any k_T. Thus, as k_T goes to infinity, we must have θ_T^* going to negative infinity, while θ_F^* and x^* remaining bounded.

Because p_T is increasing in θ_T^* and θ_F^*, which are decreasing in k_T, we have $\mathrm{d}p_T/\mathrm{d}k_T < 0$. This proves part (a). Further,

$$\frac{\mathrm{d}\theta_T^*}{\mathrm{d}k_T} - \frac{\mathrm{d}\theta_F^*}{\mathrm{d}k_T} = \left(\frac{\mathrm{d}\theta_T^*}{\mathrm{d}x} - \frac{\mathrm{d}\theta_F^*}{\mathrm{d}x}\right)\frac{\mathrm{d}x^*}{\mathrm{d}k_T} + \frac{\partial\theta_T^*}{\partial k_T}.$$

Because x^* is decreasing but remains bounded as k_T increases, we have $\mathrm{d}x^*/\mathrm{d}k_T$ going to 0 for k_T sufficiently large. On the other hand, $\partial\theta_T^*/\partial k_T$ is negative and bounded away from 0. Therefore, $\mathrm{d}\theta_T^*/\mathrm{d}k_T < \mathrm{d}\theta_F^*/\mathrm{d}k_T$ for k_T large. The derivative of p_T/p_F

with respect to k_T has the same sign as:

$$\frac{\phi\left(\frac{\theta_T^*-\mu}{\sigma}\right)}{\Phi\left(\frac{\theta_T^*-\mu}{\sigma}\right)}\frac{\mathrm{d}\theta_T^*}{\mathrm{d}k_T}-\frac{\phi\left(\frac{\theta_F^*-\mu}{\sigma}\right)}{\Phi\left(\frac{\theta_F^*-\mu}{\sigma}\right)}\frac{\mathrm{d}\theta_F^*}{\mathrm{d}k_T}.$$

The function $\phi(\cdot)/\Phi(\cdot)$ is decreasing, and $\mathrm{d}\theta_T^*/\mathrm{d}k_T < \mathrm{d}\theta_F^*/\mathrm{d}k_T < 0$. Therefore p_T/p_F decreases in k_T for k_T sufficiently large. This proves part (b). As k_T goes to infinity, p_T/p_F approaches $1 - H(\pi V)$. The limit value of π_{t+1} follows from the Bayes' formula. ■

References

Allen, F. and D. Gale (2000). "Financial contagion." *Journal of Political Economy 108*(1): 1–33.

Angeletos, G.-M., C. Hellwig, and A. Pavan (2006). "Signaling in a global game: Coordination and policy traps." *Journal of Political Economy 114*(3): 452–484.

Angeletos, G.-M., C. Hellwig, and A. Pavan (2007). "Dynamic global games of regime change: Learning, multiplicity, and timing of attacks." *Econometrica 75*(3): 711–756.

Banerjee, A. V. (1992). "A simple model of herd behavior." *Quarterly Journal of Economics 107*(3): 797–817.

Bernhardt, D. and M. Shadmehr (2011). "Collective action with uncertain payoffs: Coordination, public signals and punishment dilemmas." *American Political Science Review 105*(4), 829–851.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of Political Economy 100*(5): 992–1026.

Bordo, M. D. and A. P. Murshid (2000). "Are financial crises becoming more contagious? What is the historical evidence on contagion?" Pp. 367–404 in S. Claessens and K. Forbes, eds., *International Financial Contagion*. Boston: Kluwer.

Calvo, G. A. and E. G. Mendoza (2000). "Rational contagion and the globalization of securities markets." *Journal of International Economics 51*(1): 79–113.

Chamley, C. (1999). "Coordinating regime switches." *Quarterly Journal of Economics 114*(3), 869–905.

Chen, H., Y. Lu, and W. Suen (2014). "The power of whispers: A theory of rumor, communication and revolution." *International Economic Review*, forthcoming.

Dasgupta, A., R. Leon-Gonzalez, and A. Shortland (2011). "Regionality revisited: An examination of the direction of spread of currency crises." *Journal of International Money and Finance 30*(5), 831–848.

Diamond, D. W. and P. H. Dybvig (1983). "Bank runs, deposit insurance, and liquidity." *Journal of Political Economy 91*(3): 401–419.

Drazen, A. (2000). "Political contagion in currency crises." Pp. 47–70 in P. Krugman, ed., *Currency Crises*. Chicago: University of Chicago Press.

Edmond, C. (2013). "Information manipulation, coordination, and regime change." *Review of Economic Studies 80*(4): 1422–1458.

Eichengreen, B., A. Rose, and C. Wyplosz (1996a). "Contagious currency crises: First tests." *Scandinavian Journal of Economics 98*(4): 463–484.

Eichengreen, B., A. K. Rose, and C. Wyplosz (1996b). "Speculative attacks on pegged exchange rates: An empirical exploration with special reference to the European Monetary System." Pp. 191–228 in M. Canzoneri, W. Ethier and V. Grilli, eds., *The New Transatlantic Economy*. Cambridge: Cambridge University Press.

Eichengreen, B. and C. Wyplosz (1993). "The unstable EMS." *Brookings Papers on Economic Activity 1993*(1): 51–143.

Francisco, R. A. (1993). "Theories of protest and the revolutions of 1989." *American Journal of Political Science 37*(3): 663–680.

Gerlach, S. and F. Smets (1995). "Contagious speculative attacks." *European Journal of Political Economy 11*(1): 45–63.

Goldstein, I., E. Ozdenoren, and K. Yuan (2011). "Learning and complementarities in speculative attacks." *Review of Economic Studies 78*(1): 263–292.

Goldstein, I. and A. Pauzner (2005) "Demand-deposit contracts and the probability of bank runs." *Journal of Finance 60*(3): 1293–1327.

Hellwig, C., A. Mukherji, and A. Tsyvinski (2006). "Self-fulfilling currency crises: The role of interest rates." *American Economic Review 96*(5): 1769–1787.

Jalil, A. J. (2013). "A new history of banking panics in the United States, 1825–1929: Construction and implications." *American Economic Journal: Macroeconomics*, forthcoming.

Katz, M. N. (1997). *Revolutions and Revolutionary Waves*. New York: St. Martin's Press.

Kindleberger, C. P. and R. Z. Aliber (2011). *Manias, Panics and Crashes: A History of Financial Crises*, 6th ed. Hants: Palgrave Macmillan.

Krugman, P. (1999). "Currency crises." Pp. 420–440 in M. Feldstein, ed., *International Capital Flows*. Chicago: University of Chicago Press.

Kuran, T. (1989). "Sparks and prairie fires: A theory of unanticipated political revolution." *Public Choice 61*(1): 41–74.

Kuran, T. (1997). *Private Truths, Public Lies: The Social Consequences of Preference Falsification*. Cambridge: Harvard University Press.

Masson, P. R. (1999). "Contagion: Macroeconomic models with multiple equilibria." *Journal of International Money and Finance 18*(4): 587–602.

Morris, S. and H. S. Shin (1998). "Unique equilibrium in a model of self-fulfilling currency attacks." *American Economic Review 88*(3): 587–597.

Obstfeld, M. (1986). “Rational and self-fulfilling balance-of-payments crises.” *American Economic Review 76*(1): 72–81.

Obstfeld, M. (1994). “The logic of currency crises.” *Cahiers Economiques et Monetaires* (43): 189–213.

Ortoleva, P. (2012). “Modeling the change of paradigm: Non-Bayesian reactions to unexpected news.” *American Economic Review 102*(6): 2410–2436.

Page, S. E. (2006). “Essay: Path dependence.” *Quarterly Journal of Political Science 1*: 87–115.

Pesenti, P. and C. Tille (2000). “The economics of currency crises and contagion: An introduction.” *Economic Policy Review 6*(3): 3–16.

Reinhart, C. M. and K. Rogoff (2009). *This Time Is Different: Eight Centuries of Financial Folly*. Princeton: Princeton University Press.

Wald, A. and J. Wolfowitz (1940). “On a test whether two samples are from the same population.” *Annals of Mathematical Statistics 11*: 147–162.

Walt, S. M. (1992). “Revolution and war.” *World Politics 44*(3): 321–368.