On the Epidemic of Financial Crises

Nikolaos Demiris
Athens University of Economics & Business, Athens, Greece

Theodore Kypraios
University of Nottingham, Nottingham, UK

L. Vanessa Smith
University of Cambridge, Cambridge, UK

June 16, 2013

Abstract

This paper proposes a framework for modelling financial contagion that is based on SIR (Susceptible-Infected-Recovered) transmission models from epidemic theory. This class of models addresses two important features of contagion modelling, which are a common shortcoming of most existing empirical approaches, namely the direct modelling of the inherent dependencies involved in the transmission mechanism, and an associated canonical measure of crisis severity. The proposed methodology naturally implies a control mechanism, which is required when evaluating prospective immunisation policies that intend to mitigate the impact of a crisis. It can be implemented not only as a way of learning from past experiences, but also at the onset of a contagious financial crisis. The approach is illustrated on a number of currency crisis episodes, using both historical final outcome and temporal data. The latter require the introduction of a novel hierarchical model that we call the Hidden Epidemic Model (HEM), and which embeds the stochastic financial epidemic as a latent process. The empirical results suggest, among other, an increasing trend for global transmission of currency crises over time.

Keywords: Financial crisis, contagion, stochastic epidemic model, random graph, MCMC

1 Introduction

A financial crisis originating in one country can travel within and beyond its original neighbourhood spreading among countries like a contagious disease. Loosely speaking, this phenomenon, which was a common feature of the major recent crises, is referred to by economists as 'contagion'. Contagion by definition can only occur if there are interactions among subjects. These interactions can materialise at different levels and through different channels simultaneously, some channels being more important during particular events than others. The literature typically distinguishes between fundamentals-based contagion, a transmission of a crisis from one country to another through real and financial linkages, also known as spillovers, and pure contagion where a crisis might trigger additional crises elsewhere for reasons unexplained by fundamentals. As financial crises spread across countries, they affect nations with apparently healthy fundamentals and sound policies. The better understanding we have of their propagation mechanism, the better we are positioned in proposing policy interventions which can most effectively reduce their contagious spread.

The economics and finance literature includes a number of theoretical models that aim to explain the contagious spread of crises, emphasising trade linkages (bilateral or third party, Gerlach and Smets, 1995; Corsetti et al., 1998), financial linkages (Allen and Gale, 2000), as well as models on information asymmetries and investor behaviour (Calvo and Mendoza, 2000; Kodres and Pritsker, 2002), among others. More recently, recognition of the inherent complexities and interconnections associated with financial systems has advocated the need to consider methods from other disciplines such as ecology, epidemiology, biology and engineering in studying financial networks (May et al. 2008, Haldane, 2009). Indeed, researchers have turned to alternative methods in modelling financial contagion, which are largely based on numerical simulations. For example, Gai and Kapadia (2010) study a percolation-type process on a weighted network of banks as a model of contagion. (Bond) percolation processes are equivalent to SIR epidemics given appropriate model specification (Newman, 2010). May and Arinaminpathy (2010) pursue recent advances in the area of complex ecological systems in a study similar in spirit to Gai and Kapadia, though they employ a mean field approximation rather than resorting to simulations. Amini et al. (2010) analyse distress propagation in a network of banks, via a cascading process, and derive the asymptotic magnitude of contagion. Caporale et al. (2009) make use of agent-based

simulation models to investigate the dynamics of financial contagion.

On the empirical front considerable effort has been devoted to documenting the existence of contagion.Various tests have been proposed in the literature for this purpose including among others the Forbes and Rigobon (2002) test based on correlation coefficients and the Pesaran and Pick (2007) threshold test. Dungey et al. (2005a, 2005b) provide a review and comparison of alternative tests of contagion. The bulk of the studies suggest that there is evidence of contagion. For an overview of the empirical evidence of contagion see Dornbusch et al. (2000) and Pericoli and Sbracia (2003). Less emphasis, however, has been placed on modelling the propagation mechanism of financial crises. A class of models that features prominently in the empirical contagion literature is that of cross country probit-type regressions of a binary crisis indicator on variables representing potentially important transmission channels. Such channels have been identified as trade links (Glick and Rose, 1999), financial links and the common creditor (Caramazza et al., 2004, and Kaminsky and Reinhart, 2000), neighborhood effects (De Gregorio and Valdes, 2001), and macroeconomic similarities (Sachs et al., 1996). More recently, Dungey and Martin (2007) have proposed capturing financial market linkages during crises through the use of common factors, while Aït-Sahalia et al. (2010) suggest modelling financial contagion using mutually exciting processes.

In this paper we propose a framework for modelling financial contagion which is based on SIR (Susceptible-Infected-Recovered) transmission models from epidemic theory (see for example, Bailey, 1975). This class of models addresses two important features of contagion modelling which are a common shortcoming of most existing empirical approaches, namely the direct modelling of the inherent dependencies involved in the transmission mechanism, and an associated canonical measure of crisis severity. At the same time, it allows to incorporate features that reflect relevant theoretical and empirical evidence in the literature through the inclusion of appropriate covariates. The proposed methodology naturally implies a control mechanism, which is required when evaluating prospective immunisation policies that intend to mitigate the impact of a crisis. This control mechanism can be implemented not only as a way of learning from past experiences, but also in real-time as a contagious financial crisis unfolds.

Specifically, the approach is based on a stochastic epidemic transmission process where the population of countries is explicitly structured, and a crisis may be propagated both locally and globally. Having identified all countries that suffered during a particular crisis

episode, the focus is on how crises spread from an initially affected (or 'infected') country to the rest of the countries by considering their interaction at the local and global level. We explicitly model regional and global contagious 'contacts' and infer the rate of their transmission. These transmission rates allow us to directly quantify the severity of the crisis, in contrast to conventional approaches in the literature where severity is measured by a composite index of macroeconomic indicators, see Kaminsky and Reinhart (1998) and Kaminsky (2006). The methodology also delivers an estimate of the number of countries to financially support in order to prevent a major crisis.

In the proposed framework a country may experience a crisis not only because of its direct links to the originally infected (ground zero) country, but also due to local and/or global contacts with countries that are already infected. Thus, we allow for the so called 'cascading effect' following the terminology of Glick and Rose (1999), which is typically ignored in the literature. In fact, the proposed approach effectively considers the set of all possible transmission channels of a crisis. Furthermore, it naturally accounts for an increase in the likelihood of a crisis in a particular country given that there is a crisis elsewhere. This is explored in Eichengreen et al. (1996) and Kaminsky and Reinhart (2000) as being consistent with the existence of contagion.

The methodology is illustrated on a number of currency crisis episodes, using both historical final outcome and temporal data. The former provides information on the number of initially 'healthy' countries, that were ever affected or otherwise by a particular financial crisis. The information provided by the latter allows to estimate the 'infection' and/or 'recovery' times of the crisis, in addition to the trasmission rates. The use of temporal data is necessary for performing real-time analysis of a crisis spread. It requires the introduction of a novel hierarchical model, that we call the Hidden Epidemic Model (HEM), and which embeds the stochastic financial epidemic as a latent process. Our findings, among others, point to an increasing trend for global transmission of currency crises over time.

The paper is structured as follows. In Section 2 we introduce the modelling framework and describe the inference procedure relevant to final outcome data. Section 3 describes the hidden epidemic model associated with the use of temporal data. In Section 4 the proposed approach is illustrated with an empirical application to several currency crisis episodes. Section 5 discusses the relevance of the methodology for policy analysis. Section 6 contains some further discussion.

2 Modelling Framework

Consider the onset of a financial crisis, for example, a currency crisis or a stock market crash, with contagious effects to a given population of, say, countries. Our interest is in modelling the process associated with the propagation of the crisis. The approach can be equally applied to populations of firms or banks, among others.

2.1 The Model

Consider a closed population of N countries, partitioned into regions of varying sizes. Specifically we will assume that the population contains r_j regions/groups of size j, where $r = \sum_{j=1}^{S} r_j$ is the total number of regions and S denotes the size of the largest group. The total number of countries is then $N = \sum_{j=1}^{S} j r_j$. The crisis originates in a typically small number of countries. At the outset, all countries are deemed susceptible to the particular crisis episode. As the crisis unfolds, each country is considered to belong to any one of three states: susceptible, infected (and 'infectious') or recovered. Such processes are often referred to as SIR models in epidemic theory. In the SIR model a susceptible country is in a 'normal' state and can be affected by the crisis in question. An infective country is in a state of crisis and may trasmit it to other countries. At the end of its infectious period an individual country is considered recovered, in the sense that it plays no further role in the propagation of the crisis. Thus, we effectively assume that a country may not experience multiple recoveries within a single crisis, which appears reasonable for such applications.

A country, say j, remains in crisis for a positive random time, denoted by I_j, which is allowed to follow any specified distribution. The periods for which the individual countries remain in crisis are assumed to be a-priori independent. While infectious, a country may transmit the crisis to each country within its region at times given by the points of a Poisson process of rate λ_L. Additionally, each country may trasmit the crisis to any given country worldwide according to a Poisson process with rate λ_G/N. That is, there are two levels of mixing between countries, namely at the local and global level. The contact processes within this construction are assumed to be a-priori independent. This formulation of the model effectively assumes that each country has local contacts with rate $\lambda_L + \lambda_G/N$. Alternative parameterisations of the epidemic model can be obtained using the superposition and splitting of the Poisson process (see Kingman, 1993). The crisis ends when there are

no remaining infectious countries in the population. Notice the different scaling of the two transmission rates. This is a common assumption in two-level-mixing models and it implies that there will be more infectious contacts locally if the local population grows, while this is not the case globally.

The model described above is a stochastic epidemic two level mixing model introduced in Ball et al. (1997), where a detailed discussion of related work can be found. While the model does not explicitly assume a latent period for the crisis to unfold once a particular country has been affected, but rather assumes immediate infection, the distribution of the final outcome is invariant to very general assumptions concerning a latent period (see Ball et al., 1997; Andersson and Britton, 2000). The two level mixing model encompasses the so called generalised stochastic epidemic model (GSE), which arises by setting $(\lambda_L = 0, \lambda_G = \lambda)$, that is, assuming all countries mix homogeneously with rate λ. The particular case where the infectious period is exponentially distributed renders the model Markov. As the population size increases, it can be shown (e.g. Ethier and Kurtz 1986) that the solution of this Markov process converges to that of the corresponding system of ODEs known as the deterministic general epidemic.

Epidemic models in their simple deterministic form are not foreign to the economics literature. They typically feature in studies of treatment and control, see for example Geoffard and Philipson (1997), Gersovitz and Hammer (2004) and Toxvaerdy (2010). However, for a complex, highly non-linear phenomenon such as contagion, the stochastic model described above is more appropriate compared to any simple deterministic analog. It is particularly advantageous in that it is suitable for small populations, it can account for the regional/global nature of a crisis, while model complexity and realism can be naturally expanded in different directions, some of which will be discussed below.

Figure 1 illustrates a potential configuration of a crisis spread in the case of a small population of five countries denoted by {1,2,...,5}. The five countries are partitioned into two regions, represented by the two circles. Consider 1 as the country originally affected by the crisis in question. The grey links denote contacts made locally between countries within a region. The black links correspond to global contacts. The solid directed links represent contacts that resulted in infection while the dashed links represent lack of contact in both directions. Thus, in this particular configuration the crisis is transmitted from country 1 to countries 2 and 4.

Figure 1: Example of a crisis spead over a small population of countries

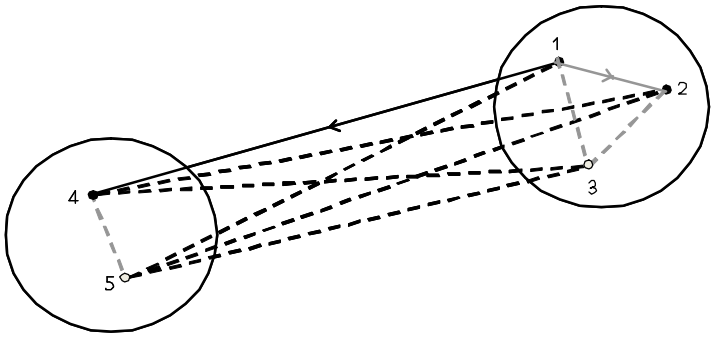

In the specification described above the population of countries is assumed to represent a complete network i.e. any country can have contacts with any other locally and globally. This is not a restrictive assumption as epidemics upon random networks can be reparameterised to correspond to epidemics on complete graphs (e.g. Neal, 2006, Newman, 2010), and thus can be incorporated into the proposed framework. Furthermore, the population is assumed to be homogeneous, that is, all countries posses similar infectivity and susceptibility properties. This assumption will be relaxed in Section 4.1.1.

Further detail can be added to the model by disaggregating the population according to individual attributes and/or relationship networks based, for example, on economic information. These details can be incorporated through the transmission rates leading to weighted networks. The study of a dynamic process on a weighted network is, essentially, equivalent to the inclusion of a covariate. This can be easily seen by considering the moment generating function of the transmission rate, say $\phi(\lambda)$, where $\phi(\lambda) = E(\exp(-\lambda I))$. In the case of a constant infectious period this is simply $\exp(-\lambda I)$, the probability of avoiding infection. When information about a particular variable, say X, becomes available, this can be incorporated using a Cox-type (log-linear) model: $\lambda = \lambda_0 \exp(\alpha X)$. Then $\phi(\lambda)$ becomes $\phi(\lambda) = \exp[-I\lambda_0 \exp(\alpha X)] = w \exp(-\lambda_0 I) = w\phi(\lambda_0)$. In other words, the inclusion of covariates using the proposed methodology is equivalent to the simultaneous estimation of the weights on our network. Dealing with multiple covariates (equivalently, weighting based on more than one variable) is straightforward as will be illustrated in Section 4.2.2. If the focus is, for example, on the contagious default of banks these covariates could represent balance sheet information and/or information on interbank relations. In the case of a currency crisis, they could represent trade and financial linkages between countries. Additional covariates accounting for macroeconomic conditions could also be included.

2.2 Threshold Parameter

An appealing feature of the proposed modelling framework is that it embodies a threshold parameter which can be utilised as a severity measure of the crisis, thus providing an inherent control mechanism. Stochastic models such as epidemic and branching processes typically generate bimodal realisations where an epidemic may or may not die out quickly, depending on the value of the threshold parameter R_0, often referred to as the basic reproduction number. This is the most important parameter in epidemic theory and is defined as the expected number of infections generated by a typical infective in an infinite susceptible population. Generally, in deterministic epidemic models, R_0 is calculated as the largest eigenvalue of the so called next generation matrix (Diekmann and Heesterbeek, 2000).

In the case of a homogeneously mixing stochastic epidemic, it holds that $R_0 = \lambda E(I)$, where λ is the contact rate and I is the infectious period. R_0 may be interpreted as a threshold parameter since its value determines whether or not a major crisis can occur. In particular, if $R_0 > 1$ a positive proportion of the infinite population will be affected by the crisis with positive probability, while if $R_0 \leq 1$ only a finite number of susceptibles will ever become infected and thus the crisis spread will swiftly die out. Hence, most control measures aim to reduce R_0 below unity. Ball and Donelly (1995) rigorously established the threshold behaviour of the general stochastic epidemic by coupling the early stages of the epidemic with a suitable branching process.

For the two-level mixing model the threshold parameter, denoted by R_*, is defined in a similar way to R_0. In this case, Ball et al. (1997) couple the epidemic with a branching process defined on groups. They show that

$$R_* = \lambda_G E(I) \mathrm{v}(\lambda_L), \tag{1}$$

where $\mathrm{v}(\lambda_L)$ is the average group final size if only local infections are permitted. Specifically, $\mathrm{v}(\lambda_L) = \frac{1}{g}\sum_{j=1}^{S} j\mu_j\pi_j$, where μ_j is the final size within a group of size j, $\pi_j = r_j/r$ is the proportion of groups of size j and g is the mean group size. It should be noted that while R_* is linear in λ_G, it is nonlinear in λ_L but linear in $\mathrm{v}(\lambda_L)$, with the latter being an increasing function of λ_L. Details of the calculation of $\mathrm{v}(\lambda_L)$ can be found in the Supplement.

Most statistical analyses require the assumption of supercriticality, that is $R_* > 1$ so that a major crisis spread is possible (e.g. Rida, 1991; Demiris and O'Neill, 2005b). Since

prophylactic measures aim to achieve $R_* \leq 1$, assuming that $R_* > 1$ is not desirable while it also results in underestimating the variability of the model parameters. The inference procedure we will employ does not condition upon $R_* > 1$.

2.3 Inference Using Final Outcome Data

2.3.1 Final Size Distribution

For a given crisis episode, when only information on the final infected number of countries is available, i.e. final outcome data, it is in principle possible to write down a system of recursive linear equations, the solution of which delivers the probability mass function required for the likelihood. However, solving such a set of equations can be numerically unstable even for small population sizes (see for example Andersson and Britton, 2000, p.18). This problem is greatly amplified for more complex cases like the two level mixing model considered here. The remainder of this section is concerned with the description of an inference procedure that surmounts this complication.

2.3.2 Data and Likelihood

For the two level mixing model, the data are of the form $\mathbf{x} = \{x_{ij}\}$ where x_{ij} denotes the number of regions containing j initially susceptible countries of which i ever suffered the crisis in question. We will describe a Bayesian inference procedure for the two infection rates λ_L and λ_G, given $\mathbf{x}$. By Bayes' Theorem, the posterior density, $\pi(\lambda_L, \lambda_G \mid \mathbf{x})$, satisfies $\pi(\lambda_L, \lambda_G \mid \mathbf{x}) \propto \pi(\mathbf{x} \mid \lambda_L, \lambda_G)\pi(\lambda_L, \lambda_G)$, where $\pi(\mathbf{x} \mid \lambda_L, \lambda_G)$ denotes the likelihood function and $\pi(\lambda_L, \lambda_G)$ the joint prior of (λ_L, λ_G). The numerical problems mentioned in Section 2.3.1 indicate that the likelihood is analytically and numerically intractable for all but a very small number of countries. We surpass this difficulty by following Demiris and O'Neill (2005a) in augmenting the parameter space using an appropriate random directed graph (digraph) as considered below.

2.3.3 Random Digraph

The characterisation of the final outcome of a stochastic epidemic in terms of random graphs is well known in epidemic theory and has been exploited by Ludwig (1975) and Barbour and Mollison (1990) among others. It has been considered for the purposes of statistical inference

by Demiris and O'Neill (2005a); see also O'Neill (2009). In short, let each country correspond to a vertex in the random digraph with a directed edge denoting a potential infection. The edges are drawn from vertex j with probability $1 - \exp(-I_j\lambda_L)$ $(1 - \exp(-I_j\lambda_G/N))$ for local (global) links. It then holds that the distribution of the final outcome of the crisis is equivalent to the distribution of the 'giant component' of the graph, that is, the random set of vertices that are connected to the initially affected country (or countries) through directed edges. Additional edges may exist in the digraph, but only the routes emanating from the initially infected can correspond to actual 'infections'. Flexible models for the infectious periods I_j such as disjoint intervals in real time, representing for example stock market opening times in the case of a stock market crash, can be easily incorporated into this framework as they only appear through the edge probabilities. To maximise computational efficiency we only consider the graph on 'infected' vertices, while all other data contributions are accounted for through the likelihood.

2.3.4 Augmented Likelihood and Posterior

Suppose that the total number of countries that are ever affected by a particular crisis is $n = \sum_i \sum_j i x_{ij}$, labeled $1, \ldots, n$, and define G as the random digraph on these n vertices. For $j = 1, \ldots, n$ let I_j denote the infectious period corresponding to vertex j and $\boldsymbol{I} = (I_1, \ldots, I_n)$. We assume that initially there is one known country in crisis, although it is trivial to consider any number of initial infectives, possibly of unknown identity. In fact, it turns out that knowledge of the initial infective is not crucial in the applications to follow**.**

The augmented posterior density may be written as

$$\pi(\lambda_L, \lambda_G, \boldsymbol{I}, G \mid \mathbf{x}) \propto \pi(\mathbf{x} \mid \lambda_L, \lambda_G, \boldsymbol{I}, G)\pi(G \mid \lambda_L, \lambda_G, \boldsymbol{I})\pi(\boldsymbol{I})\pi(\lambda_L, \lambda_G). \tag{2}$$

In (2), $\pi(\lambda_L, \lambda_G)$ and $\pi(\boldsymbol{I})$ denote the priors while the first two terms effectively represent the augmented likelihood, given by $L(\mathbf{x} \mid \lambda_L, \lambda_G, \boldsymbol{I}, G)$. Provided that G is compatible with the data, $L(\mathbf{x} \mid \lambda_L, \lambda_G, \boldsymbol{I}, G)$ is evaluated as the probability of the edges in G times the probability of no edges between the n vertices in G and the remaining $N - n$ vertices. Otherwise, $L(\mathbf{x} \mid \lambda_L, \lambda_G, \boldsymbol{I}, G)$ is set to zero.

Let ℓ_j^L (ℓ_j^G) denote the number of local (global) links emanating from vertex j and N_j^L the number of countries in region j. Further, let $G \sim \mathbf{x}$ denote the event that G is compatible

with the data. Then the augmented likelihood can be written as

$$L(\mathbf{x} \mid \lambda_L, \lambda_G, \boldsymbol{I}, G) := \pi(\mathbf{x} \mid \lambda_L, \lambda_G, \boldsymbol{I}, G)\pi(G \mid \lambda_L, \lambda_G, \boldsymbol{I}) = \mathbf{1}_{G \sim \mathbf{x}} \prod_{j=1}^{n} \Big[\{1 - \exp(-\lambda_L I_j)\}^{\ell_j^L}$$
$$\exp\left(-\lambda_L I_j \left(N_j^L - \ell_j^L\right)\right) \left\{1 - \exp\left(-\frac{\lambda_G I_j}{N}\right)\right\}^{\ell_j^G} \exp\left(-\frac{\lambda_G I_j (N - \ell_j^G)}{N}\right)\Big], \tag{3}$$

where $\mathbf{1}_C$ denotes the indicator function which takes the value of one when the event C materialises, and zero otherwise. The MCMC algorithm used to sample from (2) is largely similar to that in Demiris and O'Neill (2005a) where additional details can be found. In brief, standard random walk Metropolis samplers are sufficient for updating the infection rates, while updating G requires some attention. Specifically, G is a discrete random object with an enormous number of possible configurations, the overwhelming majority of which have negligible posterior probability. Hence, simple strategies like adding and deleting one edge at a time are preferable, as samplers based on more complex proposals can exhibit poor convergence.

3 A Hidden Epidemic Model

Thus far, analysis of the proposed model has been discussed for the case where only final outcome data are available. Next we illustrate how one can obtain additional information about the propagation of a financial crisis, for example the times of entry to and/or exit from the crisis, by making use of temporal data. This way a more complete characterisation of the crisis can be obtained, including its duration for each country. Contrary to a communicable disease where it is common to observe the times at which individuals developed symptoms or recovered, this is not typically the case for a financial crisis. The use of temporal data in the current context necessitates the introduction of a novel hierarchical model that we call the Hidden Epidemic Model (HEM). The HEM embeds the stochastic financial epidemic as a latent process that governs the behaviour of the observed temporal variables. One of the important advantages of inferring the transmission parameters using the hidden epidemic model is that it can be used for real-time analysis of a financial crisis spread and hence for the evaluation of prospective immunisation policies. Hierarchical models, specifically Hidden Markov Models, have been used in the recent literature to analyse contagion, see for example

Gravelle et al. (2006) and Mandilaras and Bird (2010). However, the focus of these studies is on detecting contagion with the former focusing on changes in volatilities and the latter on correlations across different crisis and non-crisis periods. As such, they do not contain any transmission features associated with the highly complex phenomenon of contagion.

3.1 Model Structure

Consider the observed variable Y_{it} for country i at time t where $i = 1, \ldots, N$ and $t = 1, \ldots, T$. Let S_{it} denote the state of country i at time t, and E_i and R_i the times of entry to and recovery from the crisis. The state variable S_{it} is defined as

$$S_{it} = \begin{cases} 1, & t < E_i \quad \text{(susceptible)} \\ 2, & E_i < t < R_i \quad \text{(infected)} \\ 3, & t > R_i \quad \text{(recovered).} \end{cases} \tag{4}$$

For any country i not affected by the crisis $E_i = R_i = \infty$. Hence, the total final size $l = \sum_i^N \mathbf{1}_{(E_i \neq \infty)}$.

The conditional distribution of $Y_{it}|S_{it}$ is assumed to follow a Student's t distribution with ν degrees of freedom, with mean and variance that depend on S_{it} in the following manner:

$$Y_{it}|S_{it} \sim \begin{cases} t_\nu(\mu_1, \sigma_{i1}), & \text{if} \quad S_{it} = 1, \\ t_\nu(\mu_2, \sigma_{i2}), & \text{if} \quad S_{it} = 2, \\ t_\nu(\mu_3, \sigma_{i3}), & \text{if} \quad S_{it} = 3. \end{cases}$$

The choice of the t distribution was based on preliminary results of the properties of the temporal variable used in the subsequent empirical application. In principle it can be specified as any distribution. We also make the following assumption which characterises the dependence structure of the model: $Y_{1t}, Y_{2t}, \ldots, Y_{Nt}$ conditionally on $S_{1t}, S_{2t}, \ldots, S_{Nt}$ are independent for any t, i.e., $f\left(Y_{1t}, Y_{2t}, \ldots Y_{Nt}|S_{1t}, S_{2t}, \ldots, S_{Nt}\right) = f\left(Y_{1t}|S_{1t}\right) f\left(Y_{2t}|S_{2t}\right) \cdots f\left(Y_{Nt}|S_{Nt}\right), \quad t = 1, \ldots T$, where $f(\cdot)$ denotes the density function of a Student's t distribution. That is, while the state of each country (and thus the log-returns) depends on the state of all other countries in a complex non-linear stochastic manner, conditional upon the state of the country we assume independent variation in the log-return Y_{it}. To capture the inherent dependencies involved in the transmission mechanism of a crisis, we model the state variables

$\mathbf{S} = (S_{it}, i = 1, \ldots, N, t = 1, \ldots, T)$ by the stochastic epidemic two level of mixing model described in Section 2. As in the standard Hidden Markov Models used in the literature to identify currency crises, see Martinez-Peria (2002) and Brunetti et al. (2008) among others, the above specification can be extended to allow for an autoregressive structure in the mean of the observed series, as well as conditional heteroskedasticity. The latter can be accommodated by allowing the volatilities to follow, for example, a standard GARCH form. Such extensions will not be pursued here.

3.2 Inference

Without loss of generality we assume that $\mu_\ell = 0$, for $\ell = 1, 2, 3$ as would be reasonable, for example, in the case of financial (log)return data. We are interested in inferring the parameters governing the crisis transmission, and as a by-product the volatilities $\boldsymbol{\sigma} = (\sigma_{i\ell}, i = 1, \ldots, N; \ell = 1, 2, 3)$ and times of entry to and recovery from the crisis, $\mathbf{E} = (E_1, \ldots, E_N)$ and $\mathbf{R} = (R_1, \ldots, R_N)$ respectively.

Augmented Likelihood

Let $\mathbf{Y} = (Y_{11}, \ldots, Y_{1T}; Y_{21}, \ldots, Y_{2T}, \ldots; Y_{N1}, \ldots, Y_{NT})$. The likelihood of the observed data $(\mathbf{Y})$ given the infection rates λ_G and λ_L and the volatilities $(\boldsymbol{\sigma})$ can be expressed as

$$L(\mathbf{Y}|\lambda_G, \lambda_L, \boldsymbol{\sigma}) = \int_{\mathbf{S}} f(\mathbf{Y}, \mathbf{S}|\lambda_G, \lambda_L, \boldsymbol{\sigma}) \mathrm{d}\mathbf{S}. \tag{5}$$

If the times of entry to and recovery from the crisis for all the countries in the population were known, then parameter estimation would be straightforward. This information, however, is typically not directly observable. In this case, the likelihood given by (5) is intractable due to the requirement of computing the high dimensional integral $\int_{\mathbf{S}}(\cdot)$. To surmount this problem we augment the parameter space with the sets $\mathbf{E}$ and $\mathbf{R}$, and propose a Bayesian data-augmentation framework for estimation (see, for example, O'Neill and Roberts, 1999). This approach will enable us to treat the times $E_i, R_i, i = 1, \ldots N$ as additional parameters to be estimated simultaneously with λ_L, λ_G and $\boldsymbol{\sigma}$.

Posterior Distribution

The posterior distribution of the unknown parameters of interest given the observed data and assuming independent priors for λ_L, λ_G and $\boldsymbol{\sigma}$, is expressed as

$$\begin{aligned}\pi(\lambda_G, \lambda_L, \boldsymbol{\sigma}, \mathbf{E}, \mathbf{R}|\mathbf{Y}) &\propto \pi(\mathbf{Y}|\mathbf{E}, \mathbf{R}, \boldsymbol{\sigma}, \lambda_G, \lambda_L)\pi(\mathbf{E}, \mathbf{R}|\boldsymbol{\sigma}, \lambda_L, \lambda_G)\pi(\boldsymbol{\sigma}, \lambda_L, \lambda_G) \\ &= \pi(\mathbf{Y}|\mathbf{E}, \mathbf{R}, \boldsymbol{\sigma})\pi(\mathbf{E}, \mathbf{R}|\lambda_L, \lambda_G)\pi(\boldsymbol{\sigma})\pi(\lambda_G)\pi(\lambda_L), \qquad (6)\end{aligned}$$

where $\pi(\mathbf{Y}|\mathbf{E}, \mathbf{R}, \boldsymbol{\sigma}) = \prod_{i=1}^{N}\prod_{t=1}^{T} f(Y_{it}|S_{it})$.

The second term in (6) is the contribution from the epidemic model and following, for example, Jewell et al. (2009) is given by

$$\pi(\mathbf{E}, \mathbf{R}|\lambda_L, \lambda_G) \propto \left(\prod_{j\neq k} M_j\right) \times \exp\left\{-\int_{E_k}^{T}\sum_{i,j}\mu_{ij}(t)\, dt\right\} \times \mathcal{R}, \qquad (7)$$

where M_j denotes the (infectious) pressure that an infected country is subjected to just before it enters the crisis. Thus we have that $M_j = \sum_{i\in\mathcal{C}}\mu_{ij}$ where μ_{ij} is the instantaneous rate at which country i exerts (infectious) pressure on country j just before country j enters the crisis with $\mu_{ij} = \frac{\lambda_G}{N} + \lambda_L \mathbf{1}_{(i,j \text{ in the same region})}$. The set $\mathcal{C}$ denotes the set of countries affected by the crisis at time t. The country that first entered the crisis is labelled by k and T is the time at which the crisis is assumed to be over.

The term $\mathcal{R}$ in (7) denotes the contribution of the recovery process to the likelihood of the epidemic model. Assuming independent (random) infectious periods I_i, $i = 1, \ldots, n$, this contribution is $\mathcal{R} = \prod_{i=1}^{n} f_I(I_i)$, where $f_I(\cdot)$ denotes the arbitrarily specified distribution governing the infectious period. The infectious periods are assumed to be distributed according to a Gamma distribution with (hyper) parameters γ and δ (and mean γ/δ), so then $\mathcal{R} = \prod_{i=1}^{n} \frac{\delta^\gamma}{\Gamma(\gamma)} I_i^{\gamma-1}\exp\{-I_i\delta\}$.

Priors

The terms $\pi(\boldsymbol{\sigma})$, $\pi(\lambda_G)$ and $\pi(\lambda_L)$ denote the prior distributions assigned to the parameters $\boldsymbol{\sigma}$, λ_L and λ_G respectively. Slowly varying exponential priors are assigned to λ_L and λ_G and Inverse-Gamma distributions are assigned to the volatilities. If the parameters associated with the infectious period distributions, γ and δ, are assumed to be unknown then

prior distributions need to also be assigned to these. An obvious choice for weakly informative priors would be to assume that both parameters are a-priori independent and follow exponential priors similar to λ_L, λ_G. As we do not make any assumption about which country entered the crisis first, we consider a joint prior distribution for k and E_k as follows: $\pi(k, E_k) = \pi(E_k|k)\pi(k)$ with $\pi(k) = 1/|\mathcal{C}|$ where $|\cdot|$ denotes the cardinality of a set and $\pi(E_k|k) \equiv U(-\infty, R_k)$.

3.3 Incorporating Additional Covariates

Explanatory variables are, in principle, straightforward to incorporate within a HEM through the transmission rates. In the most general case one can allow for different covariates affecting local and global rates as follows:

$$\lambda_{L,ij} = \lambda_{L0} \exp(\textstyle\sum_k \alpha_{Lk} X_{L,ijk}), \quad \lambda_{G,ij} = \lambda_{G0} \exp(\textstyle\sum_k \alpha_{Gk} X_{G,ijk}),$$

where $\lambda_{L,ij}$ is the local crisis trasmission rate from the infected country i to the susceptible country j, λ_{L0} is the local baseline transmission rate, $\mathbf{X}_{L,ij} = (X_{L,ij1}, ..., X_{L,ijK})$ are the covariates of interest, and $\boldsymbol{\alpha}_L = (\alpha_{L1}, ..., \alpha_{LK})$ are the associated coefficients; similarly for the global counterparts. An example of a covariate, $X_{L,ijk}$ ($X_{G,ijk}$) could be the exports from country i to country j, at the local (global) level or some proxy for financial linkages as will be considered in our empirical application. It is then of interest to infer the effect of the different covariates on the transmission rates by estimating the parameters $\boldsymbol{\alpha}_L$ and $\boldsymbol{\alpha}_G$. This will allow to evaluate the importance of different channels in the propagation of a crisis across countries. Note that for $\boldsymbol{\alpha}_L = \boldsymbol{\alpha}_G = \mathbf{0}$, the baseline hidden epidemic model with two levels of mixing is recovered. The posterior distribution of interest then becomes

$$\begin{aligned} \pi(\lambda_G, \lambda_L, \boldsymbol{\alpha}_L, \boldsymbol{\alpha}_G, \boldsymbol{\sigma}, \mathbf{E}, \mathbf{R}|\mathbf{Y}) &= \pi(\mathbf{Y}|\mathbf{E}, \mathbf{R}, \boldsymbol{\sigma})\pi(\mathbf{E}, \mathbf{R}|\lambda_L, \lambda_G) \\ &\times \pi(\boldsymbol{\sigma})\pi(\lambda_G)\pi(\lambda_L)\pi(\boldsymbol{\alpha}_L)\pi(\boldsymbol{\alpha}_G). \end{aligned} \tag{8}$$

We employ a Metropolis-within-Gibbs algorithm to draw samples from the posterior distribution given by (8). Apart from the volatilities which are updated using a Gibbs sampler, all the other parameters in the model, including the times of entry to and exit from the crisis, are updated using a Metropolis-Hastings algorithm with Gaussian proposal

distributions.

Remark 1 *Unlike when only final size data are available where the SIR model is invariant to a latent period, this is not true for temporal data. A straightforward extension to consider in the latter case is a Susceptible-Exposed-Infected-Recovered (SEIR) model which assumes that once a country enters the crisis state, it can only trasmit the crisis (i.e. become infective) after a certain time period elapses.*

Remark 2 *In the case of real-time analysis, the methodology can be extended to conduct inference for the countries more likely to enter the crisis, given the current state of the process. For this purpose, a trans-dimensional MCMC algorithm (Green, 1995) is required which accounts for the variable dimension of the parameter space.*

Remark 3 *A threshold parameter similar to $R^\star$ (see Section 2.2) can also be defined for the hidden epidemic model, taking into account any additional covariate information. In this case, definitions of typical infective and susceptible countries based upon averaging over their covariate values could be used for inclusion in the $v(\lambda_L)$ term in (1).*

4 Empirical Application

To illustrate the proposed methodology we analyse the contagious spead of five currency crisis episodes considered by Glick and Rose (1999), namely the breakdown of the Bretton Woods System, 1971, the collapse of the Smithsonian Agreement, 1973, the EMS Crisis, 1992, the Mexican meltdown and Tequila Effect, 1994, and the Asian Flu, 1997. Caution is required when interpreting the results for the 1973 episode, as this date coincides with the oil crisis of the same year, and may not be entirely appropriate as an example of contagion. However, we decided to include it in our analysis for illustrative purposes. The final outcome dataset is compiled from the cross-sectional binary data of Glick and Rose for 160 developed and developing countries, where a country that ultimately experienced the crisis in question takes the value of one and zero otherwise. Our choice to analyse currency crises was primarily driven by the evidence of contagion associated with these crises as reported in the literature. Glick and Rose identify the binary crisis variable from journalistic and academic histories of the various crisis episodes. An alternative approach of constructing the binary crisis indicator relies on some pre-selected crisis-identification threshold, which transforms a continuous

variable or combination of variables into the desired indicator. See Jacobs et al. (2005) for a review of the literature on identification of crises.

We begin our analysis by dividing the countries into groups based on geographical location. A detailed list of the countries and the composition of regions, as well as the countries affected by each of the crisis episodes can be found in Tables A1 and A2 respectively in the Appendix. The regional decomposition is based on the United Nations classification. Alternative population structures such as overlapping groups or networks based on the trade pattern, financial links or macroeconomic similarities between countries could also be considered. Linkages of this type will be incorporated in what follows through the use of covariates as discussed in Section 3.3.

4.1 Final Size Data

In what follows we present posterior summary estimates of the local (λ_L) and global (λ_G) trasmission rates for each of the five currency crisis episodes. Since no information is available in final outcome data regarding the mean length of the period for which a country remains in crisis, we set the infectious period distribution in advance of the data analysis. This should be taken into account when interpreting λ_L and λ_G, while R_* will not be affected as can be seen from equation (1). In the results that follow, without loss of generality the infectious period is set to unity. Slowly varying exponential distributions with rate 0.01 are used as priors.

The results, summarised in Table 1, give the mean posterior estimates and corresponding 95% credible intervals for the local and global contact rates, as well as the ratio of global (between region) to total (within region) per-country contact rate $\frac{\lambda_G/N}{\lambda_L+\lambda_G/N} = \frac{\lambda_G}{N\lambda_L+\lambda_G}$, for each crisis episode. Since the infectious period is set to unity, this ratio also represents the ratio of between to within country-to-country infectious contacts. The results indicate an increase in the global trasmission rate over time across the different currency crises with a simultaneous decrease observed in the local rate. This is particularly apparent from the ratio $\frac{\lambda_G}{N\lambda_L+\lambda_G}$, which is very small for the crises of the seventies and increases to almost a half in the case of the 1997 crisis, with the length of the corresponding intervals also increasing. The shift from local to global spread over time potentially reflects the significant increase of global financial linkages across countries in recent decades. While empirical evidence in the

literature suggests that currency crises tend to be mostly regional, only a model allowing for both local and global effects can truly assess the relative importance of each of these dimensions in the transmission of crises. However, it must be pointed out that while the 1992 crisis was mainly confined to European countries, the results in Table 1 do not bring out such a feature. Indeed, a further refinement of the analysis is required to shed light on the true nature of the EMS crisis episode. This will be undertaken in Section 4.1.1 where we consider including categorical covariates in the analysis reflecting certain characteristics of the countries involved.

Table 1: Posterior summary estimates for the two-level mixing model

Crisis Episode	λ_L	λ_G	$\frac{\lambda_G}{N\lambda_L+\lambda_G}$
1971	0.445(0.239,0.738)	0.297(0.093,0.618)	0.005(0.001,0.013)
1973	0.349(0.164,0.574)	0.446(0.169,0.868)	0.010(0.003,0.027)
1992	0.018(0.000,0.066)	1.019(0.487,1.731)	0.381(0.073,0.932)
1994	0.013(0.000,0.047)	1.043(0.516,1.777)	0.456(0.102,0.955)
1997	0.011(0.000,0.041)	1.068(0.617,1.616)	0.491(0.127,0.958)

The above results were obtained using the MCMC algorithm described in Section 2.3.4. The computational complexity is mainly determined by the size of the random graph. The small final sizes involved in this application imply highly efficient computations, that is approximate running times of 2-3 minutes per 10^6 iterations. An additional attractive feature of the considered methodology is related to the convergence of the algorithm. Specifically, it is possible to create a measure of approximate monotonicity by considering summary measures of the random graph such as the total number of links. In the case of a homogeneous population (GSE) and constant infectious period the number of links can be thought of as a sufficient statistic; in more general models it is approximately sufficient. Hence, this number can be used as an approximate convergence diagnostic of the MCMC algorithm. For all the reported results we started the algorithm from the two 'extremes' in the graph space, that is, a complete graph where all countries have local and global links with all others, and a sparse tree-like graph. Rapid convergence towards the 'high probability' region was observed for all the algorithm implementations. This region lies between the two extremes, with the 'likely' graphs being of relatively sparse form.

Additional information from the graph output can also be obtained if required. An empirical measure of the longest path and the number of generations can easily be extracted

from the digraph. Combined with knowledge relating to the duration of the crisis, this would deliver inference for the time-length of each generation of affected countries. Theoretical bounds for the longest path in the case of a homogeneous random digraph have been derived in Foss and Konstantopoulos (2003).

4.1.1 Multitype Processes

Motivated by the evidence that crises in developing economies are of a different nature than those in developed ones (Kaminsky, 2006), we consider further partitioning our population of countries into developed and developing, based on the world factbook IMF classification. This distinction allows for the possibility of these two groupings having different local and/or global rates, rather than assuming equal susceptibility and infectivity rates across all countries. It therefore enables to distinguish whether a particular crisis episode affected mostly developed and/or developing countries, avoiding potential misjudgement that may arise from overlooking the degree of economic development of the countries under study. In addition, an analysis that considers a population of multiple types can be useful in uncovering important transmission channels that may go unobserved in the context of a single type model, as will become clear in what follows. This kind of extension essentially involves introducing categorical covariates into our model specification that can be handled using multitype epidemic models. Table 2 below gives the proportions of the affected developed and developing countries across the different crisis episodes.

Table 2: Proportions of affected countries across the different crisis episodes

Crisis Episodes	1971	1973	1992	1994	1997
Developed	$\frac{19}{28}$	$\frac{19}{28}$	$\frac{10}{29}$	$\frac{2}{29}$	$\frac{4}{29}$
Developing	$\frac{0}{113}$	$\frac{0}{113}$	$\frac{0}{130}$	$\frac{9}{131}$	$\frac{13}{131}$

Consider the case of the 1992 EMS crisis where 10 out of 29 developed countries were affected, while no developing country suffered the crisis in question. Moreover, all affected countries were in Europe (see Table A2 in the Appendix) so we would expect a high local transmission rate among developed countries, which could not have been deduced from Table 1.

We consider a two type, two level mixing stochastic epidemic model to characterise the spread of crises among developed and developing countries. Multitype analysis involves a further decomposition of the local and global rates relative to the single type framework,

reflecting the data provided in Table 2. To see this let $\Lambda_L = \{\lambda_{ij,L}\}$ and $\Lambda_G = \{\lambda_{ij,G}\}$, $i, j = 1, 2$, denote the 2×2 matrices of local and global infectious rates, respectively. Subscript 1 (2) refers to the developed (developing) countries, so that for example $\lambda_{12,L}$ denotes the local rate at which the crisis affecting a developed country is transmitted to a developing country. It is evident that the number of infection rates to be estimated has now increased considerably relative to the single type model. Certain modelling restrictions are therefore necessary for identifiability, unavoidably for the λ_G's and possibly for the λ_L's depending on the dataset under consideration, see Britton (1998) and Britton and Becker (2000) for a detailed discussion.

We consider the following restriction: we allow for distinct local rates when the data contain sufficiently rich information for identification purposes. This is not always the case. Based on the data provided in Table 2, for the earlier crises of 1971, 1973 and 1992 no developing countries were affected. Hence, the data contain no information with respect to the local and global transmission of these crises from one developing country to another, that is $\lambda_{22,L}$ and $\lambda_{22,G}$, respectively, and hence results will be largely informed by the prior. In addition, convergence of the MCMC sampler may be problematic so special care is required in efforts to overcome issues of parameter redundancy. In such cases, for the local rates we assume equal susceptibility and distinct infectivity, which implies $\lambda_{11,L} = \lambda_{21,L}$, and $\lambda_{12,L} = \lambda_{22,L}$, that is developed countries have different potential in transmitting the crisis compared to developing ones, while all countries are equally susceptible. Globally, it is well known (e.g. Ball et al., 2004) that at most two transmission rates are identifiable. Thus, in keeping with our local assumptions, we assume common global susceptibility to a crisis, but distinct global infectivity. We consider this choice to be reasonable within the present context. Scenarios with alternative susceptibility levels, potentially defined based on economic information, are also possible. In fact, within our Bayesian framework it is possible to consider any structure for Λ_L and Λ_G subject to the required identifiability restrictions.

We begin by presenting in Table 3 the posterior means and standard deviations of the two type model when equal local and global vulnerability is considered for all crisis episodes. For the earlier crises of 1971, 1973 and 1992, both local and global rates are smaller for the developing countries compared to the developed countries. The results are less homogeneous for the later crises and largely reflect their relative severity, with the developing countries being generally more affected, especially globally.

Table 3: Posterior summaries (means and standard deviations) for the two type model with equal local and global vulnerability for each grouping

Crisis Episode	$\Lambda_L = \begin{pmatrix} \lambda_{11,L} & \lambda_{12,L} \\ \lambda_{21,L} & \lambda_{22,L} \end{pmatrix}$		$\Lambda_G = \begin{pmatrix} \lambda_{11,G} & \lambda_{12,G} \\ \lambda_{21,G} & \lambda_{22,G} \end{pmatrix}$	
1971	0.475(0.211)	0.091(0.089)	0.761(0.356)	0.053(0.053)
1973	0.284(0.130)	0.075(0.076)	1.028(0.417)	0.054(0.053)
1992	0.170(0.086)	0.037(0.037)	0.645(0.390)	0.100(0.099)
1994	0.115(0.093)	0.065(0.032)	0.194(0.135)	0.559(0.238)
1997	1.017(0.394)	0.070(0.030)	0.138(0.098)	0.454(0.186)

Note: The matrix Λ_L (Λ_G) contains the local (global) estimated trasmission rates, where subscripts 1 and 2 refer to the developed and developing countries, respectively. Standard deviations are reported in brackets. The restrictions $\lambda_{11,L} = \lambda_{21,L}$ and $\lambda_{12,L} = \lambda_{22,L}$ are imposed on Λ_L, and $\lambda_{11,G} = \lambda_{21,G}$ and $\lambda_{12,G} = \lambda_{22,G}$ are imposed on Λ_G.

Table 4 presents results under distinct local and equal global vulnerability. Results for this case are only reported for the more recent crises, 1994 and 1997, owing to the lack of information in identifying the local transmission rate, $\lambda_{22,L}$, as explained earlier. For both these episodes, allowing for distinct local susceptibility shows that the local transmission of the crisis is largely from developing to developed countries, as would be expected from the ratios reported in Table 2. A reasonably high local transmission rate is also observed between developed countries. Transmission rates at the global level are similar to those reported in Table 3.

Table 4: Posterior summaries (means and standard deviations) for the two type model with distinct local and equal global vulnerability for each grouping

Crisis Episode	$\Lambda_L = \begin{pmatrix} \lambda_{11,L} & \lambda_{12,L} \\ \lambda_{21,L} & \lambda_{22,L} \end{pmatrix}$		$\Lambda_G = \begin{pmatrix} \lambda_{11,G} & \lambda_{12,G} \\ \lambda_{21,G} & \lambda_{22,G} \end{pmatrix}$	
1994	0.299(0.235)	0.264(0.262)	0.191(0.137)	0.567(0.251)
	0.657(0.317)	0.070(0.035)	0.191(0.137)	0.567(0.251)
1997	0.366(0.183)	0.189(0.153)	0.147(0.103)	0.450(0.185)
	0.924(0.325)	0.064(0.036)	0.147(0.103)	0.450(0.185)

Note: The matrix Λ_L (Λ_G) contains the local (global) estimated trasmission rates, where subscripts 1 and 2 refer to the developed and developing countries, respectively. Standard deviations are reported in brackets. The restrictions $\lambda_{11,G} = \lambda_{21,G}$ and $\lambda_{12,G} = \lambda_{22,G}$ are imposed on Λ_G.

4.2 Temporal Data

We follow the definition of a currency crisis by Frankel and Rose (1996), namely a successful speculative attack that manifests itself through a large nominal depreciation of the currency.

As such, we consider daily log changes in nominal exchange rates as our observed data, that can be used to estimate the entry to and exit from the crisis, in addition to the transmission rates and volatility parameter. To capture the well documented importance of trade and financial linkages in the spread of contagious currency crises, we include related variables as covariates in the hidden epidemic model. Since the relative importance of the trasmission rates has already been extensively discussed, our focus will be on assessing the significance of these covariates.

4.2.1 Exchange Rates

We use daily log exchange rate returns for those countries affected by the crisis, for each crisis episode. While in principle exchange rate data should be considered for all countries in the sample, we proceed by making use of the knowledge of those countries that were affected by the crisis for consistency with the analysis in Section 4.1. The Deutsche Mark is used as the reference currency for the 1992 episode, given that this was the anchor currency within the European Monetary System at the time, and the US dollar for the 1994 and 1997 crises. The returns were originally computed over a five year window, including two years preceding and two years following the crisis year. However, to avoid overlapping crises, the final estimation samples were chosen as, 01/01/1990-15/03/1994, 16/03/1994-31/12/1996, and 01/01/1997-13/03/1998 for the 1992, 1994, and 1997 crises, respectively. Further details regarding the data are available in the Supplement. Among the infected countries for the 1994 and 1997 crises, Argentina and Hong Kong maintained an exchange rate peg and were therefore excluded from the sample for the purpose of the HEM analysis.

4.2.2 Covariates

Trade routes of contagion in the empirical literature typically rely on direct and indirect measures of trade based on exports. Financial routes focus primarily on bank lending channels and in particular the 'common lender effect'. The common lender effect argues that a country that relies on the same source of financing as a crisis country (being highly indebted to the same lender, as well as being highly represented in the lender's portfolio) can increase the country's financial vulnerability. A loss or underperforming loans in a country in crisis induces international creditors to reduce the amount of the loan and withdraw their financial

assets from other countries. Countries affected by a crisis rely heavily on borrowings from the common lender. About 10% or more of the common bank liabilities are normally held in an infected country (Caramazza et al. 2004). We consider trade covariates for the 1992, 1994 and 1997 crisis episodes. Financial covariates are considered only for the latter two, due to data non-availability. The crises of the seventies are excluded from the analysis for the same reason.

Trade Covariates Glick and Rose (1999) consider a number of aggregate measures of direct and indirect trade, centered on the ground zero (initial infective) country and focusing exclusively on exports. These and other similar measures are considered in a number of studies, for example Van Rijcheghem and Weder (2001), Dasgupta et al. (2011), and Haile and Pozo (2008), that use probit type models.

We use the disaggregate matrix of bilateral trade relationships between the infected countries and all trading partners directly in our model. To adjust for the varying size of countries in our sample we consider trade shares, constructed as the ratio of bilateral trade between the infected country i and each of its trading partners j within the total trade of country i. Specifically, we consider two sets of trade covariates, one based exclusively on annual export figures from the IMF Direction of Trade Statistics (DOTS) and the other based on the average of import and export figures. The latter allows to assess the importance of imports, which is ignored in the analysis of Glick and Rose (1999) as they admittedly point out.

Both trade covariates are based on annual figures averaged across the three years typically immediately preceding each crisis episode. See the supplement for more details. Countries where data for more than 20% of trading activity were not available were dropped from the sample for each crisis episode. This resulted in sample sizes of 99, 107, 106 countries for the 1992, 1994 and 1997 crises, respectively, for exports only, and 124, 133 and 132, respectively, for imports and exports. For the purpose of sensitivity analysis we further excluded countries where more than 30% of trading activity was not available and the results were very similar.

We model the instantaneous rate at which country i exerts (infectious) pressure on country j, just before country j enters the crisis as $\mu_{ij} = \left(\frac{\lambda_G}{N} + \lambda_L \mathbf{1}_{(i,j \text{ in the same region})}\right) \times e^{\alpha X_{ij}}$, where X_{ij} denotes the acting trade covariate, that is the trade share matrix of the infected countries with all countries in the sample, α is the corresponding covariate effect, and

λ_L, λ_G denote the local and global transmission rates, respectively. Although in principle one could use a more general model and allow the covariate X_{ij} to separately affect the local and global rates, as discussed in Section 3.3, we opted for a simpler model in the interest of parsimony. We assign a Gaussian prior with large variance to the parameter α as for all covariate effects in the illustrations that follow.

For a given infected country, the stronger its trade linkages with the other countries the higher the rate of transmitting the crisis to these countries, other things being equal. This interpretation is different from that of Glick and Rose (1999) and other authors who focus on the effect of trade on the probability of a country becoming infected, rather than on the probability of the infected country transmitting the crisis to other countries. In our context, the former probability can be evaluated as a function of the transmission rates and the infectious pressure exerted on that country.

Table 5 reports the mean and median posterior estimates, and the 90% credible intervals, for the coefficient of the trade covariate across the different crises, in the case of imports and exports, and exports only. For the 1992 crisis episode, a significant amount of the probability mass of the posterior distributions is concentrated around positive values of the trade coefficient, while for the 1994 and 1997 crises a significant amount of probability mass is associated with negative values (see the Supplement for graphs of the posterior densities). A negative coefficient implies that trade confers a 'protective' effect against the spread of these crises. This is not counterintuitive after inspecting the percentage of trade among the infected countries for each crisis episode given in Table A3 in the Appendix. The figures show that, for the 1994 and 1997 crises, trading activity (based on imports and exports) among infected countries is smaller than that between infected countries and their non-infected trading partners. For the 1992 crisis, trading activity among infected countries is much higher. Similar figures are observed in the case of exports only (results not shown). According to these results trade, whether based on imports and exports or on exports only, appears to be a statistically significant channel in the transmission of the 1994 currency crisis. There is no such evidence in the case of the 1992 and 1997 crises.

Table 5: Posterior summary estimates for the coefficient of the trade covariates

	1992	1994	1997
	Imports and Exports		
Mean	6.19	-30.95	-4.07
Median	6.32	-25.15	-3.02
90% CI	(-1.66,13.94)	(-74.52,-2.08)	(-17.16,5.13)
	Exports		
Mean	8.51	-48.87	-9.17
Median	8.36	-36.53	-4.76
90% CI	(-0.85,18.09)	(-125.43,-1.95)	(-38.48,5.10)

Financial Covariates Financial covariates for the 1994 and 1997 crises are constructed based on two sets of bi-annual banking data from the Bank of International Settlements (BIS) beginning in 1988, namely international lending from the international consolidated statistics and international borrowing from the locational international banking statistics. As in the case of trade, we constructed three year averages based on those years typically immediately preceding the crisis date. This resulted in 124 and 136 countries for the 1994 and 1997 crisis, respectively. We employ three commonly used measures for financial linkages between the common lender and crisis countries which are described below (see among others Van Rijckeghem and Weder, 2001, Caramazza et al., 2004, Dasgupta et al., 2011 and Haile and Pozo, 2008). USA and Japan are the common lenders for the 1994 and 1997 crises, respectively.

(i) A measure of the importance of the common lender to the crisis country, that is the proportion of borrowing from a common lender, given by $F1_i = b_{i,c}/b_i$, where $b_{i,c}$ is total borrowing of country i from the common lender c and b_i is total borrowing of country i.

(ii) A measure of the importance of an affected country to a common lender, that is the proportion of borrowings of an affected country in the lending portfolio of a common lender, expressed as $F2_i = l_{i,c}/l_c$, where $l_{i,c}$ is lending of a common lender to country i and l_c is total lending of a common lender.

(iii) A measure of competition for funds, that is the extent to which country i competes for borrowings from the same common lender as the ground zero country, given by

$$FC_i = \sum_c \frac{x_{0,c} + x_{i,c}}{x_0 + x_i} \left(1 - \frac{|x_{i,c} - x_{0,c}|}{x_{i,c} + x_{0,c}}\right), \tag{9}$$

where 0 denotes the ground zero country, c is the common lender, $x_{i,c}$ is total lending from

the common lender to county i, and x_i is total borrowings of country i. The higher the value of FC_i, the greater the competition for funds between countries i and 0.

The transmission rate from an 'infectious' country i to a 'susceptible' country j, μ_{ij}, is given by $\mu_{ij} = \left(\frac{\lambda_G}{N} + \lambda_L \mathbf{1}_{(i,j \text{ in the same region})}\right) \times e^{\alpha Z_i}$, where Z_i denotes the acting financial covariate either $F1, F2$, or FC, and α the associated coefficient.

Table 6 reports the mean, median and 90% credible interval for the posterior distribution of the coefficient α associated with the different financial covariates, for the 1994 and 1997 crises. These coefficients appear to be largely negative as would be expected. For a given infected country, the greater the proportion of borrowing from a common lender ($F1$) or in the lending portfolio of a common lender ($F2$) the more 'protected' that country becomes and so the smaller the probability of transmitting the crisis to other countries. The results point to a significant role for the common lender effect in the 1994 crisis, when measured by the F1 covariate.

Table 6: Posterior summary estimates for the coefficient of the various financial covariates

	1994			1997		
	Mean	Median	90% CI	Mean	Median	90% CI
F1	-14.14	-13.10	(-30.75,-0.99)	-0.72	-0.16	(-8.89,4.94)
F2	-15.40	-13.60	(-39.93,4.63)	-9.73	-5.91	(-34.33,5.50)
FC	-7.98	-7.53	(-19.94,2.66)	-2.10	-0.35	(-15.71,4.77)

Trade and Financial Covariates Next we consider including both trade and financial covariates together in the model so that we have $\mu_{ij} = \left(\frac{\lambda_G}{N} + \lambda_L \mathbf{1}_{(i,j \text{ in the same region})}\right) \times e^{\alpha_1 X_{ij} + \alpha_2 Z_i}$, where X_{ij} and Z_i denote the acting trade and financial covariates defined as above, with α_1 and α_2 being the corresponding effects.

When both trade and financial covariates are included in the model simultaneously, the posterior distributions are wider, as confirmed by the posterior estimates and associated credible intervals presented in Table 7. This table shows the estimates for the posterior distributions of each pair of coefficients, α_1 and α_2, associated with the corresponding trade and financial covariate. Trade refers to the trade shares computed based on the average of imports and exports. The results based on exports only were qualitatively similar. The results indicate that, as in the case where a financial covariate only is included in the model, the common lender effect continues to be significant in the 1994 crisis, not only when measured as the proportion of borrowing of an affected country from the common lender, $F1$,

but also as the proportion of borrowings of an affected country in the lending portfolio of the common lender, $F2$. This result appears to overshadow the significant effect found for the trade channel in the 1994 crisis, when only a trade covariate is included in the model. No significant effect is observed for the 1997 crisis, in line with the earlier findings.

Table 7: Posterior summary estimates for the coefficients associated with the trade and financial covariates

	1994			1997		
	Mean	Median	90% CI	Mean	Median	90% CI
Trade	-26.83	-18.30	(-85.96,2.35)	5.13	5.96	(-7.87,14.95)
$F1$	-15.87	-13.81	(-36.75,-2.70)	0.16	0.81	(-6.64,5.09)
Trade	-34.05	-25.62	(-98.49,0.94)	1.73	3.59	(-16.87,13.16)
$F2$	-53.03	-39.47	(-152.13,-1.00)	-14.56	-4.84	(-60.24,7.22)
Trade	-23.18	-17.97	(-59.28,0.37)	3.51	4.86	(-10.50,13.76)
FC	-10.45	-8.95	(-27.90,0.85)	-2.25	-1.16	(-12.80,4.26)

Note: Trade refers to the shares of the average of imports and exports.

5 Policy Relevance

Understanding how financial crises spread, through models such as the one proposed here, can enhance in the design and evaluation of immunisation policies to reduce the risks and manage the impact of contagion. So far, at the domestic level, the need for policies aimed at reducing financial fragility has been emphasised; at the international level, the role of better financial standards and of the international lender of last resort has been discussed, see Chang and Majnoni (2001). The latter can provide liquidity to crisis countries to withstand pressures of contagion, as witnessed in the recent economic crises, in the form of rescue packages.

Among the questions increasingly raised on the policy front are how severe is a particular crisis and how to distribute the available resources to sustain countries that have experienced or may be experiencing negative outcomes (how many and which countries to support). The proposed approach can offer answers to such questions, as we illustrate below, by drawing on the past experience of the five currency crises episodes analysed earlier. Lessons from these episodes can offer valuable insight to policy makers for future difficulties.

A number of immunisation policies have been developed, mostly with reference to epidemics among humans or animals (Anderson and May, 1991), though examples exist in other fields like computer science (Balthrop et al., 2004). The threshold theorem described in Sec-

tion 2.2 provides a natural formulation for evaluating this kind of strategies. In particular, having obtained the parameter estimates of our model, that is the local and global transmission rates, we can compute the severity measure R_* given by (1) for the individual crisis episodes. Recall that if $R_* > 1$ a large number of countries will suffer from the crisis in question with positive probability, while if $R_* \leq 1$ only a small number of susceptible countries will ever be affected. Control measures, therefore, typically aim to reduce R_* below unity with high probability, which we set to 0.95 in the results that follow. The situation where a major 'outbreak' is unlikely to occur is often referred to as 'herd immunity'. Having computed R_* we can then sample from its posterior density, and estimate the percentage (or number) of countries to support, denoted by $\mathrm{p_v}$, in order to prevent a major crisis. In other words, we will estimate the smallest $\mathrm{p_v}$ such that $P(R_* < 1) = 0.95$.

We will examine here a simple support strategy where the countries to be supported are chosen uniformly at random. Recall that the threshold parameter without support is given by $R_* = \frac{\lambda_G E(I)}{\nu} \sum_j j\mu_j\pi_j$. When a proportion $\mathrm{p_v}$ of the countries receives support the threshold reduces to $R_*(\mathrm{p_v}) = \frac{\lambda_G E(I)}{\nu} \sum_j j\mu_j \sum_{z\geq j} \pi_r \binom{z}{j} (1-\mathrm{p_v})^j \mathrm{p_v}^{z-j}$, since the number of countries prone to the crisis reduces. The percentage $\mathrm{p_v}$ is computed as the solution to the equation $R_*(\mathrm{p_v}) = 1$. Thus, $\mathrm{p_v}$ and R_* are non-linear functionals of the basic model parameters, at least when the population structure is non-random and known. Inference for this type of functionals has been considered for two level mixing models in Britton and Becker (2000) assuming for simplicity that the outcome within each region is independent of the fate of other regions. Ball et al. (2004) also consider estimating R_* and $\mathrm{p_v}$ using asymptotic arguments and assuming $R_* > 1$. The methods considered in this paper are free from these assumptions. Moreover, assuming $R_* > 1$ is not particularly satisfactory in our examples, as will soon become apparent. Alternative support policies, like supporting whole regions chosen at random, are also possible, in which case $\mathrm{p_v} = 1 - 1/R_*$, a well known result in epidemic theory, see for example Anderson and May (1982), and Ball and Lyne (2006) for the case of structured stochastic models. In the subsequent results we assume that financial support to a country confers complete protection from the crisis, or at least it effectively prevents a particular country from becoming affected.

Table 8 gives estimates of the posterior mean of R_*, $P(R_* < 1)$ and $\mathrm{p_v}$. A graph of the posterior distribution of R_* for all crisis episodes can be found in the supplement. The mean posterior estimate for R_* is greater than one for all episodes. In particular, the posterior

mean of the 1971 and 1973 crises is 2.19 and 2.95 respectively, while that of the nineties crises is around 1.2. These results corroborate our earlier empirical findings relating to the increasing importance of λ_G, which is now reflected in the decreasing length of the 95% credible intervals for R_*. The corresponding probability of a crisis being subcritical ('minor'), given under the heading $P(R^* < 1)$, is smaller for the earlier crises compared to the later. In accordance, the associated percentage, $\mathrm{p_v}$, of countries to support in order to avoid a major crisis with a 0.95 probability decreases over time as noted by the figures in the last column. The results suggest that the crises of the seventies were of greater severity compared to those of the nineties.

Table 8: Estimated control measures for the various crisis episodes

Crisis Episode	R^*	$P(R^*<1)$	$\mathrm{p_v}$
1971	2.186(0.674,4.610)	0.094	0.291(0.034,0.523)
1973	2.946(0.602,3.632)	0.115	0.265(0.029,0.513)
1992	1.242(0.525,2.425)	0.325	0.240(0.016,0.508)
1994	1.194(0.568,2.170)	0.356	0.230(0.015,0.504)
1997	1.200(0.665,1.938)	0.293	0.210(0.011,0.443)

Note: R^* measures the severity of the crisis, $P(R^* < 1)$ is the probability that only subcritical ('minor') crises can occur and $\mathrm{p_v}$ is the critical protection coverage i.e. the percentage of countries to support in order to avoid a major crisis with a 0.95 probability. Figures in parentheses are the 95% credibles intervals.

The threshold parameter R_* and associated control measures are also available in the multitype context, though the calculations are somewhat more involving. In the two type case, R_* is defined as the largest eigenvalue of a matrix $M = \{m_{ij}\}$, $i, j = 1, 2$, where, crudely speaking, m_{12} describes the mean number of contacts from countries of type 1 (developed) to countries of type 2 (developing). More details regarding the computation of the m_{ij} elements can be found in the Supplement.

In addition to the support policy considered above, alternative policies can be implemented that provide partial coverage to more countries as opposed to offering complete protection to a smaller number. This approach is more involved and additional optimisation techniques are required. It should be noted that if the local transmission rate is small, as may be the case in a globalised economy, there will not be a significant difference between the various adopted policies, at least for homogeneous individuals. Britton and Becker (2000) consider immunisation strategies within a simpler model allowing the level of susceptibility to the crisis to vary among individuals. Such extensions can be naturally accommodated within the framework of multitype epidemics discussed earlier.

As demonstrated from the above analysis, the proposed methodology offers a set of valuable control measures that can be used to analyse contagious financial crises. When implemented in real-time, preliminary estimates of the control measures can be computed to evaluate prospective immunisation policies and reduce the countries' vulnerability to international contagion, see Fraser et al. (2009) for a recent example related to the 2009 influenza pandemic.

6 Discussion

Crises have been a recurrent feature of financial markets for a long time, though it was primarily the crises of the nineties that sparked increased interest in the study of their dynamics and propagation. The experience of the recent crises has rekindled such interest, emphasising the complexity of the financial system, and the need for a better understanding and monitoring of systemic risk and contagious effects.

This paper proposed a modelling framework to analyse financial contagion based on a stochastic epidemic process. The proposed approach directly accounts for the inherent dependencies involved in the propagation of financial crises. In particular, it allows for local and global transmissions through direct and indirect links to the originally affected country. Perhaps more importantly, the framework provides an implicit control mechanism, which includes a canonical measure of the crisis severity along with an estimate of the number of countries to financially support, in order to prevent a major crisis.

The suggested framework is general in nature, in that it can be applied to any contagious financial crisis not only as a way to learn from past experiences, but also as the crisis unfolds, enabling the evaluation of prospective immunisation policies. At the same time the approach is flexible, as it allows to incorporate features that reflect relevant theoretical and empirical evidence in the literature through the inclusion of appropriate covariates. Another notable feature is that it is simple enough to be amenable to model analysis and statistical inference procedures. This is in contrast to more complex stochastic models which are typically explored through simulations. Relating the latter to real data is far from trivial, unless simplifying assumptions are considered. It should be recognised, however, that the assessment of epidemic model fit is not a trivial matter. This is partly due to the bimodal realisations of such model outcomes. In the preceding application, we performed an informal

assessment by simulating 1000 epidemics from the estimated posterior predictive distribution of the stochastic epidemic two level mixing model for each crisis episode. The results showed larger outbreaks for the seventies and smaller epidemics for the nineties, indicating reasonable model fit.

In addition, the suggested framework naturally caters for temporal data, through the use of the proposed novel hidden epidemic model, contrary to the percolation-based approach. From a methodological viewpoint, it also has several appealing features. It dispenses with approximations such as infinite population, independent regions or supercriticality as is often assumed in epidemic modelling. Moreover, the proposed approach enables one to obtain additional information relating to the propagation of a crisis, not discussed thus far. This includes the set of potential infectors and their corresponding probabilities. The estimation of such crisis characteristics will be more precise in the case of temporal data. Inference regarding 'the most likely path' of the crisis spread can also be conducted, resulting in a set of the most probable pathways and their corresponding probabilities. Related work is presented in Shah and Zaman (2011). On the control side, an estimated cost can be assigned to each country and the total cost of the crisis in question can be calculated.

Further extensions can also be accommodated. The present set-up assumes a-priori independence of the recovery mechanism across countries, i.e. independent infectious periods. This can be relaxed and alternative processes can be incorporated into the framework, tailored to the application under consideration. A simple autoregressive structure could be a reasonable starting point to characterise the potential dependence of infectious periods. More involved density-dependent processes, can also be considered. For example, the model could be combined with models for the spread of rumours. The latter are similar to epidemic models but the stifling (recovery) of an individual depends on the recovery of others. In fact, within the economics literature, rumour models have proved useful in the study of the effects of diffusion of beliefs on market outcomes, see Banerjee (1993) and Kosfeld (2005). However, the use of classical rumour models alone (for example the Daley-Kendall and Maki-Thompson models, see Daley and Gani, 2000, for a review) would be inappropriate in the current context, as they imply that a major crisis will always occur. An interesting combination of epidemic and rumour models is presented in Nekovee et al. (2007) using interacting Markov chains. They demonstrate that the phase transition of classical epidemics is sufficient in characterising their model behaviour. They show that in the case of inhomo-

geneous networks the threshold parameter should be multiplied by $E(k^2)/E(k)$ where k is the degree of connections and E denotes the expectation operator. The $E(k^2)/E(k)$ factor accounts for non-homogeneous contact networks, see Anderson and May (1991) for an early discussion in the context of HIV modelling. In practice, an extension of this kind would only affect the recovery component of the likelihood associated with the hidden epidemic model, while the infection mechanism and subsequently the control mechanism, would remain unaffected. Mathematically this holds true since the spreading mechanism is of second order and its effect on the initial stages of spreading is negligible.

The superposition of poisson contact processes used in this paper may not be entirely appropriate for banking applications since a number of studies show that the degree distribution of the interbank market network follows a power law (see for instance Arinaminpathy et al., 2012, and references therein). However, it should be noted that by using a non-homogeneously mixing population we do allow for a more realistic (small world-type) topology of the underlying interaction network along which a crisis may spread. As such, the methodology of this paper can be used to conduct inference for a number of related models. For example, there are connections to the work of Arinaminpathy et al. (2012) where a bank's individual health could be determined by the local transmission rate while their confidence indicator may be thought of as corresponding to the global rate. In addition, the individual banks could be assigned to one of two types: large or small. Alternatively, one could adopt a non-poisson contact process, see Streftaris and Gibson (2012) for recently developed models where alternative renewal processes are used. Allowing for time-varying intensities could further be of interest. These can be accommodated by replacing λI_i with $\int_0^\infty \lambda_i(t)dt$ for 'suitable' infectiousness functions $\lambda_i(t)$. Such extensions and variants thereof, are left for future work.

Acknowledgements

We thank the Joint Editor, Associate Editor and two referees for valuable comments and suggestions which improved the paper. We are grateful to Yacine Aït-Sahalia, Frank Diebold, Mardi Dungey, Petra Geraats, Pete Neal, Phil O'Neill, Hashem Pesaran, Chris Rogers and Til Shuermann for useful comments on a preliminary version of this work. We further thank seminar and conference participants at the University of Cambridge, the Catholic University

of Leuven, Columbia University, Erasmus University, LSE, the Universitat Pompeu Fabra, USC, the New York Fed, the World Bank, the Second Cambridge/Penn conference, the Fifth Cambridge-Princeton conference, the conference on Interconnections in Financial Markets: Theoretical and Empirical Perspectives, in Cambridge, and the European Meeting of the Econometric Society in Malaga. Lyudmyla Hvozdyk, Siobhan McNamara and M. Tugrul Vehbi provided excellent research assistance in preparing the datasets used for the empirical illustration. The first author gratefully aknowledges financial support from the AUEB research funding for excellence and extroversion.

Appendix

Table A.1: List of countries and composition of regions

Northern America	**Southern Europe**	**Eastern Asia**	**Southern Africa**
US†	Bosnia	China	Botswana
Canada†	Croatia	Hong Kong†	Lesotho
Greenland*	Greece†	Japan†	South Africa
Central America	Italy†	Korea†	Swaziland
Belize	FYROM	Macao	**Eastern Africa**
Costa Rica	Malta†	Mongolia	Ethiopia
El Salvador	Portugal†	Chinese Taipei**	Kenya
Guatemala	Slovenia†	**South East Asia**	Madagascar
Honduras	Spain†	Cambodia	Malawi
Mexico	Yugoslavia	Indonesia	Mauritius
Nicaragua	**Eastern Europe**	Laos	Mozambique
Panama	Belarus	Malaysia	Reunion*
Southern America	Bulgaria	Myanmar	Rwanda
Argentina	Czech Republic	Philippines	Tanzania
Bolivia	Hungary	Singapore†	Uganda
Brazil	Moldova	Thailand	Zambia
Chile	Poland	Vietnam	Zimbabwe
Colombia	Romania	**Western Asia**	**Western Africa**
Ecuador	Russia	Armenia	Benin
French Guiana*	Slovak Republic	Azerbaijan	Burkina Faso
Guyana	Ukraine	Bahrain	Gambia
Paraguay	**Western Europe**	Cyprus	Ghana
Peru	Austria†	Georgia	Guinea-Bissau
Suriname	Belgium†	Iraq	Ivory Coast
Uruguay	France†	Israel†	Liberia
Venezuela	Germany†	Jordan	Mali
Caribbean	Netherlands†	Kuwait	Mauritania
Bahamas	Switzerland†	Lebanon	Niger
Barbados	**Central Asia**	Oman	Nigeria
Dominican Republic	Kazakhstan	Qatar	Senegal
Guadeloupe*	Kyrgyzstan	Saudi Arabia	Sierra Leone
Haiti	Tajikistan	Syria	Togo
Jamaica	Turkmenistan	Turkey	**Middle Africa**
Martinique*	Uzbekistan	United Arab Emirates	Angola
Trinidad	**Southern Asia**	Yemen	Cameroon
Northern Europe	Afghanistan	**Northern Africa**	Central Africa Republic
Denmark†	Bangladesh	Algeria	Congo
Estonia	India	Egypt	Gabon
Finland†	Iran	Libya	Guinea
Iceland†	Pakistan	Morocco	Zaire
Ireland†	Sri Lanka	Sudan	**Oceania**
Latvia		Tunisia	Australia†
Lithuania			Fiji
Norway†			New Caledonia
Sweden†			New Zealand†
UK†			Papua New Guinea

Notes: The regional decomposition is based on the United Nations classification. † denotes countries classified as developed based on the world factbook IMF classification, while the rest of the countries belong to the developing group. * and ** denote the countries that are not available in the BIS consolidated and locational banking statistics and the IMF Direction of Trade Statistics (DOTS) databases, respectively, used to demonstrate the introduction of covariates into the analysis. Guinea is also known as French Guinea. Guyana is also known as British Guiana. Zaire is what is currently known as the Democratic Republic of the Congo. Chinese Taipei is otherwise known as Taiwan. Yugoslavia is what used to be Yugoslavia SFR before its break up.

Table A.2: Crisis episodes and affected countries

Crisis	Year	Countries Affected
Break Down of Bretton Woods	1971	US, Denmark, Finland, Ireland, Norway, Sweden, UK, Greece, Italy, Portugal, Spain, Austria, Belgium, France, Germany‡, Netherlands, Switzerland, Australia, New Zealand
Collapse of the Smithsonian Agreement	1973	US, Denmark, Finland, Iceland, Norway, Sweden, UK, Greece, Italy, Portugal, Austria, Belgium, France, Germany‡, Netherlands, Switzerland, Japan, Australia, New Zealand
EMS Crisis	1992	Denmark, Finland‡, Ireland, Sweden, UK, Italy, Portugal, Spain, Belgium, France
Mexican Meltdown and Tequila Effect	1994	Canada, Mexico‡, Argentina, Brazil, Peru, Venezuela, Hungary, Hong Kong, Indonesia, Philippines, Thailand
Asian Flu	1997	Mexico, Argentina, Brazil, Czech Republic, Hungary, Poland, Pakistan, Hong Kong, Korea, Taiwan, Indonesia, Malaysia, Philippines, Singapore,Thailand‡, Vietnam, South Africa

Note: ‡ Denotes the country where the crisis originated. Czech Republic and Russia Federation are included in the 1971 and 1973 episodes corresponding to Czechoslovakia and USSR. The following countries: Armenia, Azerbaijan, Belarus, Bosnia, Croatia, Estonia, FYROM, Georgia, Kazakhstan, Kyrgyzstan, Latvia, Lithuania, Moldova, Slovak Republic, Slovenia, Tajikistan, Turkmenistan, Ukraine, Uzbekistan, did not exist at the time of the 1971 and 1973 crisis. Slovak Republic did not exist at the time of the 1992 crisis.

Table A.3: Percentage of trade among infected countries for the different crisis episodes

1992 Crisis Infected countries	% of trade among infected countries	1997 Crisis Infected countries	% of trade among infected countries
Denmark	0.42	Mexico	0.03
Finland	0.47	Brazil	0.09
Ireland	0.61	Czech Republic	0.07
Sweden	0.39	Hungary	0.07
UK	0.37	Poland	0.07
Italy	0.36	Pakistan	0.13
Portugal	0.55	Korea	0.15
Spain	0.48	Indonesia	0.23
Belgium	0.40	Malaysia	0.31
France	0.43	Philippines	0.19
1994 Crisis Infected countries	% of trade among infected countries	Singapore	0.33
		Thailand	0.22
Canada	0.03	Vietnam	0.46
Mexico	0.04	South Africa	0.04
Brazil	0.08		
Peru	0.14		
Venezuela	0.08		
Hungary	0.01		
Indonesia	0.04		
Philippines	0.05		
Thailand	0.04		

Note: Trade refers to the shares of the average of imports and exports.

References

Aït-Sahalia, Y., Cacho-Diaz, J.A. and R.J.A. Laeven (2010). Modeling financial contagion using mutually exciting processes. NBER Working Paper No. 15850.

Allen, F. and D. Gale (2000). Financial contagion. Journal of Political Economy 108, 1-33.

Amini, H., Cont, R. and A. Minca (2010). Resilience to contagion in financial networks. Mimeo.

Anderson, R.M. and R. M. May (1982). Population biology of infectious diseases. Report of the Dahlem Workshop, Berlin, 14th-19th March 1982, Springer-Verlag, Berlin.

Anderson, R. M., and R. M. May (1991). Infectious diseases of humans. Oxford, U.K.: Oxford University Press.

Andersson, H. and T. Britton (2000). Stochastic epidemic models and their statistical analysis. Lecture Notes in Statistics 151, Springer, New York.

Arinaminpathy, N., Kapadia, S., and R.M. May (2012). Proceedings of the National Academy of Sciences of the United States of America, 109, 18338–18343.

Ball, F. G. (1986). A unified approach to the distribution of total size and total area under the trajectory of infectives in epidemic models. Adv. Appl. Probab. 18, 289–310.

Ball, F. G., Britton, T. and O. D. Lyne (2004). Stochastic multitype epidemics in a community of households: Estimation of threshold parameter R_* and secure vaccination coverage. Biometrika, 91, 345-362.

Ball, F. and P. J. Donelly (1995). Strong approximations for epidemic models. Stochastic Processes and Applications, 55, 1-21.

Ball, F. G. and O. D. Lyne (2006). Optimal vaccination schemes for epidemics among a population of households, with application to variola minor in Brasil. Statistical Methods in Medical Research, 15, 481-487.

Ball, F.G., Mollison, D. and G. Scalia-Tomba (1997). Epidemics with two levels of mixing. Annals of Applied Probability 7, 46-89.

Balthrop, J., Forrest, S., Newman, M.E.J. and M.M. Williamson (2004). Technological networks and the spread of computer viruses. Science, 304, 527-529.

Banerjee, A.V. (1993). The economics of rumors. Review of Economic Studies 60, 309–327.

Barbour, A.D. and D. Mollison (1990). Epidemics and random graphs. In Stochastic Processes in epidemic theory, eds. Gabriel J. P. and Lefevre, C., Lecture notes in Biomathematics 86, 86-89.

Britton, T. (1998) Estimation in multitype epidemics. J. R. Statist. Soc. B 60, 993–679.

Britton, T. and N.G. Becker (2000). Estimating the immunity coverage required to prevent epidemics in a community of households. Biostatistics, 1, 389-402.

Brunetti, C., Scotti, C., Mariano, R. S. and A.H.H. Tan (2008). Markov switching GARCH models of currency turmoil in Southeast Asia. Emerging Markets Review 9, 104–128.

Calvo, G.A. and E.G. Mendoza (2000). Rational contagion and the globalization of securities markets. Journal of International Economics 55, 79-113.

Caporale, G.M., Serguieva, A. and H. Wu (2009). Financial contagion: evolutionary optimization of a multinational agent-based model. Intelligent Systems in Accounting, Finance and Management. Intell. Sys. Acc. Fin. Mgmt. 16, 111-125.

Caramazza, F., Ricci, L. and R. Salgado (2004). International financial contagion in currency crises. Journal of International Money and Finance 23, 51–70.

Chang, R. and G. Majnoni (2001). International Contagion: Implications for Policy. In: Claessens, S., Forbes, K. (Eds), International Financial Contagion. Kluwer Academic Publishers.

Corsetti, G., Pesenti, P., Roubini, N. and C. Tille (1998). Structural links and contagion effects in the Asian crisis: A welfare based approach. New York, New York University.

Daley, D.J. and J. Gani (2000) Epidemic modelling: An introduction. Cambridge University Press, Cambridge, UK.

Dasgupta, A., Leon-Gonzalez, R. and A. Shortland (2011), Regionality Revisited: An Examination of the Direction of Spread of Currency Crises. Journal of International Money and Finance 30, 831–848

De Gregorio, J. and R. Valdes (2001). Crisis transmission: Evidence from the debt, tequila and Asian flu crises. The World Bank Economic Review 15, 289-314.

Demiris, N. and P.D. O'Neill (2005a). Bayesian inference for stochastic multitype epidemics in structured populations via random graphs. Journal of the Royal Statistical Society Series B 67, 731-745.

Demiris, N. and P.D. O'Neill (2005b). Bayesian inference for epidemic models with two levels of mixing. Scand. J. Statist. 32, 265–280.

Diekmann, O. and J.A.P, Heesterbeek (2000). Mathematical epidemiology of infectious diseases: model building, analysis and interpretation. New York: Wiley.

Dornbusch, R., Park, Y.C. and S. Claessens (2000). Contagion: Understanding how it spreads. The World Bank Research Observer, 15, 177-197.

Dungey, M., Fry, R., González-Hermosillo, B. and V.L. Martin (2005a). A comparison of alternative tests of contagion with applications. In Dungey, M. and D. Tambakis (eds), Identifying international financial contagion: Progress and challenges, Ch. 3, Oxford University Press, New York.

Dungey, M., Fry, R., González-Hermosillo, B. and V.L. Martin (2005b). Sampling properties of contagion tests. Mimeo.

Dungey, M. and V.L. Martin (2007). Unravelling Financial Market Linkages During Crises. Journal of Applied Econometrics, 22, 89-119.

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

Ethier, S.N. and T.G. Kurtz (1986). Markov Processes: Characterization and Convergence. Wiley, New York.

Forbes, K.J. and R. Rigobon (2002). No contagion, only interdependence: Measuring stock market comovements. Journal of Finance, 5, 2223-2261.

Foss, and T. Konstantopoulos. (2003) Extended renovation theory and limit theorems for stochastic ordered graphs. Markov Proc. Rel. Fields, 9, 413-468.

Frankel, J.A. and A.K. Rose (1996). Currency crashes in emerging markets: An empirical treatment. Journal of International Economics, 41, 351-366.

Fraser, C., Donnelly, C.A., Cauchemez, S., Hanage, W.P., Van Kerkhove, M.D., Hollingsworth,

T.D., Griffin, J., Baggaley, R.F., Jenkins, H.E., Lyons, E.J., et al. (2009). Pandemic Potential of a Strain of Influenza A (H1N1) : Early Findings. Science, 324, 1557-1561.

Gai, P. and S. Kapadia (2010). Contagion in financial networks. Proc. Roy. Soc. A, vol. 466, no. 2120, 2401-2423.

Geoffard, P.-Y. and T. Philipson (1997). Disease eradication: Private versus public vaccination. American Economic Review, 87, 222-230.

Gerlach, S. and F. Smets (1995). Contagious speculative attacks. European Journal of Political Economy 11, 5–63.

Gersovitz, M. and J. S. Hammer (2004). The economical control of infectious diseases. Economic Journal, 114, 1-27.

Glick, R. and A.K. Rose (1999). Contagion and trade: Why are currency crises regional? Journal of International Money and Finance 18, 603-617.

Gravelle, T., Kichian, M. and J. Morley (2006). Detecting shift-contagion in currency and bond markets. Journal of International Economics, 68, 409-423.

Green, P.J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. Biometrika, 82, 711-732.

Haldane, A. (2009). Rethinking the Financial Network. www.bankofengland.co.uk/publications/speeches/2009/speech386.pdf.

Haile, F. and S. Pozo (2008). Currency Crisis Contagion and the Identification of Transmission Channels. International Review of Economics and Finance 17, 572-588.

Jacobs, J.P.A.M., Kuper, G.H. and Lestano (2005). Identifying financial crises. In Dungey, M. and D. Tambakis (Eds), Identifying International Financial Contagion: Progress and Challenges, Oxford University Press, Oxford, 86-110.

Jewell, C.P., Kypraios, T., Neal, P. and G.O. Roberts (2009). Bayesian analysis for emerging infectious diseases. Bayesian Analysis, 4, 465 - 498.

Kaminsky, G.L. (2006). Currency crises: Are they all the same? Journal of International Money and Finance 25, 503-527.

Kaminsky, G.L and C.M. Reinhart (1998). Financial crises in Asia and Latin America: Then

and now. American Economic Review 88, 444-448.

Kaminsky, G.L. and C.M. Reinhart (2000). On crises, contagion and confusion. Journal of International Economics 51, 145-168.

Kingman, J.F.C. (1993) Poisson processes. Oxford University Press.

Kodres, L.E. and M. Pritsker (2002). A rational expectations model of financial contagion. Journal of Finance, 57, 769–799.

Kosfeld, M. (2005). Rumours and markets. Journal of Mathematical Economics 41, 646–664.

Ludwig, D. (1975) Final size distributions for epidemics. Mathematical Biosciences 23, 33-46.

Mandilaras, A. and G. Bird (2010). A Markov switching analysis of contagion in the EMS. Journal of International Money and Finance, 29, 1062-1075.

Martinez Peria, M. S. (2002). A regime-switching approach to the study of speculative attacks: A focus on EMS crises. Empirical Economics, 27, 299-334.

May, R. M. and N. Arinaminpathy (2010). Systemic risk: the dynamics of model banking systems. J. R. Soc. Interface 7, 823-838.

May, R.M., Levin, S.A. and G. Sugihara (2008). Complex systems: Ecology for bankers, 451, 893-895.

Neal, P.J. (2006). Multitype randomised Reed-Frost epidemics and epidemics upon random graphs. Annals of Applied Probability 16, 1166-1189.

Nekovee, M., Moreno, Y., Bianconi, G. and M. Marsili (2007). Theory of rumour spreading in complex social networks. Physica A 374, 457-470.

Newman, M.E.J. (2010). Networks: An Introduction. Oxford University Press.

O'Neill, P.D. (2009). Bayesian inference for stochastic multitype epidemics in structured populations using sample data. Biostatistics, 10, 779-791.

O'Neill, P.D. and G.O. Roberts (1999). Bayesian inference for partially observed stochastic epidemics. Journal of the Royal Statistical Society Series A 162, 121-129.

Pericoli, M. and M. Sbracia (2003). A primer on financial contagion. Journal of Economic Surveys, 17, 571-608.

Pesaran, M.H. and A. Pick (2007). Econometric issues in the analysis of contagion. Journal of Economic Dynamics and Control 31, 1245-1277.

Rida, W.N. (1991). Asymptotic properties of some estimators for the infection rate in the general stochastic epidemic model. J. Roy. Statist. Soc. B, 53, 269–283.

Sachs, J., Tornell, A. and A. Velasco (1996). Financial crises in emerging markets: The lessons from 1995. Brookings Paper 27, 147-199.

Scalia-Tomba G. (1985) Asymptotic final size distribution for some chain-binomial processes. Adv. Appl. Prob. 17, 477-495.

Shah, D. and T. Zaman (2011). Rumors in a network: Who's the culprit? IEEE Transactions on Information Theory, 57, 5163-5181.

Streftaris, G. and Gibson, G. J. (2012). Nonexponential tolerance to infection in epidemic systems – modelling, inference and assessment. Biostatistics, 13, 580-593.

Toxvaerdy, F. (2010). Infection, Acquired Immunity and Externalities in Treatment. Unpublished Manuscript. CEPR Discussion Paper, 8112.

Van Rijckeghem, C., and B. Weder (2001), Sources of Contagion: Is it Finance or Trade?, Journal of International Economics, 54, 293-308.