Speed of Financial Contagion and Optimal Timing for Intervention*

Diana G. Mikhail†

Draft, December 31, 2019

Abstract

What constitutes timely intervention during a systemic crisis? Intervention that is 'too early' may not be appropriately designed due to the uncertainty surrounding the systemic nature of the crisis and may trigger panic. Intervention that is 'too late' may exacerbate the severity and duration of the crisis and ultimately prove ineffective in limiting financial contagion. In this paper, I develop a model to study optimal timing for intervention in a *Core-Periphery* interbank network of the financial system. Optimal timing for intervention during a cascade depends on two trade-offs: resilience of nodes against maturity of liabilities ('speed of financial contagion') and welfare of defaulting nodes against the cost of a systemic bailout. I find that faster contagion necessitates more immediate intervention and there is a threshold of speed beyond which immediate intervention becomes optimal. A 'too-interconnected-to-fail' effect arises endogenously in my model where a systemic bailout is warranted earlier when core nodes default even when it is more expensive. This finding is robust even when core nodes contribute less welfare to the financial system.

Keywords: Networks, Financial Fragility and Contagion, Financial Interconnectedness, Systemic Risk, Core-periphery, Bailouts

JEL Codes: D85, G01, G38, H81

*I am grateful to my advisors Burton Hollifield and Selman Erol for their invaluable input and guidance. I would also like to thank Chris Telmer, Brent Glover, Bryan Routledge, David Childers, Shu Lin Wee and participants at the CMU Finance Seminar for helpful comments and suggestions.

†Carnegie Mellon University. Email: dmikhail@andrew.cmu.edu.

1 Introduction

Optimal intervention arguably depends on distinguishing fundamental from self-fulfilling crises. Intervention is meant to avert self-fulfilling crises – policy tools such as lender of last resort (LOLR) and deposit insurance are designed to prevent coordination failures. Nonetheless, it is difficult to distinguish between fundamental and self-fulfilling crises. As a result, intervention has been the usual outcome. There is a growing debate in literature on the optimality of such intervention decision. The dominant view is that intervention may avert bad equilibria in a self-fulfilling crisis but may only delay a necessary adjustment in a fundamental crisis.

Optimal timing is necessary for optimal intervention. However, the literature fails to consider the optimality of intervention timing. Policy intervention during a systemic crisis is time-sensitive. Empirical evidence suggests that the timing of such policy interventions vary by 3-41 months from the crisis start date. Immediate intervention may be optimal in a self-fulfilling crisis but is it also optimal in a fundamental-based crisis? The government typically delays intervention due to its uncertainty about the specific cause of crisis and its systemic nature. In turn, premature intervention may be inappropriately designed or unnecessary which can trigger panic. Late intervention will 'miss' the crisis, exacerbate crisis' severity and duration, and ultimately prove ineffective in limiting financial contagion.

In this paper, I build a new model to study optimal timing for intervention in a financial network setting. The methodology is applicable to any network topology. The economy consists of a financial network of banks and a regulator. The set of possible linkages between banks is fixed over the horizon of the model but each period, a random subset of these linkages are due (i.e., maturity of liabilities is random). As a result, the economy may be in a cascade state or a normal state each period. In a normal state, all nodes are able to meet their obligations given their capital buffers. In a cascade state, however, some nodes are unable to meet their obligations and intervention may be warranted. Intervention is in the form of a one-time systemic bailout of all defaulting nodes in a period.

Two trade-off forces influence optimal timing for intervention. The first trade-off is the maturity of liabilities against the resilience (i.e., capital buffers) of nodes. The speed of financial contagion captures this trade-off and is measured as the number of nodes in default per period. It represents the extent of shock propagation over time. Shorter maturity and/or lower capital buffers accelerate contagion which may necessitate earlier intervention. The second trade-off is the welfare contribution of otherwise defaulting nodes against the cost of a systemic bailout. The cost of a systemic bailout is the shortfall amount of all defaulting nodes. Lower welfare contribution and/or higher cost may delay intervention. The regulator decides the optimal period for a systemic bailout given the above trade-offs.

Findings from a *Core-Periphery* network yield new insights into optimal timing for intervention. The speed of financial contagion has direct bearing on the immediacy of optimal intervention. Fast contagion necessitates earlier intervention and there is a threshold beyond which immediate intervention becomes optimal. A 'too-interconnected-to-fail' effect arises endogenously in the model. Earlier intervention is optimal when core nodes default relative to when periphery nodes default even when it is more expensive. Effect holds even when core nodes contribute less welfare to the financial system. The welfare-cost trade-off may delay or stop intervention. The higher the cost or lower the welfare of defaulting nodes, the less immediate and less likely the intervention. In addition, there is a lower bound on welfare below which intervention is suboptimal in some cascade states.

Model predictions have implications on the optimal design of intervention policy. First, the speed of intervention should be guided by the speed of financial contagion to be optimal. Fast contagion may necessitate earlier intervention while slow contagion may allow the government more time to learn about the nature of the crisis and design appropriate policy prior to intervention. Second, it is optimal to bailout core nodes first even when it is expensive as the distress of core nodes accelerates the speed of financial contagion.

This paper alerts regulators to the fact that effective intervention should consider timing. Timely intervention can minimize disruption to the financial system through reducing the economic costs incurred by financial institutions and accelerating their recovery from the crisis. Moreover, it can substantially reduce the fraction of institutions that face financial distress or default during a systemic crisis. The 2008-2009 financial crisis highlights that financial system stability is integral to the stability of the global economy. Timely intervention that appropriately addresses vulnerabilities in the financial system will have the critical macroeconomic consequences of limiting the extent of financial contagion which will ultimately reduce the severity and duration of systemic crises.

1.1 Literature Review

This paper relates to three strands of the literature: (1) financial networks and the extent of financial contagion, (2) optimality of bailouts, and (3) financial crises.

The link between interconnectedness and the extent of financial contagion is well-established in the literature. In contrast with this literature, this paper measures the speed of financial contagion (i.e., the extent of financial contagion over time) and use it to identify an optimal time for intervention. Allen and Gale (2000) develop a network model of liquidity shocks where they show that increased interconnectedness improves the resilience of the financial system by increasing risk-sharing. Acemoglu et al. (2015a) find that below some threshold of severity, increased interconnectedness indeed provide the documented risk-sharing benefits but beyond some threshold, facilitates financial contagion and fragility. Elliott et al. (2014) study how two important features of the structure of the financial network - diversification and integration - influence the extent of default cascades.

Literature shows that the anticipation of government intervention results in suboptimal consequences with respect to network formation and bank risk-taking. Cordella and Yeyati (2003), Freixas and Rochet (2013), Gorton and Huang (2004) among others document the moral hazard consequences that takes the form of excessive risk taking behavior in the presence of bailouts. Acemoglu et al. (2015b) and Erol and Vohra (2014) study endogenous network formation in the presence of systemic risk. In this paper, I consider the role of intervention timing in the optimality of intervention.

In addition, this paper is motivated by and contributes to the literature on financial crises. Laeven and Valencia (2008, 2010, 2012, 2013) develop a comprehensive systemic crises database that spans 1970-2011. The database contains the containment and resolution measures for a subset of 65 systemic crises. The database illustrates that it takes on average 5 months from the date of the crisis for the crisis to be deemed systemic and 3-41 months from the date of the crisis to observe the first form of government intervention. My paper contributes to this strand of literature since optimal timing of intervention matters for the duration and severity of systemic crises.

The paper proceeds as follows. Section 2 presents a model of optimal timing for a systemic bailout in a financial network setting. Section 3 applies the model to a *Core-Periphery* network and describes the findings. Section 4 concludes the paper.

2 Model: Optimal Timing of Systemic Bailouts

In this section, I develop a model to study optimal timing for a systemic bailout in a financial network setting. The economy consists of a financial network with n banks and the regulator. The regulator has uncertainty about the maturity of liabilities. The model has T periods. Each period, the economy may be in a normal state or a cascade state. In a cascade state, some nodes default and a systemic bailout may be warranted. The regulator decides the optimal period for a systemic bailout given two trade-offs: resilience of nodes against the maturity of liabilities and welfare of defaulting nodes against the cost of a systemic bailout. In 2.1, I describe the financial network setting and its dynamics for the optimal timing problem. In 2.2, I illustrate the trade-off forces that the regulator faces in the model when deciding an optimal timing for intervention. In 2.3, I present the methodology for solving an optimal stopping problem in this dynamic financial network setting.

2.1 Financial Network with n Banks

The set of possible linkages among n banks is fixed over the duration of the model. Let L_{ij} denote an $n \times n$ matrix that maps nominal obligations due from each node i to node j. Let $\bar{k}_i$ denote the total obligations of node i to all other nodes in the financial system.

$$\bar{k}_i = \sum_{j=1}^{n} L_{ij}$$

Each period, a random subset of these linkages are effective (i.e., maturity of liabilities is random). Let k_t denote an $n \times 1$ vector of liabilities due in a given period. In any given period t, a node may owe some fraction of its outstanding liabilities to a subset of its links, owe its total outstanding liabilities $\bar{k}_i$ or have no liabilities due. Node obligations in a given period cannot be deferred to future periods. Let L_{ij}^t denote an $n \times n$ matrix that captures the map of liabilities due in period t and let Π_{ij}^t denote an $n \times n$ matrix of liabilities proportional to total obligations due in period t.

$$\Pi_{ij}^t = \begin{cases} L_{ij}^t / k_t^i & \text{if } k_t^i > 0 \\ 0 & \text{otherwise} \end{cases} \qquad \text{where } \forall i, \sum_{j=1}^{n} \Pi_{ij}^t = 1$$

Assumption 1. *(No obligations deferral): Liabilities due in any given period cannot be deferred to future periods.*

For each node, there is a mapping from obligations that may be due each period to the nodes/links to which obligations are due. Such construction ensures that the realization of k_t provides a mapping to liabilities owed L_{ij}^t and proportional liabilities due Π_{ij}^t, where liabilities that are not owed in a given period are replaced with zero entries. Consider an example where a core node n_1 owes \$1 to n_2, n_3 and n_4 where n_2 and n_3 are core nodes and n_4 is a periphery node. Suppose that the mapping each period is such that a node may owe core or periphery but not both. Then, if the maturity realization of $k_t^1 = \$2$ then this means that n_1 only owes n_2 and n_3 (core nodes) in this period. If, however, the maturity realization is $k_t^1 = \$1$ then the node only owes n_4 (periphery node) in this period. In this example, a maturity realization of \$1 for n_1 entails zero entries for n_2 and n_3 in L_{ij}^t and similarly, a maturity realization of \$2 entails a zero entry for n_4 in L_{ij}^t.

Assumption 2. *(No recurrence of obligations to a node): Liabilities due in any given period must satisfy obligations to a subset of nodes. This ensures that a node cannot owe the same node more than once over the duration of the model.*

Defaults in the financial network depend on a node's cash flow relative to its liabilities. Let e_t denote an $n \times 1$ vector of cash flows used to pay liabilities owed within a period. If a node's cash flow e_t^i exceeds liabilities owed k_t^i, the node pays the liabilities it owes to the network. If, however, a node's cash flow e_t^i falls below liabilities owed k_t^i, the node defaults and pays its cash flow to the network. Claimant nodes are paid in proportion to the amount they were owed by the defaulting node. Let $\mathbf{D}(k_t, e_t)$ denote the set of defaulting nodes i such that $k_t > e_t$. Defaults are pinned down by k_t and e_t before any payments are made.[1] Let p_t denote an $n \times 1$ vector of payments at any given period t. Payments determined by the realization of k_t and e_t today are cleared next period. The law of motion for a node's cash flow e_t depends on its previous cash flow e_{t-1}, previous period payment to the network p_{t-1} and previous period payments received from the network $\Pi_{ij}^{t-1} p_{t-1}$.

$$p_t = min(k_t, e_t)$$

$$e_t = e_{t-1} - p_{t-1} + \Pi_{ij}^{t-1} p_{t-1}$$

Assumption 3. *(Proportionality): Claimant nodes are paid by the defaulting node in proportion to the original amount owed.*

Assumption 4. *(Next period payment clearing): Payments due in the current period are cleared next period.*

The state of the economy depends on the realization of k_t and e_t. At any given period t, the economy can either be in a cascade state or in a normal state. If the realization of k_t is less than e_t for all nodes n in period t, all nodes are able to meet their obligations and the economy is in a normal state. If, however, the realization of k_t for at least one node i is greater then e_t, there is at least one observed default and the economy is in a cascade state.

[1] If there is a payment clearing mechanism each period, the order with which payments are cleared would change the sequencing and number of defaults.

Let c_t denote the state of the economy in period t.

$$\text{State of the Economy:} \begin{cases} k_t > e_t & c_t = 1 \text{ (cascade state)} \\ k_t \leq e_t & c_t = 0 \text{ (normal state)} \end{cases}$$

2.2 Regulator

Since k_t is stochastic and e_t is deterministic, the regulator cannot foresee whether the degree of resilience of the financial system is sufficient to avert a cascade state next period. Each period, the regulator observes the state of the economy and optimally decides whether or not to intervene. Intervention is in the form of a systemic bailout where all nodes in default at the time of intervention are bailed out. The regulator decides the optimal period for a systemic bailout (i.e., may only intervene once over the horizon of the model).

Assumption 5. *(Systemic or untargeted bailout): In the case of intervention, the regulator must bailout all nodes in default in a given period. This assumption implicitly entails that the regulator has enough budget to do so.*

Assumption 6. *(One-shot bailout): Regulator can only intervene once over the horizon of the model.*

Identifying an optimal period for a systemic bailout depends on two trade-off forces in the model. The first trade-off is the maturity of liabilities against the resilience of nodes. This trade-off influences how quickly distress propagates through the financial system in a cascade state. I refer to this trade-off as the speed of financial contagion. Keeping maturity k_t fixed, the greater the resilience e_t, the slower the propagation of the shock. Keeping resilience e_t fixed, the shorter the maturity of liabilities k_t, the faster the cascade.

Speed of financial contagion critically depends on both maturity k_t and resilience e_t as a node with little or no cash flow will not default unless it has outstanding obligations in a given period. Therefore, I define a measure for speed that accounts for the interaction of k_t and e_t – the number of nodes in default per period. Slow contagion allows the regulator more time to learn about the systemic nature of the crisis and design appropriate policy for later intervention. Fast contagion, on the other hand, may necessitate earlier intervention. The speed of financial contagion thus constitutes the first important trade-off in identifying an optimal time for intervention.

The second important trade-off in identifying an optimal time for intervention is the welfare-cost of a systemic bailout. Given the realization of k_t and e_t, the regulator can identify nodes in default and incur the shortfall needed to restore the system to a normal state. The shortfall for each node i takes a value of $k_t^i - e_t^i$ when a node defaults and zero otherwise. The cost of a systemic bailout is the shortfall amount of all nodes in the default set in each period t. The benefit of a systemic bailout lies in the welfare contribution of otherwise defaulting nodes. Let w denote an $n \times 1$ vector that specifies the welfare contribution of each node to the economy. Keeping cost fixed $(k_t - e_t)$, the higher the welfare contribution, the earlier the intervention. Keeping welfare contribution fixed (w), the higher the cost, the later the intervention.

The cost and benefit of a systemic bailout are endogenous to the state of the economy as the set of defaulting nodes is pinned down by each pair (k_t, e_t). The set of defaulting nodes can identify the exact cost of a systemic bailout since it depends of the shortfall amount $(k_t - e_t)$ of all defaulting nodes. It also identifies the benefit of a systemic bailout since it depends of the welfare contribution of all otherwise defaulting nodes.

$$\text{Cost of Systemic Bailout} = \sum_{i \in \mathbf{D}} (k_t^i - e_t^i)$$

$$\text{Benefit of Systemic Bailout} = \sum_{i \in \mathbf{D}} w_i$$

Assumption 7. *(Welfare contribution of each node is not time-varying): The welfare contribution of each node to the financial system is constant over the horizon of the model.*

2.3 Optimal Stopping Problem

Using the above setting, I develop a finite horizon stochastic optimal control model for optimal intervention timing during a cascade.

State Space (k_t, e_t)**.** For each node, a state space for k_t^i is defined such that it encompasses possible liabilities due in a period. Maturity of liabilities is random. Each period, a node may have no obligations due, owe some fraction of its obligations or all of its obligations. The state space for each node is constructed such that it can identify the subset of nodes that are due payment in a period. Using the state space defined for each node, we can generate all possible permutations of the k_t vector and its corresponding mapping to L_{ij}^t and Π_{ij}^t. Let K denote the number of *unique* permutations of the k_t vector based on the defined state space for each node.

The cash flow vector e_t depends on the realization of maturity k_t. At each time step, the new cash flow is the old cash flow less payments made to the network plus payments received from the network. In a cascade state, each node pays the minimum of what it has and what it owes. In a normal state, each node pays only what it owes. Let e_0 denote an $n \times 1$ vector with the initial cash flow for each node in the financial system. Pairing e_0 with each of the possible K realizations of k_t yields K realizations of e_1. Repeating this process for each of the K realizations of e_1 yields $K \times K$ new realizations of e_2 and so on. The state space for e_t expands with each transition to include the new unique vectors of e_{t+1}. The computational cost substantially increases with the number of transitions. For tractability, I limit the state space of e_t to the first transition. Let E denote the number of *unique* vectors in the state space of e_t.

$$e_{t+1} = \begin{cases} e_t - p_t + \Pi_{ij}^{t}{}' p_t & c_t = 1 \ \textit{where} \ p_t = min\{k_t, e_t\} \\ e_t - k_t + \Pi_{ij}^{t}{}' k_t & c_t = 0 \end{cases}$$

We can now construct the state space for the problem $KE = K \times E$ where each state is a pair of vectors (k_t, e_t).

Action (Control) Space (X)**.** In each state, the regulator observes (k_t, e_t) and optimally decides whether a systemic bailout is warranted. Let x_t denote the regulator's decision.

$$x_t \in X = \{\text{bailout}(1), \text{no bailout}(0)\}$$

.

Reward Function (k_t, e_t, x_t)**.** For each pair (k_t, e_t), we can identify the set of defaulting nodes (if any) and compute the reward from intervention. In the case of intervention, the reward is the welfare contribution of otherwise defaulting nodes less the cost of a systemic bailout $(k_t - e_t)$. There is no immediate reward from the regulator's decision to wait. Let R denote a $KE \times 2$ matrix that contains the reward from intervention vs. waiting in each (k_t, e_t) state.

$$f(k, e, x) = \begin{cases} 0 & x = 0, \quad c_t = 1 \\ \sum\limits_{i \in \mathbf{D}} w - \sum\limits_{i \in \mathbf{D}} (k - e) & x = 1, \quad c_t = 1 \end{cases}$$

Joint Transition Map (k_t, e_t)**.** A probability transition matrix is defined over the liabilities state space for each node such that liabilities due in the current state cannot reoccur in the future. If the current maturity state is no liabilities owed, then all states tomorrow are equally likely (i.e., assigned equal probability). Otherwise, a positive probability is assigned to states other than the current state and a transition probability of zero is assigned to the current state. We can now construct an $K \times K$ probability transition matrix where each entry in the matrix is an $n \times 1$ vector that maps the transition probability of each element from k_t to k_{t+1} according to the probability transition matrix defined for each node. The maturity realizations across nodes are independent and therefore, the joint probability can be computed as the product of vector elements. This simplifies each entry of the probability transition matrix from an $n \times 1$ vector of element-wise probabilities corresponding to the maturity of each node to a scalar that represents the transition probability for k_t as a vector. This yields an $K \times K$ probability transition matrix (K_{ij}) for k_t.

We can now construct the probability transition matrix for the problem $KE \times KE$. For each pair (k_t, e_t), we compute e_{t+1} based on the law of motion defined for the cash flow vector. If a pair's e_{t+1} falls within the defined state space for e_t, we parse out the k^{th} row corresponding to the pair's k_t from the K_{ij} matrix into the respective e_{t+1} columns. If, however, the computed e_{t+1} does not fall within the state space, then we assign a transition probability of 1 to the state that corresponds to the current state. Essentially, this creates an absorbing state for any pair with an e_{t+1} that falls outside the state space for e_t. All other entries in the matrix are zero.

Value Function (k, e)**.** Let $V_t(k, e)$ denote the value of a systemic bailout for a financial system with liabilities-resilience pair $(\vec{k_t}, \vec{e_t})$. In a cascade state, the payoff from intervention is the welfare gained from otherwise defaulting nodes less the amount of the shortfall needed to bailout all nodes in default. When the regulator decides to wait, there is no immediate value but regulator retains the option to intervene in the next state where with some probability the cascade stops or continues. The option to step in expires at the terminal period and therefore, the post-terminal value is zero.

$$V_t(k, e) = max\{\sum_{i \in \mathbf{D}} w - \sum_{i \in \mathbf{D}} (k - e),\ 0 + \delta \mathbb{E} V_{t+1}(k', e')\}$$

$$V_{T+1}(k, e) = 0$$

3 Optimal Timing for a Systemic Bailout in a Core-Periphery Network

In this section, I apply the framework from section 2 to study optimal timing for a systemic bailout in a core-periphery network with $n = 6$ nodes and $T = 5$ periods. In 3.1, I describe the environment. In 3.2, I apply the methodology for constructing the state space and transition map given the network topology. In 3.3, I show the parameter calibration for the problem. In 3.4, I present the findings for this core-periphery network.

3.1 Environment

Suppose that there are five periods, $T = 5$. The regulator decides the optimal period for a systemic bailout. Intervention in $t = 1$ is deemed immediate.

Financial Network. Consider an economy with 6 banks where n_1, n_2, and n_3 are *Core* nodes and n_4, n_5, and n_6 are *Periphery* nodes. In this economy, *Core* nodes are systemically important not due to their size but due to their degree of interconnectedness with the financial system. Each core node is linked to two core nodes and one periphery node. Each periphery node is linked to one core node. These linkages represent the set of possible interconnections between nodes over the horizon of the model. Suppose that each node owes \$1 to each link it forms with a counterparty. We can now construct the map of obligations due (L_{ij}) from each node i to node j and compute the total obligations due ($\bar{k}$) for each node i. Figure 1 illustrates the map of liabilities owed L_{ij} and total obligations due $\bar{k}$ for this network.

Figure 1: A Core-Periphery Network

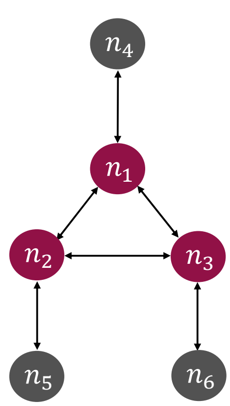

$$
\boldsymbol{L_{ij}} = \begin{array}{c} \begin{array}{cccccc} n_1 & n_2 & n_3 & n_4 & n_5 & n_6 \end{array} \\ \begin{bmatrix} 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \end{array} \begin{array}{c} \\ n_1 \\ n_2 \\ n_3 \\ n_4 \\ n_5 \\ n_6 \end{array} \qquad \bar{\boldsymbol{k}} = \begin{bmatrix} 3 \\ 3 \\ 3 \\ 1 \\ 1 \\ 1 \end{bmatrix} \begin{array}{c} n_1 \\ n_2 \\ n_3 \\ n_4 \\ n_5 \\ n_6 \end{array}
$$

Based on the set of possible linkages, the maximum amount each node may owe to the financial system is \$3 for a core node and \$1 for a periphery node. The uncertainty about maturity of liabilities entails that some fraction of these liabilities may become due to a subset of links/nodes each period. If the cash flow in the financial system is not sufficient to cover the liabilities due, some nodes default and the economy enters a cascade state. The regulator then optimally decides when to intervene.

3.2 Problem State Space and Transition Map (k_t, e_t)

State Space (k_t, e_t). I describe below the steps required to construct the state space for the problem. Steps 1 and 2 construct the state space for k_t. Steps 3 and 4 use the state space for k_t to pin down the state space for e_t. Step 5 uses the state space for k_t and e_t to construct the final state space for the problem which is composed of $25,272$ (k_t, e_t) pairs.

Step 1: Define a liabilities state space for each node and its mapping to linkages. Each element in the state space represents liabilities owed to a subset of possible linkages. As a result, the state space has a mapping to L_{ij}^t where liabilities that are not due in period t are replaced with zero entries. I define the state space and mapping to linkages to be identical for each core and periphery node. For each core node, the state space for $k_t^{i=1:3}$ is $\{0, 1, 2\}$ where a realization of $\{0\}$ indicates no liabilities due, a realization of $\{1\}$ indicates liabilities due to the periphery node and a realization of $\{2\}$ indicates liabilities due to core nodes. For each periphery node, the state space for $k_t^{i=4:6}$ is $\{0, 1\}$ where a realization of $\{0\}$ indicates no liabilities due and a realization of $\{1\}$ indicates liabilities due to the core node.

Figure 2: Node State Space and Mapping to Linkages

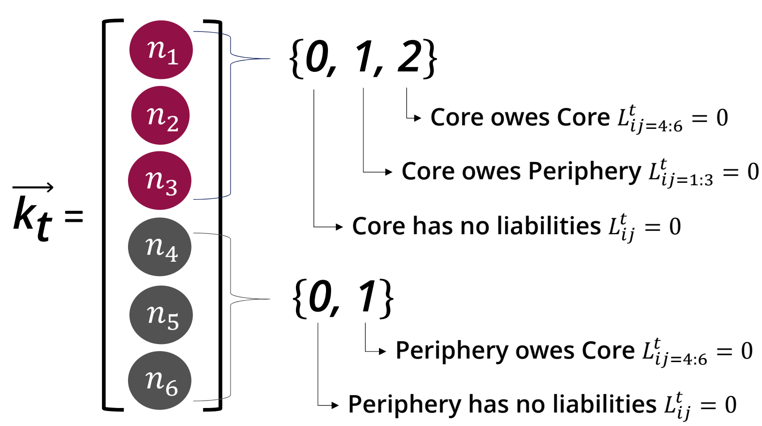

Step 2: Use the defined state space for each node to generate permutations of the k_t vector. Given that the defined state space is $\{0, 1, 2\}$ for each core node and $\{0, 1\}$ for each periphery node, we can generate all possible permutations of the k_t vector. Each k_t vector (6×1) specifies what each node in the financial system owes the network in a given period.

This step results in 216 unique vectors for k_t.

Step 3: Initialize the cash flow vector e_0. In order to generate the state space for e_t, we must initialize the cash flow vector. I assume that each node has an initial cash flow of \$1 (i.e., $e_0 = \vec{1}$).

Step 4: Compute e_{t+1} *for each unique permutation of the* k_t *vector and keep unique* e_t *vectors.* Use e_0 and the law of motion for the cash flow vector to compute e_1 for each unique permutation of the k_t vector. Recall that the law of motion is $e_{t+1} = e_t - p_t + \Pi_{ij}^{t}{}' p_t$ where $p_t = min\{k_t, e_t\}$. This step results in 216 vectors for e_1. Note that pairing each of the e_1 vectors with 216 possible k_t vectors yields 216×216 possible e_2 vectors and so on. For tractability, I limit the state space of e_t to e_0 and the 216 e_1 vectors from the first transition. Each e_t vector (6×1) specifies the cash flow that each node in the financial system has to meet its obligations in a given period. This step yields 117 unique vectors for e_t.

Step 5: Construct the joint state space (k_t, e_t) *from unique* k_t *and* e_t *vectors.* Using the 216 unique vectors for k_t from Step 2 and the 117 unique vectors for e_t from Step 4, we can now construct the final state space for the problem. The problem state space is composed of all possible (k_t, e_t) pairs (i.e., 216×117 states). This step results in $25,272$ states where each state is a pair of vectors (k_t, e_t).

Figure 3: Problem State Space: Selected Steps

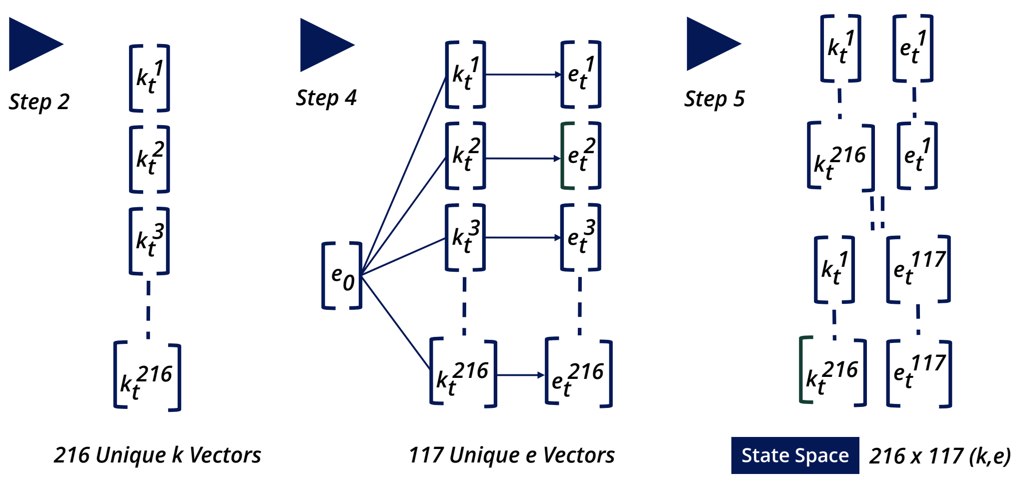

Using the problem state space, we can now compute the reward from intervention for each pair (k_t, e_t). The reward is the welfare contribution of otherwise defaulting nodes less the cost of a systemic bailout. The cost of a systemic bailout is the shortfall of all defaulting nodes (i.e., $k_t - e_t$). The benefit from intervention is the welfare contribution of bailed out nodes. I assign equal welfare contribution to all nodes and compare findings using two calibrations:

high welfare ($w = \vec{20}$) and low welfare ($w = \vec{2}$). Let R denote a $25,272 \times 2$ matrix that contains the regulator's reward in each (k_t, e_t) state from intervention vs. waiting.

Joint Transition Map (k_t, e_t)**.** I describe below the steps required to construct the joint probability transition matrix for the problem. Steps 1 and 2 construct the probability transition matrix for k_t. Steps 3 and 4 construct the final transition map for the problem $25,272 \times 25,272$.

Step 1: Define a probability transition matrix over the liabilities state space for each node. The transition probabilities are assigned such that liabilities owed this period cannot reoccur in the future. I define an identical probability transition matrix for each core and periphery node. For a core node, the defined mapping to linkages is such that a node may have liabilities due to core $\{2\}$, liabilities due to periphery $\{1\}$ or no liabilities due $\{0\}$ in any given state. Therefore, a core node that owes \$2 in the current state can transition to states $\{1\}$ or $\{0\}$ with equal probability. Similarly, a core node that owes \$1 in the current state can transition to states $\{2\}$ or $\{0\}$ with equal probability. For a periphery node, the defined mapping to linkages is such that a node may have liabilities due to core $\{1\}$ or no liabilities due $\{0\}$ in any given state. Therefore, a periphery node that owes \$1 in the current state will transition to state $\{0\}$ with probability 1 since it satisfied all its obligations. A core or periphery node with no liabilities due in the current state can transition to all states with equal probability in the next state.

Figure 4: Node Probability Transition Matrix

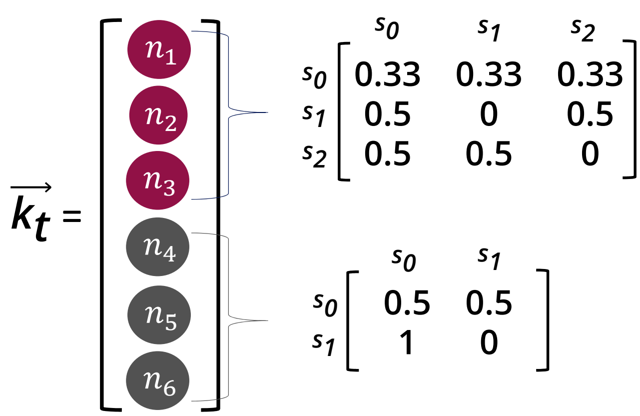

Step 2: Construct a probability transition matrix for k_t. Use the defined probability transition matrix for each node to map the probability of each element in each permutation of k_t to transition to k_{t+1}. This step yields 216×216 vectors where each vector contains the probability for each node to transition from k_t^i to k_{t+1}^i. Next, compute the joint probability

that each k_t vector transitions to k_{t+1}. This step simplifies each transition vector to a scalar (product of vector elements) assuming independence of realizations across nodes. As a result, the probability transition matrix for k_t simplifies from 216×216 vectors to 216×216 elements where each element represents the probability that k_t transitions to k_{t+1}.

Figure 5: Probability Transition Matrix for k_t

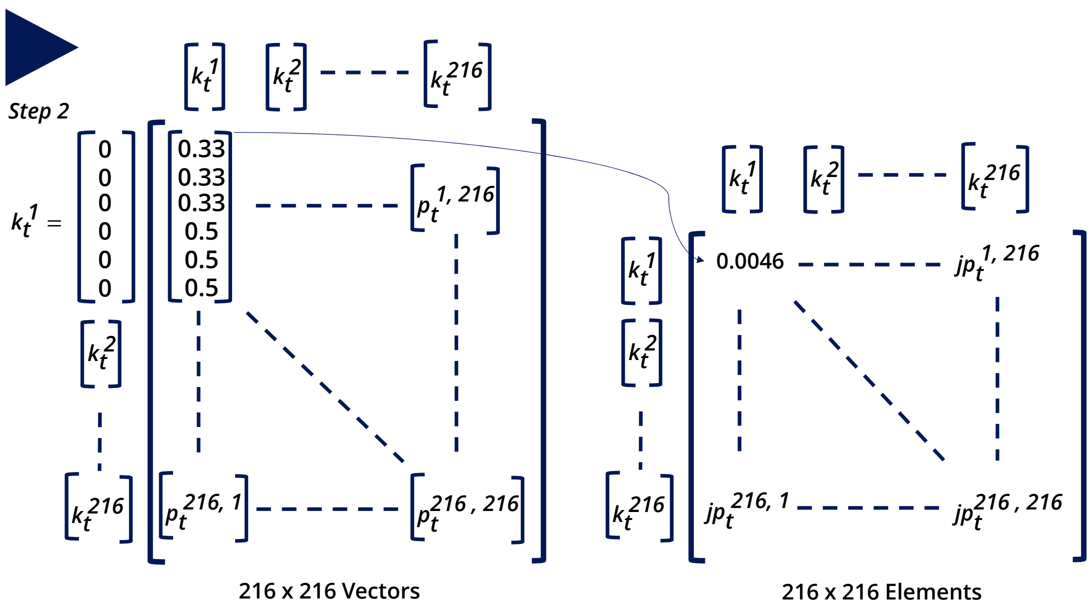

Step 3: Compute e_{t+1} for each (k_t, e_t) pair in the state space. Given a pair (k_t, e_t), e_{t+1} is known with certainty while k_{t+1} is not. Recall that the only source of uncertainty in this economy stems from the realization of k_{t+1} (Step 2). In order to identify the transition probability from (k_t, e_t) to (k_{t+1}, e_{t+1}), we first compute e_{t+1} for each (k_t, e_t) pair. This step allows us to limit possible transition states for (k_t, e_t) to pairs with e_{t+1}.

Step 4: Construct a joint probability transition matrix for (k_t, e_t). For (k_t, e_t) pairs where e_{t+1} falls within the state space for e_t, parse out the row corresponding to the pair's k_t from Step 2 to the columns corresponding to e_{t+1}. For (k_t, e_t) pairs where e_{t+1} falls outside the state space for e_t, assign a transition probability of 1 to the current state. This step yields $25,272 \times 25,272$ matrix for the joint transition probability of (k_t, e_t). Let P denote this probability transition matrix.

3.3 Calibration

Using the below calibration, this finite horizon stochastic optimal control problem is solved numerically using dynamic programming.

Parameters		
Variable Name	Description	Value
T	Horizon	5
δ	Discount Factor	0.96
w	Welfare Contribution	$\vec{20}$, $\vec{2}$

There are 5 inputs to the problem: horizon ($T = 5$), discount factor ($\delta = 0.96$), decision ($x_t = \{0, 1\}$), intervention reward (R matrix), and probability transition matrix (P matrix). I compare findings under a high welfare ($w = \vec{20}$) relative to a low welfare ($w = \vec{2}$) calibration.

3.4 Findings

Table 1 shows the state space decomposition of optimal timing decision under both welfare calibrations. Immediate intervention is optimal in 74.9% and 72.3% of all states under high and low welfare calibrations, respectively. Intervention delay is optimal in 16% and 14% of all states under high and low welfare calibrations, respectively. In a cascade state, intervention is always optimal under a high welfare calibration but is suboptimal in 4.3% of all states under a low welfare calibration. As would be expected, intervention is strictly suboptimal when the economy is in a normal state under both welfare calibrations.

Table 1: State Space Decomposition and Timing of Intervention
Number of States (in % of all states)

		Low Welfare		*High Welfare*	
		Cascade	*Normal*	*Cascade*	*Normal*
Timing of Intervention	*Immediate t=1*	18,263 (72.3%)		18,925 (74.9%)	
	Wait until t=2	126 (0.5%)		34 (0.1%)	
	Wait until t=3	332 (1.3%)		77 (0.3%)	
	Wait until t=4	1,288 (5.1%)		1,380 (5.5%)	
	Wait until t=5	1,850 (7.3%)		2,531 (10.0%)	
	No Intervention	1,088 (4.3%)	2,325 (9.2%)		2,325 (9.2%)

3.4.1 There is a Threshold of Speed Beyond Which Intervention is Optimal

The higher the number of defaults per period (speed), the more immediate the intervention. Defaults of two or more nodes within a period accounts for 77.6% and 75.9% of states where immediate intervention is optimal under high and low welfare calibrations, respectively. Table 1 shows that there is a threshold of speed beyond which immediate intervention is optimal. Under high welfare calibration, immediate intervention is optimal when three or more nodes default. Under low welfare calibration, on the other hand, immediate intervention is optimal when four or more nodes default. This indicates that the greater the welfare, the sooner we approach the threshold at which immediate intervention becomes optimal.

Table 2: Speed of Contagion and Timing of Intervention
Number of States (in % of states with same # of defaults)

High Welfare

Speed of Contagion (Fast ←→ Slow)

Timing of Intervention	Defaults per Period: *5*	*4*	*3*	*2*	*1*
Immediate t=1	15 (100%)	492 (100%)	4,317 (100%)	9,859 (99.92%)	4,242 (51.4%)
Wait until t=2				4 (0.04%)	30 (0.4%)
Wait until t=3				4 (0.04%)	73 (0.9%)
Wait until t=4					1,380 (16.7%)
Wait until t=5					2,531 (30.7%)

Low Welfare

Speed of Contagion (Fast ←→ Slow)

Timing of Intervention	Defaults per Period: *5*	*4*	*3*	*2*	*1*
Immediate t=1	15 (100%)	492 (100%)	4,247 (98.4%)	9,114 (92.4%)	4,395 (53.2%)
Wait until t=2			6 (0.1%)	40 (0.4%)	80 (1.0%)
Wait until t=3			11 (0.3%)	94 (1.0%)	227 (2.7%)
Wait until t=4			29 (0.7%)	295 (3.0%)	964 (11.7%)
Wait until t=5			16 (0.4%)	216 (2.2%)	1,618 (19.6%)

3.4.2 The Higher the Cost, the Less Immediate and Less Likely the Intervention

The cost of intervention is considered 'high' if the amount of the shortfall is greater than 4, 'medium' if the amount is greater than 2 and less than or equal to 4, and 'low' if the amount is less than 2. Table 2 shows how optimal timing varies with the cost of intervention. Under high welfare calibration, the likelihood of intermediate intervention when the cost is 'high' is 1.2% relative to 38.3% and 60.4% under 'medium' and 'low' cost, respectively. Under low welfare calibration, the likelihood of immediate intervention when the cost is 'high' is 0.9% which indicates that intervention becomes more suboptimal under high cost and low welfare. Under both calibrations, we can also observe that the regulator is less likely to step in with a delay under 'high' cost relative to 'medium' or 'low' cost.

Table 3: Cost and Timing of Intervention

Number of States (in % of states with same intervention decision)

High Welfare

Timing of Intervention		Cost of Intervention: High	Medium	Low	
	Immediate t=1	236 (1.2%)	7,249 (38.3%)	11,440 (60.4%)	**18,925 (100%)**
	Wait until t=2		3 (8.8%)	31 (91.2%)	**34 (100%)**
	Wait until t=3		3 (3.9%)	74 (96.1%)	**77 (100%)**
	Wait until t=4			1,380 (100.0%)	**1,380 (100%)**
	Wait until t=5			2,531 (100.0%)	**2,531 (100%)**

Low Welfare

Timing of Intervention		Cost of Intervention: High	Medium	Low	
	Immediate t=1	173 (0.9%)	6,578 (36.0%)	11,512 (63.0%)	**18,263 (100%)**
	Wait until t=2	6 (4.8%)	25 (19.8%)	95 (75.4%)	**126 (100%)**
	Wait until t=3	11 (3.3%)	70 (21.1%)	251 (75.6%)	**332 (100%)**
	Wait until t=4	22 (1.7%)	263 (20.4%)	1,003 (77.9%)	**1,288 (100%)**
	Wait until t=5	16 (0.9%)	211 (11.4%)	1,623 (87.7%)	**1,850 (100%)**

3.4.3 Immediate Intervention is More Likely When a Core Node Defaults Even at the Expense of Higher Cost

Although the calibration assigns equal welfare to core and periphery nodes, I observe that intervention is more immediate when a core node defaults than when a periphery node defaults. Table 3 shows that, under high welfare calibration, the default of three, two and one core nodes corresponds to an immediate intervention likelihood of 5.7%, 35.4% and 47.4%, respectively. This compares to only 1%, 12.4% and 42.8% when three, two and one periphery nodes default, respectively. Similar finding holds under low welfare calibration. This emphasizes the systemic importance of core nodes to the network topology and is consistent with the 'too-interconnected-to-fail' effect.

Table 4: Systemic Importance and Likelihood of Immediate Intervention
Number of States (in % of states with immediate intervention)

High Welfare

# of Core Nodes in Default	# of Periphery Nodes in Default: 3	2	1	0	
3	0 (0.0%)	6 (0.0%)	168 (0.9%)	906 (4.8%)	**1,080 (5.7%)**
2	9 (0.0%)	261 (1.4%)	2,187 (11.6%)	4,233 (22.4%)	**6,690 (35.4%)**
1	63 (0.3%)	1,107 (5.8%)	4,652 (24.6%)	3,144 (16.6%)	**8,966 (47.4%)**
0	117 (0.6%)	974 (5.1%)	1,098 (5.8%)	0 (0.0%)	
	189 (1.0%)	**2,348 (12.4%)**	**8,105 (42.8%)**		

Low Welfare

# of Core Nodes in Default	# of Periphery Nodes in Default: 3	2	1	0	
3	0 (0.0%)	6 (0.0%)	168 (0.9%)	844 (4.6%)	**1,018 (5.6%)**
2	9 (0.0%)	261 (1.4%)	2,179 (11.9%)	3,746 (20.5%)	**6,195 (33.9%)**
1	63 (0.3%)	1,107 (6.1%)	4,418 (24.2%)	3,106 (17.0%)	**8,694 (47.6%)**
0	117 (0.6%)	950 (5.2%)	1,289 (7.1%)	0 (0.0%)	
	189 (1.0%)	**2,324 (12.7%)**	**8,054 (44.1%)**		

Table 4 illustrates how the cost of intervention influences the likelihood of immediate intervention when core nodes default relative to when periphery nodes default. The likelihood of immediate intervention under 'high' cost when two core nodes default is 0.8% relative to 0.2% when two periphery nodes default. Similarly, under high welfare calibration, the likelihood of immediate intervention under 'medium' cost when three and two core nodes default is 4.2% and 19.6% relative to only 1% and 7.1% when three and two periphery nodes default, respectively. The same result holds under low welfare calibration. This suggests that more immediate intervention is optimal when two or more core nodes default even when it comes at the expense of higher cost which further supports the 'too-interconnected-to-fail' effect. Notice, however, that this effect does not hold when only one core node defaults. The likelihood of immediate intervention under 'medium' cost when only one periphery node defaults is 17.4% and 16.9% relative to 13.9% and 13.3% when one core node defaults under high and low welfare calibrations, respectively.

Table 5: Systemic Importance and Cost of Intervention

Number of States (in % of states with immediate intervention)

High Welfare

Cost of Intervention

# of Core Nodes in Default	High	Medium	Low
3	92 (0.5%)	790 (4.2%)	198 (1.0%)
2	144 (0.8%)	3,708 (19.6%)	2,838 (15.0%)
1		2,634 (13.9%)	6,332 (33.5%)

Cost of Intervention

# of Periphery Nodes in Default	High	Medium	Low
3		189 (1.0%)	
2	36 (0.2%)	1,338 (7.1%)	974 (5.1%)
1	120 (0.6%)	3,288 (17.4%)	4,697 (24.8%)

Low Welfare

Cost of Intervention

# of Core Nodes in Default	High	Medium	Low
3	36 (0.2%)	784 (4.3%)	198 (1.1%)
2	137 (0.8%)	3,250 (17.8%)	2,808 (15.4%)
1		2,427 (13.3%)	6,267 (34.3%)

Cost of Intervention

# of Periphery Nodes in Default	High	Medium	Low
3		189 (1.0%)	
2	36 (0.2%)	1,338 (7.3%)	950 (5.2%)
1	113 (0.6%)	3,080 (16.9%)	4,861 (26.6%)

Robustness: What if Core Nodes Contribute Less Welfare?

Table 5 shows the results under two welfare calibrations where core nodes contribute less welfare relative to periphery nodes. Under a calibration of $w_c = 15$ and $w_p = 20$, we observe that the same finding holds. Under a calibration of $w_c = 1$ and $w_p = 200$, however, the result holds when two and three core nodes default but not when one core node defaults. Note that this calibration assigns the lowest possible welfare contribution to core nodes. The default of one core node may not trigger more immediate intervention due to the very low welfare contribution. When more than one core node defaults, however, the systemic ramifications on the network prevail and it becomes optimal to step in more immediately despite the low welfare contribution.

Table 6: Robustness: Systemic Importance and Likelihood of Immediate Intervention
Number of States (in % of states with immediate intervention)

Welfare: C=15, P=20

# of Core Nodes in Default	# of Periphery Nodes in Default: 3	2	1	0	
3	0 (0.0%)	6 (0.0%)	168 (0.9%)	906 (4.7%)	**1,080 (5.6%)**
2	9 (0.0%)	261 (1.3%)	2,187 (11.3%)	4,229 (21.9%)	**6,686 (34.6%)**
1	63 (0.3%)	1,107 (5.7%)	4,653 (24.1%)	3,133 (16.2%)	**8,956 (46.3%)**
0	117 (0.6%)	975 (5.0%)	1,526 (7.9%)	0 (0.0%)	
	189 (1.0%)	**2,349 (12.1%)**	**8,534 (44.1%)**		

Welfare: C=1, P=200

# of Core Nodes in Default	# of Periphery Nodes in Default: 3	2	1	0	
3	0 (0.0%)	6 (0.0%)	168 (1.2%)	173 (1.2%)	**347 (2.4%)**
2	9 (0.1%)	261 (1.8%)	2,187 (15.0%)	1,073 (7.4%)	**3,530 (24.3%)**
1	63 (0.4%)	1,107 (7.6%)	4,651 (32.0%)	1,124 (7.7%)	**6,945 (47.8%)**
0	117 (0.8%)	975 (6.7%)	2,624 (18.0%)	0 (0.0%)	
	189 (1.3%)	**2,349 (16.2%)**	**9,630 (66.2%)**		

3.4.4 There is a Welfare Lower Bound Above Which Intervention is Optimal

Table 7 shows how optimal timing varies with welfare. We can observe that above some threshold of welfare, intervention becomes optimal in all cascade states. When $w = 0$, intervention is suboptimal in all cascade states. When $w = 1$, intervention is suboptimal in $16,131$ (70.3%) of all cascade states. This drops to only 4.7% of all cascade states when $w = 2$ and to 0% thereafter. When $w \geq 2$, all cascade states warrant intervention yet at varying points in time. This suggests that the welfare lower bound in this example is $w = 2$.

Table 7: Welfare and Optimal Intervention
Number of States (in % of all cascade states)

		Welfare				
		w=0	w=1	w=2	w=3	w=20
Intervention Decision	Immediate		6,712 (29.3%)	18,263 (79.6%)	18,835 (82.1%)	18,925 (82.5%)
	With Delay		104 (0.5%)	3,596 (15.7%)	4,112 (17.9%)	4,022 (17.5%)
	No Intervention	22,947 (100.0%)	16,131 (70.3%)	1,088 (4.7%)		

4 Conclusion

When is it optimal to bailout the financial system during a systemic crisis? This paper builds a model to study optimal timing for a systemic bailout in any financial network setting. The regulator decides the optimal period for a systemic bailout. There are two trade-off forces that influence optimal timing for intervention. The first trade-off is the maturity of liabilities against the resilience of nodes which captures how quickly distress can propagate through the financial system (i.e., the speed of financial contagion). The second trade-off is the welfare benefit against the shortfall cost of a systemic bailout.

Applying the model to a *Core-Periphery* network setting yields new insights into optimal timing for intervention. I show that the speed of financial contagion has critical bearing on the immediacy of intervention as it can identify the point beyond which immediate intervention becomes optimal. Moreover, I find that earlier intervention is optimal in states where core nodes default relative to states where periphery nodes default even when it is more expensive. This finding is robust to a calibration where core nodes contribute less welfare to the financial system and is consistent with the 'too-interconnected-to-fail' effect.

Policy intervention during a systemic crisis should incorporate optimal timing to be effective. Immediate intervention may not be appropriately designed due to the uncertainty surrounding the nature of the crisis which may trigger panic. Late intervention, on the other hand, will incur substantial economic costs and ultimately prove ineffective in containing financial contagion. The model illustrates the trade-off forces that influence optimal timing for intervention. Findings introduce a sense of urgency in resolving problems faced by systemically important nodes. They, further, validate the current macroprudential regulation approach which calls for greater scrutiny over the most interconnected players.

There are several possible extensions to the model. One extension would be to allow the set of possible linkages to change over the horizon of the model. This would provide a more accurate representation of reality where banks attempt to limit their exposure to a shock during a cascade by changing or diversifying their counterparties period-to-period. Other extensions may include allowing for targeted bailouts where the regulator can optimally choose which nodes to bailout in any given period and introducing a dynamic payment clearing mechanism.

References

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015a). Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608.

Acemoglu, D., Ozdaglar, A. E., and Tahbaz-Salehi, A. (2015b). Systemic risk in endogenous financial networks. *SSRN Electronic Journal*.

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

Cordella, T. and Yeyati, E. L. (2003). Bank bailouts: moral hazard vs. value effect. *Journal of Financial Intermediation*, 12(4):300–330.

Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review*, 104(10):3115–3153.

Erol, S. and Vohra, R. (2014). Network formation and systemic risk. *SSRN Electronic Journal*.

Freixas, X. and Rochet, J.-C. (2013). Taming systemically important financial institutions. *Journal of Money, Credit and Banking*, 45:37–58.

Gorton, G. and Huang, L. (2004). Liquidity, efficiency, and bank bailouts. *American Economic Review*, 94(3):455–483.

Laeven, L. and Valencia, F. (2008). Systemic banking crises: A new database. *IMF Working Papers*, 08(224):1.

Laeven, L. and Valencia, F. (2010). Resolution of banking crises: The good, the bad, and the ugly. *IMF Working Papers*, 10(146):1.

Laeven, L. and Valencia, F. (2012). Systemic banking crises database: An update. *IMF Working Papers*, 12(163):33.

Laeven, L. and Valencia, F. (2013). Systemic banking crises database. *IMF Economic Review*, 61(2):225–270.