Regulating Financial Networks Under Uncertainty

CARLOS RAMÍREZ*

This version: October 21, 2019.

*Board of Governors of the Federal Reserve System. I thank Co-Pierre George, Nathan Foley-Fisher, Daniel Opolot for detailed readings of early versions of the paper. I also thank Celso Brunetti, Julio Deride, Michael Gofman (discussant), Sebastian Infante, Zafer Kanik, José Lopez (discussant), Shawn Mankad, Borghan Narajabad, David Rappoport, João Santos (discussant), Stéphane Verani, Vladimir Yankov, Ke Wang, and participants at the 2019 ESEM-EEA, the 2019 EFA, the Bank of Portugal, the Fifth Annual Conference on Network Science and Economics, the 2018 LACEA-LAMES, WU, the Federal Reserve Board, the 6th Annual CIRANO-Walton Workshop on Networks, and the 2018 Federal Reserve Research Scrum for their suggestions. All remaining errors are my own. This article represents the views of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or other members of its staff. E-mail: carlos.ramirez@frb.gov.

Regulating Financial Networks Under Uncertainty

ABSTRACT

I study the problem of regulating a network of interdependent financial institutions that is prone to contagion when there is uncertainty regarding its precise structure. I show that such uncertainty reduces the scope for welfare-improving interventions. While improving network transparency potentially reduces this uncertainty, it does not always lead to welfare improvements. Under certain conditions, regulation that reduces the risk-taking incentives of a small set of institutions can improve welfare. The size and composition of such a set crucially depend on the interplay between (i) the (expected) susceptibility of the network to contagion, (ii) the cost of improving network transparency, (iii) the cost of regulating institutions, and (iv) investors' preferences.

Keywords: Financial networks, contagion, policy design under uncertainty.
JEL classification: C6, E61, G01.

The financial crisis that began in 2007 underscored the relevance of interdependencies among financial institutions—e.g., commercial banks, money market funds, investment banks, and insurance companies—in the functioning of modern economies. While, in normal times, these interdependencies—in the form of contractual obligations or common exposures—can be beneficial, as they help institutions manage liquidity or diversity risk, they can also create channels through which shocks propagate in times of economic stress. These channels might cause problems at one institution to spread to others, potentially leading to cascades of distress with economy-wide implications.

In light of the potential harmful side effects of these interdependencies, policymakers across the globe implemented responses that directly or indirectly take into account the interconnected nature of modern financial systems so as to preserve the benefits of interdependencies while managing their unintended negative consequences. When designing these responses, however, policymakers are confronted with an inconvenient truth: it is hard to determine the precise structure of the network of exposures among financial institutions because of the opacity, complexity, and multifaceted nature of their linkages. Importantly, this problem becomes particularly acute in times of economic stress, as spirals of fire sales may become relevant. A natural question then arises: How can policymakers regulate a network of interdependent financial institutions when those policymakers are fundamentally uncertain about its precise structure? Despite its importance, this question has been overlooked by most of the literature. This paper partially fills this gap by developing a model to study the behavior of such policymakers.

The main results are as follows. First, I show that uncertainty about the precise architecture of the network can reduce the scope for welfare-improving interventions, as such uncertainty gives rise to difficulties in determining the likelihood of systemic events. However, this lack of certainty does not necessarily justify a non-interventionist policy, considering the negative consequences of cascades of distress. Second, while improving network transparency could help policymakers overcome forecasting limitations, it does not necessarily lead to welfare improving interventions as improving network transparency is also costly. Because confidentiality might be valuable to institutions, improving network transparency might compromise their market position and potentially decrease efficiency. Third, under certain conditions, regulation that reduces the risk-taking incentives of a set of institutions can increase welfare. Importantly, the size and composition of that set is determined by the interplay between (i) the (expected) susceptibility of the network to contagion, (ii) the cost of improving network transparency, (iii) the cost of regulating institutions, and (iv) investors' preferences.

The model is motivated by an economy in which financial institutions (banks, for short) are interconnected through an exogenous network of opaque exposures, on either the asset side or the liability side, that cannot be mitigated through contractual protections. In times of economic stress, some of these exposures (henceforth referred to as contagious exposures) function as propagation mechanisms, as banks become more vulnerable to distress affecting related banks (henceforth referred to as neighbors). Cascades of distress may occur as a result of contagion, as the distress affecting one bank could cause distress to that bank's neighbors, which, in turn, may cause distress to the neighbors' neighbors, and so on.

To capture policymakers' inability to ascertain the precise architecture of the network in times of economic stress, I assume that the set of contagious exposures is unknown when designing interventions. Because banks fail to internalize the consequences of their actions on the spread of distress, introducing regulation potentially leads to a Pareto improvement. A planner seeks to maximize welfare by imposing preemptive liquidity restrictions on a set of banks. While liquidity restrictions decrease banks' likelihood of distress, they are not costless as they limit banks' ability to allocate funds toward more productive investment opportunities, thereby introducing resource misallocation. Although the planner is uncertain which exposures may propagate distress, she can improve network transparency at a cost. By improving network transparency, the planner can strategically target banks with the highest number of contagious exposures first to limit the spread of distress more effectively, avoiding losses associated with regulating an excessively large number of banks.

I first analyze the behavior of the planner when the distribution of contagious exposures across banks is known. I show that the optimal policy—which is jointly determined by a choice of network transparency and a selection of restricted banks—is shaped by the interplay between the distribution of contagious exposures, the cost of improving network transparency, and the costs of restricting banks. If the network of contagious exposures exhibits a highly asymmetric structure (that is, there is high variation in the number of contagious exposures across banks), then a handful of banks play an important role in the propagation of distress. Learning the identity of those banks becomes critical to adequately avoid contagion, as regulating them effectively deters the emergence of cascades of distress. As a result, improving network transparency tends to be optimal. However, if the network of contagious exposures exhibits a highly symmetric structure (that is, there is small variation in the number of contagious exposures across banks), every bank is likely to play a similar role in the propagation of distress when conditions deteriorate. Thus, improving network transparency tends not to be optimal as more information regarding the precise structure of the network does not necessarily allow the planner to uncover the most contagious banks. Finally, higher costs of restricting banks lead to a smaller fraction of banks that can be restricted.

Next, I analyze the behavior of the planner when the distribution of contagious exposures is unknown. In this case, the planner faces model uncertainty as she unsure about the model that describes how the economy behaves in times of economic stress. The optimal intervention is then affected by investors' attitudes toward ambiguity and their beliefs regarding the susceptibility of the network to contagion. Under certain conditions, small changes in beliefs generate significant changes in the optimal set of restricted banks. When investors are sufficiently ambiguity averse, they worry the number of restricted banks might not be sufficiently large to prevent cascades of distress. As a result, more banks might need to be restricted as network uncertainty increases. Importantly, the value of improving network transparency now also depends on the extent of network uncertainty. In many cases, improving network transparency tends to be more valuable as network uncertainty increases. This is because more transparency allows the planner to hedge the risk of implementing non-optimal policies ex-ante as a result of not knowing the distribution of contagious exposures.

The first set of results informs the ongoing debate regarding the optimal design of preemptive macroprudential regulations. While post-crisis reforms with a macroprudential dimension have focused principally on large financial institutions, my results underscore that the architecture of the financial system (and not just the size of institutions) matters for policy design. In addition, these results provide a rationale for regulation that seeks to improve network transparency and, in particular, improve information disclosure, as institution-level information may be critical to effectively limit the impact of cascades of distress. More broadly, these results highlight the importance of developing privacy-preserving methods for sharing financial exposures. The second set of results highlights that an appropriate macroprudential regulatory framework must be mindful of the uncertainty regarding the pattern of interdependencies among institutions.

Related literature. This paper contributes to two strands of the literature. First, this paper adds to a body of work that explores how network features of the financial system affect the likelihood of contagion. An incomplete list includes Rochet and Tirole (1996), Allen and Gale (2000), Freixas et al. (2000), Eisenberg and Noe (2001), Lagunoff and Schreft (2001), Dasgupta (2004), Leitner (2005), Nier et al. (2007), Allen and Babus (2009), Haldane and May (2011), Allen et al. (2012), Amini et al. (2013), Cont et al. (2013), Georg (2013), Zawadowski (2013), Cabrales et al. (2014), Elliott et al. (2014), Glasserman and Young (2015, 2016), Acemoglu et al. (2015), and Castiglionesi et al. (2019). Unlike these papers, my paper explicitly focuses on the planner's problem in the presence of spillovers and uncertainty regarding the pattern of linkages among institutions. Second, my paper adds to recent research that explores how policy interventions affect the mechanism through which shocks propagate (see, for example, Beale et al. (2011), Gai et al. (2011), Battiston et al.

(2012), Goyal and Vigier (2014), Alvarez and Barlevy (2015), Adrian et al. (2015), Aldasoro et al. (2017), Erol and Ordoñez (2017), Gofman (2017), and Galeotti et al. (2018)). While my paper also focuses on how contagion varies with different interventions, it provides a tractable framework in which optimal policies can be determined under uncertainty regarding the economy's connectivity structure.

The rest of the paper is organized as follows. The next section provides a motivating example which illustrates the key ideas of the paper. Section II introduces a more flexible framework to see the extent to which these ideas can be generalized. Section III explores how regulation affects expected total output. Section IV characterizes optimal interventions. Section V concludes. All derivations appear in the Appendix.

I. Motivating Example

Consider a two-period economy with three banks. Time is indexed by $t \in \{t_O, t_F\}$, with $t_O < t_F$. There is a representative investor who owns all assets in the economy. Although she is risk-neutral, she cares about model uncertainty as her preferences that can be characterized by the smooth ambiguity model of Klibanoff et al. (2005).[1] While it is commonly acknowledged that banks' payoffs are linked through the network of exposures depicted in figure 1(a), the precise position of banks in 1(a) is unknown.

At $t = t_O$, a planner imposes preemptive liquidity restrictions on a set of banks so as to maximize the representative investor's smooth ambiguity certainty equivalent. Liquidity restrictions simply force restricted banks to hold more liquid portfolios. As it becomes clear in section II, this intervention possibly leads to a Pareto improvement as banks take more risk than is socially optimal. At $t = t_F$, economic conditions deteriorate, payoffs are realized,

[1]In a broad sense, these preferences capture circumstances in which investors are uncertain about the "true model" that determines the behavior of the economy. Given the uncertainty about the model, investors may exhibit aversion to (or preference for) that uncertainty. For example, if investors are averse to such uncertainty, they worry about making non-optimal decisions ex ante because they do not know the "true model." Importantly, with these preferences, investors' tastes for risk and model uncertainty can be separated in a simple form. For more details, see Klibanoff et al. (2005) and Maccheroni et al. (2013).

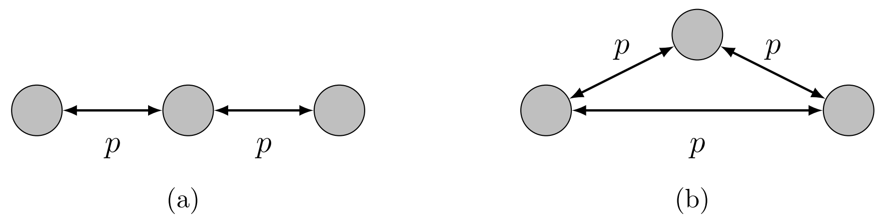

Figure 1. Two network architectures among three banks.

and consumption occurs. When economic conditions deteriorate, cascades of liquidity shocks might occur as a result of contagion. In particular, the following events occur before payoffs are realized:

- One bank (chosen uniformly at random) faces an adverse liquidity shock.
- If that bank is affected, the shock can propagate to others via randomly selected exposures. Each exposure is contagious (independently of others) with probability $0 < p < 1$. Bank i faces a liquidity shock if (1) there is a sequence of contagious exposures between i and the first bank that faces the liquidity shock, and (2) every bank within that sequence is affected by the liquidity shock.

At a basic level, cascades of liquidity shocks can be broadly interpreted as liquidity-driven crises in which liquidity shocks affecting certain banks induce liquidity shocks for some of their neighbors (as in Diamond and Rajan (2011), Caballero and Simsek (2013), and Stein (2013)). In times of stress, those neighbors may face a run due to solvency concerns, which, in turn, potentially causes solvency concerns about some of the neighbors' neighbors, possibly generating cascades of runs. Consequently, cascades of liquidity shocks could also be interpreted as crises of confidence (as in Zhou (2018)).

The random selection of contagious exposures serves as a metaphor for market participants and regulators having difficulty assessing how exposures react in times of economic stress. This difficulty makes the planner fundamentally uncertain which banks are more prone to propagate shocks when economic conditions deteriorate, motivating the assumption that the exact position of banks in the network is unknown.[2]

A bank's output is set at zero if such a bank is affected by a liquidity shock at $t = t_F$; otherwise, its output is set at one dollar. As regulation forces banks to hold more liquid portfolios, I assume restricted banks absorb, rather than amplify, liquidity shocks. This assumption has two implications. First, restricted banks become resilient to liquidity shocks. Second, restricted banks do not propagate shocks within any sequence of contagious exposures as they become resilient to shocks, thereby decreasing the likelihood of distress of their (direct and indirect) neighbors. As a result, regulation reshapes the way that shocks spread. This is because, from the perspective of shock propagation, imposing restrictions on a bank can be represented by the removal of such a bank and its exposures from any realized network of contagious exposures.

While restrictions generate positive externalities, they are not costless as they prevent banks from investing in illiquid (and possibly more profitable) assets, potentially deterring bank lending and/or introducing resource misallocation. To capture this idea in a simple form, I assume that imposing restrictions on any given bank entails paying a cost of c dollars, with $0 < c < 1$.

I now focus on two cases. First, in section I.A, I characterize the optimal intervention when p is known. Second, in section I.B, I characterize the optimal intervention when p is unknown. Importantly, when p is unknown, there is another layer of uncertainty as the planner is now also unsure about the model that describes how the economy behaves in times of stress. This type of uncertainty seeks to capture situations wherein regulators are also uncertain about the process determining how shocks spread when a crisis manifests.

[2] Of course, this random selection process—which is similar to the one used in Ramírez (2017)—provides a crude approximation of how liquidity shocks propagate in times of stress. Yet, it allows me to provide a tractable analysis of cascades of shocks within a more general framework (see section II). The main results continue to hold if a small set of banks is initially affected by liquidity shocks. Conditional on banks i, j, and k being connected via contagious exposures, the existence of a contagious exposure between i and j is independent of the existence of a contagious exposure between j and k. A richer model would include local dependencies among such events so that the effect that a single distressed neighbor has on a bank depends critically on whether other neighbors face liquidity shocks. For a model that introduces such dependencies, see Watts (2002). If one introduces such dependencies, the basic trade-off behind the main results should continue to appear.

A. p is known

In this case, there is no model uncertainty as the shock propagation mechanism at $t = t_F$ is commonly known. Consequently, the representative investor's smooth ambiguity certain equivalent equals expected total output, $\mathbb{E}_p[\text{TO}]$, as she is risk neutral. Thus, regulation is selected so as to maximize $\mathbb{E}_p[\text{TO}]$.

How many banks should be regulated? In the absence of any further information, the planner is unable to determine the position of banks in the network. Hence, the planner imposes restrictions at random. Let x denote the number of restricted banks. For any given p, expected total output is

$$\begin{aligned} \mathbb{E}_p\left[\text{TO}|x=0\right] = \frac{2}{3}(1-p)(3+p) \quad &\text{and} \quad \mathbb{E}_p\left[\text{TO}|x=1\right] = (1-c) + \frac{4}{9}(3-p), \\ \mathbb{E}_p\left[\text{TO}|x=2\right] = 2(1-c) + \frac{2}{3} \quad &\text{and} \quad \mathbb{E}_p\left[\text{TO}|x=3\right] = 3(1-c). \end{aligned} \tag{1}$$

Consequently, the optimal number of restricted banks, $x^*_p(c)$, is

$$x^*_p(c) \quad = \quad \begin{cases} 3, \text{ if} & c \leq \frac{1}{3} \\ 2, \text{ if} & \frac{1}{3} < c \leq \frac{1}{3} + \frac{4}{9}p \\ 1, \text{ if} & \frac{1}{3} + \frac{4}{9}p < c \leq \frac{1}{3} + \frac{8}{9}p + \frac{2}{3}p^2 \\ 0, \text{ if} & \frac{1}{3} + \frac{8}{9}p + \frac{2}{3}p^2 < c. \end{cases} \tag{2}$$

That is, when c is sufficiently small, it is optimal to regulate as many banks a possible as the marginal cost of regulation is negligible. As c increases, the marginal cost of regulation becomes material, and, thus, it is optimal to regulate fewer banks. Importantly, the optimal intervention, $x^*_p(c)$, hinges on the interplay between the marginal cost of regulation, c, and the susceptibility of the economy to contagion, captured by p.

Improving network transparency. Now suppose that before implementing restrictions, the planner could learn the identity of the bank in the middle of figure 1(a) by paying a cost of κ dollars, with $0 < \kappa \leq \frac{4}{3}$. Parameter κ can be broadly interpreted as the cost of designing and implementing policies to improve information disclosure and transparency regarding the network of contagious exposures.[3] This cost could arise in economies wherein more transparency might decrease banks' confidentiality. As confidentiality is valuable to banks, improving transparency could compromise banks' market position, reducing their incentives to lend and potentially decreasing market efficiency.[4]

Why would the planner pay κ? If she decides not to pay, then banks are ex-ante identical from her point of view. By paying κ, however, the planner can improve the effectiveness of her intervention by targeting the bank in the middle of figure 1(a) first. Naturally, the planner's decision will depend on how much this information helps her to mitigate contagion more effectively. It directly follows from the way shocks propagate

$$\mathbb{E}_p\left[\text{TO}|\text{pay } \kappa \text{ and restrict the bank in the middle of figure } 1(a)\right] = (1-c) + \frac{4}{3} - \kappa. \quad (3)$$

By comparing equations (3) and (1), I determine the set of pairs (κ, c) for which paying κ and then targeting the bank in the middle generates higher expected output than regulating banks at random,

$$T_p \equiv \left\{(\kappa, c)\middle| \left(\frac{1}{3} \leq (c-\kappa) \text{ and } \frac{1}{3} \leq c \leq \frac{1}{3} + \frac{4}{9}p\right) \text{ or } \left(\kappa \leq \frac{4}{9}p \text{ and } \frac{1}{3} + \frac{4}{9}p \leq c \leq \frac{1}{3} + \frac{8}{9}p + \frac{2}{3}p^2\right) \cdots \right.$$
$$\left. \text{or } \left(\kappa + c \leq \frac{1}{3} + \frac{12}{9}p + \frac{2}{3}p^2 \text{ and } \frac{1}{3} + \frac{8}{9}p + \frac{2}{3}p^2 \leq c\right)\right\}.$$

The definition of T_p underscores that the decision to improve network transparency depends

[3] These policies may allow regulators to uncover banks that play an important role in the transmission of shocks when a crisis materializes. Two important examples of such policies are the Comprehensive Liquidity Assessment and Review (CLAR) and the Dodd-Frank Act supervisory stress test, run annually by the Federal Reserve. In these programs, regulators evaluate the liquidity risk profile of bank holding companies (BHCs) through a range of metrics and project whether BHCs would be vulnerable during times of weak economic conditions. Other examples include programs implemented by the SEC such as forms N-MFP and PF. Form N-MFP requires registered money market funds to report their portfolio holdings and other information on a monthly basis, while form PF requires private funds to report assets under management.

[4] Thinking of banks as secret keepers is also consistent with this idea; see Dang et al. (2017) for more details.

on the interaction between (i) the marginal cost of regulation, c, (ii) the cost of improving network transparency, κ, and (iii) the susceptibility of the economy to contagion, p. Figure 2 illustrates this result by depicting T_p for different values of p.

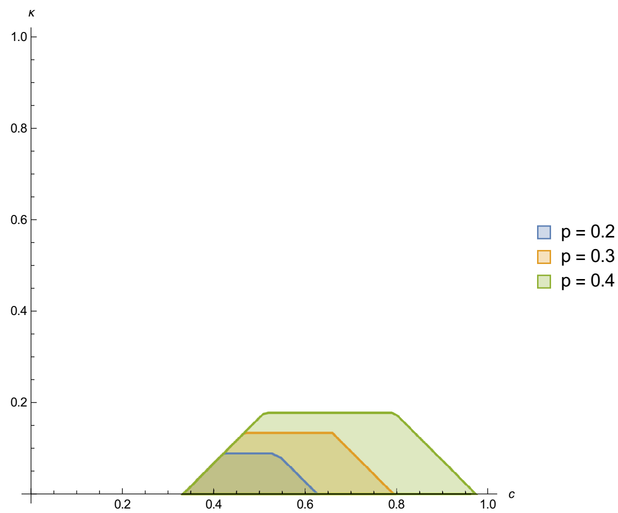

Figure 2. T_p when $p \in \left\{\frac{2}{10}, \frac{3}{10}, \frac{4}{10}\right\}$.

What would happen if the network architecture was different? Now suppose that it is commonly acknowledged that banks are linked as in figure 1(b). For now, assume that the position of banks in the network is unknown. For any given p, expected total output is

$$\begin{aligned} \mathbb{E}_p\left[\mathrm{TO}|x=0\right] &= 2(1-p)^2(1+p) \quad \text{and} \quad \mathbb{E}_p\left[\mathrm{TO}|x=1\right] = (1-c)+\frac{2}{3}(2-p), \\ \mathbb{E}_p\left[\mathrm{TO}|x=2\right] &= 2(1-c)+\frac{2}{3} \quad \text{and} \quad \mathbb{E}_p\left[\mathrm{TO}|x=3\right] = 3(1-c). \end{aligned} \tag{4}$$

Consequently, the optimal number of regulated banks, $\widehat{x_p}(c)$, is given by

$$\widehat{x_p}(c) = \begin{cases} 3, \text{ if} & c \leq \frac{1}{3} \\ 2, \text{ if} & \frac{1}{3} < c \leq \frac{1}{3}+\frac{2}{3}p \\ 1, \text{ if} & \frac{1}{3}+\frac{2}{3}p < c \leq \frac{1}{3}+\frac{4}{3}p+2p^2(1-p) \\ 0, \text{ if} & \frac{1}{3}+\frac{4}{3}p+2p^2(1-p) < c. \end{cases} \tag{5}$$

The comparison between expressions (2) and (5) uncovers two results. First, when c is sufficiently small, it is optimal to regulate as many banks a possible regardless of the underlying network architecture. Second, as c increases, regulating fewer banks becomes optimal. But, how many banks should be regulated? The comparison between x_p^* and $\widehat{x}_p$ highlights that the nature of the network architecture also has first order effects for the design of interventions. This is because the susceptibility of the economy to contagion depends on the interplay between p and the underlying network architecture. For example, an increase in p increases the susceptibility of the economy to contagion relatively more if the network is captured by figure 1(b) rather than by figure 1(a).

Should network transparency be improved? If banks are linked as depicted in figure 1(b), all banks are ex-ante identical from the perspective of shock propagation. Therefore, the effectiveness of interventions cannot be increased by learning the position of banks in the network. Hence, it is never optimal to pay κ. Consequently, the decision to improve network transparency also depends on the underlying network architecture.

B. p is unknown (robust interventions)

Now suppose p is unknown. Assume that it is commonly acknowledged that p can take two values $\{p_L, p_H\}$, with $p_L < p_H$, and $p = p_L$ with probability ϕ, with $0 < \phi < 1$. To highlight the importance of model uncertainty, assume banks' payoffs are linked as in figure 1(a). To facilitate exposition, hereinafter assume $p_L = \frac{1}{5}$, $p_H = \frac{4}{5}$, and $c = \frac{5}{9}$.

How many banks should be regulated? Suppose the planner is unable to determine the position of banks in the network and, thus, she restricts banks at random. Let x_A denote the number of restricted banks. The representative investor's smooth ambiguity certain equivalent, $\mathrm{SCE}(x_A)$, is then

$$\mathrm{SCE}(x_A = 0) \equiv \mathbb{E}_{\bar{p}}[\mathrm{TO}|x_A = 0] - \left(\frac{\theta}{2}\right)\mathbb{V}\left[\mathbb{E}_p(\mathrm{TO}|x_A = 0)\right]$$

$$
\begin{aligned}
&= \frac{2}{3}(1-\bar{p})(3+\bar{p}) - \left(\frac{\theta}{2}\right)\left(\frac{6}{25}\right)^2 \phi(1-\phi)(25+\phi(1-\phi)), \\
\text{SCE}(x_A = 1) &\equiv \mathbb{E}_{\bar{p}}\left[\text{TO}|x_A = 1\right] - \left(\frac{\theta}{2}\right)\mathbb{V}\left[\mathbb{E}_p\left(\text{TO}|x_A = 1\right)\right] \\
&= \left(1-\frac{5}{9}\right) + \frac{4}{9}(3-\bar{p}) - \left(\frac{\theta}{2}\right)\left(\frac{4}{15}\right)^2 \phi(1-\phi), \\
\text{SCE}(x_A = 2) &\equiv \mathbb{E}_{\bar{p}}\left[\text{TO}|x_A = 2\right] - \left(\frac{\theta}{2}\right)\mathbb{V}\left[\mathbb{E}_p\left(\text{TO}|x_A = 2\right)\right] \\
&= 2\left(1-\frac{5}{9}\right) + \frac{2}{3}, \\
\text{SCE}(x_A = 3) &\equiv \mathbb{E}_{\bar{p}}\left[\text{TO}|x_A = 3\right] - \left(\frac{\theta}{2}\right)\mathbb{V}\left[\mathbb{E}_p\left(\text{TO}|x_A = 3\right)\right] \\
&= 3\left(1-\frac{5}{9}\right),
\end{aligned}
$$

where $\bar{p} = \phi p_L + (1-\phi)p_H = \frac{4-3\phi}{5}$, θ is a non-negative coefficient capturing the representative investor's attitude toward model uncertainty, and $\mathbb{V}[\mathbb{E}_p(\cdot|x_A)] \equiv \sum_{p\in\{p_L,p_H\}}(\mathbb{E}_p(\cdot|x_A) - \mathbb{E}_{\bar{p}}(\cdot|x_A))^2\mathbb{P}(p)$.

To illustrate how the optimal intervention varies with model uncertainty, figure 3(a) depicts the number of restricted banks selected so as to maximize SCE, $x_A^*(\phi,\theta)$. Suppose θ is held fixed. As ϕ increases, the economy is less susceptible to contagion with higher probability. Thus, as regulating banks is costly, $x_A^*(\phi,\theta)$ is a weakly decreasing function of ϕ. Now suppose ϕ is held fixed. Given the values for p_L, p_H, and c, it is always optimal to regulate two banks when $\phi \leq \frac{1}{2}$. Thus, changes in θ do not alter the optimal intervention. A similar idea applies when $0.85 \approx \underline{\phi} < \phi < \bar{\phi} \approx 0.97$, as it is optimal to regulate only one bank. However, when $\frac{1}{2} < \phi < \underline{\phi}$ or $\bar{\phi} \leq \phi$, changes in θ could alter the optimal intervention. As θ increases, the representative investor exhibits more aversion to model uncertainty, and, thus, the planner acts as if she dislikes making non-optimal decisions ex-ante more because she is unsure about the precise value of p. Consequently, it becomes optimal to restrict two banks rather than one if $\frac{1}{2} < \phi < \underline{\phi}$, for sufficiently large values of θ—or to restrict one bank rather than no bank if $\bar{\phi} \leq \phi$. Thus, $x_A^*(\phi,\theta)$ is a weakly increasing function of θ.

Improving network transparency. Now suppose the planner could learn the identity

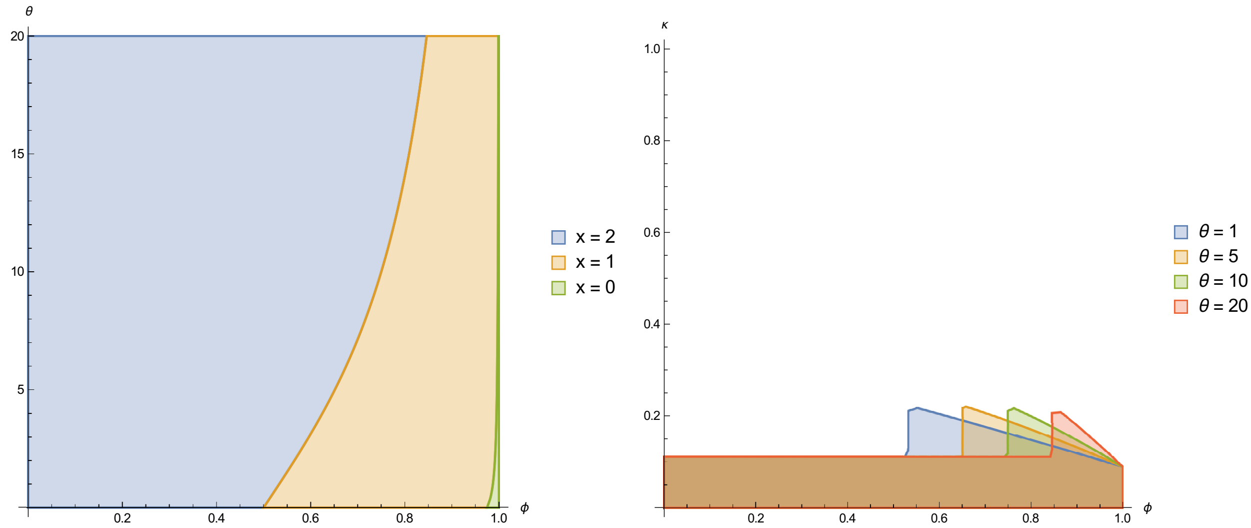

(a) $x_A^*(\phi, \theta)$. Random targeting.

(b) (ϕ, κ) for which improving network transparency is optimal.

Figure 3. Robust optimal interventions.

of the bank in the middle by paying κ. Under what conditions would the planner pay κ? Figure 3(b) answers this question by depicting the set of pairs (ϕ, κ) for which it is optimal to learn this information so as to regulate the bank in the middle first. Importantly, such a set is reshaped by θ. Intuitively, more transparency allows the planner to hedge the risk of making non-optimal decisions ex-ante as a result of not knowing p. As θ increases, the aversion for such risk increases, which, in turn, makes network transparency more valuable.

C. Summary of findings and challenges.

Without model uncertainty, the optimal intervention is characterized by the interplay between three characteristics of the economy: (1) its susceptibility to contagion, (2) the marginal cost of regulation, and (3) the cost of improving network transparency. When model uncertainty is incorporated, beliefs regarding the nature of the network architecture reshape this interplay as they alter the expected susceptibility of the economy to contagion. Additionally, the optimal intervention is affected by investors' attitude toward model uncertainty.

While the motivating example characterizes the optimal intervention, it is hard to know what to make of results which rely on an economy with three banks and a specific network architecture. In particular, it is not clear what happens if the size of the economy grows large and how different network architectures can be incorporated into the analysis. The next sections show that these findings continue to be valid for economies with arbitrary sizes and network architectures as long as contagious exposures are randomly determined.

II. General Framework

In this section, I pose the problem in a more flexible framework in a way that parallels the motivating example.

Environment. Take the motivating example and extend it along five dimensions. First, consider n banks, each endowed with one dollar; with n being potentially large. Write B_n for the set that contains all banks and $\mathcal{P}(B_n)$ for the power set of B_n. Second, suppose the network of exposures has an arbitrary architecture. While banks may differ in their number of exposures, they are ex ante identical in other respects, such as size and leverage.

Third, instead of assuming that exposures are contagious with probability p, assume that the resulting distribution of contagious exposures across banks can be characterized by a distribution $\{p_k\}_{k=0}^{n-1}$, where p_k denotes the probability that a randomly chosen bank has k contagious exposures at $t = t_F$. The fact that contagious exposures are randomly determined is critical for the analysis that follows. Importantly, as the network architecture is arbitrary, p_k is allowed to take any functional form.

Fourth, assume that if the planner pays κ, she is able to rank banks based on their future number of contagious exposures. Then, she restricts banks using that ranking; restricting banks with the highest number of contagious exposures first so as to prevent contagion more effectively. While, in reality, the planner might collect different pieces of information, and, thus, decide between intermediate levels of network transparency, this assumption keeps computations tractable.

Fifth, generalize payoffs and describe banks' optimization problem. To do so, introduce an interim period t_I, with $t_O < t_I < t_F$, where banks are allowed to react to regulation. In particular, for each dollar invested at $t = t_I$, suppose the (random) payoff of bank i at $t = t_F$, π_i, equals

$$\pi_i(\omega_i) \quad = \quad \omega_i \times R_L + (1 - \omega_i) \times R_I - \beta_{\omega_i} \times \varepsilon_i, \tag{6}$$

where ω_i denotes the fraction of bank i's portfolio invested in liquid assets; R_L and R_I denote the (random) payoff of liquid and illiquid assets, respectively. Liquid assets are (exogenously determined) investment opportunities that can be easily converted into cash. Illiquid assets can be broadly interpreted as (long-term) projects seeking financing. Random variable ε_i equals 1 if bank i faces an adverse liquidity shock; otherwise, $\varepsilon_i = 0$. The term β_{ω_i} captures the effect of liquidity shocks on bank i's payoff.

For simplicity, I assume ω_i can take two values, ω_L or ω_H, with $0 \leq \omega_L < \omega_H \leq 1$. Additionally, illiquid assets are assumed to yield a higher expected payoff than liquid assets; hence, $\mathbb{E}[R_L] < \mathbb{E}[R_I]$. Finally, I assume that the liquidity of bank i's portfolio, ω_i, alters the effect of ε_i on bank i's payoff. In particular,

$$\beta_{\omega_i} \quad = \quad \begin{cases} 0, \text{ if } \omega_i = \omega_H \\ \omega_L \times R_L + (1 - \omega_L) \times R_I, \text{ otherwise.} \end{cases} \tag{7}$$

Equations (6) and (7) impose parallelism between payoffs within the general framework and the motivating example. Banks with more liquid portfolios are not affected by liquidity shocks when economic conditions deteriorate, while banks with more illiquid portfolios fail when facing a liquidity shock. Consequently, when choosing ω_i, bank i faces the following trade-off: the more liquid its portfolio, the higher its resilience to liquidity shocks, but potentially the lower its future payoff.

As in the motivating example, restrictions take a simple form: the planner forces restricted banks to hold more liquid portfolios. As a result, the optimal portfolio allocation of bank i, ω_i^*, solves

$$\begin{aligned} \max_{\omega_i \in \{\omega_L, \omega_H\}} \quad & \mathbb{E}_i \left[\pi_i(\omega_i) \right] \\ \text{s.t.} \quad & \omega_H \times e_i \leq \omega_i \text{ (regulatory constraint)}, \end{aligned} \tag{8}$$

where e_i equals 1 if bank i is restricted at $t = t_O$ and 0 otherwise; operator $\mathbb{E}_i$ emphasizes that bank i chooses ω_i^* based on its available information and subjective beliefs. Problem (8) underscores the idea that restricted banks optimize within the confines of regulatory constraints.

Additional Assumptions. For tractability, I make two further assumptions. First, I assume banks underestimate the likelihood of being affected by cascades of liquidity shocks (in a sense properly specified in Appendix A.A). Therefore, from banks' perspective, investing in illiquid assets is more lucrative than storing funds in liquid assets.[5] Consequently, bank i chooses $\omega_i^* = \omega_L$, unless the planner imposes restrictions on i. Second, I assume there is at least one bank that if restricted would cause the representative investor's smooth certainty equivalent to increase.

The first assumption ensures there is space for regulation as the market equilibrium is not efficient. The second assumption ensures that regulation potentially leads to a Pareto improvement. Subsection II.A provides more details.

Information Structure. As in the motivating example, in the absence of any further information, neither banks nor the planner can determine the position of banks within the network of contagious exposures.

Timeline. Figure 4 depicts the timeline of events within the general framework.

[5]This assumption is consistent with the "underestimated risks" factor highlighted by the IGM Forum (2017) as one of the most prominent factors contributing to the 2007–2009 financial crisis as well as banks' lack of appreciation of downside risks highlighted by Gennaioli and Shleifer (2018).

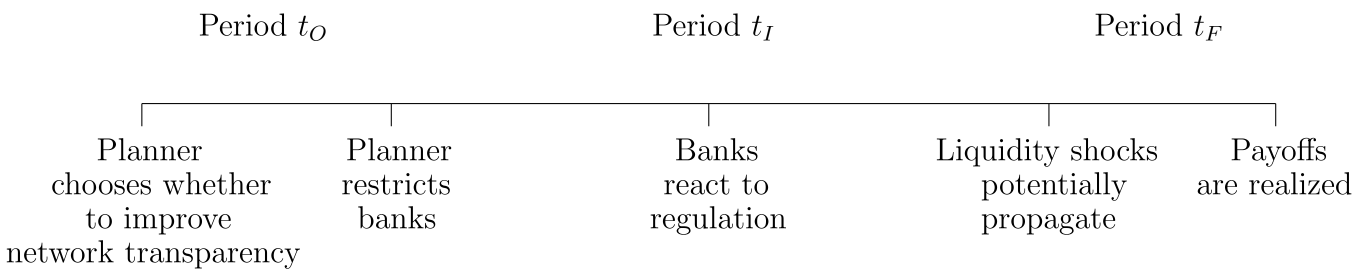

Figure 4. Model timeline

A. The need for regulation

To better understand the reason the market equilibrium is not efficient and interventions possibly lead to a Pareto improvement, this subsection compares expected total output under two cases. In the first, I compute expected total output under the market equilibrium. In the second, I compute expected total output if one bank is restricted. In what follows, $\{p_k\}_k$ is assumed to be known.

Market Equilibrium. Without regulation, every bank holds a fraction ω_L of its portfolio in liquid assets. As a result, the expected total output generated in the market equilibrium, $\mathbb{E}\left[TO_E\right]$, equals

$$\mathbb{E}\left[\mathrm{TO_E}\right] \;=\; \mathbb{E}\left[\sum_{i=1}^{n}\pi_i\right] = n\left(\omega_L\mathbb{E}[R_L] + (1-\omega_L)\mathbb{E}[R_I]\right) - \left(\sum_{i=1}^{n}\mathbb{E}_0[\beta_{\omega_i}\varepsilon_i]\right)$$

where $\mathbb{E}_0[\beta_{\omega_i}\varepsilon_i] = \mathbb{E}[\beta_{\omega_i}\varepsilon_i|\omega_i = \omega_{-i} = \omega_L]$.

Introducing regulation. To appreciate the potential benefits of regulation, suppose only one bank, say bank i, is forced to hold a fraction ω_H of its portfolio in liquid assets. Let $\mathbb{E}\left[TO_i\right]$ denote expected total output in this case. Then, $\mathbb{E}\left[TO_i\right]$ equals

$$\mathbb{E}\left[\mathrm{TO}_i\right] = \overbrace{\omega_H\mathbb{E}[R_L] + (1-\omega_H)\mathbb{E}[R_I] - \mathbb{E}'[\beta_{\omega_i}\varepsilon_i]}^{\text{output derived from bank } i\text{'s response to regulation}} \;+$$

$$\underbrace{(n-1)\left(\omega_L\mathbb{E}[R_L]+(1-\omega_L)\mathbb{E}[R_I]\right)-\left(\sum_{j\neq i}^{n}\mathbb{E}''[\beta_{\omega_j}\varepsilon_j]\right)}_{\text{output derived from other banks' actions}},$$

with $\mathbb{E}'[\beta_{\omega_i}\varepsilon_i]=\mathbb{E}[\beta_{\omega_i}\varepsilon_i|\omega_i=\omega_H \text{ and } \omega_{-i}=\omega_L]$ and $\mathbb{E}''[\beta_{\omega_j}\varepsilon_j]=\mathbb{E}[\beta_{\omega_j}\varepsilon_j|\omega_i=\omega_H \text{ and } \omega_{-i}=\omega_L]$. Thus, the difference $(\mathbb{E}\left[\mathrm{TO}_i\right]-\mathbb{E}\left[\mathrm{TO}_E\right])$ describes the welfare effects of imposing liquidity restrictions on bank i. This difference can be written as

$$\underbrace{(\mathbb{E}_0[\beta_{\omega_i}\varepsilon_i]-\mathbb{E}'[\beta_{\omega_i}\varepsilon_i])+\sum_{j\neq i}^{n}(\mathbb{E}_0[\beta_{\omega_j}\varepsilon_j]-\mathbb{E}''[\beta_{\omega_j}\varepsilon_j])}_{\text{benefit}} \quad - \quad \underbrace{\Delta\omega\mathbb{E}[\Delta R]}_{\text{cost}}.$$

That is, the benefit of increasing the liquidity of bank i's portfolio is composed of two terms. The first term captures the increase in bank i's resilience to liquidity shocks. The second term captures the increase in the resilience of bank i's neighbors (and the neighbors of those neighbors, and so on), as bank i no longer propagates shocks when conditions deteriorate. Importantly, bank i fails to internalize this term when choosing ω_i. The cost of increasing bank i's liquidity captures the decrease in the expected payoff of bank i, as illiquid assets yield a higher expected payoff than liquid assets.

One of the aforementioned assumptions ensures that there exists at least one bank, say bank l, that if restricted, the following inequality is satisfied

$$(\mathbb{E}_0[\beta_{\omega_l}\varepsilon_l]-\mathbb{E}'[\beta_{\omega_l}\varepsilon_l])+\sum_{j\neq l}^{n}(\mathbb{E}_0[\beta_{\omega_j}\varepsilon_j]-\mathbb{E}''[\beta_{\omega_j}\varepsilon_j]) \quad > \quad \Delta\omega\mathbb{E}[\Delta R].$$

Namely, the increase in resilience of bank l, its neighbors, and the neighbors of those neighbors, $(\mathbb{E}_0[\beta_{\omega_l}\varepsilon_l]-\mathbb{E}'[\beta_{\omega_l}\varepsilon_l])+\sum_{j\neq l}^{n}(\mathbb{E}_0[\beta_{\omega_j}\varepsilon_j]-\mathbb{E}''[\beta_{\omega_j}\varepsilon_j])$, more than compensates the losses associated with regulating bank l, $\Delta\omega\mathbb{E}[\Delta R]$. Consequently, the market equilibrium is not efficient and interventions are potentially welfare-improving.

III. Welfare Effects of Regulation

To better understand the welfare effects of regulation in a more general context, this section studies how arbitrary regulation alters expected total output. Because restricted banks are forced to change their risk-taking behavior, regulation reshapes the way that liquidity shocks propagate, thereby modifying the distribution of total output when a crisis manifests.

Without loss of generality, suppose the planner restricts every bank within an arbitrary set $\mathcal{R}_x$. Hereinafter, subscripts indicate the relative size of a set (compared to the size of the economy). Namely, the cardinality of $\mathcal{R}_x$ is $n \times x$, where $x = \frac{s}{n}$ and $s \in \{0, 1, \cdots, n\}$. Then,

$$\begin{aligned}\left(\frac{1}{n}\right)\text{TO} &= \left(\frac{1}{n}\right)\left(\sum_{i \in \mathcal{R}_x} \pi_i\right) + \left(\frac{1}{n}\right)\left(\sum_{i \notin \mathcal{R}_x} \pi_i\right) \\ &= x(R_I - \omega_H \Delta R) + \left(\frac{1}{n}\right)\left(\sum_{i \notin \mathcal{R}_x} \pi_i\right),\end{aligned}$$

where $\Delta\omega \equiv \omega_H - \omega_L$ and $\Delta R \equiv R_I - R_L$; $\left(\frac{1}{n}\right)$ is a normalization term used in case the economy grows large. To determine the distribution of $\left(\frac{1}{n}\right)\left(\sum_{i \notin \mathcal{R}_x} \pi_i\right)$, it is illustrative to analyze how shocks propagate when economic conditions deteriorate. Because the bank that initially faces a liquidity shock is selected uniformly at random, the probability that such a bank was restricted at $t = t_O$ is x. In this case, contagion is prevented from its onset, as restricted banks do not propagate liquidity shocks. However, if such a bank was not restricted, then at least one bank faces a liquidity shock and that shock might spread. Therefore,

$$\left(\frac{1}{n}\right)\sum_{i \notin \mathcal{R}_x} \pi_i = \begin{cases} (1-x)[R_I - \omega_L \Delta R], \text{ with probability } x \\ \left(1 - x - \frac{m}{n}\right)[R_I - \omega_L \Delta R], \text{ with probability } (1-x)\phi_m^{\mathcal{R}_x} \text{ with } m = 1, \cdots, n(1-x), \end{cases}$$

where $\phi_m^{\mathcal{R}_x}$ denotes the probability that m banks are affected by the liquidity shock at $t = t_F$ once every bank in $\mathcal{R}_x$ has been restricted at $t = t_O$. After some algebra, it can be shown that (see Appendix A.B)

$$\left(\frac{1}{n}\right)\mathbb{E}[\mathrm{TO}|\mathcal{R}_x] \;=\; \nu - \underbrace{\nu(1-x)\frac{\langle\phi^{\mathcal{R}_x}\rangle}{n}}_{\text{costs of contagion}} - \underbrace{x\Delta\omega\mathbb{E}[\Delta R],}_{\text{losses associated with regulation}} \tag{9}$$

where $\nu \equiv \mathbb{E}[R_I] - \omega_L\mathbb{E}[\Delta R]$. The term $\langle\phi^{\mathcal{R}_x}\rangle \equiv \left(\sum_{m=1}^{n(1-x)} m\phi_m^{\mathcal{R}_x}\right)$ denotes the expected number of banks affected by liquidity shocks at $t = t_F$ after every bank in $\mathcal{R}_x$ has been restricted.

It directly follows from (9) that increasing the relative size of $\mathcal{R}_x$, x, increases the losses arising from liquidity restrictions—captured by $x\Delta\omega\mathbb{E}[\Delta R]$. Intuitively, once a fraction x of banks are restricted, banks' expected payoffs decrease by $x\Delta\omega\mathbb{E}[\Delta R]$, as restricted banks invest a higher fraction of their portfolio in assets with lower expected returns. Importantly, increasing x not only increases the aforementioned losses but also alters the costs arising from the spread of liquidity shocks—captured by $\nu(1-x)\frac{\langle\phi^{\mathcal{R}_x}\rangle}{n}$. While increasing x decreases $(1-x)$, it is not clear whether increasing x will increase or decrease $\langle\phi^{\mathcal{R}_x}\rangle$ as such an increase changes the composition of banks within $\mathcal{R}_x$.

It is then pivotal to determine probabilities $\left\{\phi_m^{\mathcal{R}_x}\right\}_m$ to study how $\langle\phi^{\mathcal{R}_x}\rangle$ varies with x. While computing these probabilities is challenging even for small and medium-sized economies—as liquidity shocks may propagate in intricate ways—the following two lemmas show that, within the model, these probabilities can be computed for economies with arbitrary sizes and network architectures.

Before computing probabilities $\left\{\phi_m^{\mathcal{R}_x}\right\}_m$, I present the following result, which characterizes how random and strategic targeting affect the distribution of contagious exposures among non-restricted banks.

LEMMA 1 (Probabilities $\left\{\theta_k^{\mathcal{R}}\right\}_k$): *For a given distribution $\{p_k\}_k$, let $\langle k\rangle = \sum_{k=0}^{n-1} kp_k$ denote the expected number of contagious exposures per bank. Let $\mathcal{R}_{x_r}$ denote the set of restricted banks chosen after the planner decides not to pay κ. Let $\mathcal{R}_{x_t}$ denote the set of restricted banks* *chosen after the planner decides to pay κ so as to strategically target banks. Write $\theta_k^{\mathcal{R}}$ for the probability that a non-restricted bank shares k contagious exposures with other non-restricted banks, with $k \geq 0$, after every bank in $\mathcal{R}$ has been restricted.*

- *If $\mathcal{R} = \mathcal{R}_{x_r}$, then*

$$\theta_k^{\mathcal{R}} \;=\; \begin{cases} \sum_{j=k}^{n-1} p_j \binom{j}{k} \left(1 - x_r\right)^k x_r^{j-k} & \textit{if } k = \{0, \cdots, n(1-x_r) - 1\} \\ 0 & \textit{otherwise.} \end{cases}$$

- *If $\mathcal{R} = \mathcal{R}_{x_t}$, then*

$$\theta_k^{\mathcal{R}} \;=\; \begin{cases} \sum_{j=k}^{k_t} p_j \binom{j}{k} (1 - \kappa_t)^k \kappa_t^{j-k} & \textit{if } k = \{0, \cdots, k_t\} \\ 0 & \textit{otherwise,} \end{cases}$$

  where k_t and κ_t satisfy the system of equations

$$x_t \;=\; 1 - \sum_{k=0}^{k_t} p_k \quad \textit{and} \quad \kappa_t \;=\; 1 - \frac{1}{\langle k \rangle} \left(\sum_{k=0}^{k_t} k p_k \right).$$

Lemma 1 states that the way the planner targets banks matters as it changes the distribution of contagious exposures among non-restricted banks. With strategic targeting, the planner restricts the most contagious banks first, resulting in the removal of a larger fraction of contagious exposures than when restricting at random (from the perspective of shock propagation). As a result, the way the planner targets banks alters the susceptibility of the economy to contagion. With this result at hand, the next lemma characterizes probabilities $\left\{\phi_m^{\mathcal{R}_x}\right\}_m$.

LEMMA 2 (Probabilities $\left\{\phi_m^{\mathcal{R}_x}\right\}_m$): *For a given distribution $\{p_k\}_k$, define probabilities $\theta_k^{\mathcal{R}_x}$*

as in lemma 1. Suppose all banks within the set $\mathcal{R}_x$ are restricted. Then

$$\phi_m^{\mathcal{R}_x} = \begin{cases} \frac{\langle \theta^{\mathcal{R}_x} \rangle}{(m-1)!} \left(\frac{d^{m-2}}{dz^{m-2}} \left[g\left(z, \left\{\theta_k^{\mathcal{R}_x}\right\}_k\right)^m \right] \right) \Big|_{z=0} & , \text{ with } m = \{2, \cdots, n(1-x)\} \\ \theta_0^{\mathcal{R}_x} & , \text{ with } m = 1, \end{cases}$$

where $\left(\frac{d^{m-2}}{dz^{m-2}} \left[g\left(z, \left\{\theta_k^{\mathcal{R}_x}\right\}_k\right)^m \right] \right) \Big|_{z=0}$ denotes the $(m-2)$ derivative of $g\left(z, \left\{\theta_k^{\mathcal{R}_x}\right\}_k\right)^m$ evaluated at $z = 0$, with

$$g(z, \left\{\theta_k^{\mathcal{R}_x}\right\}_k) = \sum_{k=0}^{n(1-x)-2} \left(\frac{(k+1)\theta_{k+1}^{\mathcal{R}_x}}{\langle \theta^{\mathcal{R}_x} \rangle} \right) z^k \quad \text{and} \quad \langle \theta^{\mathcal{R}_x} \rangle = \sum_{k=0}^{n(1-x)-1} k\theta_k^{\mathcal{R}_x}.$$

Lemmas 1 and 2 allow me to compute $\langle \phi^{\mathcal{R}_x} \rangle$ for any distribution $\{p_k\}_k$. Notably, the computation of $\left\{\phi_m^{\mathcal{R}_x}\right\}_m$ requires calculating sums and derivatives, which can be done numerically for any finite n. However, for certain families of distributions $\{p_k\}_k$, it is possible to derive closed-form expressions, as example 1 shows.

EXAMPLE 1 (Poisson): *Suppose $\{p_k\}_{k=0}^{n-1}$ follows a Poisson distribution with parameter α, that is, $p_k = e^{-\alpha} \frac{\alpha^k}{k!}$.*

- *If $\mathcal{R}_x = \mathcal{R}_{x_r}$, then*

$$\phi_m^{\mathcal{R}_x} = \frac{e^{-(1-x_r)\alpha m} \left((1-x_r)\,\alpha m\right)^{m-1}}{m!}, \; m = \{1, \cdots, n(1-x_r)\}.$$

- *If $\mathcal{R}_x = \mathcal{R}_{x_t}$, then*

$$\phi_m^{\mathcal{R}_x} = \frac{e^{-(1-\kappa_t)\alpha m} \left((1-\kappa_t)\,\alpha m\right)^{m-1}}{m!}, \; m = \{1, \cdots, n\,(1-x_t)\}.$$

  where $\kappa_t = 1 - \frac{1}{\alpha} \left(e^{-\alpha} \sum_{k=0}^{k_x} \frac{k\alpha^k}{k!} \right)$ and $x_t = e^{-\alpha} \sum_{k=k_t}^{n-1} \frac{\alpha^k}{k!}$.

For illustration, suppose $\{p_k\}_k$ follows a Poisson distribution with parameter α. Importantly, when $\{p_k\}_k$ follows a Poisson distribution, α captures the first two moments of

$\{p_k\}_k$. Figure 5 illustrates that $\langle \phi^{\mathcal{R}_x} \rangle$ can be a non-monotonic mapping of x. Figure 5(a) depicts $\langle \phi^{\mathcal{R}_x} \rangle$ as a function of x if the planner restricts banks at random. When $\alpha = 1$, the susceptibility of the economy to contagion is small. This is because both the average number of contagious exposures per bank and the variation of contagion exposures across banks are small. Thus, $\langle \phi^{\mathcal{R}_x} \rangle$ is a weakly decreasing function of x, as increasing x effectively tilt the distribution $\{\phi_m^{\mathcal{R}_x}\}_m$ (see figure 6(a)), making cascades of liquidity shocks relatively less likely, thereby decreasing $\langle \phi^{\mathcal{R}_x} \rangle$. However, when $\alpha > 1$, the variation of contagious exposures across banks might be large enough to allow contagion to be far-reaching. Consequently, $\langle \phi^{\mathcal{R}_x} \rangle$ becomes a non-monotonic function of x. For small values of x, increasing x isolates banks with only few contagious exposures with high probability, making cascades relatively more likely (see figure 6(b)), which, in turn, increases $\langle \phi^{\mathcal{R}_x} \rangle$. However, when x is relatively large, increasing x isolates a sufficiently large number of banks, decreasing the average number of contagion exposures enough to tilt the distribution $\{\phi_m^{\mathcal{R}_x}\}_m$, curbing the likelihood of cascades, thereby decreasing $\langle \phi^{\mathcal{R}_x} \rangle$ (see figure 6(c)).

Figure 5(b) depicts $\langle \phi^{\mathcal{R}_x} \rangle$ as a function of x if the planner restricts banks based on their future number of contagious exposures. As figure 5(b) shows, $\langle \phi^{\mathcal{R}_x} \rangle$ continue to vary with x, but not necessarily in a continuous fashion. Thus, the dependence of $\langle \phi^{\mathcal{R}_x} \rangle$ on x not only hinges on the susceptibility of the economy to contagion—which is encoded in $\{p_k\}_k$—but also hinges on how restricted banks are selected.

Improving network transparency. It follows from (9) that a planner's decision to learn the identities of the most connected banks depends on how much that information helps reduce $\langle \phi^{\mathcal{R}_x} \rangle$. If only a few banks play a key role in the propagation of shocks, targeting those banks substantially reduces $\langle \phi^{\mathcal{R}_x} \rangle$, and, hence, collecting information about their identities may be worth the cost. Consequently, improving network transparency has an intrinsic value to the extent that it allows the planner to dampen cascades of liquidity shocks more effectively. As the next section shows, this value dictates the optimal choice of network transparency and it is determined by the interplay between the distribution $\{p_k\}_k$ and the marginal cost of regulation, $\Delta\omega\mathbb{E}[\Delta R]$.

IV. Optimal Interventions

This section characterizes optimal interventions—which are jointly determined by a choice of network transparency and a selection of restricted banks—and explores how these policies vary with the primitives of the model. Mirroring the motivating example, I now focus on two cases. Section IV.A describes the optimal intervention when $\{p_k\}_k$ is known, paralleling the analysis of section I.A. Section IV.B mirrors the analysis of section I.B, studying economies wherein $\{p_k\}_k$ is unknown.

A. $\{p_k\}_k$ is known

When $\{p_k\}_k$ is known, there is no model uncertainty. Consequently, regulation is selected so as to maximize expected total output.

For ease of exposition, I first study the optimal selection of restricted banks, given a choice of network transparency. I then study the optimal choice of network transparency.

A.1. Selecting the optimal set of restricted banks

The next proposition shows that the planner's problem has a solution.

PROPOSITION 1 (Existence): *Take the choice of network transparency as given. For any distribution $\{p_k\}_k$, there exists a solution of the planner's problem, that is,*

$$\max_{\mathcal{R}\,\in\,\mathcal{P}(B_n)} \quad \mathbb{E}\left[\left(\frac{1}{n}\right) TO\middle|\mathcal{R}\right] - \kappa\times\mathbb{1}_\kappa$$

where $\mathbb{1}_\kappa$ is an indicator function that equals 1 if the planner pays κ and 0 otherwise.

The above proposition guarantees that the planner's problem can be solved. The basic strategy behind the proof is to show that $\mathbb{E}\left[\left(\frac{1}{n}\right)\mathrm{TO}\middle|\mathcal{R}\right]$ is bounded from above. This is true because $0 \leq \frac{\langle \phi^{\mathcal{R}} \rangle}{n} \leq \left(1 - \frac{|\mathcal{R}|}{n}\right)$ for any $\mathcal{R} \in \mathcal{P}(B_n)$.

The next proposition shows that, under somewhat more restrictive conditions, I can uniquely determine the relative size of the optimal set of restricted banks. To pose those conditions in a sharper way, I first introduce the following definitions.

DEFINITION 1: *A simple intervention is a set of restricted banks.*

That is, a simple intervention is any element of the power set of B_n, $\mathcal{P}(B_n)$.

DEFINITION 2: *For a given distribution $\{p_k\}_k$, a collection of simple interventions $\mathcal{C}_R \equiv \mathcal{C}_R(\{p_k\}_k) = \{\mathcal{R}_x | \mathcal{R}_x$ is a simple intervention$\} \subseteq \mathcal{P}(B_n)$ is said to be size differentiable if the mapping $g : \mathcal{C}_R \to [0,1]$, defined as $g(\mathcal{R}_x) \equiv \langle \phi^{\mathcal{R}_x} \rangle$, is differentiable for all x, with $0 \leq x \leq 1$.*

While the above definition might seem abstract, there are several instances in which a planner might focus on size differentiable interventions. For instance, suppose $\{p_k\}_k$ follows a Poisson distribution. If the planner decides not to pay κ, she restricts banks at random. Figure 5(a) shows that the set of feasible interventions is size differentiable. However, if the planner decides to pay κ so as to learn the identity of the most connected banks, figure 5(b) shows that the set of feasible interventions is no longer size differentiable.

PROPOSITION 2 (Optimal size): *Take the choice of network transparency as given. For a given distribution $\{p_k\}_k$, suppose the planner only focuses on simple interventions within a size differentiable collection, $\mathcal{C}_R(\{p_k\}_k)$. Take any three elements of $\mathcal{C}_R$, say $\mathcal{R}_z$, $\mathcal{R}_{x_0}$, and $\mathcal{R}_{x_1}$, with $z = (1-\alpha)x_0 + \alpha x_1 \in \left(\frac{1}{n}, 1\right)$, $\frac{1}{n} \leq x_0 < x_1 \leq 1$ and $\alpha \in (0,1)$. If*

$$\frac{(1-\alpha)(1-x_0)\langle \phi^{\mathcal{R}_{x_0}} \rangle + \alpha(1-x_1)\langle \phi^{\mathcal{R}_{x_1}} \rangle}{(1-\alpha)(1-x_0) + \alpha(1-x_1)} \quad > \quad \langle \phi^{\mathcal{R}_z} \rangle,$$

then the relative size of the optimal intervention, $\mathcal{R}_{x^*} \in \mathcal{C}_R(\{p_k\}_k)$, approximately solves

$$\underbrace{\nu \left(\frac{\langle \phi^{\mathcal{R}_{x^*}} \rangle}{n} - (1-x^*) \frac{\partial}{\partial x} \left(\frac{\langle \phi^{\mathcal{R}_x} \rangle}{n} \right) \Bigg|_{x=x^*} \right)}_{\textit{marginal benefit}} \quad = \quad \underbrace{\Delta \omega \mathbb{E}[\Delta R],}_{\textit{marginal cost}} \tag{10}$$

as n grows large.

Simply put, the conditions in the above proposition ensure that expected total output is a strictly concave mapping of the relative size of simple interventions, and, thus, I can determine the optimal fraction of restricted banks, x^*, by solving the first order condition of the planner's problem (represented by equation (10)).

At the fundamental level, equation (10) highlights the planner's trade-off when selecting the set of restricted banks: x^* is deliberately chosen so as to limit the spread of liquidity shocks while at the same time avoiding excessive losses from liquidity restrictions. The optimal intervention ensures that the benefits of restricting the last bank are equal to the losses associated with restricting such a bank.

A.2. Value of network transparency

Before implementing restrictions, the planner decides whether to learn the identities of the most contagious banks. By learning this information, the planner can strategically target them first so as to confine the spread of liquidity shocks at a lower cost. Naturally, her decision will depend on how much this information helps mitigate contagion more effectively.

Let $\mathcal{R}_{x_t}$ and $\mathcal{R}_{x_r}$ denote the set of restricted banks chosen if the planner strategically targets banks or restricts them at random, respectively. Define $\Delta x \equiv (x_r - x_t)$. The social value of improving network transparency, SVI, is then

$$\begin{aligned} \text{SVI} &\equiv \left(\frac{1}{n}\right)\left(\mathbb{E}[\text{TO}|\mathcal{R}_{x_t}] - \mathbb{E}[\text{TO}|\mathcal{R}_{x_r}]\right) \\ &= \Delta x \Delta\omega \mathbb{E}[\Delta R] + \nu\left((1-x_r)\frac{\langle\phi^{\mathcal{R}_{x_r}}\rangle}{n} - (1-x_t)\frac{\langle\phi^{\mathcal{R}_{x_t}}\rangle}{n}\right). \end{aligned} \tag{11}$$

The left-hand side of equation (11) highlights the two benefits of improving network transparency. The first term, $\Delta x \Delta\omega \mathbb{E}[\Delta R]$, captures the fact that more transparency potentially allows the planner to restrict fewer banks. The second term, $\nu\left((1-x_r)\frac{\langle\phi^{\mathcal{R}_{x_r}}\rangle}{n} - (1-x_t)\frac{\langle\phi^{\mathcal{R}_{x_t}}\rangle}{n}\right)$, captures the fact that more transparency possibly allows the planner to decrease the spread of liquidity shocks more effectively. Even when $x_r = x_t$, $\langle \phi^{\mathcal{R}_{x_r}} \rangle \geq \langle \phi^{\mathcal{R}_{x_t}} \rangle$ as the planner strategically chooses $\mathcal{R}_{x_t}$ so as to curb contagion more efficiently than when restricting banks at random.

Importantly, the value of network transparency depends on the pattern of contagious exposures across banks, as SVI varies with $\{p_k\}_k$, because x_r, x_t, $\langle \phi^{\mathcal{R}_{x_r}} \rangle$, and $\langle \phi^{\mathcal{R}_{x_t}} \rangle$ are implicit functions of $\{p_k\}_k$.

Optimal rule. Improving network transparency is optimal if and only if

$$\Delta x \Delta \omega \mathbb{E}[\Delta R] + \nu \left((1 - x_r) \frac{\langle \phi^{\mathcal{R}_{x_r}} \rangle}{n} - (1 - x_t) \frac{\langle \phi^{\mathcal{R}_{x_t}} \rangle}{n} \right) \quad \geq \quad \kappa.$$

As a result, the planner's decision of whether to improve network transparency is closely linked to how useful that transparency is to limit the propagation of liquidity shocks.

A.3. Optimal Interventions in Large Economies

This section provides closed form characterizations of optimal interventions to better illustrate how interventions vary with the primitives of the model. To do so, I focus on the limiting case $n \to \infty$, as it considerably facilitates computations. When $n \to \infty$, cascades of finite size have no material effect. Consequently, in what follows, I focus on large cascades of liquidity shocks defined as events in which a finite fraction of banks in an economy of infinite size face a liquidity shock as a result of any one bank initially facing a liquidity shock.

Appendix A.D shows that when $n \to \infty$, the optimal fraction of restricted banks, x^*, equals

$$x^* \quad = \quad \begin{cases} x_r, \;\text{ if } \Delta\omega\mathbb{E}[\Delta R] \leq \min\left\{\frac{\nu}{x_r}, \frac{\kappa}{\Delta x}\right\} \\ x_t, \;\text{ if } \min\left\{\frac{\nu}{x_r}, \frac{\kappa}{\Delta x}\right\} < \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu - \kappa}{x_t} \\ 0, \qquad\quad \text{otherwise.} \end{cases} \tag{12}$$

where x_t and x_r denote the smallest fraction of banks that must be restricted to prevent large cascades of liquidity shocks if either the planner strategically targets banks or restricts them at random.

Intuitively, when $\Delta\omega\mathbb{E}[\Delta R] \leq \min\left\{\frac{\nu}{x_r}, \frac{\kappa}{\Delta x}\right\}$, the marginal cost of regulation is so small that the losses associated with random targeting are negligible. It is then optimal not to improve network transparency, and, thus, $x^* = x_r$. When $\min\left\{\frac{\nu}{x_r}, \frac{\kappa}{\Delta x}\right\} < \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu-\kappa}{x_t}$, less efficient interventions due to random targeting become sufficiently costly, as they involve restricting an excessively large fraction of banks. It is then optimal to strategically select banks, and, thus, $x^* = x_t$. Finally, when $\Delta\omega\mathbb{E}[\Delta R] > \frac{\nu-\kappa}{x_t}$, the marginal cost of regulation is too large, and, thus, a non-interventionist policy is optimal; that is, $x^* = 0$.

Appendix A.D also shows that when $n \to \infty$ the social value of improving network transparency, SVI, equals

$$\text{SVI} = \begin{cases} \Delta x \Delta\omega\mathbb{E}[\Delta R], \text{ if } \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu}{x_r} \\ \nu - x_t\Delta\omega\mathbb{E}[\Delta R], \text{ if } \frac{\nu}{x_r} < \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu}{x_t} \\ 0, \qquad \text{otherwise.} \end{cases} \tag{13}$$

That is, when the marginal cost of regulation is small, i.e., $\Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu}{x_r}$, the value of improving network transparency is proportional to Δx, as more transparency possibly allow the planner to restrict fewer banks. Importantly, the value of improving transparency increases with $\Delta\omega\mathbb{E}[\Delta R]$, as strategically targeting directly reduces the losses associated with ineffective regulation due to random targeting. When the marginal cost of regulation is intermediate, i.e., $\frac{\nu}{x_r} < \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu}{x_t}$, it becomes optimal to improve network transparency as long as κ is sufficiently small. In this case, the value of improving transparency decreases with $\Delta\omega\mathbb{E}[\Delta R]$ because, as $\Delta\omega\mathbb{E}[\Delta R]$ increases, any intervention becomes more costly to begin with. For sufficiently large values of $\Delta\omega\mathbb{E}[\Delta R]$, a non-interventionist policy is optimal, and, thus, network transparency adds no efficiency gains from a policy perspective; therefore, $\text{SVI} = 0$.

Importantly, expressions (12) and (13) underscore that the optimal intervention crucially depends on the interplay between (a) the susceptibility of the economy to contagion, through x_t and x_r, (b) the marginal cost of regulation, $\Delta\omega\mathbb{E}[\Delta R]$, and (c) the cost of improving network transparency, κ. Thus, at the fundamental level, the findings presented in section I.A can be extended to economies with larger sizes and arbitrary network architectures.

To better illustrate how $\{p_k\}_k$ alters the optimal intervention, I now focus on two distinct families of distributions $\{p_k\}_k$: Poisson and Power-laws.

Poisson networks. Suppose $\{p_k\}_k$ follows a Poisson distribution with parameter $\alpha > 0$. Then,

$$x_r = 1 - \frac{1}{\alpha} \quad \text{and} \quad x_t = x_r - \frac{e^{-\alpha}\alpha^{K_\alpha}}{K_\alpha!},$$

where K_α solves $\frac{1}{\alpha} = \sum_{j=0}^{(K_\alpha-2)} \frac{e^{-\alpha}\alpha^j}{j!}$.

Substituting x_r and x_t into (12) fully characterizes $x^*(\alpha)$. This characterization helps in analyzing how interventions vary with the underlying network architecture. Figure 8(a) illustrates $x^*(\alpha)$ for different values of $\Delta\omega\mathbb{E}[\Delta R]$ when $\kappa = 1/10$. When α is small, both contagious exposures are less frequent and the variation of contagious exposures across banks is small. Thus, if $\Delta\omega\mathbb{E}[\Delta R]$ is not sufficiently large, the planner has little incentive to improve network transparency as no single bank plays a determinant role in the spread of liquidity shocks. Thus, $x^* = x_r$. When α increases, however, the economy becomes more prone to contagion, because (1) more banks exhibit a higher number of contagious exposures and (2) there is more variation in contagious exposures across banks. As a result, the planner now has more incentives to identify the most contagious banks. Consequently, unless $\Delta\omega\mathbb{E}[\Delta R]$ is sufficiently small, $x^* = x_t$. When $\Delta\omega\mathbb{E}[\Delta R]$ is sufficiently large, the losses associated with regulation are considerable, and, thus, $x^* = 0$.

Substituting x_r and x_t into (13) characterizes SVI. Figure 8(b) illustrates $\text{SVI}(\Delta\omega\mathbb{E}[\Delta R], \alpha)$ through a contour plot. Curves connect pairs $(\Delta\omega\mathbb{E}[\Delta R], \alpha)$ where network transparency yields the same value (which are indicated in the figure). When $\Delta\omega\mathbb{E}[\Delta R]$ is sufficiently small, the costs associated with regulation are negligible, and, thus, restricting as many banks as possible is optimal regardless of the network architecture. Hence, network transparency has no value. Similarly, when $\Delta\omega\mathbb{E}[\Delta R]$ is sufficiently large, network transparency has no value, as a non-interventionist policy is always optimal.

For intermediate values of $\Delta\omega\mathbb{E}[\Delta R]$, however, the analysis becomes more involved, as improving network transparency may be optimal depending on the interplay between κ and the susceptibility of the economy to contagion, captured by α. Suppose $\Delta\omega\mathbb{E}[\Delta R]$ is held fixed at one of these values. When α is small, contagion is an almost zero-probability event. As a consequence, network transparency has no value as no regulation is required. As α increases, the susceptibility of the economy to contagion increases, thereby increasing the planner's incentives to identify the most contagious banks. Thus, the value of network transparency increases. When α is sufficiently large, contagion tends to happen anyway. Thus, network transparency adds no value.

Power-law networks. Suppose $\{p_k\}_k$ follows a Power-law distribution with parameter $\alpha > 2$—that is, $p_k \propto k^{-\alpha}$. Assume further that the minimum number of contagious exposures per bank equals one. Then,

$$\begin{aligned} x_r &= \begin{cases} 1 - \left(\left(\frac{2-\alpha}{3-\alpha}\right) - 1\right)^{-1} & \text{if } \alpha > 3 \\ 1 & \text{if } 2 < \alpha \leq 3. \end{cases} \\ x_t &= K_\alpha^{(1-\alpha)} \end{aligned}$$

where K_α satisfies $K_\alpha^{2-\alpha} - 2 = \left(\frac{2-\alpha}{3-\alpha}\right)(K_\alpha^{3-\alpha} - 1)$.

As before, substituting x_r and x_t into expressions (12) and (13) characterizes x^* and SVI. To appreciate the differences between Poisson and Power-law distributions, figure 7 illustrates $x_r(\alpha)$ and $x_t(\alpha)$ under both distributions. Notably, Δx is visibly larger in the Power-law than in the Poisson case; especially when $2 < \alpha < 3$. The reason is simple. When $\{p_k\}_k$ follows a

Power-law and $2 < \alpha < 3$, only the first moment of $\{p_k\}_k$ is finite while higher order moments are infinite. In this case, the high variation in the number of contagious exposures across banks ensures that liquidity shocks affecting one bank almost surely affect a non-negligible fraction of them through propagation, making the economy highly susceptible to contagion. Because only an extremely small fraction of banks exhibit an excessively large number of contagious exposures, the planner is likely to miss such banks if targeting banks at random. Consequently, network transparency tends to prove more helpful in the Power-law than in the Poisson case.

Figure 9(a) depicts $x^*(\alpha)$ for different values of $\Delta\omega\mathbb{E}[\Delta R]$ when $\{p_k\}_k$ follows a Power-law. Consistent with the previous analysis, when $\alpha \leq 3$, the planner decides to improve network transparency so as to minimize the losses associated with excessive regulation; hence, $x^* = x_t$. However, for larger values of α and sufficiently small values of $\Delta\omega\mathbb{E}[\Delta R]$, $x^* = x_r$.

Figure 9(b) illustrates $\mathrm{SVI}(\Delta\omega\mathbb{E}[\Delta R], \alpha)$ through a contour plot. As before, when $\Delta\omega\mathbb{E}[\Delta R]$ is either sufficiently small or large, improving network transparency adds no value. For intermediate values of $\Delta\omega\mathbb{E}[\Delta R]$, improving network transparency may be optimal. Importantly, the analysis in the Power-law case differs fundamentally from the Poisson case. When $\alpha < 3$, large cascades cannot be prevented without improving network transparency (see Appendix A.D.1). As a consequence, learning the identity of the most contagious banks is considerably valuable, and, thus, it tends to be optimal.

Discussion. To sum up, the marked differences between the connectivity structures of Poisson and Power-law networks underscore three important findings. First, different network architectures exhibit different susceptibility to contagion, imposing distinct challenges to planners when trying to mitigate cascades of liquidity shocks. Second, the planner's ability to prevent cascades in certain networks heavily depends on how restricted banks are selected, and, thus, the extent of network transparency may be critical. Third, the scope for welfare-improving interventions is intimately linked to the degree of symmetry exhibited by the network of contagious exposures, as such degree alters the susceptibility of the economy to contagion. In particular, when $\{p_k\}_k$ is somewhat symmetric—in the sense that the number of contagious exposures does not vary too much across banks—banks do not significantly differ in their role spreading liquidity shocks. Consequently, network transparency brings limited efficiency gains. Yet, when $\{p_k\}_k$ is somewhat asymmetric, a small fraction of banks drives the spread of liquidity shocks, and, thus, there are considerable efficiency gains from improving network transparency.

B. $\{p_k\}_k$ *is unknown*

Mirroring the analysis of subsection I.B, this subsection studies economies wherein $\{p_k\}_k$ is unknown. Here, the planner faces model uncertainty, as she is unsure about the generating process of contagious exposures, which ultimately determines how shocks propagate when a crisis manifests.[6]

To better appreciate the implications of model uncertainty within the general framework, suppose it is commonly acknowledged that $\{p_k\}_k$ comes from a known family of distributions. Assume further that each distribution within that family can be identified by a parameter α (such as in the Poisson or Power-law case). Now suppose α is unknown. Let $\{\{p_k^\alpha\}_k\}_{\alpha\in\mathcal{A}}$ denote such a family of distributions where $\mathcal{A}$ denotes the set of plausible values for α. For a given value of α, the previous analysis shows that if every bank in $\mathcal{R}_x$ is restricted, then

$$\left(\frac{1}{n}\right)\mathbb{E}_\alpha[\mathrm{TO}|\mathcal{R}_x] \;=\; \nu-\nu(1-x)\left(\sum_{m=1}^{n(1-x)}\frac{m}{n}\phi_m^{\mathcal{R}_x}(\alpha)\right)-x\Delta\omega\mathbb{E}[\Delta R]$$

where probabilities $\phi_m^{\mathcal{R}_x}$ are written as $\phi_m^{\mathcal{R}_x}(\alpha)$ to emphasize their dependence on the precise value of α. Importantly, if α is random, probabilities $\phi_m^{\mathcal{R}_x}(\alpha)$ are now random variables, which, in turn, makes expected total output a random variable.

Given the representative investor's preferences, the planner takes into account the extent

[6] See Routledge and Zin (2009) and Easley and O'Hara (2010) for models that connect liquidity and model uncertainty. See Ruffino (2014) for a discussion of some implications of model uncertainty for regulation.

of model uncertainty when designing interventions, as she imposes restrictions so as to maximize the representative investor's smooth ambiguity certainty equivalent. Suppose it is commonly acknowledged that α is distributed over $\mathcal{A}$ according to a distribution f, with finite barycenter $\bar{\alpha} \equiv \int_{\alpha\in\mathcal{A}} \alpha df$. Then, given a choice of network transparency, the planner chooses a set of restricted banks, $\mathcal{R}_x$, to solve

$$\max_{\mathcal{R}_x \in \mathcal{P}(B_n)} \quad \mathbb{E}_{\bar{\alpha}}\left(\frac{1}{n}\mathrm{TO}\middle|\mathcal{R}_x\right) - \left(\frac{\theta}{2}\right) \times \mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_\alpha\left(\mathrm{TO}\middle|\mathcal{R}_x\right)\right) - \kappa \times \mathbb{1}_\kappa, \tag{14}$$

where operator $\mathbb{E}_{\bar{\alpha}}(\cdot)$ denotes the expectation when $\alpha = \bar{\alpha}$. Operator $\mathbb{V}_f(\cdot)$ denotes the variance of expected total output, computed using distribution f. As in the motivating example, θ captures the representative investor's attitude toward ambiguity. Notably, the analysis in this section is equivalent to the analysis in the previous section when $\mathcal{A}$ is singleton or $\theta = 0$. When $\mathcal{A}$ is singleton, there is no model uncertainty, as the exact value of α is known. When $\theta = 0$, the representative investor is ambiguity neutral. Thus, the planner does not mind not knowing α and implements restrictions as if $\alpha = \bar{\alpha}$.

B.1. Selecting the optimal set of restricted banks

The next proposition shows that problem (14) has a solution.

PROPOSITION 3 (Existence under Ambiguity): *Take the choice of network transparency as given. Then, for any distributions $\{p_k\}_k$ and f, there exists a solution of problem (14).*

The above proposition guarantees that the planner's problem can be solved when α is unknown. The basic strategy behind the proof is to show that $\mathbb{E}_{\bar{\alpha}}\left[\left(\frac{1}{n}\right)\mathrm{TO}\middle|\mathcal{R}_x\right]$ is bounded from above while $\mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_\alpha\left(\mathrm{TO}\middle|\mathcal{R}_x\right)\right)$ is bounded from below. This is true because $0 \leq \frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_x}\rangle}{n} \leq (1-x)$ while $\mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_\alpha\left(\mathrm{TO}\middle|\mathcal{R}_x\right)\right) \geq 0$ for any $\mathcal{R}_x \in \mathcal{P}(B_n)$, where $\langle\phi_{\bar{\alpha}}^{\mathcal{R}_x}\rangle \equiv \sum_{m=1}^{n(1-x)} m\phi_m^{\mathcal{R}_x}(\bar{\alpha})$.

For conciseness, I defer to proposition 4 in Appendix A.C the discussion of conditions under which the relative size of restricted banks can be uniquely determined when α is unknown. For now, suppose those conditions are satisfied. To illustrate the planner's trade-off in this environment, I now rewrite the first order condition of problem (14) as

$$\underbrace{\nu\left(\frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_{x^a}}\rangle}{n}-(1-x^a)\frac{\partial}{\partial x}\left(\frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_x}\rangle}{n}\right)\Bigg|_{x=x^a}\right)}_{\text{marginal benefit}} = \underbrace{\Delta\omega\mathbb{E}[\Delta R] + \left(\frac{\theta}{2}\right)\nu^2\frac{\partial}{\partial x}\left((1-x)^2\int_{\alpha\in\mathcal{A}}\left(\frac{\langle\phi_{\alpha}^{\mathcal{R}_x}\rangle}{n}-\frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_x}\rangle}{n}\right)^2 df(\alpha)\right)\Bigg|_{x=x^a}}_{\text{marginal cost}}$$

To appreciate the importance of model uncertainty, it is illustrative to emphasize the similarities between the above equation and equation (10). While the marginal benefit in both equations is similar, the marginal cost now has an extra component—the second term in the RHS of the above equation. Importantly, this component is unrelated to losses arising from liquidity restrictions. This cost arises solely from the fact that (a) α is unknown, and (b) the representative investor exhibits aversion to ambiguity. Consequently, the optimal fraction of restricted banks, x^a, now also hinges on distribution f, and the representative investor's attitudes toward ambiguity.

B.2. Value of network transparency.

Let $\mathcal{R}_{x_t^a}$ and $\mathcal{R}_{x_r^a}$ denote the set of restricted banks chosen if the planner strategically target banks or restricts them at random, respectively. Define $\Delta x^a \equiv (x_r^a - x_t^a)$. The social value of improving network transparency under model uncertainty, SVI_a, is then

$$\begin{aligned}
\text{SVI}_a \;\equiv\; & \left(\frac{1}{n}\right)\left(\mathbb{E}_{\bar{\alpha}}[\text{TO}|\mathcal{R}_{x_t^a}]-\mathbb{E}_{\bar{\alpha}}[\text{TO}|\mathcal{R}_{x_r^a}]\right)-\left(\frac{\theta}{2}\right)\left(\mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_{\alpha}\left(\text{TO}|\mathcal{R}_{x_t^a}\right)\right)-\mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_{\alpha}\left(\text{TO}|\mathcal{R}_{x_r^a}\right)\right)\right)\\
&= \Delta x^a\Delta\omega\mathbb{E}[\Delta R]+\nu\left((1-x_r^a)\frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_{x_r^a}}\rangle}{n}-(1-x_t^a)\frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_{x_t^a}}\rangle}{n}\right)\\
&+\left(\frac{\theta}{2}\right)\nu^2\left((1-x_r^a)^2\left[\int_{\alpha\in\mathcal{A}}\left(\frac{\langle\phi_{\alpha}^{\mathcal{R}_{x_r^a}}\rangle}{n}-\frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_{x_r^a}}\rangle}{n}\right)^2 df(\alpha)\right]-(1-x_t^a)^2\left[\int_{\alpha\in\mathcal{A}}\left(\frac{\langle\phi_{\alpha}^{\mathcal{R}_{x_t^a}}\rangle}{n}-\frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_{x_t^a}}\rangle}{n}\right)^2 df(\alpha)\right]\right).
\end{aligned}$$

Thus, SVI now has three components. The first two terms in the RHS of the above equation capture ideas similar to the two components described in the RHS of equation (11). The first term arises from the fact that, on average, transparency helps decrease losses generated from excessive regulation. The second term captures the fact that, on average, transparency allows the planner to limit the spread of liquidity shocks better. However, the third term is new and captures the idea that the representative investor is ambiguity averse and α is unknown. Hence, she dislikes making non-optimal decisions ex-ante as a result of not knowing α. Thus, this term can be broadly interpreted as the extent to which transparency allows the planner to hedge the risk of implementing non-optimal policies ex-ante as a result of not knowing α. Importantly, the perception of such risk is intimately linked to the underlying family of networks $\{\{p_k^\alpha\}_k\}_{\alpha\in\mathcal{A}}$ as well as the distribution f.

Optimal rule. The optimal choice of network transparency follows a simple rule. If $\mathrm{SVI}_a \geq \kappa$, the social benefit of increased network transparency outweighs its cost, and, thus, it is optimal to improve network transparency. Notably, the extent of network uncertainty now plays a key role in this decision.

B.3. Optimal Interventions in Large Economies

To appreciate the importance of network uncertainty for policy design, I now explore (via numerical solutions) how the optimal intervention varies with (a) f, and (b) the representative investor's aversion to ambiguity, θ. To facilitate comparison between subsections IV.A and IV.B, I focus on Poisson and Power-law networks in what follows. For concreteness, I assume hereafter that $\nu = 1$, $\Delta\omega\mathbb{E}[\Delta R] = 2$, $\kappa = 1/10$, and $n = 50$, while f follows a truncated normal distribution with mean $\bar{\alpha} = 3$, variance $\sigma^2 > 0$, and $\mathcal{A} = [2, 4]$. Hereinafter, I capture the extent of model uncertainty through variation in σ^2.

Poisson networks. Figure 10 highlights the implications for policy design of changes in distribution f and θ. Figure 10(a) depicts the optimal fraction of restricted banks, x^*, as a function of σ^2. Unless θ is sufficiently large, variation in distribution f does not generate variation in x^*. When the representative investor does not experience sufficiently large disutility from making non-optimal decisions ex ante, it is optimal for the planner to choose her policy as if $\alpha = \bar{\alpha}$. Because $\bar{\alpha} = 3$ and $\Delta\omega\mathbb{E}[\Delta R] = 2$, a non-interventionist policy is optimal regardless of the extent of network uncertainty (see figure 8(a)).

However, when the representative investor exhibits sufficiently high aversion to ambiguity, the optimal policy can be heavily affected by changes in model uncertainty, as figure 10(a) shows. As σ^2 increases, the extent of network uncertainty increases. When facing high network uncertainty, the planner worries the fraction of restricted banks may not be sufficiently large to prevent large cascades. As a result, it tends to be optimal to increase the fraction of restricted banks as σ^2 increases. In this case, the lack of certainty about the network is not a justification for inaction, but rather the opposite, considering the considerable negative consequences of large cascades. Importantly, when the representative investor is sufficiently averse to ambiguity, figure 10(b) shows that the value of network transparency is non-negative for a large set of pairs $(\Delta\omega\mathbb{E}[\Delta R], \sigma^2)$ as it is extremely costly to implement non-optimal interventions ex ante. As a result, improving network transparency is valuable even though the network of contagious exposures might exhibit a symmetric architecture.

Power-law networks. Figure 11 shows that differences in the nature of the network architecture have important implications for policy making under network uncertainty. Consistent with the previous results, figure 11(a) shows that x^* is heavily affected by changes in network uncertainty when the representative investor exhibits sufficiently high aversion to ambiguity. Notably, figure 11(b) shows that, for relatively small values of σ^2, the value of information increases as $\Delta\omega\mathbb{E}[\Delta R]$ increases. Intuitively, as $\Delta\omega\mathbb{E}[\Delta R]$ increases, it becomes more costly to make mistakes (by restricting an excessively large fraction of banks). Consequently, as $\Delta\omega\mathbb{E}[\Delta R]$ increases, the higher the planner's incentives to identify banks that drive the propagation of shocks, and, thus, the greater the value of network transparency.

V. Conclusion

My primary goal has been to show that it is possible to define an optimal solution of the problem of regulating a network of interdependent financial institutions under uncertainty regarding its precise structure. By incorporating results from the literature on random graphs, my model makes it possible to compute optimal interventions in economies with arbitrary sizes and network architectures. When the process that generates the network architecture is also unknown, the model characterizes optimal interventions by drawing insights from the literature on decision-making under ambiguity. As the size of the economy grows large, interventions that prevent large cascades of distress can be analytically determined.

While the proposed framework does not capture the economic incentives underlying the formation of interdependencies among institutions or the reasons some of them may be more prone to propagating shocks than others, it provides a simple, yet general, approximation of the problem faced by policymakers nowadays, where the lack of detailed information and the high complexity of interactions among institutions besets the regulation and supervision of financial networks. In doing so, the proposed framework provides a benchmark to which other models can be compared to.

My emphasis on the relevance of network uncertainty should not be understood as downplaying the role that leverage, size, and short-term funding play in the design of optimal policies. As the network structure interacts with these variables, regulation should be mindful of such an interaction so as to take into consideration how financial (and non financial) institutions react to regulation and how such reactions contribute to financial stability.

Finally, network uncertainty is not only a problem for regulators as it also gives rise to uncertainty for market participants, especially in times of economic stress. In doing so, network uncertainty itself can be a source of cascades of liquidity shocks. Future research in this area is certainly called for.

REFERENCES

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, 2015, Systemic risk and stability in financial networks, *American Economic Review* 105, 564–608.

Adrian, Tobias, Daniel Covitz, and Nellie Liang, 2015, Financial stability monitoring, *Annual Review of Financial Economics* 7, 357–395.

Aldasoro, Iñaki, Domenico Delli Gatti, and Ester Faia, 2017, Bank networks: Contagion, systemic risk and prudential policy, *Journal of Economic Behavior and Organization* 142, 164 – 188.

Allen, Franklin, and Ana Babus, 2009, Networks in finance, in Paul R. Kleindorfer, and Yoram Wind, eds., *In The Network Challenge: Strategy, Profit, and Risk in an Interlinked World* (Wharton School Publishing).

Allen, Franklin, Ana Babus, and Elena Carletti, 2012, Asset commonality, debt maturity and systemic risk, *Journal of Financial Economics* 104, 519–534.

Allen, Franklin, and Douglas Gale, 2000, Financial contagion, *Journal of Political Economy* 108, 1–33.

Alvarez, Fernando, and Gadi Barlevy, 2015, Mandatory disclosure and financial contagion, Working Paper 21328, National Bureau of Economic Research.

Amini, Hamed, Rama Cont, and Andreea Minca, 2013, Resilience to contagion in financial networks, *Mathematical Finance* .

Battiston, Stefano, Domenico Delli Gatti, Mauro Gallegati, Bruce Greenwald, and Joseph Stiglitz, 2012, Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk, *Journal of Economic Dynamics and Control* 36, 1121–1141.

Beale, Nicholas, David G. Rand, Heather Battey, Karen Croxson, Robert M. May, and Martin A. Nowak, 2011, Individual versus systemic risk and the regulator's dilemma, *Proceedings of the National Academy of Sciences* 108, 12647–12652.

Caballero, Ricardo J., and Alp Simsek, 2013, Fire sales in a model of complexity, *Journal of Finance* 68, 2549–2587.

Cabrales, Antonio, Piero Gottardi, and Fernando Vega-Redondo, 2014, Risk sharing and contagion in networks, *CESifo Working Paper* .

Castiglionesi, Fabio, Fabio Feriozzi, and Guido Lorenzoni, 2019, Financial integration and liquidity crises, *Management Science* 65, 955–975.

Cohen, Reuven, Keren Erez, Daniel ben Avraham, and Shlomo Havlin, 2000, Resilience of the internet to random breakdowns, *Physical Review Letters* 85, 4626–4628.

Cohen, Reuven, Keren Erez, Daniel ben Avraham, and Shlomo Havlin, 2001, Breakdown of the internet under intentional attack, *Physical Review Letters* 86, 3682–3685.

Cont, Rama, Amal Moussa, and Edson Santos, 2013, Network structure and systemic risk in banking systems, in Jean-Pierre Fouque, and Joe Langsam, eds., *Handbook of Systemic Risk* (Cambridge University Press).

Dang, Tri Vi, Gary Gorton, Bengt Holmström, and Guillermo Ordoñez, 2017, Banks as secret keepers, *American Economic Review* 107, 1005–1029.

Dasgupta, Amil, 2004, Financial contagion through capital connections: A model of the origin and spread of bank panics, *Journal of the European Economics Association* 2, 1049–1084.

Diamond, Douglas W., and Raghuram G. Rajan, 2011, Fear of fire sales, illiquidity seeking, and credit freezes, *The Quarterly Journal of Economics* 126, 557–591.

Easley, David, and Maureen O'Hara, 2010, Liquidity and valuation in an uncertain world, *Journal of Financial Economics* 97, 1 – 11.

Eisenberg, Larry, and Thomas Noe, 2001, Systemic risk in financial systems, *Management Science* 47, 236–249.

Elliott, Matthew, Benjamin Golub, and Matthew Jackson, 2014, Financial networks and contagion, *American Economic Review* 104, 3115–3153.

Erol, Selman, and Guillermo Ordoñez, 2017, Network reactions to banking regulations, *Journal of Monetary Economics* 89, 51 – 67, Carnegie-Rochester-NYU Conference Series on the Macroeconomics of Liquidity in Capital Markets and the Corporate Sector.

Freixas, Xavier, Bruno M. Parigi, and Jean-Charles Rochet, 2000, Systemic risk, interbank relations, and liquidity provision by the central bank, *Journal of Money, Credit and Banking* 32, 611–638.

Gai, Prasanna, Andrew Haldane, and Sujit Kapadia, 2011, Complexity, concentration and contagion, *Journal of Monetary Economics* 58, 453–470.

Galeotti, Andrea, Benjamin Golub, and Sanjeev Goyal, 2018, Targeting interventions in networks, *Unpublished manuscript* .

Gennaioli, Nicola, and Andrei Shleifer, 2018, *A Crisis of Beliefs Investor Psychology and Financial Fragility* (Princeton University Press).

Georg, Co-Pierre, 2013, The effect of the interbank network structure on contagion and common shocks, *Journal of Banking and Finance* 37, 2216–2228.

Glasserman, Paul, and H. Peyton Young, 2015, How likely is contagion in financial networks?, *Journal of Banking and Finance* 50, 383–399.

Glasserman, Paul, and H. Peyton Young, 2016, Contagion in financial networks, *Journal of Economic Literature* 54, 779–831.

Gofman, Michael, 2017, Efficiency and stability of a financial architecture with too-interconnected-to-fail institutions, *Journal of Financial Economics* 124, 113–146.

Goyal, Sanjeev, and Adrien Vigier, 2014, Attack, defense, and contagion in networks, *Review of Economic Studies* 81, 1518–1542.

Haldane, Andrew G., and Robert M. May, 2011, Systemic risk in banking ecosystems, *Nature* 469.

IGM Forum, 2017, Factors contributing to the 2008 global financial crisis.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2005, A smooth model of decision making under ambiguity, *Econometrica* 73, 1849–1892.

Lagunoff, Roger, and Stacey L. Schreft, 2001, A model of financial fragility, *Journal of Economic Theory* 99, 220 – 264.

Leitner, Yaron, 2005, Financial networks: Contagion, commitment, and private sector bailouts, *Journal of Finance* 60, 2925–2953.

Maccheroni, Fabio, Massimo Marinacci, and Doriana Ruffino, 2013, Alpha as ambiguity: Robust mean-variance portfolio analysis, *Econometrica* 81, 1075–1113.

Molloy, Michael, and Bruce Reed, 1995, A critical point for random graphs with a given degree sequence, *Random Structures and Algorithms* 6, 161–180.

Molloy, Michael, and Bruce Reed, 1998, The size of the giant component of a random graph with a given degree sequence, *Combinatorics, Probability and Computing* 7, 295–305.

Newman, M. E. J., 2007, Component sizes in networks with arbitrary degree distributions, *Phys. Rev. E* 76, 045101.

Nier, Erlend, Jing Yang, Tanju Yorulmazer, and Amadeo Alentorn, 2007, Network models and financial stability, *Journal of Economic Dynamics and Control* 31, 2033–2060.

Ramírez, Carlos, 2017, Firm networks and asset returns, *Finance and Economics Discussion Series. Board of Governors of the Federal Reserve System* 14.

Rochet, Jean-Charles, and Jean Tirole, 1996, Interbank lending and systemic risk, *Journal of Money, Credit, and Banking* 28, 733–762.

Routledge, Bryan R., and Stanley E. Zin, 2009, Model uncertainty and liquidity, *Review of Economic Dynamics* 12, 543 – 566.

Ruffino, Doriana, 2014, Some implications of knightian uncertainty for finance and regulation, *FEDS Notes* .

Stein, Jeremy C., 2013, Liquidity regulation and central banking, Remarks at the "Finding the Right Balance" 2013 Credit Markets Symposium sponsored by the Federal Reserve Bank of Richmond, Charlotte, North Carolina.

Watts, Duncan J., 2002, A simple model of global cascades on random networks, *Proceedings of the National Academy of Sciences* 99, 5766–5771.

Zawadowski, Adam, 2013, Entangled financial systems, *Review of Financial Studies* 26, 1291–1323.

Zhou, Zhen, 2018, Systemic bank panics in financial networks, *Unpublished Manuscript* .

For Online Publication: Appendix for "Regulating Financial Networks Under Uncertainty"

Appendix A Mathematical Derivations

This section contains explanations of definitions and derivations of examples, lemmas, and propositions mentioned in the body of the paper.

A Understanding banks' beliefs and their behavior

The payoff of bank i is given by

$$\pi_i = \omega_i R_L + (1-\omega_i) R_I - \beta_{\omega_i} \varepsilon_i,$$

where

$$\beta_{\omega_i} \quad = \quad \begin{cases} (\omega_L R_L + (1-\omega_L) R_I) \text{ with probability } p_i \text{ if } \ \omega_i = \omega_L, \\ 0, \qquad\qquad\qquad\qquad\qquad \text{otherwise,} \end{cases}$$

with $p_i = \mathbb{P}[i \text{ faces a liquidity shock} | \omega_i = \omega_L]$. Given how liquidity shocks propagate among banks,

$$p_i(\mathcal{C}_i) = \frac{1}{n} + \left(1 - \frac{1}{n}\right) \frac{\mathbb{E}[|\mathcal{C}_i|]}{n}$$

where $\mathcal{C}_i$ denotes the set of banks (directly or indirectly) connected to bank i —via a sequence of contagious exposures at $t = t_F$—whose portfolio contains a fraction ω_L in liquid assets. From bank i's perspective,

$$\mathbb{E}_i[\pi_i | \omega_i = \omega_H] - \mathbb{E}_i[\pi_i | \omega_i = \omega_L] \quad = \quad \nu \left[\frac{1}{n} + \left(1 - \frac{1}{n}\right) \frac{\mathbb{E}_i[|\mathcal{C}_i|]}{n} \right] - \Delta\omega \mathbb{E}[\Delta R].$$

DEFINITION 3: *Bank i is said to underestimate the likelihood of being affected by cascades of liquidity shocks if $\mathbb{E}_i[|\mathcal{C}_i|] = o(n)$.*

Consequently, if bank i underestimates the likelihood of being affected by cascades of liquidity shocks, then

$$\lim_{n \to \infty} \left(\mathbb{E}_i[\pi_i | \omega_i = \omega_H] - \mathbb{E}_i[\pi_i | \omega_i = \omega_L] \right) \quad < \quad 0.$$

Namely, as n grows large, investing in illiquid assets is more lucrative than storing funds in cash from bank i's perspective.

B Welfare Effects of Regulation

Suppose the planner restricts all banks in $\mathcal{R}$, with $|\mathcal{R}| = nx$. Then

$$\left(\frac{1}{n}\right)\mathbb{E}[\mathrm{TO}|x] \quad = \quad x(\mathbb{E}[R_I] - \omega_H \mathbb{E}[\Delta R]) + \left(\frac{1}{n}\right)\sum_{i\notin\mathcal{R}} \mathbb{E}[\pi_i]$$

It is worth noting that

$$\left(\frac{1}{n}\right)\sum_{i\notin\mathcal{R}} \pi_i \quad = \quad \begin{cases} (1-x)[R_I - \omega_L \Delta R], \text{ with probability } x \\ \left(1 - x - \frac{m}{n}\right)[R_I - \omega_L \Delta R], \text{ with probability } (1-x)\phi_m^x \text{ with } m = 1, \cdots, n(1-x). \end{cases}$$

where ϕ_m^x denotes the probability that m nonrestricted banks are affected by liquidity shocks, once $n \times x$ banks have been restricted. Consequently,

$$\begin{aligned}
\left(\frac{1}{n}\right)\mathbb{E}[\mathrm{TO}|x] \quad &= \quad x(\mathbb{E}[R_I] - \omega_H \mathbb{E}[\Delta R]) \\
&+ \quad (\mathbb{E}[R_I] - \omega_L \mathbb{E}[\Delta R]) \left(x(1-x) + (1-x)^2 \underbrace{\left(\sum_{m=1}^{n(1-x)} \phi_m^x\right)}_{=1} - \frac{(1-x)}{n} \underbrace{\sum_{m=1}^{n(1-x)} m\phi_m^x}_{=\langle\phi^x\rangle} \right) \\
&= \quad x(\mathbb{E}[R_I] - \omega_H \mathbb{E}[\Delta R]) + (1-x)(\mathbb{E}[R_I] - \omega_L \mathbb{E}[\Delta R]) \left(1 - \frac{\langle\phi^x\rangle}{n}\right) + \\
&= \quad (\mathbb{E}[R_I] - \omega_L \mathbb{E}[\Delta R]) - (1-x)\left(\mathbb{E}[R_I] - \omega_L \mathbb{E}[\Delta R]\right)\frac{\langle\phi^x\rangle}{n} - x\Delta\omega\mathbb{E}[\Delta R] \\
&= \quad \nu - \nu(1-x)\frac{\langle\phi^x\rangle}{n} - x\Delta\omega\mathbb{E}[\Delta R]
\end{aligned}$$

where $\langle\phi^x\rangle$ denotes the expected number of nonrestricted banks affected by liquidity shocks once $n \times x$ banks have been restricted.

Proof of Lemma 1. There are two cases.

(a) $\mathcal{R} = \mathcal{R}_{x_r}$. Here, banks are ex ante identical from the point of view of the planner, as she is unable to identify whether some banks will exhibit more contagious exposures than others. Hence, when solving her problem, the planner effectively acts as if she restricts banks uniformly at random.
It is then illustrative to explore the distribution of contagious exposures among non-restricted banks after imposing restrictions on a fraction $\frac{r}{n}$. Consider a bank with k_0 contagious exposures. After imposing restrictions, that bank may have k contagious exposures, with $k \leq k_0$, as some of its neighbors may be restricted. Additionally, the probability that a subset of k neighbors is not restricted is $\left(1 - \frac{r}{n}\right)^k$, whereas the probability that the remaining neighbors are restricted is $\left(\frac{r}{n}\right)^{k_0-k}$. Because there are $\binom{k_0}{k}$ different subsets of k neighbors, the distribution of contagious exposures among

nonrestricted banks is

$$\theta_k^r = \begin{cases} \sum_{j=k}^{n-1} p_j \binom{j}{k} \left(1-\frac{r}{n}\right)^k \left(\frac{r}{n}\right)^{j-k} & \text{if } k=\{0,\cdots,n-r-1\} \\ 0 & \text{otherwise.} \end{cases} \tag{A1}$$

In other words, θ_k^r captures the probability that a nonrestricted bank shares k contagious exposures with other nonrestricted banks, once r banks are restricted.

(b) $\mathcal{R} = \mathcal{R}_{x_t}$. Here, the planner is able to identify which banks will exhibit the highest number of contagious exposures. Then, she can use that information and restrict such banks first to prevent contagion more efficiently than restricting at random, as such intervention results in the removal of a larger fraction of contagious exposures.
Suppose the planner imposes restrictions on all banks with more than k_x contagious exposures, with $k_x \geq s^*$. Restricting those banks is equivalent to restricting a fraction x of banks with the highest number of contagious exposures. The relationship between x and k_x is given by

$$x = \sum_{k_x<k} p_k \implies x = 1 - \sum_{k=0}^{k_x} p_k. \tag{A2}$$

Implementing the above policy results in an approximate random removal of contagious exposures from nonrestricted banks, as contagious exposures of restricted banks no longer propagate liquidity shocks. The probability κ_x that a contagious exposure leads to a restricted bank equals

$$\kappa_x = \sum_{k_x<k} \frac{kp_k}{\sum_k kp_k} = \frac{1}{\langle k \rangle}\left(\sum_{k_x<k} kp_k\right) \implies \kappa_x = 1 - \frac{1}{\langle k \rangle}\left(\sum_{k=0}^{k_x} kp_k\right). \tag{A3}$$

It is important to note that the network of contagious exposures that remains after implementing the above policy is equivalent to a network in which the maximum number of contagious exposures per bank is k_x and a fraction κ_x of banks is restricted uniformly at random. It follows from the previous analysis that the probability that a nonrestricted bank has k contagious exposures once a fraction x of banks has been restricted, φ_k^x, is given by

$$\varphi_k^x = \begin{cases} \sum_{j=k}^{k_x} p_j \binom{j}{k} (1-\kappa_x)^k \kappa_x^{j-k} & \text{if } k=\{0,\cdots,k_x\} \\ 0 & \text{otherwise.} \end{cases} \tag{A4}$$

□

Proof of Lemma 2. It is worth noting that $\{\theta_k^r\}_{k=0}^{n-r-1}$ represents the degree distribution of a randomly generated network among $(n-r)$ banks. Consequently, ϕ_m^r is equivalent to the probability that a randomly chosen nonrestricted bank belongs to a connected subgraph of size m. Therefore, one can directly apply results in Newman (2007) and show that

$$\phi_m^r = \begin{cases} \frac{\langle\theta^r\rangle}{(m-1)!}\left(\frac{d^{m-2}}{dz^{m-2}}\left[g\left(z,\{\theta_k^r\}_k\right)^m\right]\right)\Big|_{z=0} & \text{, with } m=\{2,\cdots,n-r\} \\ \theta_0^r & \text{, with } m=1. \end{cases}$$

where $\langle\theta^r\rangle$ denotes the average number of contagious exposures among nonrestricted banks and $\left(\frac{d^{m-2}}{dz^{m-2}}\left[g\left(z,\{\theta_k^r\}_k\right)^m\right]\right)\Big|_{z=0}$ denotes the $(m-2)$ derivative of $g\left(z,\{\theta_k^r\}_k\right)^m$ evaluated at $z=0$, where $g(z,\{\theta_k^r\}_k)$ represents the excess degree distribution function of $\{\theta_k^r\}_k$, defined as

$$g(z,\{\theta_k^r\}_k) \equiv \sum_{k=0}^{n-r-2}\left(\frac{(k+1)\theta_{k+1}^r}{\langle\theta^r\rangle}\right)z^k.$$

□

REMARK 1 (Numerical solutions): *When solving the model numerically, the* $(m-2)$ *derivative of* $g\left(z,\{\theta_k^r\}_k\right)^m$ *evaluated at* $z=0$*, can be approximated by*

$$\frac{1}{\epsilon^{m-2}}\left[\sum_{j=0}^{m-2}(-1)^{m-2-j}\binom{m-2}{j}g\left(j\epsilon,\{\theta_k^r\}_k\right)^m\right]$$

with $\epsilon>0$ *sufficiently small.*

Proof of Example 1. There are two cases.

(a) $\mathcal{R}=\mathcal{R}_{x_r}$. When $\{p_k\}_{k=0}^{n-1}$ follows a Poisson distribution with parameter α, $\{\theta_k^r\}_{k=0}^{n-r-1}$ approximately follows a Poisson distribution of parameter $\left(1-\frac{r}{n}\right)\alpha$. As a result,

$$g(z,\{\theta_k^r\}_{k=0}^{n-r-1}) = e^{\left(1-\frac{r}{n}\right)\alpha(z-1)},$$

and, thus,

$$\phi_m^r = \frac{e^{-\left(1-\frac{r}{n}\right)\alpha m}\left(\left(1-\frac{r}{n}\right)\alpha m\right)^{m-1}}{m!}, \ m=\{1,\cdots,n-r\}.$$

(b) $\mathcal{R}=\mathcal{R}_{x_t}$. Here, $p_k=e^{-\alpha}\frac{\alpha^k}{k!}$, with $k=\{0,n-1\}$. The result follows directly from (a) after substituting p_k into

$$\kappa_x = \sum_{k_x<k}\frac{kp_k}{\sum_k kp_k} = \frac{1}{\langle k\rangle}\left(\sum_{k_x<k}kp_k\right) \implies \kappa_x = 1-\frac{1}{\langle k\rangle}\left(\sum_{k=0}^{k_x}kp_k\right).$$

□

C Optimal Interventions

Proof of Proposition 1. For a given distribution $\{p_k\}_k$, $\mathbb{E}\left[\left(\frac{1}{n}\right)\mathrm{TO}\middle|\mathcal{R}\right]$ is bounded from above because $0 \leq \frac{\langle\phi^{\mathcal{R}}\rangle}{n} \leq \left(1-\frac{|\mathcal{R}|}{n}\right)$ for any $\mathcal{R} \in \mathcal{P}(B_n)$. Thus, there exists a set $\mathcal{R} \in \mathcal{P}(B_n)$ where $\mathbb{E}\left[\left(\frac{1}{n}\right)\mathrm{TO}\middle|\mathcal{R}\right]$ attains its maximum. □

Proof of Proposition 2. For a given distribution $\{p_k\}_k$, proposition 1 ensures the planner's problem has a solution. Now suppose the planner only focuses on simple interventions within a size differentiable collection, $\mathcal{C}_R(\{p_k\}_k)$. Additionally, for any three elements of $\mathcal{C}_R$, say $\mathcal{R}_z$, $\mathcal{R}_{x_0}$, and $\mathcal{R}_{x_1}$,

$$\begin{aligned}
\left(\frac{1}{n}\right)\mathbb{E}[\mathrm{TO}|z] &= \eta-(1-z)\left(\frac{\nu}{n}\right)\langle\phi^z\rangle - z\Delta\omega\mathbb{E}[\Delta R] \\
&= (1-\alpha)\left(\eta-(1-x_0)\left(\frac{\nu}{n}\right)\langle\phi^z\rangle - x_0\Delta\omega\mathbb{E}[\Delta R]\right) \\
&+ \alpha\left(\eta-(1-x_1)\left(\frac{\nu}{n}\right)\langle\phi^z\rangle - x_1\Delta\omega\mathbb{E}[\Delta R]\right) \\
&= (1-\alpha)\left(\eta-(1-x_0)\left(\frac{\nu}{n}\right)\langle\phi^{x_0}\rangle - x_0\Delta\omega\mathbb{E}[\Delta R]\right) \\
&+ \alpha\left(\eta-(1-x_1)\left(\frac{\nu}{n}\right)\langle\phi^{x_1}\rangle - x_1\Delta\omega\mathbb{E}[\Delta R]\right) \\
&+ \left(\frac{\nu}{n}\right)\left((1-\alpha)(1-x_0)\langle\phi^{x_0}\rangle+\alpha(1-x_1)\langle\phi^{x_1}\rangle-[(1-\alpha)x_0+\alpha(1-x_1)]\langle\phi^z\rangle\right) \\
&> \left(\frac{1}{n}\right)\left((1-\alpha)\mathbb{E}[\mathrm{TO}|x_0]+\alpha\mathbb{E}[\mathrm{TO}|x_1]\right)
\end{aligned}$$

where the last inequality holds if and only if

$$\frac{(1-\alpha)(1-x_0)\langle\phi^{x_0}\rangle+\alpha(1-x_1)\langle\phi^{x_1}\rangle}{(1-\alpha)(1-x_0)+\alpha(1-x_1)} \quad > \quad \langle\phi^z\rangle.$$

As a result, $\left(\frac{1}{n}\right)\mathbb{E}[\mathrm{TO}]$ is a strictly concave function of x, and thus, x^* approximately satisfies

$$\frac{\partial}{\partial x}\left(\left(\frac{1}{n}\right)\mathbb{E}[\mathrm{TO}]\right)\bigg|_{x=x^*} \quad = \quad 0 \tag{A5}$$

as n grows large.[7] Equation (A5) can be rewritten as

$$\nu\left(\frac{\langle\phi^{x^*}\rangle}{n}-(1-x^*)\frac{\partial}{\partial x}\left(\frac{\langle\phi^x\rangle}{n}\right)\bigg|_{x=x^*}\right) \quad = \quad \Delta\omega\mathbb{E}[\Delta R]$$

□

[7]The fraction of restricted banks must be a rational number because n is a natural number. However, the solution of equation (A5) could be an irrational number. Nonetheless, the optimal intervention x^* gets arbitrarily close to the solution of equation (A5) as n grows large.

Proof of Proposition 3. For any distribution $\{p_k\}_k$ and finite barycenter $\bar{\alpha}$, $\mathbb{E}_{\bar{\alpha}}\left[\frac{1}{n}\mathrm{TO}\right]$ is bounded from above. Because $\mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_\alpha\left(\mathrm{TO}\middle|\mathcal{R}_x\right)\right) \geq 0$ for any $\mathcal{R}_x \in \mathcal{P}(B_n)$ and distribution f, the objective function of the planner's problem is bounded from above. Hence, it attains its maximum for certain set $\mathcal{R} \in \mathcal{P}(B_n)$. $\square$

Mirroring the ideas in proposition 2, the next proposition shows that, under certain conditions, I can uniquely determine the size of the optimal set of restricted banks when $\{p_k\}_k$ is unknown. To pose these conditions in a sharper way, I first introduce a definition

DEFINITION 4: *Given distributions $\{p_k\}_k$ and f, a collection of simple interventions $\mathcal{C}_R\left(\{p_k\}_k, f\right)$ is said to be size differentiable under ambiguity if the mapping $h : \mathcal{C}_R \to [0,1]$, defined as $h(\mathcal{R}_x) \equiv \frac{\langle\phi^{\mathcal{R}_x}\rangle}{n} - \left(\frac{\theta}{2}\right) \times \mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_\alpha\left(TO\middle|\mathcal{R}_x\right)\right)$ is differentiable for all x, $0 \leq x \leq 1$.*

PROPOSITION 4 (Interior Solution): *Take the choice of network transparency as given. Given distributions $\{p_k\}_k$ and f, suppose the planner only focuses on simple interventions within a collection that is size differentiable under ambiguity, $\mathcal{C}_R(\{p_k\}_k, f)$. Assume further that*

$$\mathbb{E}_{\bar{\alpha}}\left(\frac{1}{n}TO\middle|\mathcal{R}_x\right) - \left(\frac{\theta}{2}\right) \times \mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_\alpha\left(TO\middle|\mathcal{R}_x\right)\right)$$

is a strictly concave function of x, for any set $\mathcal{R}_x \in \mathcal{C}_R(\{p_k\}_k, f)$. Then, the relative size of the optimal intervention, $\mathcal{R}_{x^a} \in \mathcal{C}_R(\{p_k\}_k, f)$, approximately solves

$$\begin{aligned}\nu\left(\frac{\langle\phi_{\bar{\alpha}}^{x^a}\rangle}{n} - (1-x^a)\frac{\partial}{\partial x}\left(\frac{\langle\phi_{\bar{\alpha}}^{x}\rangle}{n}\right)\Bigg|_{x=x^a}\right) &= \Delta\omega\mathbb{E}[\Delta R] \\ &+ \left(\frac{\theta}{2}\right)\nu^2\frac{\partial}{\partial x}\left((1-x)^2\int_{\alpha\in\mathcal{A}}\left(\frac{\langle\phi_{\alpha}^{x}\rangle}{n} - \frac{\langle\phi_{\bar{\alpha}}^{x}\rangle}{n}\right)^2 df(\alpha)\right)\Bigg|_{x=x^a}.\end{aligned}$$

as n grows large.

Proof. If

$$\mathbb{E}_{\bar{\alpha}}\left(\frac{1}{n}\mathrm{TO}\middle|\mathcal{R}_x\right) - \left(\frac{\theta}{2}\right) \times \mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_\alpha\left(\mathrm{TO}\middle|\mathcal{R}_x\right)\right)$$

is a strictly concave function of x, for all $\mathcal{R}_x \in \mathcal{C}_R(\{p_k\}_k, f)$, then the solution of the planner's problem is interior. Notably,

$$\begin{aligned}\mathbb{E}_{\bar{\alpha}}\left(\frac{1}{n}\mathrm{TO}_\alpha\middle|\mathcal{R}_x\right) &= \nu - (1-x)\nu\frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_x}\rangle}{n} - x\Delta\omega\mathbb{E}(\Delta R) \\ \mathbb{V}_f\left(\frac{1}{n}\mathbb{E}_\alpha\left(\mathrm{TO}\middle|\mathcal{R}_x\right)\right) &= (1-x)^2\nu^2\int_{\alpha\in\mathcal{A}}\left(\frac{\langle\phi_{\alpha}^{\mathcal{R}_x}\rangle}{n} - \frac{\langle\phi_{\bar{\alpha}}^{\mathcal{R}_x}\rangle}{n}\right)^2 df(\alpha).\end{aligned}$$

As a result, the first order condition of the planner's problem can be rewritten as

$$\begin{aligned}\nu\left(\frac{\langle\phi_{\bar{\alpha}}^{x^a}\rangle}{n}-(1-x^a)\frac{\partial}{\partial x}\left(\frac{\langle\phi_{\bar{\alpha}}^{x}\rangle}{n}\right)\Bigg|_{x=x^a}\right) &= \Delta\omega\mathbb{E}[\Delta R] \\ &+ \left(\frac{\theta}{2}\right)\nu^2\frac{\partial}{\partial x}\left((1-x)^2\int_{\alpha\in\mathcal{A}}\left(\frac{\langle\phi_{\alpha}^{x}\rangle}{n}-\frac{\langle\phi_{\bar{\alpha}}^{x}\rangle}{n}\right)^2 df(\alpha)\right)\Bigg|_{x=x^a}.\end{aligned}$$

Because the fraction of restricted banks must be a rational number, x^* gets arbitrarily close to the solution of the above equation as n gets large. □

D *Optimal Interventions in Large Economies*

When designing optimal interventions, it is pivotal to understand under which conditions cascades of liquidity shocks are non-negligible. As I focus on a system of infinite size, any cascade of finite size becomes negligible as the economy grows large. The following section provides a detailed definition and analysis of cascades of liquidity shocks in large economies.

D.1 Large Cascades of Liquidity Shocks

The Rise of Large Cascades of Liquidity Shocks. To fix notation, let $\mathcal{G}_n$ denote a network of contagious exposures among n banks and $\{\mathcal{G}_n\}_{n\in\mathbb{N}}$ denote a sequence of such networks, indexed by the number of banks n. Let $\mathcal{S}(\mathcal{G}_n)$ denote the largest subset of connected banks in $\mathcal{G}_n$, and let $|\mathcal{S}(\mathcal{G}_n)|$ denote the cardinality of such a set. To determine the condition under which large cascades of liquidity shocks arise, one can use the following idea, similar to the one proposed by Molloy and Reed (1995) and Cohen et al. (2000). Let n_0 denote a large natural number. Suppose there are two banks belonging to each element in the subsequence $\{\mathcal{S}(\mathcal{G}_n)\}_{n\geq n_0}$—say, i and j, which are directly connected. If bank i (or j) is also directly connected to another bank—and loops of contagious exposures can be ignored—then the size of the largest sequence of connected banks is proportional to the size of the system—i.e., $\lim_{n\to\infty}\frac{\mathbb{E}|\mathcal{S}(\mathcal{G}_n)|}{n}>0$—and, thus, large cascades of liquidity shocks occur; otherwise, the largest sequence of connected banks is fragmented, and, thus, $\lim_{n\to\infty}\frac{\mathbb{E}|\mathcal{S}(\mathcal{G}_n)|}{n}=0$.[8] Therefore, the condition that determines the emergence of large cascades of liquidity shocks is given by

$$\lim_{n\to\infty}\mathbb{E}_n\left[k_i|i\leftrightarrow j\right] \;=\; \lim_{n\to\infty}\sum_{k_i}k_i\mathbb{P}_n\left[k_i|i\leftrightarrow j\right] \;\leq\; 2, \tag{A6}$$

where $\mathbb{P}_n\left[k_i|i\leftrightarrow j\right]$ denotes the probability that bank i has k_i contagious exposures, given that i and j are connected via one contagious exposure. It follows from Bayes' rule that

$$\mathbb{P}_n\left[k_i|i\leftrightarrow j\right] \;=\; \frac{\mathbb{P}_n\left[i\leftrightarrow j|k_i\right]\mathbb{P}_n\left[k_i\right]}{\mathbb{P}_n\left[i\leftrightarrow j\right]}.$$

[8]As n grows large, loops of contagious exposures can be ignored for $\frac{\mathbb{E}_n(k^2)}{\mathbb{E}_n(k)}<2$. For more details, see Cohen et al. (2000).

Because contagious exposures are randomly determined,

$$\mathbb{P}_n\left[i \leftrightarrow j\right] = \frac{\mathbb{E}_n[k]}{n-1} \quad \text{and} \quad \mathbb{P}_n\left[i \leftrightarrow j | k_i\right] = \frac{k_i}{n-1}.$$

Thus, equation (A6) is equivalent to

$$\lim_{n\to\infty} \frac{\mathbb{E}_n[k^2]}{\mathbb{E}_n[k]} \;\; \leq \;\; 2, \tag{A7}$$

It is important to note that the derivation of equation (A7) does not rely on the functional form of $\mathbb{P}_n[k]$ and applies to any distribution of links in which banks are randomly connected to each other. Equation (A7) establishes that if, in the limit, there is enough variation in the number of contagious exposures among banks, liquidity shocks affecting one bank almost surely affects a non-negligible fraction of them. High variation in the number of contagious exposures makes the economy more prone to contagion, as banks with a large number of contagious exposures can effectively reach a large fraction of banks.

Preventing large cascades of liquidity shocks. Because restricting a bank not only precludes that bank from facing liquidity shocks, but also precludes that bank from propagating liquidity shocks, restricting a sufficiently large fraction of banks can potentially prevent the emergence of large cascades of liquidity shocks. When x exceeds a certain threshold, x^*, large cascades of liquidity shocks can be prevented, as the network of contagious exposures disintegrates into smaller and disconnected parts, keeping liquidity shocks locally confined. Importantly, the value of x^* critically depends on how restricted bank are selected, as the ratio in (A7) varies across policies.

First, suppose a fraction x of bank are restricted uniformly at random. After imposing restrictions, a bank with k_0 contagious exposures may only have k contagious exposures, with $k \leq k_0$, as some of its neighbors may be restricted. In addition, the probability that a subset of k neighbors is not restricted is $(1-x)^k$, whereas the probability that the remaining neighbors are restricted is x^{k_0-k}. Because there are $\binom{k_0}{k}$ different subsets of k neighbors, the distribution of contagious exposures among nonrestricted banks is

$$\mathbb{P}'_n\left(k\right) \;\; = \;\; \sum_{k\geq k_0} p^{\alpha}_{k_0} \binom{k_0}{k} (1-x)^k x^{k_0-k},$$

and, thus,

$$\mathbb{E}'_n[k] = \langle k \rangle (1-x) \quad \text{and} \quad \mathbb{E}'_n[k^2] = \langle k^2 \rangle (1-x)^2 + \langle k \rangle x(1-x), \tag{A8}$$

where expectations with superscript prime denote expectations after implementing restrictions. After banks are restricted, large cascades of liquidity shocks arise if and only if

$$\lim_{n\to\infty} \frac{\mathbb{E}'_n[k^2]}{\mathbb{E}'_n[k]} \;\; \leq \;\; 2. \tag{A9}$$

It then directly follows from substituting equation (A8) into equation (A9) that x^* must satisfy

$$x^* = 1 - \frac{\langle k \rangle}{\langle k^2 \rangle - \langle k \rangle}.$$

Now, suppose banks with the highest number of contagious exposures are restricted. The following computations closely follow the ideas in Cohen et al. (2001). Restricting banks with more than $K(x^*)$ contagious exposures is approximately equivalent to restricting a fraction x^* of banks, where x^* satisfies

$$x^* = 1 - \sum_{k=0}^{K(x^*)} p_k^\alpha.$$

Take a bank with k contagious exposures. The fraction of contagious exposures attached to all banks with k contagious exposures equals $\frac{kp_k^\alpha}{\langle k \rangle}$. As a consequence, the fraction of contagious exposures attached to restricted banks is

$$s(x^*) \;=\; \frac{1}{\langle k \rangle} \left(\sum_{k=K(x^*)+1}^{n-1} kp_k^\alpha \right) \;=\; 1 - \frac{1}{\langle k \rangle} \left(\sum_{k=0}^{K(x^*)} kp_k^\alpha \right)$$

Because imposing restrictions on a set of banks can be represented by the removal of such banks and their exposures, the optimal policy x^* must satisfy

$$s(x^*) \;=\; x^* \;=\; 1 - \frac{\langle k(x^*) \rangle}{\langle k(x^*)^2 \rangle - \langle k(x^*) \rangle}.$$

Therefore,

$$\begin{aligned}
1 - s(x^*) &= \frac{\langle k(x^*) \rangle}{\langle k(x^*)^2 \rangle - \langle k(x^*) \rangle} \\
\frac{1}{\langle k \rangle} \left(\sum_{k=0}^{K(x^*)} kp_k^\alpha \right) &= \frac{\sum_{k=0}^{K(x^*)} kp_k^\alpha}{\sum_{k=0}^{K(x^*)} k^2 p_k^\alpha - \sum_{k=0}^{K(x^*)} kp_k^\alpha} \\
\frac{1}{\langle k \rangle} \left(\sum_{k=0}^{K(x^*)} kp_k^\alpha \right) &= \frac{\sum_{k=0}^{K(x^*)} kp_k^\alpha}{\sum_{k=0}^{K(x^*)} k(k-1) p_k^\alpha} \\
\langle k \rangle &= \sum_{k=0}^{K(x^*)} k(k-1) p_k^\alpha. \qquad \text{(A10)}
\end{aligned}$$

which determines the condition under which large cascades of liquidity shocks emerge.

Preventing large cascades in Poisson networks. If the planner cannot identify whether some banks will exhibit more contagious exposures than others, then

$$x^* \;=\; 1 - \left(\frac{1}{\alpha}\right).$$

However, if the planner is able to identify which banks will exhibit the highest number of contagious exposures, then

$$x^* \;=\; 1 - e^{-\alpha}\left(\sum_{k=0}^{K(x^*)} \frac{\alpha^k}{k!}\right), \text{ where } K(x^*) \text{ satisfies } \alpha \;=\; e^{-\alpha}\left(\sum_{k=2}^{K(x^*)} \frac{\alpha^k}{(k-2)!}\right).$$

The derivation of the above equations uses the following arguments. For a Poisson distribution with parameter α, the first two moments are given by $\langle k\rangle = \alpha$ and $\langle k^2\rangle = \alpha^2 + \alpha$. Thus, when restricting at random, the optimal policy is given by $x^* = 1 - \frac{1}{\alpha}$. Provided that $p_k^\alpha = e^{-\alpha}\frac{\alpha^k}{k!}$, a direct application of condition (A10) yields

$$x^* \;=\; 1 - e^{-\alpha}\left(\sum_{k=0}^{K(x^*)} \frac{\alpha^k}{k!}\right), \text{ where } K(x^*) \text{ satisfies } \alpha \;=\; e^{-\alpha}\left(\sum_{k=2}^{K(x^*)} \frac{\alpha^k}{(k-2)!}\right).$$

Preventing large cascades in Power-law networks. If the planner cannot identify whether some banks will exhibit more contagious exposures than others, then

$$x^* \;=\; \begin{cases} 1 - \left(\left(\frac{2-\alpha}{3-\alpha}\right)k_0 - 1\right)^{-1} & \text{if } \alpha > 3 \\ 1 & \text{if } 1 \leq \alpha \leq 3. \end{cases}$$

However, if the planner is able to identify which banks will exhibit the highest number of contagious exposures, then

$$x^* \;=\; 1 - \sum_{k=0}^{K(x^*)} k^{-\alpha}.$$

where $K(x^*)$ satisfies

$$\left(\frac{K(x^*)}{k_0}\right)^{2-\alpha} - 2 \;=\; \left(\frac{2-\alpha}{3-\alpha}\right)k_0\left(\left(\frac{K(x^*)}{k_0}\right)^{3-\alpha} - 1\right).$$

The derivation of the above equations uses the following arguments. A continuous Power-law distribution with parameter α, minimal value k_0, and maximum value K, satisfies

$$\langle k\rangle = k_0^{\alpha-1}K^{2-\alpha}\left(\frac{\alpha-1}{\alpha-2}\right) \text{ and } \langle k^2\rangle \;=\; k_0^{\alpha-1}K^{3-\alpha}\left(\frac{\alpha-1}{\alpha-3}\right) \quad \text{if } 1<\alpha<2$$

$$\langle k\rangle = k_0\left(\frac{\alpha-1}{\alpha-2}\right) \text{ and } \langle k^2\rangle \;=\; k_0^{\alpha-1}K^{3-\alpha}\left(\frac{\alpha-1}{\alpha-3}\right) \quad \text{if } 2<\alpha<3$$

$$\langle k \rangle = k_0 \left(\frac{\alpha - 1}{\alpha - 2} \right) \quad \text{and} \quad \langle k^2 \rangle \quad = \quad k_0^2 \left(\frac{\alpha - 1}{\alpha - 3} \right) \qquad \text{if } 3 < \alpha$$

As a consequence, when K grows large,

$$\langle k \rangle = k_0 \left(\frac{\alpha - 1}{\alpha - 2} \right) \text{ if } \alpha > 2 \quad \text{and} \quad \langle k^2 \rangle = k_0^2 \left(\frac{\alpha - 1}{\alpha - 2} \right) \text{ if } \alpha > 3$$

and they diverge in all other cases.

Now, consider the case when the planner cannot differentiate among banks before implementing her policy. Using the above equations and the condition that determines the emergence of large cascades of liquidity shocks, it is easy to show that

$$x^* \quad = \quad \begin{cases} 1 - \left(\left(\frac{2-\alpha}{3-\alpha} \right) k_0 - 1 \right)^{-1} & \text{if } \alpha > 3 \\ 1 & \text{if } 1 \leq \alpha \leq 3. \end{cases}$$

as n grows large.

When the planner can identify banks with the highest number of contagious exposures, the following equation

$$\sum_{k=k_0}^{k_x} k(k-1)p_k \quad = \quad \langle k \rangle$$

determines the emergence of large cascades of liquidity shocks. Because the network follows a Power-law distribution with parameter α, the above equation is equivalent to

$$(\alpha - 1)k_0^{\alpha - 1} \left(\frac{k_x^{3-\alpha} - k_0^{3-\alpha}}{3 - \alpha} - \frac{k_x^{2-\alpha} - k_0^{2-\alpha}}{2 - \alpha} \right) \quad = \quad k_0 \left(\frac{\alpha - 1}{\alpha - 2} \right)$$

which is equivalent to

$$\left(\frac{k_0}{3 - \alpha} \right) \left(\left(\frac{k_x}{k_0} \right)^{3-\alpha} - 1 \right) - \left(\frac{1}{2 - \alpha} \right) \left(\left(\frac{k_x}{k_0} \right)^{2-\alpha} - 2 \right) \quad = \quad 0,$$

and, thus, k_x can be derived from α.

D.2 Interventions in Large Economies

Suppose the planner restricts all banks in $\mathcal{R}$, with $|\mathcal{R}| = nx$. Then

$$\left(\frac{1}{n} \right) \mathbb{E}[\text{TO}|x] \quad = \quad x(\mathbb{E}[R_I] - \omega_H \mathbb{E}[\Delta R]) + \left(\frac{1}{n} \right) \sum_{i \notin \mathcal{R}} \mathbb{E}[\pi_i]$$

For a given information set, let x^* denote the smallest fraction of banks that must be restricted to prevent the emergence of large cascades of liquidity shocks. As the economy grows large, it is worth noting that

$$\lim_{n\to\infty}\left(\frac{1}{n}\right)\sum_{i\notin\mathcal{R}}\pi_i = \begin{cases}(1-x)[R_I-\omega_L\Delta R], \text{ with probability } x \text{ if } x<x^* \\ 0, \text{ with probability } (1-x) \text{ if } x<x^* \\ (1-x)[R_I-\omega_L\Delta R], \text{ with probability } 1 \text{ if } x\geq x^*.\end{cases}$$

Let $\mathcal{R}^c$ denote the complement set of $\mathcal{R}$. The above expression follows from the fact that if $x\geq x^*$, then the size of the largest connected component in $\mathcal{R}^c$ is almost surely of order $\rho^2\log(n)$, where ρ is the highest degree within $\mathcal{R}^c$. However, if $x<x^*$ the size of the largest connected component is of order n—and the size of the second largest connected component is of order $\log(n)$; for more details, see Molloy and Reed (1998). As a result,

$$\lim_{n\to\infty}\left(\frac{1}{n}\right)\mathbb{E}[\text{TO}|x] = \begin{cases}x(\mathbb{E}[R_I]-\omega_H\mathbb{E}[\Delta R])+(1-x)(\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]), \text{ if } x\geq x^* \\ x(\mathbb{E}[R_I]-\omega_H\mathbb{E}[\Delta R])+x(1-x)(\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]), \text{ if } x<x^*.\end{cases}$$

Define $\Delta x=(x^*-x)$. To determine the optimal policy, it is worth noting that $H(x)\equiv \lim_{n\to\infty}\left(\frac{1}{n}\right)(\mathbb{E}[\text{TO}|x^*]-\mathbb{E}[\text{TO}|x])$ equals

$$H(x) = \begin{cases}-\Delta x\Delta\omega\mathbb{E}[\Delta R], \text{ if } x\geq x^* \\ (1-x)^2(\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R])-\Delta\omega\mathbb{E}[\Delta R]\Delta x, \text{ if } x<x^*.\end{cases}$$

If $\left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right)\geq x^*$, then $H(x)\geq 0$, $\forall\ x$ in $[0,x^*]$. Thus, x^* generates higher expected total output than any other fraction $0\leq x\leq 1$, as $H(x)$ is strictly positive when $x>x^*$.

However, if $\left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right)<x^*$, then $H(0)<0$. Consequently, $H(x)<0$, when $0\leq x<x^*$, as $H(x)$ is an increasing function of x when $x<x^*$ and $H(x^*)=0$. Therefore, $x=0$ maximizes expected total output. Then, as n grows large, the optimal policy converges to the following intervention:

$$x^{\text{optimal}} = \begin{cases}x^*, \text{ if } \left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right)\geq x^* \\ 0, \qquad \text{otherwise.}\end{cases}$$

Importantly, the value of x^* does not depend on the values of $\mathbb{E}[R_I]$, ω_L, $\mathbb{E}[\Delta R]$, or $\Delta\omega$. However, x^* does depend on the distribution $\{p_k^\alpha\}_{k=1}$ and how banks are targeted. In particular, if bank-level information is not acquired, then $x^*=1-\frac{\langle k\rangle}{\langle k^2\rangle-\langle k\rangle}$. If bank-level information is acquired, then $x^*=1-\sum_{k=k_{\min}}^{K(x^*)}p_k^\alpha$—where $K(x^*)$ is the solution of the following equation $\langle k\rangle=\sum_{k=k_{\min}}^{K(x^*)}k(k-1)p_k^\alpha$ and $k_{\min}$ is the smallest number of links that a bank might have.

As in the paper, x_t and x_r denote the smallest fraction of banks that must be restricted to prevent large cascades of liquidity shocks if either the planner strategically targets banks or restricts them at random. Because $x_t\leq x_r$, then $\Delta x\geq 0$.

EXAMPLE 2 (Poisson Networks): *If the network exhibits a Poisson degree distribution of* *parameter α, then*

$$\begin{aligned} x_r &= 1-\frac{1}{\alpha} \\ x_t &= x_r-\frac{e^{-\alpha}\alpha^{K_\alpha}}{K_\alpha!} \end{aligned}$$

where K_α solves $\frac{1}{\alpha}=\sum_{j=0}^{(K_\alpha-2)}\frac{e^{-\alpha}\alpha^j}{j!}$. Thus, $\Delta x=\frac{e^{-\alpha}\alpha^{K_\alpha}}{K_\alpha!}$.

Proof. With random targeting, $x_r=1-\frac{\langle k\rangle}{\langle k^2\rangle-\langle k\rangle}$. The result follows directly from the fact a Poisson network with parameter α yields $\langle k\rangle=\alpha$ and $\langle k^2\rangle=\alpha^2+\alpha$. With strategic targeting, $x_t=1-\sum_{k=k_{\min}}^{K^*}p_k^\alpha=\sum_{k=K^*+1}^{\infty}p_k^\alpha$—where K^* is the solution of the equation $\langle k\rangle=\sum_{k=k_{\min}}^{K^*}k(k-1)p_k^\alpha$, with $p_k^\alpha=\frac{e^{-\alpha}\alpha^k}{k!}$. It is worth noting that the fraction of contagious exposures attached to restricted banks, p, equals $p=\frac{1}{\langle k\rangle}\sum_{k=K^*+1}^{\infty}kp_k^\alpha$. Importantly, if $p=x_r$, then large cascades of liquidity shocks are prevented. Because $p=\frac{1}{\langle k\rangle}\sum_{k=K^*+1}^{\infty}kp_k^\alpha=\sum_{k=K^*}^{\infty}\frac{e^{-\alpha}\alpha^k}{k!}=(\sum_{k=K^*+1}^{\infty}\frac{e^{-\alpha}\alpha^k}{k!})+\frac{e^{-\alpha}\alpha^{K^*}}{K^*!}=x_t+\frac{e^{-\alpha}\alpha^{K^*}}{K^*!}$, then $x_t=x_r-\frac{e^{-\alpha}\alpha^{K^*}}{K^*!}$. □

EXAMPLE 3 (Power-law Networks): *If the network exhibits a Power-law degree distribution of parameter α and $k_{min}=1$, then*

$$\begin{aligned} x_r &= \begin{cases} 1-\left(\left(\frac{2-\alpha}{3-\alpha}\right)-1\right)^{-1} & \text{if } \alpha>3 \\ 1 & \text{if } 1\leq\alpha\leq 3. \end{cases} \\ x_t &= K_\alpha^{(1-\alpha)} \end{aligned}$$

where K_α satisfies

$$K_\alpha^{2-\alpha}-2 = \left(\frac{2-\alpha}{3-\alpha}\right)\left(K_\alpha^{3-\alpha}-1\right).$$

As a result,

$$\Delta x = \begin{cases} 1-\left(\left(\frac{2-\alpha}{3-\alpha}\right)-1\right)^{-1}-K_\alpha^{(1-\alpha)} & \text{if } \alpha>3 \\ 1-K_\alpha^{(1-\alpha)} & \text{if } 1\leq\alpha\leq 3. \end{cases}$$

To determine the (endogenous) value of network transparency and the optimal intervention, it is illustrative to analyze how $\left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right)$ compares to x_t and x_r. First, suppose $x_r\leq\left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right)$. Then, $\lim_{n\to\infty}\left(\frac{1}{n}\right)(\mathbb{E}[\mathrm{TO}|x_t]-\mathbb{E}[\mathrm{TO}|x_r])=\Delta x\Delta\omega\mathbb{E}[\Delta R]$, which represents the value of network transparency. Consequently, if $\Delta x\Delta\omega\mathbb{E}[\Delta R]\geq\kappa$, then improving network transparency is optimal and $x^{\text{optimal}}=x_t$. Otherwise, it is optimal not to improve transparency and $x^{\text{optimal}}=x_r$. Second, suppose $x_t\leq\left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right)<x_r$. Then $\lim_{n\to\infty}\left(\frac{1}{n}\right)(\mathbb{E}[\mathrm{TO}|x_t]-\mathbb{E}[\mathrm{TO}|x=0])=(\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R])-x_t\Delta\omega\mathbb{E}[\Delta R]$ represents the value of network transparency. Thus, if $(\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R])-x_t\Delta\omega\mathbb{E}[\Delta R]\geq\kappa$, then improving transparency is optimal and $x^{\text{optimal}}=x_t$. Otherwise, improving transparency

is suboptimal and $x^{\text{optimal}} = 0$. Finally, suppose $\left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right) < x_t$. Then, $x^{\text{optimal}} = 0$ no matter what, and, thus, the value of transparency is 0. As a result, improving network transparency is sub-optimal.

As a consequence, the optimal intervention is given by

$$x^{\text{optimal}} = \begin{cases} x_r, \text{ if } x_r \leq \min\left\{\left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right), x_t + \frac{\kappa}{\Delta\omega\mathbb{E}[\Delta R]}\right\} \\ x_t, \text{ if } \frac{\kappa}{\Delta\omega\mathbb{E}[\Delta R]} + x_t \leq x_r \leq \left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right) \text{ or} \\ \qquad x_t \leq \left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]-\kappa}{\Delta\omega\mathbb{E}[\Delta R]}\right) \text{ and } \left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right) < x_r \\ 0, \qquad \text{otherwise,} \end{cases} \tag{A11}$$

which can be rewritten as

$$x^{\text{optimal}} = \begin{cases} x_r, \text{ if } \Delta\omega\mathbb{E}[\Delta R] \leq \min\left\{\frac{\nu}{x_r}, \frac{\kappa}{\Delta x}\right\} \\ x_t, \text{ if } \min\left\{\frac{\nu}{x_r}, \frac{\kappa}{\Delta x}\right\} < \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu-\kappa}{x_t} \\ 0, \qquad \text{otherwise.} \end{cases} \tag{A12}$$

The endogenous value of network transparency is

$$\text{SVI} = \begin{cases} \Delta x\Delta\omega\mathbb{E}[\Delta R], \text{ if } x_r \leq \left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right) \\ (\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]) - x_t\Delta\omega\mathbb{E}[\Delta R], \text{ if } x_t \leq \left(\frac{\mathbb{E}[R_I]-\omega_L\mathbb{E}[\Delta R]}{\Delta\omega\mathbb{E}[\Delta R]}\right) \leq x_r \\ 0, \qquad \text{otherwise,} \end{cases} \tag{A13}$$

which is equivalent to

$$\text{SVI} = \begin{cases} \Delta x\Delta\omega\mathbb{E}[\Delta R], \text{ if } \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu}{x_r} \\ \nu - x_t\Delta\omega\mathbb{E}[\Delta R], \text{ if } \frac{\nu}{x_r} < \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu}{x_t} \\ 0, \qquad \text{otherwise.} \end{cases} \tag{A14}$$

It directly follows from the above analysis

EXAMPLE 4 (Optimal Intervention and Value of Network Transparency in Poisson Networks): *If the network exhibits a Poisson degree distribution of parameter α, then*

$$x^{optimal} = \begin{cases} \left(1-\frac{1}{\alpha}\right), \ \textit{if } \Delta\omega\mathbb{E}[\Delta R] \leq \min\left\{\frac{\nu\alpha}{\alpha-1}, \frac{\kappa K_\alpha!}{e^{-\alpha}\alpha^{K_\alpha}}\right\} \\ \left(1-\frac{1}{\alpha}\right) - \left(\frac{e^{-\alpha}\alpha^{K_\alpha}}{K_\alpha!}\right), \ \textit{if } \min\left\{\frac{\nu\alpha}{\alpha-1}, \frac{\kappa K_\alpha!}{e^{-\alpha}\alpha^{K_\alpha}}\right\} < \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu-\kappa}{\left(1-\frac{1}{\alpha}\right)-\left(\frac{e^{-\alpha}\alpha^{K_\alpha}}{K_\alpha!}\right)} \\ 0, \qquad \textit{otherwise.} \end{cases}$$

where K_α solves $\frac{1}{\alpha} = \sum_{j=0}^{(K_\alpha - 2)} \frac{e^{-\alpha}\alpha^j}{j!}$. The endogenous value of network transparency is

$$SVI = \begin{cases} \left(\frac{e^{-\alpha}\alpha^{K_\alpha}}{K_\alpha!}\right)\Delta\omega\mathbb{E}[\Delta R], \ \textit{if } \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu\alpha}{\alpha-1} \\ \nu - \left(\left(1-\frac{1}{\alpha}\right) - \frac{e^{-\alpha}\alpha^{K_\alpha}}{K_\alpha!}\right)\Delta\omega\mathbb{E}[\Delta R], \ \textit{if } \frac{\nu\alpha}{\alpha-1} < \Delta\omega\mathbb{E}[\Delta R] \leq \frac{\nu}{\left(1-\frac{1}{\alpha}\right)-\left(\frac{e^{-\alpha}\alpha^{K_\alpha}}{K_\alpha!}\right)} \\ 0, \qquad \textit{otherwise.} \end{cases}$$

Appendix B Figures

This section contains figures mentioned in the body of the paper.

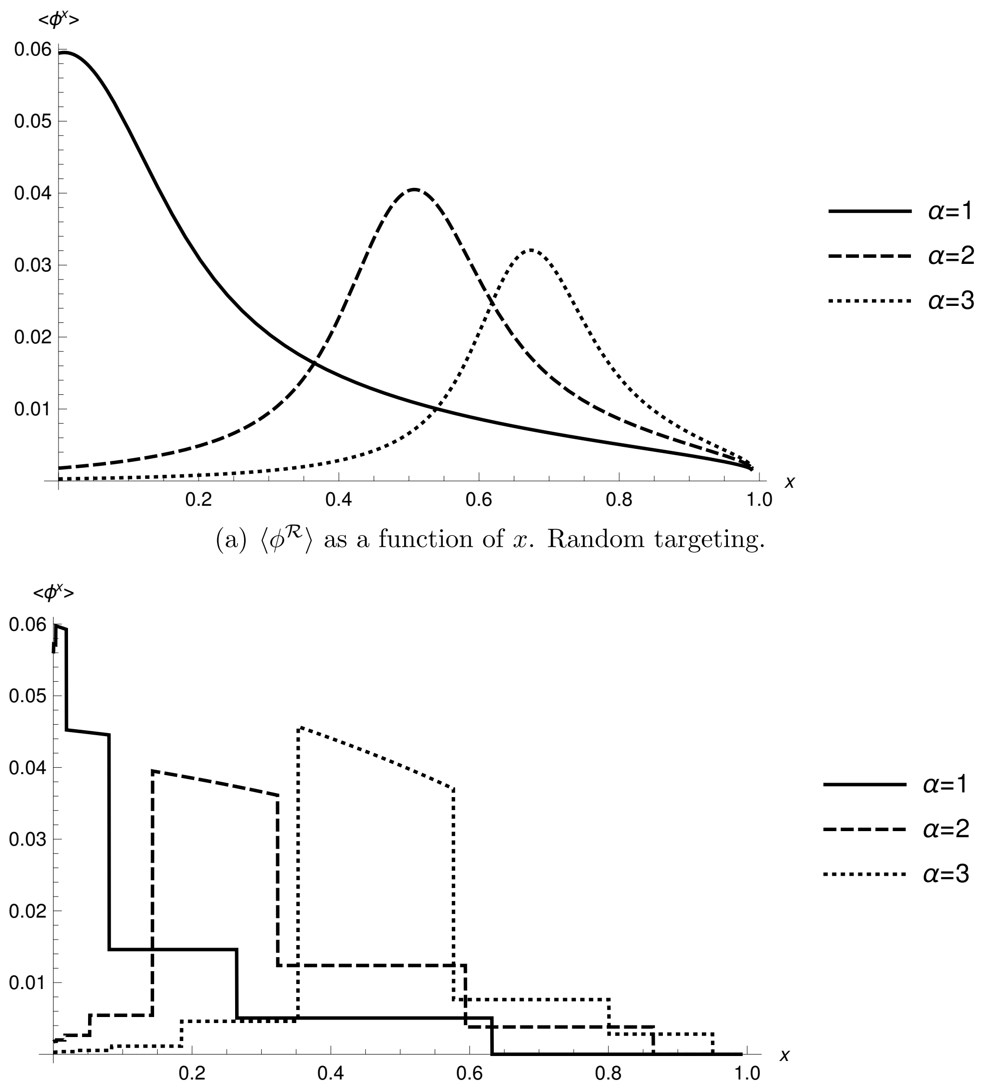

(a) $\langle\phi^{\mathcal{R}}\rangle$ as a function of x. Random targeting.

(b) $\langle\phi^{\mathcal{R}}\rangle$ as a function of x. Targeting based on banks' future number of contagious exposures.

Figure 5. $\{p_k\}_k$ follows a Poisson distribution with parameter α.

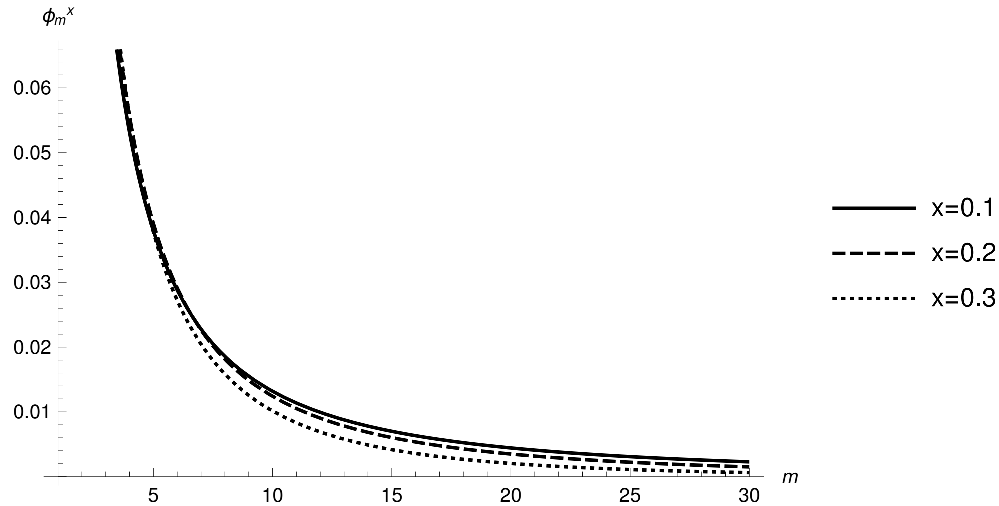

(a) $\{\phi_m^{\mathcal{R}}\}_m$ if $\alpha = 1$, and $n = 100$.

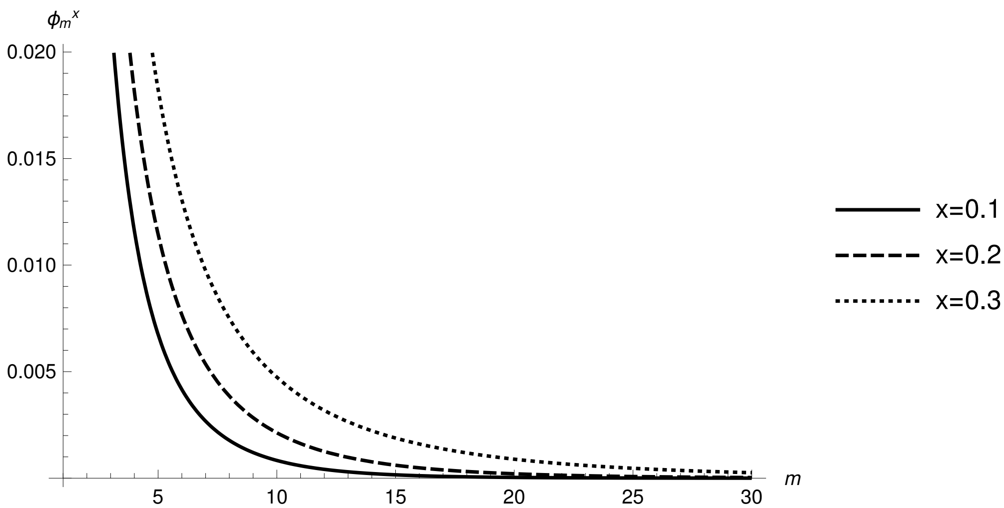

(b) $\{\phi_m^{\mathcal{R}}\}_m$ if $\alpha = 2$, and $n = 100$.

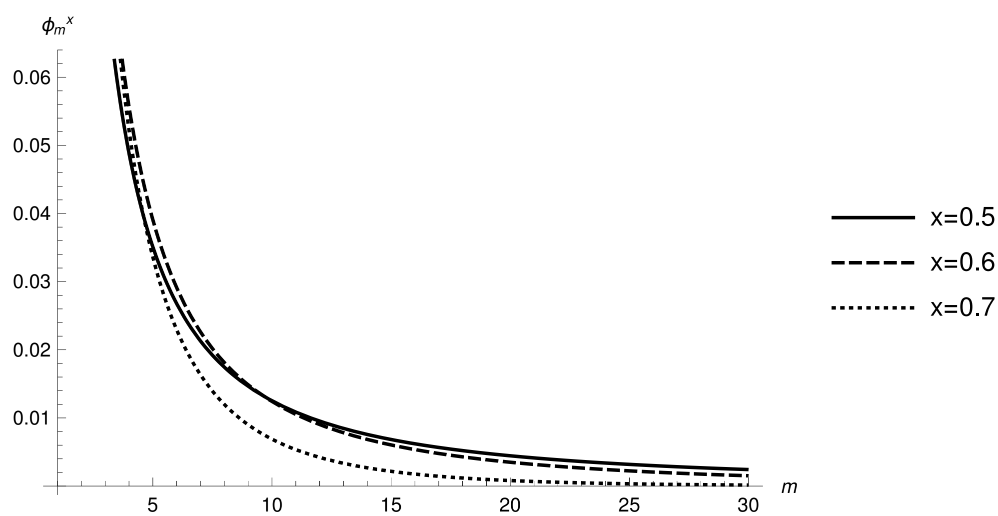

(c) $\{\phi_m^{\mathcal{R}}\}_m$ if $\alpha = 2$, and $n = 100$.

Figure 6. $\{p_k\}_k$ follows a Poisson distribution with parameter α and restricted banks are selected at random

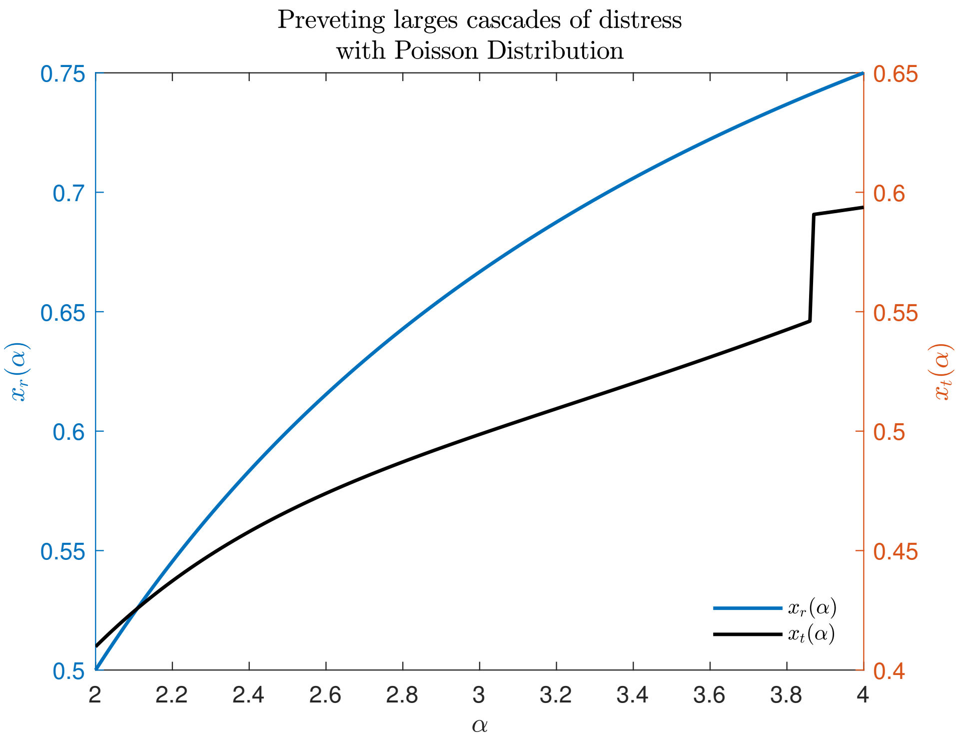

(a) Poisson

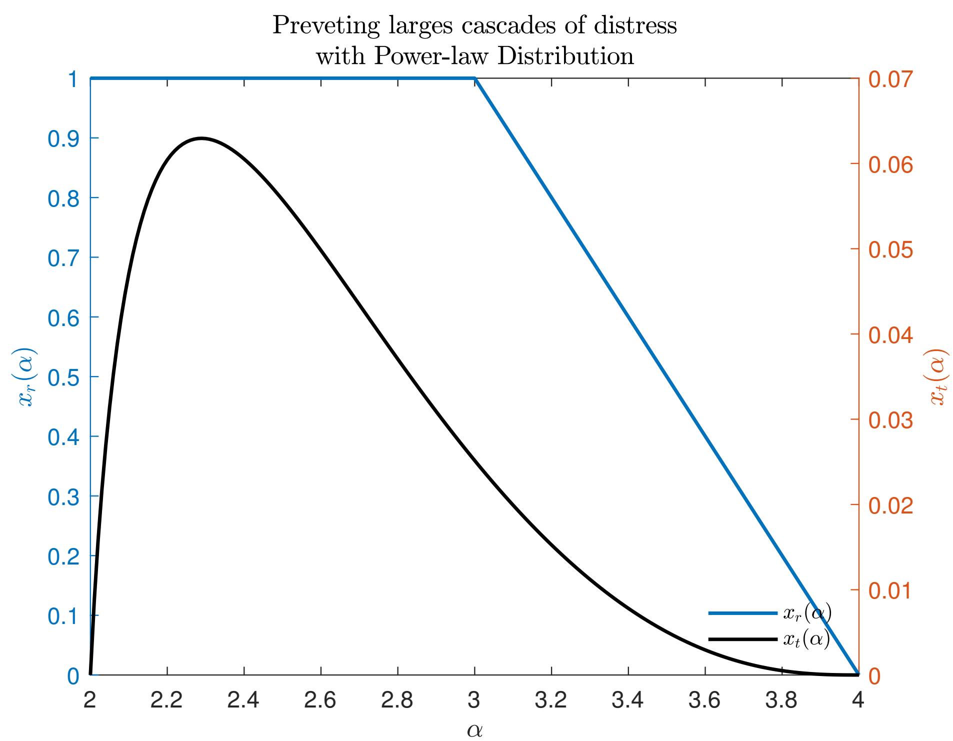

(b) Power-law

Figure 7. Preventing large cascades of distress

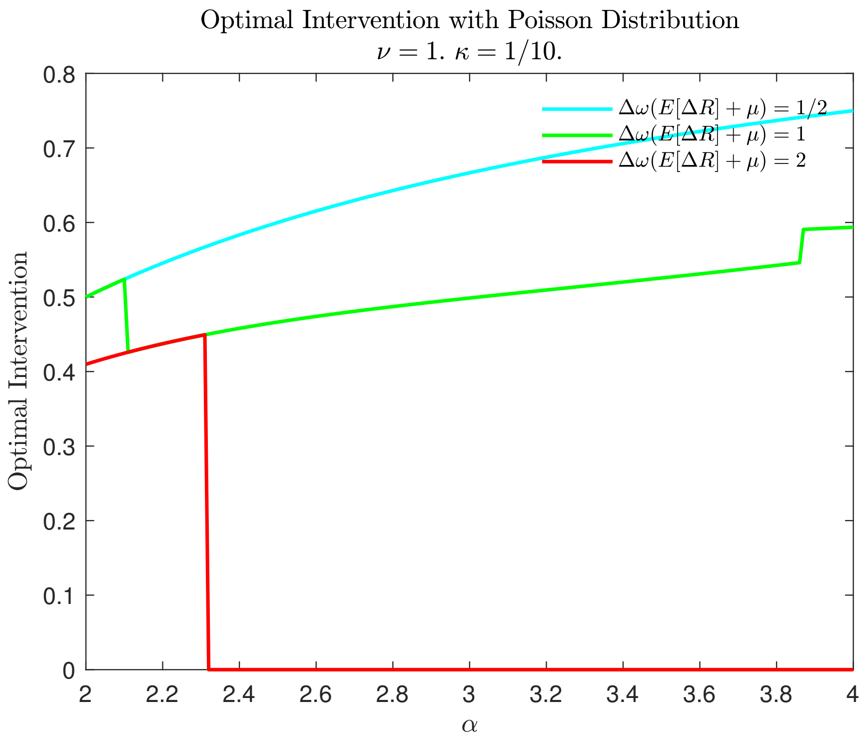

(a) Optimal fraction of restricted banks

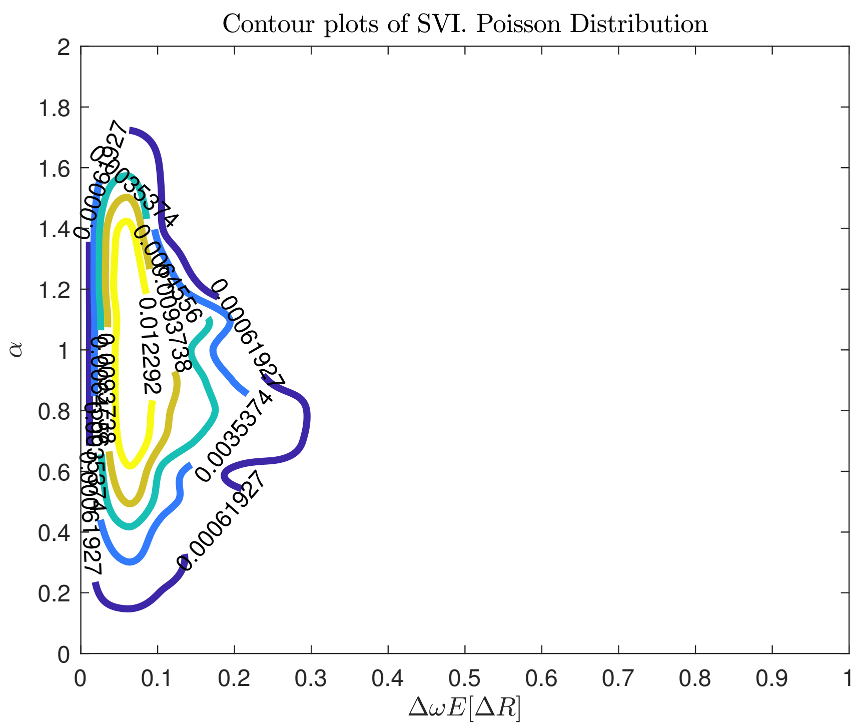

(b) Value of network transparency

Figure 8. Optimal intervention as a function of α in Poisson networks.

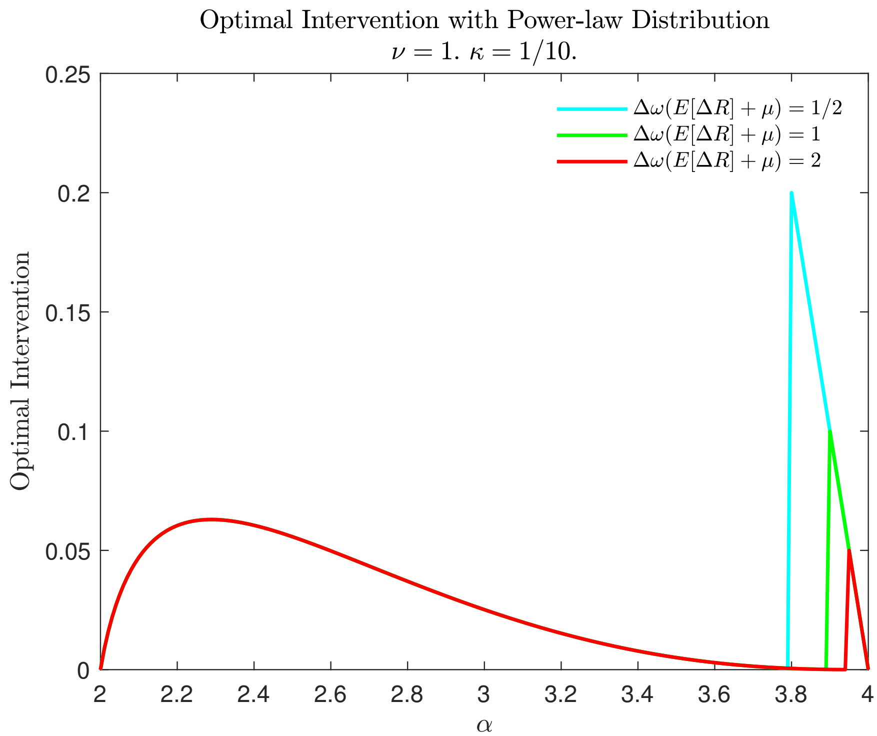

(a) Optimal fraction of restricted banks

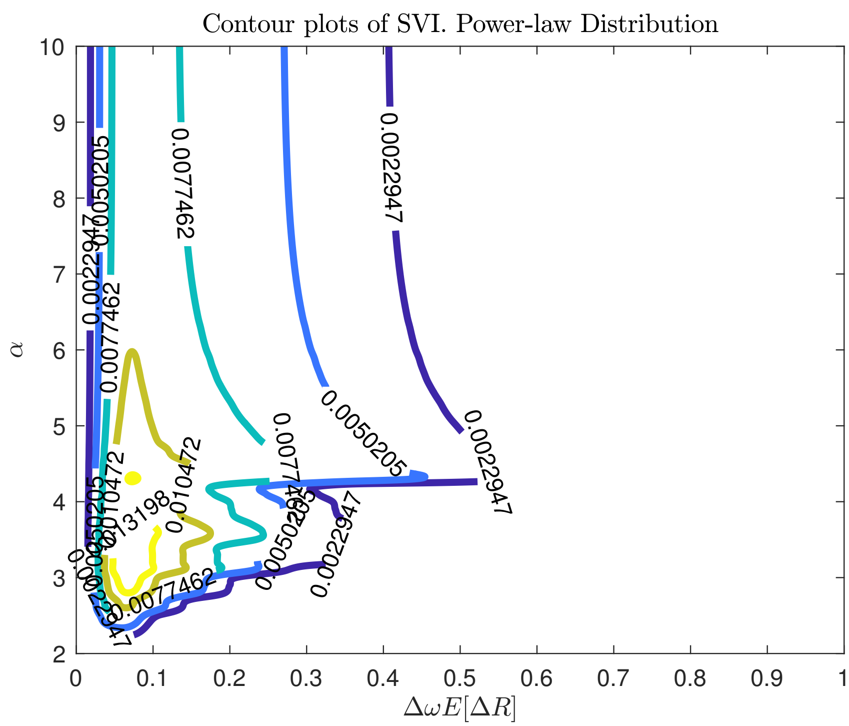

(b) Value of network transparency

Figure 9. Optimal intervention as a function of α in Power-law networks.

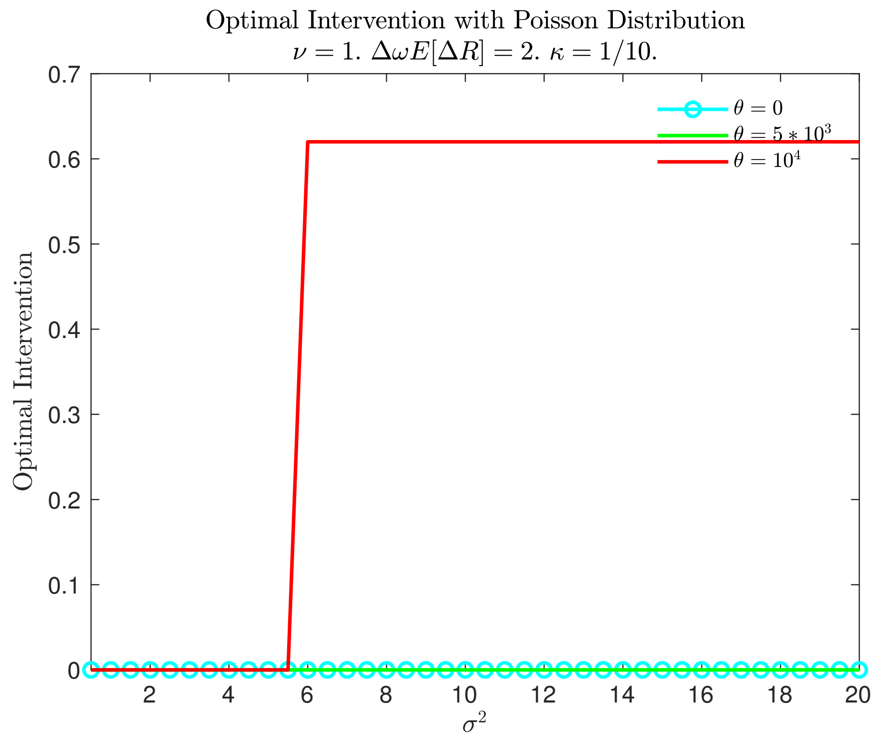

(a) Optimal fraction of restricted banks

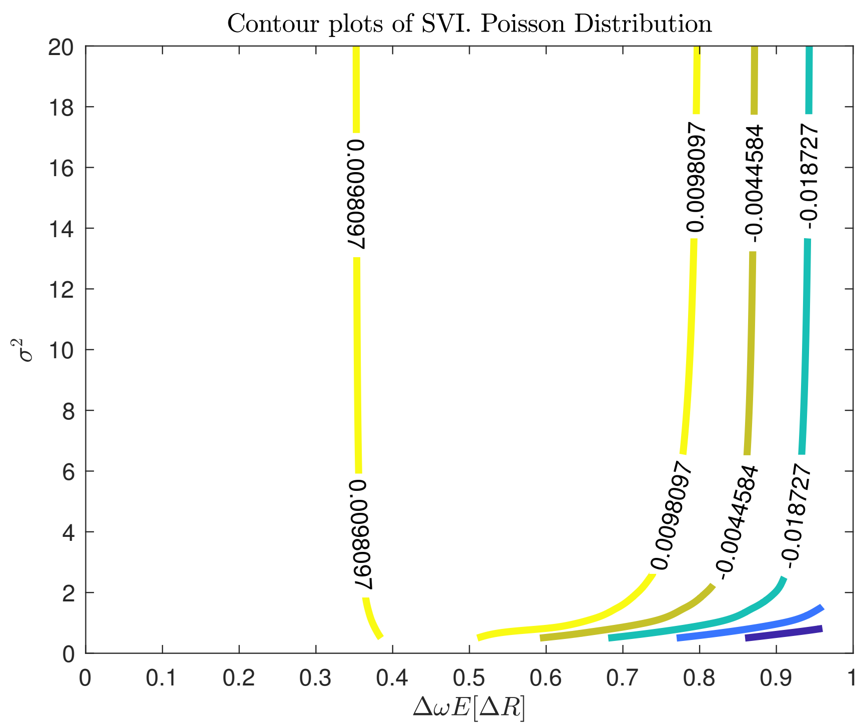

(b) Value of network transparency. $\theta = 10^4$.

Figure 10. Optimal intervention as a function of σ^2 in Poisson networks with model uncertainty.

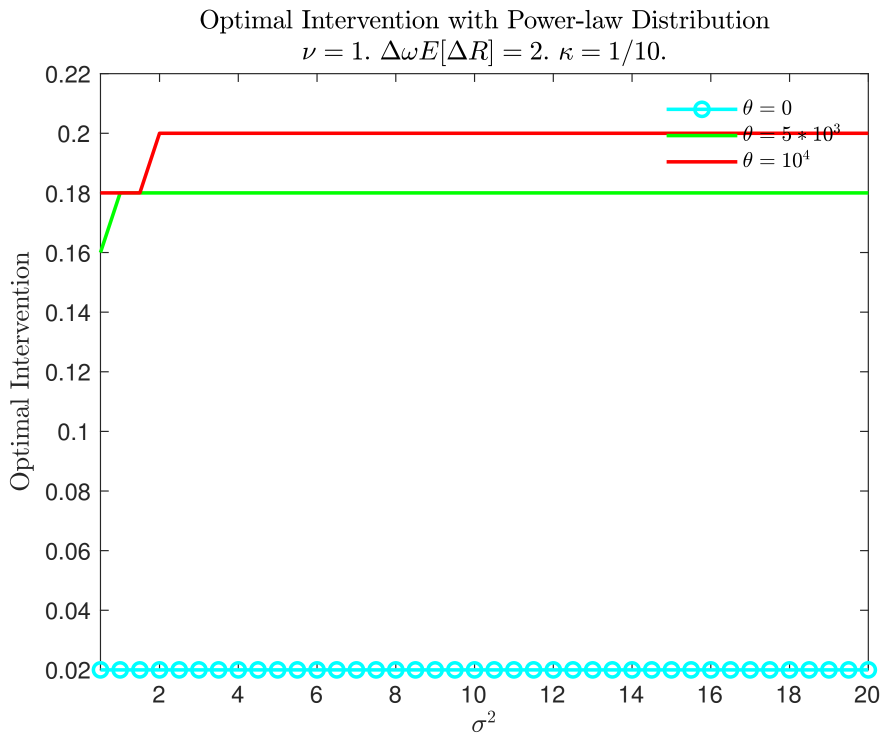

(a) Optimal fraction of restricted banks

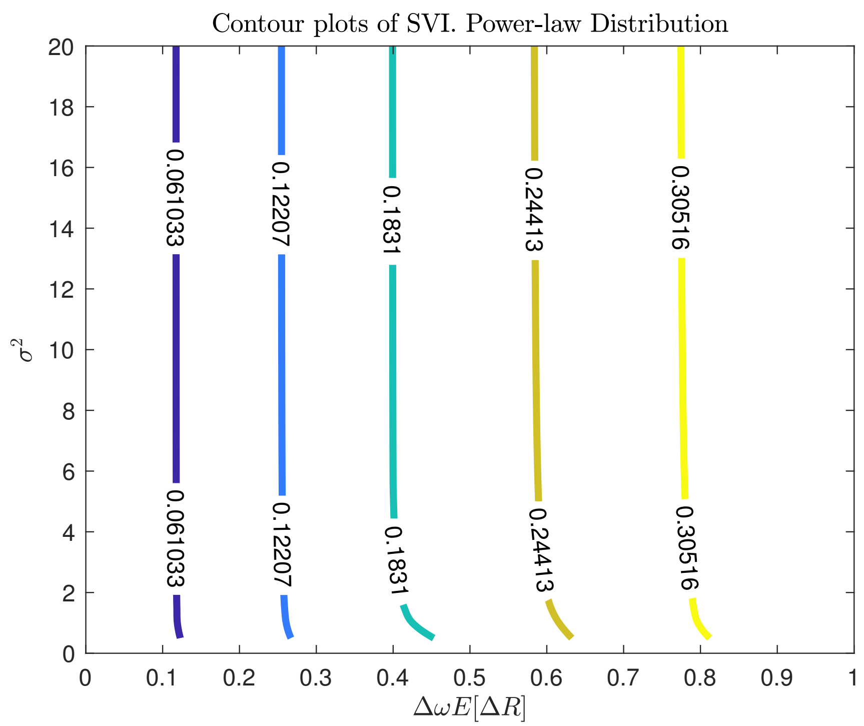

(b) Value of network transparency. $\theta = 10^4$.

Figure 11. Optimal intervention as a function of σ^2 in Power-law networks with model uncertainty.