

Fire Sales in a Model of Complexity

RICARDO J. CABALLERO and ALP SIMSEK*

ABSTRACT

We present a model of financial crises that stem from endogenous *complexity*. We conceptualize complexity as banks' *uncertainty* about the financial network of cross exposures. As conditions deteriorate, cross exposures generate the possibility of a domino effect of bankruptcies. As this happens, banks face an increasingly complex environment since they need to understand a greater fraction of the financial network to assess their own financial health. Complexity dramatically amplifies banks' perceived counterparty risk, and makes relatively healthy banks reluctant to buy risky assets. The model also features a novel *complexity externality*.

ONE OF THE MOST damaging aspects of financial crises is the enormous *uncertainty* they generate, and a central factor behind this uncertainty is the *complexity* of the linkages among modern financial institutions (banks, for short). The concern for the uncertainty and complexity combination, and the perverse fire sales that accompany it, influences private and public choices alike. Federal Reserve Chairman Ben Bernanke, in his testimony to the Senate on April 3, 2008 following the Fed's Bear Stearns intervention, captures this concern as follows:

Our financial system is extremely complex and interconnected, and Bear Stearns participated extensively in a range of critical markets. The sudden failure of Bear Stearns likely would have led to a chaotic unwinding of positions in those markets and could have severely shaken confidence. The company's failure could also have cast doubt on the financial positions of some of Bear Stearns' thousands of counterparties and perhaps of companies with similar businesses. . . . Moreover, the adverse impact of a default would not have been confined to the financial system but would

*Both the authors are from MIT and NBER. We thank Jeff Campbell, John Geanakoplos, Campbell Harvey, David Laibson, Fred Malherbe, Juan Ocampo, Katerina Smidkova, Fernando Vega-Redondo, Adam Zawadowski, an anonymous referee, and an Associate Editor for valuable comments. We also thank the seminar participants at Bilkent University, Boston College, The Federal Reserve Bank of Boston, Boston University, Harvard University, MIT, Middle East Technical University, New York University, the University of Chicago, the University of Michigan; conference participants at the AEA meetings, Brown University, CIED-Tel Aviv University, Columbia University, ECB, IRFMP, MFI, MNB-CEPR, PSU-Cornell, SAET, SED for their comments; and the NSF for financial support. This paper covers and extends the substantive issues in (and hence, replaces) "Complexity and Financial Panics," NBER WP #14997.

DOI: 10.1111/jofi.12087

have been felt broadly in the real economy through its effects on asset values and credit availability.¹

Unfortunately, Chairman Bernanke's testimony would prove prescient only a few months later during the Lehman episode, when the demise of the investment bank wrecked havoc around the world. Moreover, the concern for a repeat of such turmoil is the central reason behind the multiple recent attempts to insulate the rest of Europe from the sovereign debt problems of its periphery.

In this paper, we present a model of crises that builds upon the idea that complexity, a dormant factor during normal times, becomes acutely relevant and self-reinforcing during crises. Complexity matters in our model not directly (and in this sense this is not a "networks" paper) but through the uncertainty it generates and through the reactions of the economic agents to this uncertainty.

The basic structure of our model is a network of cross exposures between financial institutions that is susceptible to a domino effect of bankruptcies. However, we make assumptions such that domino effects are of limited size in the absence of our informational mechanism. In this context, we conceptualize complexity by banks' uncertainty about cross exposures. In particular, banks have only local knowledge of cross exposures: they understand their own exposures, but they are increasingly uncertain about cross exposures of banks that are farther away (in the network) from themselves.² During normal times, banks only need to understand the financial health of their direct counterparties. In contrast, when a surprise liquidity shock hits parts of the network, a domino effect of bankruptcies becomes possible, and banks become concerned that they might be indirectly hit. Banks' uncertainty about cross exposures, a dormant factor in normal times, suddenly becomes relevant. In particular, banks now need to understand the financial health of the counterparties of their counterparties (and their counterparties). Since banks only have local knowledge of the exposures, they cannot rule out an indirect hit. They now perceive significant counterparty risk that leads them to retrench into a liquidity conservation mode.

This structure exhibits strong interactions with secondary markets for banks' assets. Banks in distress can sell their legacy assets to meet the surprise liquidity shock. The natural buyers of the legacy assets are other banks in the financial network, which may also receive an indirect hit. When the surprise shock is small, the domino effect is small and buyers can rule out an indirect hit. In this case, buyers purchase the distressed banks' legacy assets at their "fair" prices (which reflect the fundamental value of the assets). In contrast, when the surprise shock is large, larger domino effects become possible and buyers cannot rule out an indirect hit. As a precautionary measure, they hoard liquidity and turn into sellers. The price of legacy assets plummets to fire

¹ Available at <http://www.federalreserve.gov/newsevents/testimony/bernanke20080y03a.htm>.

² In practice, banks also face many other sources of complexity (e.g., asset payoffs). Our modest goal in this paper is to focus on one source of complexity (about the structure of cross exposures), and to understand the role that this type of complexity plays during crises.

sale levels (i.e., below fundamentals), which in turn exacerbates the domino effect.

This feedback mechanism can generate multiple equilibria for intermediate levels of the surprise shock. When legacy assets garner a fair price in the secondary market, the banks in distress have access to more liquidity, and the surprise shock is contained after fewer bankruptcies. When the domino effect is smaller, the natural buyers rule out an indirect hit and demand legacy assets, which ensures that these assets trade at their fair prices. Set against this benign scenario is the possibility of a *fire sale equilibrium*, where the price of legacy assets collapses to fire sale levels. This leads to a greater number of bankruptcies and a larger domino effect. With a larger domino effect, the natural buyers become worried about an indirect hit and they sell their own legacy assets, which reinforces the collapse of asset prices.

Our model features a distinct *complexity externality*, which stems from the dependence of banks' counterparty risk on the endogenous size of the domino effect. In particular, any action that increases the size of the domino effect increases the counterparty risk perceived by banks that are uncertain about the financial network, and banks dislike this effect. Our model features two variants of this complexity externality (one nonpecuniary, one pecuniary), each of which supports different government policies. First, a bailout of the distressed banks financed by small lump-sum taxes on all banks may lead to a Pareto improvement. The market equilibrium is unable to replicate this allocation because each bank fails to take into account that its contribution to a bailout will reduce the counterparty risk faced by all other banks. Second, in the range of multiple equilibria, policies that increase asset prices may lead to a Pareto improvement by coordinating the banks on the fair price equilibrium. In this range, the fire sale equilibrium is Pareto inefficient because a bank that sells assets does not take into account the effect of its decision on other banks' counterparty risk. In particular, this bank generates a (small) reduction in asset prices, which in turn leads to a larger domino effect and greater counterparty risk, real and perceived, for all other banks.

Our paper is related to several strands of the literature. An extensive literature highlights the possibility of network failures and contagion in financial markets. An incomplete list includes Rochet and Tirole (1996), Kiyotaki and Moore (1997a), Allen and Gale (2000), Freixas, Parigi, and Rochet (2000), Eisenberg and Noe (2001), Lagunoff and Schreft (2001), Dasgupta (2004), Leitner (2005), Cifuentes, Ferrucci, and Shin (2005), Rotemberg (2008), Allen, Babus, and Carletti (2010), Acemoglu et al. (2012), and Zawadowski (2013) (see Allen and Babus (2009) for a survey). Many of these papers focus on the mechanisms through which solvency and liquidity shocks may generate a domino effect in the financial network. In contrast, we take these phenomena as the reason for the rise in banks' uncertainty and we focus on the effect of this uncertainty on banks' prudential actions. It is also worth pointing out that the uncertainty mechanism we emphasize in this paper is operational *even for a relatively small amount of network contagion*. The contagion literature is sometimes criticized because it appears unlikely that many financial institutions would be caught

up in a domino effect of bankruptcies.³ But as this paper illustrates, even *partial* domino effects can have large aggregate effects since they greatly increase counterparty risk and trigger widespread prudential actions.⁴ A caveat is that our model applies for shocks that are either unanticipated or difficult to insure *ex ante* (see Section V.A for details).

Our paper is most directly related to the literature on flight to quality and Knightian uncertainty in financial markets, as in Caballero and Krishnamurthy (2008), Easley and O'Hara (2009), Routledge and Zin (2009), and Hansen and Sargent (2010). Our contribution relative to this literature is in generating the uncertainty from the complexity of the financial network itself. Our work complements a number of recent papers that focus on other sources of uncertainty during crises. Brunnermeier and Sannikov (2012) show that exogenous uncertainty is amplified in a fire sale episode, because price uncertainty increases natural buyers' balance sheet uncertainty (which in turn feeds back into price uncertainty). Dang, Gorton, and Holmstrom (2010) show that uncertainty (and asymmetric information) in credit markets increases during crises because debt contracts become information sensitive.

In the canonical model of fire sales, fire sales happen because the natural buyers of the assets experience financial distress simultaneously with sellers (see Shleifer and Vishny (1992, 1997) and Kiyotaki and Moore (1997b)). More recently, Brunnermeier and Pedersen (2005) show that, when there are few players, unconstrained potential buyers may choose not to arbitrage fire sales in the short run because they anticipate a better deal in the future. Our model lies somewhere in between these two views: most potential buyers are unconstrained, as in Brunnermeier and Pedersen (2005), but they are fearful of going about their normal arbitrage role because of uncertainty (and in this sense they are distressed as in Shleifer and Vishny (1992)). It is the complexity of the environment that sidelines potential buyers and exacerbates the fire sales. Importantly, this mechanism works even when the number of market participants is large.⁵

The organization of this paper is as follows. In Section I we describe the environment for a benchmark case with no uncertainty about cross exposures. Section II characterizes the equilibrium for this benchmark and illustrates the

³ See Upper (2007) for a survey of the empirical literature that uses counterfactual simulations to assess the danger of contagion. Regarding this literature, Brunnermeier et al. (2009) note that "it is only with implausibly large shocks that the simulations generate any meaningful contagion."

⁴ The role of domino effects in elevating complexity and uncertainty was also highlighted by Haldane (2009), who nicely captures the mechanism when he wrote that at times of stress "knowing your ultimate counterparty's risk becomes like solving a high-dimension Sudoku puzzle."

⁵ Other papers that investigate the mechanisms for fire sales and asset price dislocations in financial markets include Allen and Gale (1994), Gromb and Vayanos (2002), Geanakoplos (2003, 2009), Lorenzoni (2008), Brunnermeier and Pedersen (2009), Acharya, Gale, and Yorulmazer (2011), Garleanu and Pedersen (2011), Stein (2012), and Diamond and Rajan (2011) (see Shleifer and Vishny (2011) for a recent survey). More broadly, this paper belongs to an extensive literature on financial crises that highlights the connection between panics and a decline in the financial system's ability to channel resources to the real economy (see, for example, Caballero and Kurlat (2008) for a survey).

mechanics of (partial) domino effects in our setting. Section III contains our main results. There, banks have only local knowledge about cross exposures, and a sufficiently large shock increases banks' counterparty risk and leads to fire sales in secondary markets. This section also highlights the interaction between counterparty risk and fire sales, and demonstrates the possibility of multiple equilibria. In Section IV we describe the complexity externality and its policy implications. Section V discusses the role of a number of key assumptions. The paper concludes with final remarks in Section VI and several appendices.

I. Basic Environment and Equilibrium

We consider an economy with three dates $\{0, 1, 2\}$ and a single consumption good (a dollar). The economy has n continuums of financial intermediaries denoted by $\{b^j\}_{j=0}^{n-1}$. Each of these continuums consists of identical banks. For simplicity, we refer to each continuum b^j as *bank b^j* , which is our unit of analysis.⁶ Banks start with a given balance sheet at date 0 (which will be described shortly), but they only consume at date 2. Banks can transfer their date 0 dollars to date 2 by investing in one of two ways. First, banks can keep their dollars in cash, which yields one dollar at the next date per dollar invested. Second, banks can invest in an asset. Each unit of the asset yields $R > 1$ dollars at date 2 (and no dollars at date 1). The asset is supplied elastically at date 0 at a normalized price of one.

While the asset yields a higher date 2 return than cash, it is completely illiquid at date 1. In particular, it is not possible to sell or borrow against the asset at date 1. (Thus, a bank cannot convert the asset to dollars at date 1.) This assumption captures the standard liquidity and return trade-off, which is prevalent in financial markets. The microfoundations that lead to this trade-off are well known (e.g., Holmstrom and Tirole (1998)). One can think of the cash in this model as the liquid securities, such as U.S. Treasuries, that yield a lower return but retain their market value at times of distress. In contrast, the asset can be thought of as illiquid securities, such as asset-backed securities, which potentially yield a higher return but lose their market value at times of distress.

Each bank initially has y dollars and $1 - y$ units of legacy assets. At date 0, which is the only meaningful decision date in our model, banks can trade legacy assets in a secondary market at an endogenous price p . This price cannot exceed one because legacy assets and new assets are identical (and the price of the latter is one). We also assume that the natural buyers of legacy assets are the other banks in the model. In particular, outside agents (lower valuation users) demand asset elastically at a discounted price $p_{\text{scrap}} < 1$. Thus, if legacy assets are sold to outside agents, they garner a price $p = p_{\text{scrap}}$. We refer to this situation as a *fire sale* of legacy assets.

⁶ The only reason for the continuum is for banks to take other banks' decisions as given.

Forward Neighbor Bank	Assets	Liabilities	Backward Neighbor Bank
	<ul style="list-style-type: none"> • short-term debt claim with face value z • 1-y legacy assets (keep or sell) • y dollars (invest in cash or asset) 	<ul style="list-style-type: none"> • short-term debt claim with face value z • equity 	

Figure 1. The initial balance sheet of a generic bank.

The basic premise of our model can now be informally described. An unexpected liquidity shock generates the possibility that banks might need dollars at date 1. This in turn shifts banks' investments at date 0 from the asset to cash (flight to quality), which has two effects. First, as banks stop buying new assets, there is a credit crunch. Second, as banks stop buying legacy assets (and as they try to sell their own legacy assets to raise dollars), there is a fire sale of legacy assets in the secondary market. The contribution of our paper is to describe the role of complexity associated with counterparty risk in generating this flight-to-quality episode. To this end, we gradually introduce the main ingredients of our model.

A. Cross Exposures and the Financial Network

At date 0, each bank has short-term debt claims worth z dollars on one other bank, which we call the *forward neighbor bank*. We assume that short-term debt cannot be rolled over and it must be paid back at date 1, which will be without loss of generality.⁷ On the liability side, the bank also has z dollars of short-term debt claims held by another bank, which we call the *backward neighbor bank*. Figure 1 illustrates the initial balance sheet of a bank.

The role of these cross-debt claims is to capture various types of unsecured cross exposures that are common in the financial system, and the counterparty risk that they bring about. One source of cross exposures is interbank loans. Upper (2007) documents that interbank loans constitute a large fraction of banks' balance sheets in many European countries.⁸ A second and potentially much larger source of cross exposures is OTC derivative contracts (such as interest rate swaps or credit default swaps (CDSs)) traded between financial institutions. The Bank for International Settlements (BIS) reports that gross credit exposures in over-the-counter (OTC) derivative markets in G10 countries and

⁷ Appendix A considers an extension of the model in which banks have the option to roll over and shows that the equilibrium is unchanged.

⁸ To give two examples, Upper (2007) notes that "at the end of June 2005 interbank credits accounted for 29% of total assets of Swiss banks and 25% of total assets of German banks."




Figure 2. Financial Network. The left panel illustrates a circle network. The right panel illustrates the template network in which slots are to be filled by banks.

Switzerland had exceeded \$5 trillion by the end of 2008.⁹ The cross-debt claims of this model can be viewed as capturing the uncollateralized portion of these exposures (although the Lehman crisis revealed that even fully collateralized repo loans can be frozen by bankruptcy courts).

In Caballero and Simsek (2009), we provide one rationale (out of many) for cross exposures from their role in facilitating bilateral liquidity insurance, as in Allen and Gale (2000). In this paper, we take the exposures as given and analyze their role in generating flight-to-quality episodes.

Banks’ cross exposures form a *financial network*. For simplicity, we assume that the network takes the form of a circle. The left panel of Figure 2 illustrates a circle network. The notation, $b^j \leftarrow b^{j+1}$, illustrates that bank b^{j+1} has debt claims on bank b^j . Note that banks are ordered around a circle, with bank b^0 having debt claims on bank b^{n-1} .

We conceptualize complexity about counterparty risk with banks’ uncertainty about the financial network. In particular, banks have only *local knowledge of cross exposures*: they understand their own exposures, but they are increasingly uncertain about cross exposures of banks that are farther away (in the network) from themselves. We capture this notion by assuming that banks have only *local knowledge about the financial network*: they know the identity of their forward neighbor bank (on which they have debt claims), but they do not know how the rest of the banks are ordered around the circle (i.e., which banks are exposed to which other banks).

To formalize this ingredient, we need to consider a larger set of networks than the above example. To this end, consider the template network illustrated in the right panel of Figure 2, which is described in terms of financial slots to be filled by banks. Next consider a permutation $\sigma : \{0, 1, \dots, n - 1\} \rightarrow \{0, 1, \dots, n - 1\}$, which assigns bank j to slot $i = \sigma(j)$. Note that the bank assigned to slot i has debt claims on the bank that is assigned to slot $i - 1$. However, the identities of these banks depend on the permutation, σ , which introduces uncertainty about the financial network. In particular, we let

$$\mathcal{N} = \{\sigma | \sigma : \{0, 1, \dots, n - 1\} \rightarrow \{0, 1, \dots, n - 1\} \text{ is a permutation}\}$$

⁹ Source: BIS semiannual OTC derivatives statistics. Gross credit exposures take into account bilateral netting between the same pair of counterparties. Gross market values of exposures, which do not take into account this netting, are much larger (more than \$20 trillion in interest rate derivatives and more than \$5 trillion in credit derivatives by the end of 2008).

denote the set of all possible financial networks. The earlier example corresponds to the particular realization in which each bank is assigned to the same slot as its identity.

We capture banks' uncertainty over the network by assuming that banks do not know the realization, σ . In particular, let $\mathcal{N}^j(\sigma) \subset \mathcal{N}$ denote the set of financial networks that bank b^j finds possible given the actual realization, σ . We refer to the collection $\{\mathcal{N}^j(\sigma)\}_{j,\sigma}$ as an *uncertainty model* for banks.¹⁰

B. Surprise Shock and Banks' Response

At date 0, the banks learn that a *rare event* (which they had not anticipated at the unmodeled date -1) has happened and one bank, b^0 , will become distressed. To remain solvent, this bank needs to make θ dollars of payment (to an outsider) at date 1.

This outside debt is senior to the short-term debt to the neighbor bank (it can be equivalently interpreted as a shock to the value of the bank's assets in a version of the model in which banks have heterogeneous assets).¹¹ Consequently, these losses might spill over to other banks via the financial network and may bring them into financial distress at date 1. To prepare for date 1, at date 0 the banks take one of the actions $A_0^j \in \{S, B\}$, which are restricted to a binary choice set for simplicity.¹² As a precautionary measure, the bank may choose $A_0^j = S$ to invest all of its y dollars in cash and sell all of its legacy assets $1 - y$ in the secondary market, keeping a completely liquid balance sheet. Alternatively, the bank may choose $A_0^j = B$ to be a potential buyer of assets. In this case, the bank retains its own legacy assets on its balance sheet and uses its dollars to buy either new or legacy assets (whichever is more profitable).

The bank chooses A_0^j to maximize its equity value at date 2, subject to meeting its debt payment at date 1. Given the rare event, a bank may not be able to pay back its debt in full (despite the precautionary measures it might take), but instead ends up paying $q_1^j \leq z$. Similarly, the value of the bank's date 2 equity may be $q_2^j \leq R$. Note that the bank is either *solvent*, pays $q_1^j = z$, and observes date 2 equity value $q_2^j \geq 0$, or is *insolvent*, pays $q_1^j < z$, and observes date 2 equity value $q_2^j = 0$.

The bank makes its decision at date 0 while facing Knightian uncertainty about the network and the counterparty risk that it might bring about. In

¹⁰ A simpler alternative to the permutations is to have banks ordered in the circle in the same order as the locations (i.e., bank 1 in location 1, bank 2 in location 2, etc.) and have the uncertainty be about the identity of the bank in distress rather than about the linkages between the banks. We chose the slightly more cumbersome route of permutations because it better aligns with the idea of complexity that we want to capture here. But mechanically, the results would be very similar to the alternative formulation.

¹¹ In practice, banks are also highly leveraged, which implies that even small reductions in the value of assets could lead to large losses of capital. Thus, the empirical counterpart of θ is the leveraged losses suffered by the bank because of a shock to its assets.

¹² This assumption does not affect any of the qualitative conclusions of the model, as demonstrated by an earlier version of this paper, Caballero and Simsek (2009).

particular, the bank considers the range of possible financial networks, $\mathcal{N}^j(\sigma)$, and chooses an action that is robust to this uncertainty. Formally, let $(q_1^j(\sigma), q_2^j(\sigma))$ denote the bank's debt and equity payment in equilibrium given the financial network, σ . We follow Gilboa and Schmeidler's 1989 Maximin expected utility representation and write the bank's optimization problem as¹³

$$\max_{A_0^j(\sigma) \in \{S, B\}} \min_{\tilde{\sigma} \in \mathcal{N}^j(\sigma)} q_2^j(\tilde{\sigma}). \tag{1}$$

As we discuss in Section V, Knightian uncertainty is not necessary for our main results (risk-aversion suffices), but it is both realistic in this context and technically convenient.

C. Secondary Market and Equilibrium

Legacy assets are traded in a centralized exchange that opens (just) at date 0. Given the legacy asset price p , the banks that choose $A_0^j = S$ sell all of their legacy assets ($1 - y$ units for each bank) while the banks that choose $A_0^j = B$ are potential buyers of legacy assets and may spend up to y (their dollars). If $p < 1$, potential buyers spend all of their y dollars on legacy assets, while if $p = 1$, they are indifferent between buying legacy or new assets. Recall that $p_{\text{scrap}} < 1$ denotes the valuation of outside agents. Thus, the market-clearing condition for legacy assets can be written as

$$(1 - y) \sum_j \mathbf{1}\{A_0^j = S\} - \frac{y}{p} \sum_j \mathbf{1}\{A_0^j = B\} \begin{cases} \leq 0, & \text{if } p = 1 \\ = 0, & \text{if } p \in (p_{\text{scrap}}, 1) . \\ \geq 0, & \text{if } p = p_{\text{scrap}} \end{cases} \tag{2}$$

The first term on the left-hand side of the equation denotes the total supply of legacy assets while the second term denotes the maximum potential demand. If the left-hand side of equation (2) is negative for each $p \in [p_{\text{scrap}}, 1]$, then legacy assets trade at their fair value one, potential buyers are indifferent between buying legacy and new assets, and they buy just enough legacy assets to clear the market. If the left-hand side of equation (2) is zero for some $p \in [p_{\text{scrap}}, 1]$, then p is the equilibrium price. If the left-hand side is positive for each $p \in [p_{\text{scrap}}, 1]$, then there is excess supply of legacy assets and their price is given by p_{scrap} .

DEFINITION 1: *An equilibrium is a collection of bank actions, debt payments, and equity values, $[\{A_0^j(\sigma), q_1^j(\sigma), q_2^j(\sigma)\}_j]_{\sigma \in \mathcal{N}}$, and a price level $p \in [p_{\text{scrap}}, 1]$ for legacy assets such that, given the realization of the financial network σ , each bank b^j chooses its actions according to the worst-case financial network*

¹³ The preferences captured by (1) are a special case of Gilboa and Schmeidler (1989) in which the set of priors consists of the convex hull of the Dirac delta measures corresponding to the networks, $\tilde{\sigma} \in \mathcal{N}^j(\sigma)$.

that it finds possible (see problem (1)) and the legacy asset market clears (see equation (2)).

To characterize the equilibrium, it is useful to define the notion of a bank's *distance* from the original distressed bank. The latter bank, b^0 , has distance $d = 0$ from itself. The backward neighbor of the original distressed bank has distance $d = 1$. Similarly, the backward neighbor of the backward neighbor has distance $d = 2$. In this way, each bank can be assigned a unique distance for each realization of the network. The distance is the only payoff-relevant variable in this economy. In particular, as we formally show in Appendix V.B, the banks' equilibrium payoffs and actions can be written as a function of their distance. That is, there exist functions $A_0[\cdot]$, $Q_1[\cdot]$, and $Q_2[\cdot]$ such that

$$(A_0^j(\sigma), q_1^j(\sigma), q_2^j(\sigma)) = (A_0[d], Q_1[d], Q_2[d]),$$

where d denotes the distance of bank b^j given the network σ .

In the next section, we demonstrate that a bank is insolvent, $Q_1[d] < z$, if and only if it is sufficiently close to the original distressed bank. Similarly, a bank chooses a precautionary action, $A_0[d] = S$, if and only if d is sufficiently small. In view of these observations, we define the following notions of a *domino effect* and a *flight to quality*, which facilitate the characterization of equilibrium.

DEFINITION 2: Consider a collection of bank actions and payoffs $\{A_0^j, q_1^j, q_2^j\}_j$.

(i) There is a domino effect of size D if banks with distance $d \leq D - 1$ are insolvent (i.e., they pay $q_1^j < z$) while banks with distance $d \geq D$ are solvent (i.e., they pay $q_1^j = z$).

(ii) There is a flight to quality of size F if banks with distance $d \leq F - 1$ choose $A_0^j = S$ while banks with distance $d \geq F$ choose $A_0^j = B$.

Note that D also corresponds to the number of banks that are insolvent, and F corresponds to the number of banks that choose the precautionary action. In subsequent sections, D and F will be useful to summarize the equilibrium in this economy.

II. Equilibrium in the Certainty Benchmark

In this section, we characterize the equilibrium with no uncertainty, which provides a benchmark for our main results with complexity (in terms of the information structure, it essentially corresponds to the Allen and Gale (2000) environment). We show that, if the number of banks is sufficiently large, then there can only be a partial domino effect and a partial flight to quality, that is, $D < n$ and $F < n$. Moreover, D and F are "proportional" to the size of the initial shock, θ . That is, when banks have perfect knowledge of the financial network, a sufficiently deep financial system is resilient to perturbations. These benign results contrast with those we obtain in the next section once we introduce complexity.

Formally, consider the (no) uncertainty model, $\mathcal{N}^j(\sigma) = \{\sigma\}$ for each j , so that each bank knows the realization of the financial network. In this benchmark, we characterize the equilibrium under the following parametric conditions:

$$ny > \lceil \theta \rceil \quad \text{and} \quad z + y + (1 - y)p_{\text{scrap}} \geq \theta. \tag{3}$$

Here, $\lceil x \rceil$ denotes the ceiling function, that is, the unique integer such that $\lceil x \rceil - 1 < x \leq \lceil x \rceil$. The first condition in (3) says that the financial system has sufficient aggregate liquidity to meet the unexpected liquidity shock, θ . The second condition simplifies the notation but does not play an essential role.

Our characterization consists of three steps. First, we characterize a generic bank’s optimal action (and solvency) taking the payoffs and actions of other banks as given. Second, we take the asset price, p , as given and we characterize the partial equilibrium corresponding to banks’ actions and payoffs. Third, we characterize the general equilibrium price and allocations.

A bank’s optimal action depends on its *liquidity need* at date 1. The liquidity need of a bank with distance d is

$$z - Q_1[d - 1] + \theta[d = 0], \tag{4}$$

where $\theta[\cdot]$ denotes the product of θ and the indicator function. The first term captures the payment the bank needs to make on its short-term debt. The second term captures the equilibrium payment the bank receives from its forward neighbor. The last term captures the additional payment that the original distressed bank needs to make. A bank with a liquidity need of zero optimally chooses the aggressive action, $A_0^j = B$, to maximize its equity value at date 2. In contrast, a bank with a strictly positive liquidity need would be insolvent (and liquidated at date 1) if it chose the aggressive action. Thus, it is optimal for this bank to choose the precautionary action, $A_0^j = S$ (see Appendix B). By taking the precautionary action, the bank keeps its y dollars in cash and sells $1 - y$ units of legacy assets in the secondary market, obtaining an *available liquidity* of

$$l(p) = y + (1 - y)p. \tag{5}$$

If $l(p)$ is greater than the bank’s liquidity need in (4), then the bank will be able to avoid insolvency by choosing the precautionary action. Otherwise, the bank will be insolvent despite taking the precautionary action.

It follows that the bank chooses the precautionary action, $A_0^j = S$, if its liquidity need is strictly positive, and it is insolvent if its liquidity need is strictly greater than $l(p)$. We next use this characterization to solve for the partial equilibrium, that is, the optimal actions and payoffs of all banks taking the price, p , as given.

PROPOSITION 1: *(Partial equilibrium in the certainty benchmark.) Suppose the price of legacy assets is fixed at $p \in [p_{\text{scrap}}, 1]$ and condition (3) holds. Then there*




Figure 3. The partial domino effect and flight to quality in the certainty benchmark.

is a domino effect of size

$$D(p) = \left\lceil \frac{\theta}{l(p)} \right\rceil - 1, \tag{6}$$

and a flight to quality of size $F = D(p) + 1$ (see Definition 2). Both the domino effect and the flight to quality are contained, that is, $D(p) < n$ and $F < n$.

Figure 3 illustrates this result, the proof of which is relegated to Appendix B. Intuitively, a number of banks that have the shortest distance to the original distressed bank have a positive liquidity need. Thus, these banks choose the precautionary action, $A_0^j = S$. By doing so, each of these banks uses its available liquidity, $l(p)$, to meet its liquidity need to some extent. In particular, as the shock propagates from one bank to another, the liquidity need decreases by $l(p)$. The initial liquidity need, θ , is fully met only by combining the available liquidities of $\lceil \frac{\theta}{l(p)} \rceil = D(p) + 1$ banks. It follows that banks with distance $d \leq D(p)$ choose the precautionary action. Moreover, all but the last one of these banks are insolvent. The last bank with distance $d = D(p)$ avoids insolvency because it is able to meet its liquidity need fully. This bank pays its backward neighbor bank in full. Consequently, the banks with distance $d > D(p)$ have zero liquidity need. These banks choose the aggressive action, $A_0^j = B$. It follows that there is a domino effect of size $D(p)$ and a flight to quality of size $D(p) + 1$.

Equation (6) shows that the size of the domino effect is “proportional” to the ratio of the size of the shock to the banks’ available liquidity, $\theta/l(p)$. A larger shock naturally leads to a larger domino effect. A reduction in available liquidity for banks also leads to a larger domino effect. Intuitively, this is because, when $l(p)$ is lower, banks that choose the precautionary action are less able to fight the shock. We next state the main result of this section, which characterizes the general equilibrium price and allocations.

PROPOSITION 2: (*General equilibrium in the certainty benchmark.*) Consider the certainty benchmark and suppose the conditions in (3) hold. Then,

- (i) The unique equilibrium price is $p = 1$ (no fire sales).
- (ii) There is a domino effect of size $\lceil \theta \rceil - 1$ and a flight to quality of size $\lceil \theta \rceil$.
- (iii) The aggregate amount of new asset purchases is $\mathcal{Y} = n\gamma - \lceil \theta \rceil$.

This result follows by combining Proposition 1 with the secondary market-clearing condition (2). Note that the banks with distance $d \leq D(p)$ choose $A_0^j = S$ and sell all of their existing assets. The remaining banks choose $A_0^j = B$, that is, they are potential buyers of assets. Condition (3) ensures that, for any price $p \in [p_{\text{scrap}}, 1]$, the demand from potential buyers exceeds the supply from distressed banks. This implies that the unique equilibrium price is $p = 1$. Given this price, the size of the domino effect is characterized by Proposition 1. The aggregate new asset purchases is calculated by considering the asset demand by potential buyers net of the legacy asset supply by distressed banks (see the proof in Appendix B).

Intuitively, if the domino effect is only partial and banks know the financial network, then there exist safe banks that will not realize losses from cross exposures and know that much. These banks do not sell assets and are ready to use their dollars to purchase assets from distressed banks. When the aggregate liquidity of the financial system is sufficiently large (see condition (3)), the demand from these potential buyers ensures that legacy assets trade at their fair price one.

Figure 4 illustrates this result by plotting the equilibrium variables as a function of the initial shock, θ . Note that the price is fixed at one, the size of the domino effect is increasing in θ , and the aggregate new asset purchases is decreasing in θ . Intuitively, as θ increases, there are more losses to be contained, which further spreads the insolvency. As the insolvency spreads, more banks keep their dollars in cash, which lowers \mathcal{Y} . Note, however, that \mathcal{Y} decreases “smoothly” with θ . These results offer a benchmark for the next section, where we show that, once auditing becomes costly, both D and \mathcal{Y} may experience large changes with small increases in θ .

III. Equilibrium with Complexity

We next introduce our key ingredient, complexity, which we model as banks’ uncertainty about cross exposures. As we will see, in this context when the shock is small, the system behaves exactly as in the benchmark. But when the shock is large, banks’ uncertainty about cross exposures becomes relevant since they need to understand distant linkages. Their inability to figure out these linkages leads to a complex environment and increases banks’ perceived counterparty risk. This increase in complexity (and associated counterparty risk) overturns the relatively benign implications of the benchmark environment.




Figure 4. Equilibrium in the certainty benchmark. The top, middle, and bottom panels, respectively, plot the size of the domino effect, the asset price, and the aggregate new asset purchases as a function of the losses in the originating bank.

Formally, we consider the uncertainty model

$$\mathcal{N}^j(\sigma) = \left\{ \tilde{\sigma} \mid \begin{array}{l} \text{such that } b^j \text{ is in slot } i \text{ and } b^{\text{forward neighbor}} \text{ is in slot } i - 1, \\ \text{where } i = \sigma(j). \end{array} \right\}. \quad (7)$$

Note that the bank knows its slot and its forward neighbor's slot, and thus, it also knows the identity of its forward neighbor bank. However, the bank is unsure about how the rest of the banks are ordered in the rest of the circle.¹⁴ We next characterize the equilibrium by repeating the analysis of Section II for this uncertainty model. The characterization similarly consists of three steps: (i) banks' optimal actions, (ii) partial equilibrium for a given p , and (iii) general equilibrium price and allocations.

Recall from Section I that in the certainty benchmark a bank (with distance d) chooses the precautionary action, $A_0^j = S$, if and only if its liquidity need is strictly positive. With uncertainty, the bank does not necessarily know its

¹⁴ In an earlier version, we assumed that each bank also knows the identity of its backward neighbor (which has claims on it). This assumption does not change any of the results. Intuitively, this information would only be relevant if the bank faced uncertainty about whether its backward neighbor will withdraw its debt claims early. Appendix A shows that in this model all banks withdraw their debt claims early. Thus, the identity of the backward neighbor is not useful information.

exact liquidity need in (4). This is because the bank does not necessarily know the amount, $Q_1[d - 1]$, it will receive from its forward neighbor. Nonetheless, Appendix C shows that the characterization of the bank's optimal action is equally simple in this case: it chooses its action, A_0^j , as if it will receive with certainty the lowest possible payment from its forward neighbor.

Using the fact that banks have only local knowledge of the network, we can further characterize their optimal actions. First, consider a bank with distance $d \leq 1$. Given the uncertainty model in (7), this bank knows its distance. Consequently, it knows the payment, $Q_1[d - 1]$, it will receive from its forward neighbor. Thus, the optimal action of this bank is characterized exactly as in the certainty benchmark.

Next consider the optimal action of a bank with distance $d \geq 2$. This bank is uncertain about its distance, and it finds possible all distances $\tilde{d} \in \{2, 3, \dots, n - 1\}$. Consequently, it does not necessarily know the payment, $Q_1[\tilde{d} - 1]$, it will receive from its forward neighbor. The worst-case scenario obtains when the bank is at the closest possible distance, $\tilde{d} = 2$. It follows that this bank chooses its optimal action as if it is at distance $\tilde{d} = 2$. Put differently, *the banks that are uncertain about their distance to the distressed bank choose their precautionary action as if they are closer to the distressed bank than they actually are.*

The following proposition, which is the analogue of Proposition 1 for this setting, characterizes the partial equilibrium.

PROPOSITION 3: *(Partial equilibrium with complexity.) Consider the economy with network uncertainty. Suppose the price of legacy assets is fixed at $p \in [p_{\text{scrap}}, 1]$ and condition (3) holds. Recall that $D(p) = \lceil \frac{\theta}{l(p)} \rceil - 1$ denotes the size of the domino effect in the certainty benchmark (see equation (6)).*

(i) *If $\theta \leq 2l(p)$ (so that $D(p) \leq 1$), then there is a domino effect of size $D(p)$ and a flight to quality of size $F = D(p) + 1$.*

(ii) *If $\theta > 2l(p)$ (so that $D(p) \geq 2$), then there is a domino effect of size $D(p)$ and a flight to quality of size $F = n$.*

Figure 5 illustrates this result by plotting the equilibrium actions (and volencies) corresponding to the two cases. The first case concerns a liquidity shock, θ , that is smaller than the available liquidity of two banks (i.e., the original distressed bank and its backward neighbor). In this case, part (i) of the proposition (and the top panel of Figure 5) shows that the partial equilibrium is the same as in the certainty benchmark. To see this, recall that banks at distance $d \geq 2$ act as if they are at distance 2. In this case, the liquidity shock is sufficiently small that the bank at distance 2 does not suffer any losses from cross exposures. Consequently, banks with distance $d \geq 2$ optimally choose the aggressive action. This leads to the same partial equilibrium as in the certainty benchmark. The proof in Appendix C formalizes this argument.

The second case concerns a liquidity shock, θ , which is greater than the available liquidity of two banks. In this case, part (ii) of the proposition (and the bottom panel of Figure 5) shows that the equilibrium features a much larger flight to quality than the certainty benchmark. In particular, all banks in the financial system choose the precautionary action, $A_0^j = S$. To see this, note that




Figure 5. The partial domino effect and the precautionary actions with network uncertainty. The top panel illustrates the first case of Proposition 3, $\theta \leq 2l(p)$. The bottom panel displays the second case, $\theta > 2l(p)$.

the liquidity shock in this case is sufficiently large to generate a domino effect of at least size 2. Thus, it is optimal for a bank at distance 2 to choose the precautionary action, $A_0^j = S$. Consequently, banks with distance $d \geq 2$ also choose the precautionary action. This leads to a flight to quality of size n .

Intuitively, if the domino effect (generated by the initial shock) is sufficiently small, the environment is *simple* in the sense that banks' uncertainty about the financial network does not generate counterparty risk. In particular, banks that are uncertain about their distance \tilde{d} can rule out an indirect hit. Hence, these banks continue to act as in the certainty benchmark despite being uncertainty averse. In contrast, if the domino effect is sufficiently large, then the environment is *complex* in the sense that banks' network uncertainty generates counterparty risk. That is, banks that are uncertain about their distance cannot rule out an indirect hit. Since they are uncertainty averse, they respond by choosing the precautionary action.

The following proposition jointly characterizes the equilibrium price and allocations.

PROPOSITION 4: (*General equilibrium with complexity.*) Consider the economy with network uncertainty and suppose the conditions in (3) hold.

(i) *Unique fair price equilibrium:* If $\theta \leq 2l(p_{\text{scrap}})$, then there is a unique equilibrium with price $p = 1$ (no fire sales). There is a domino effect of size $D(1) = \lceil \theta \rceil - 1$ and a flight to quality of size $F = \lceil \theta \rceil$. The aggregate amount of new asset purchases is $\mathcal{Y} = ny - \lceil \theta \rceil$.

(ii) *Unique fire sale equilibrium:* If $\theta > 2$, then there is a unique equilibrium with price $p = p_{\text{scrap}}$ (fire sales). There is a domino effect of size $D(p_{\text{scrap}}) = \lceil \frac{\theta}{l(p_{\text{scrap}})} \rceil - 1$ and a flight to quality of size $F = n$. The aggregate amount of new asset purchases is $\mathcal{Y} = 0$.

(iii) *Multiple equilibria:* If $\theta \in (2l(p_{\text{scrap}}), 2]$, then there is a fair price equilibrium as in part (i) and a fire sale equilibrium as in part (ii).

Figure 6 illustrates this result. There is a unique equilibrium for sufficiently small and large levels of θ , but there are multiple equilibria for intermediate levels of θ . Note also that the fair price equilibrium is the same as the equilibrium in the certainty benchmark (see Proposition 2), while the fire sale equilibrium is very different. In particular, the fire sale equilibrium features a greater flight to quality than the certainty benchmark ($F = n$ vs. $F = \lceil \theta \rceil$). This large precautionary reaction generates a fire sale in the secondary asset market ($p = p_{\text{scrap}}$). It also leads to a larger credit crunch than the certainty benchmark ($\mathcal{Y} = 0$ vs. $\mathcal{Y} = ny - \lceil \theta \rceil$).

Proposition 4 is our main result and it shows that, as the initial losses (measured by θ) increase, the equilibrium makes a very large and discontinuous jump compared to the certainty benchmark. This jump could be realized either in the region of multiple equilibria if banks coordinate on the precautionary action, or in the region of a single equilibrium if initial losses are sufficiently large. The resulting equilibrium features a flight-to-quality episode that is disproportionate to the size of the initial shock. The central ingredient for this result is complexity, that is, banks' uncertainty about the financial network.

The proof of Proposition 4 is relegated to Appendix C. Intuitively, when θ is sufficiently small, the size of the domino effect is manageable (i.e., below the critical threshold of two) regardless of the price of legacy assets. In this case, the environment is simple (i.e., banks' network uncertainty does not generate counterparty risk). In contrast, when θ is sufficiently large, the size of the domino effect is unmanageable and the environment is complex (i.e., banks' network uncertainty generates counterparty risk) regardless of the price.

For intermediate levels of θ , the interaction between the asset price and complexity of the environment generates multiple equilibria. If legacy assets trade at their fair price, then there is more market liquidity to counter the initial liquidity shock. This leads to a smaller domino effect and a simple environment. Since the environment is simple, banks that are uncertain about their distance are potential buyers in the secondary market, which ensures




Figure 6. Equilibria with network uncertainty. The panels plot various equilibrium variables as a function of the shock, θ . The top panel plots the size of the domino effect in partial equilibrium, $D(p)$, for price level $p = p_{\text{scrap}}$ (dashed line) and price level $p = 1$ (solid line). The second panel plots the general equilibrium price, p . The last panel plots the aggregate new asset purchases, Y .

that legacy assets trade at their fair price. Set against this benign scenario is the possibility of a fire sale equilibrium, in which the price of legacy assets collapses. This reduces market liquidity available to distressed banks, which leads to a larger domino effect and a complex environment. Facing a complex environment, banks that are uncertain about their distance panic and sell their legacy assets, which reinforces the collapse of asset prices.

Note also that, whenever there are multiple equilibria, the fair price equilibrium Pareto dominates the fire sale equilibrium for all banks. This observation suggests that there are externalities in our setting, which we analyze next.

IV. Complexity Externality and Policy Implications

Our model features a complexity externality, which stems from the dependence of banks' counterparty risk on the endogenous size of the domino effect. In particular, any action that exacerbates the domino effect increases the counterparty risk perceived by banks that are uncertain about the financial network,

and they dislike this effect. Our model features two variants of this complexity externality (one nonpecuniary, one pecuniary), each of which supports different types of policies. The rest of this section discusses the two variants and their policy implications.

A. *Nonprice Complexity Externality and Bank Bailouts*

To illustrate this externality, it is useful to start with a simple example. Consider an alternative economy with a continuum of (measure one) agents, $i \in I$, with utility functions

$$u(x^i) - ca^i.$$

Here, x^i denotes agent i 's endowment, $a^i \in \{0, 1\}$ denotes a costly action taken by agent i , and $u(\cdot)$ denotes a standard and strictly concave utility function. Suppose also that each x^i is a random variable with mean one and variance

$$1 - \int_I a^i di.$$

In particular, each agent can take a costly action that can (slightly) reduce the variance of endowments of all agents in this economy.

In this example, consider, respectively, the equilibrium and the planner's allocations. In equilibrium, no agent takes the costly action because she incurs a positive cost while having only a negligible effect on the variance of its own consumption. On the other hand, for sufficiently small $c > 0$, a social planner would have all agents choose $a^i = 1$. This allocation gives each agent a constant consumption at a relatively small cost (by assumption), which is a Pareto improvement over the equilibrium allocation.

In this example, the competitive equilibrium is Pareto inefficient because of a nonpecuniary externality that operates through the production technology. In particular, an agent i does not internalize the fact that her action affects the endowment variance of all other agents. By choosing $a^i = 0$, this agent exerts a negative externality on all other agents, which leads to a Pareto inefficiency.

We next describe the nonprice complexity externality of our model, which is reminiscent of the externality in this example. To this end, consider the setup of Proposition 3, that is, suppose there is network uncertainty and prices are exogenously fixed (which shuts down any pecuniary channels). Suppose also that

$$\theta \in (2l(p), 3l(p)),$$

which ensures that there is a domino effect of size two and a flight to quality of size n (see Proposition 3). In particular, all banks choose the precautionary action, $A_0^j = S$. Banks' Minimax utility at date 0 (see equation (1)) is given by

$$\begin{cases} 0 & \text{if } d < 2, \\ 3l(p) - \theta \in (0, l(p)) & \text{if } d \geq 2. \end{cases} \tag{8}$$

In this setting, consider a modification of equilibrium by introducing the possibility of a “bailout” of the distressed bank, b^0 , by other banks. In particular, each bank j can choose to contribute some of her date 0 dollars, y , to a bailout fund. Without loss of generality, suppose banks’ actions are restricted to a binary set, $\{0, \frac{\theta}{n}\}$, that is a bank either contributes zero dollars or $\frac{\theta}{n}$ dollars to the bailout fund. Note that contributing $\frac{\theta}{n}$ dollars is feasible because the banks have sufficient aggregate liquidity by condition (3). Once all contributions are made, the total amount in the fund is used to pay some (possibly all) of the liquidity need, θ , of bank b^0 . The equilibrium is then characterized as before with a potentially lower level of the shock for the original distressed bank and a lower level of date 0 dollars for the contributing banks.

In this modified equilibrium, banks optimally choose to contribute zero dollars to the bailout fund (and thus, the equilibrium remains unchanged). To see this, consider a bank with distance $d \geq 2$. By contributing to the bailout fund, this bank incurs a positive cost while receiving no benefits. This is because this bank alone is not able to change the size of the domino effect (since it is infinitesimal by assumption). On the other hand, consider a social planner that requires all banks to contribute $\frac{\theta}{n}$ dollars. With this bailout policy, the original distressed bank remains solvent and the domino effect disappears. In particular, banks’ counterparty risk also disappears. Consequently, banks choose the aggressive action, that is, they keep their legacy assets and they spend their remaining, $y - \frac{\theta}{n}$, dollars to acquire new assets. Their Minimax utility at date 0 is given by

$$(1 - y)R + \left(y - \frac{\theta}{n}\right)R = \left(1 - \frac{\theta}{n}\right)R.$$

Comparing this expression with the utility in (8) shows that this bailout policy leads to a Pareto improvement as long as n or R is sufficiently large. The fact that banks with distance $d < 2$ are better off is not remarkable because these banks are (either directly or indirectly) bailed out. However, it is remarkable that all other banks at distance $d \geq 2$ are also better off.

The equilibrium is Pareto inefficient for the same reason as in the earlier example. Each bank with distance $d \geq 2$ does not internalize that its contribution, $\frac{\theta}{n}$, would mitigate the domino effect and thus reduce the counterparty risk faced by other banks. By not contributing, this bank exerts a negative externality on other banks, which we refer to as the *nonprice complexity externality*. A bank bailout policy generates a Pareto improvement by internalizing this externality. Viewed differently, network stability (and similarly, endowment stability in the earlier example) is a public good. Each bank would like to enjoy this good because it reduces its counterparty risk. However, each bank would rather not incur the costs and free ride on other banks. The bailout could be viewed as the provision of the public good of stability, which solves the free rider problem.

We stress that the nonprice complexity externality is different from the fire sale externality that is common in the literature. In particular, in the above

setting there cannot be a fire sale externality because the asset price is fixed. We next consider the setting with endogenous asset price to illustrate the second variant of the complexity externality.

B. Price Complexity Externality and Government Asset Purchases

This externality operates through the interaction of legacy asset prices and the size of the domino effect. In particular, a bank that decides to sell assets (i.e., chooses the precautionary action, $A_0^j = S$) has a (small) negative impact on asset prices. This in turn has a (small) positive impact on the size of the domino effect. In particular, with a lower asset price, the available liquidity, $l(p)$, of each bank is lower. Thus, the crisis is contained after a greater number of insolvencies (see equation (6)). The increase in the size of the domino effect increases the counterparty risk faced by other banks and lowers their welfare, demonstrating the price complexity externality.

The price complexity externality is what leads to multiple Pareto-ranked equilibria in our setup, as we have already seen in Proposition 4. In particular, an increase in counterparty risk due to a reduction in the legacy asset price not only lowers the welfare of many banks, but also induces these banks to take extreme precautionary measures, which includes further asset sales. The sale of assets by banks in panic mode reduces asset prices further, which leads to a vicious cycle culminating in the fire sale equilibrium. In contrast, an increase in asset prices reduces the counterparty risk, which may mitigate precautionary measures and turn more sellers into buyers, leading to a virtuous spiral toward the fair price equilibrium. In particular, a social planner that puts a floor on asset prices (e.g., through an asset purchase policy) can generate a Pareto improvement by coordinating banks on the fair price equilibrium.

We stress that the price complexity externality is also different from the usual fire sale externality (e.g., in Kiyotaki and Moore (1997b) or in Lorenzoni (2008)). It is true that both externalities operate through asset prices. However, the commonalities end there because the particular channels for the two externalities are different. In a fire sale externality, the decrease in asset prices erodes the net worth of financial institutions that are *natural buyers* of this asset. This in turn tightens these institutions' borrowing constraints, which lowers their welfare and puts further downward pressure on asset prices. Instead, in the price complexity externality, the decrease in asset prices increases the counterparty risk perceived by financial institutions that are uncertain about the network. The increase in counterparty risk (as opposed to binding constraints) is what lowers the welfare of these institutions. Moreover, their precautionary reaction (as opposed to binding constraints) is what puts further downward pressure on asset prices.

This comparison also suggests that the price complexity externality could be much more potent than the fire sale externality. To see this concretely, consider a drop in the price of subprime mortgage-backed securities. From the lens of the conventional fire sale externality, this shock should mostly affect the natural buyers of these securities. In particular, it should not greatly affect

the institutions that specialize in other businesses or other asset classes (or natural buyers that happen not to hold the securities at the time of the shock). In contrast, from the lens of the price complexity externality, this shock could have a much bigger impact. In particular, suppose the shock is sufficiently large that it leads to the failure of some natural buyers and generates the possibility of domino effects. This in turn increases the counterparty risk faced by all financial institutions that are uncertain about the financial network. In practice, this includes virtually all financial institutions, illustrating the much greater scope of the price complexity externality.¹⁵

Finally, it is also worth noting that the price complexity externality amplifies the standard overleverage result, for example, in Lorenzoni (2008). The previous literature notes that banks tend to leverage and invest too much *ex ante* because they do not internalize the fire sale externalities that their deleveraging generates during crises. Our model suggests that banks might choose to overleverage and overinvest for an additional reason: because they do not take into account the increase in counterparty risk that their asset sales generate when domino effects become possible.

V. Discussion of Key Assumptions

The previous sections establish that small shocks can dramatically increase banks' counterparty risk in view of banks' uncertainty about the financial network of cross exposures. To illustrate the main mechanism, we make a number of strong assumptions. In particular, we assume that the shock is unanticipated, that the network has a particularly simple form, and that banks face Knightian uncertainty about the financial network. In this section, we discuss and support these assumptions.

A. *The Role of the Unanticipated Shock*

In our model, banks do not anticipate the surprise shock (at date -1). If banks assigned a sufficiently large probability to the shock, they would also naturally seek insurance against it at the unmodeled date -1 , and the financial system would become more resilient to the shock. However, the kind of shocks that we are trying to capture are by their very nature either impossible or very costly to insure against. One could imagine two versions of this insurance: (i) a specific insurance contract, such as "bank b^j receives some dollars if bank b^0 loses θ dollars" or (ii) a blanket insurance contract, such as "bank b^j receives some dollars if its forward neighbor is insolvent (i.e., a CDS on the forward neighbor)." We next argue that the first type of (specific) insurance is implausible, while the second type of (blanket) insurance is likely to be very costly and possibly ineffective during severe systemic events.

¹⁵ In our model, we assume for simplicity that the natural buyers of the asset are the same as banks that face network uncertainty. Instead, this discussion suggests that natural buyers are likely to be a subset of the banks that face network uncertainty. Our model could be easily modified to incorporate this feature.

Specific insurance contracts are not plausible due to the complexity of banks' balance sheets. This complexity, which is the reason for network uncertainty in our model, also naturally makes the event very difficult to describe *ex ante*. Arguably, banks would be unable to tell if the shock would hit bank b^0 , another bank, or a combination of banks. In addition, they would be unable to tell the size of the shock, and thus the size of the domino effect that it would bring about. These observations resonate well with the run-up to the recent crisis. Even though the possibility of the subprime losses was discovered in the summer of 2007, the exact form and magnitude of these losses remained not well understood until 2008. Empirical evidence in support of this point can be gleaned from the *BIS Quarterly Review* for December 2007 (Fender and Hordahl (2007)). Graph 4 of this review illustrates that the 5-year CDS spreads increased for banks on average between June 2007 and December 2007. However, the same graph also shows that the *dispersion* in large banks' spreads actually *decreased* over the same period, that is, banks' CDS spreads became closer to one another. The *BIS Review* interprets this lack of differentiation as "possibly reflecting a continuing lack of transparency about their (banks') exposures" (p. 7)

Even if specific insurance was not possible, banks could still purchase blanket insurance against all possible shocks that could hit the financial system. For example, a bank could buy a CDS on its forward neighbor, which would effectively reduce its exposure, z , and immunize it against any shock to the network. However, this type of insurance is very costly because it comes as a bundle. In particular, while it insures banks against systemic shocks that could generate domino effects, it also insures them against garden variety shocks that lead to the failure of their counterparty without a systemic component. Banks might prefer not to buy insurance against the latter types of shocks, which are always present in the financial system. If this is the case, then banks might find it optimal not to buy CDS insurance (or other types of blanket insurance) even though they recognize the benefits that this would bring about during a systemic crisis.

A separate restriction on both specific and blanket insurance concerns the counterparty risk in CDS contracts. During a systemic event, the sellers of these contracts might themselves become distressed, and thus might be unable to deliver on their promises. This point is supported by the Financial Crisis Inquiry Commission's (FCIC) January 2011 report on the causes of the financial crisis, which quotes the managing director at AIG Financial Products as saying: "And we're one of the few guys who can do that (sell disaster insurance). Because if you think about it, no one wants to buy disaster protection from someone who is not going to be around . . . That was AIGFP's sales pitch to the Street or to banks."¹⁶ Ironically, when the disaster struck AIGFP was one of the first institutions to collapse. In particular, AIGFP would be unable to deliver on its

¹⁶ Available at http://fcic-static.law.stanford.edu/edn_media/fcic_reports/fcic_final_report_full.pdf.

insurance promises absent an intervention by the Fed, which illustrates further the practical difficulty of obtaining insurance against systemic events.¹⁷

Finally, one might imagine that banks could mitigate the counterparty risk in CDS contracts by requiring sellers to post sufficient collateral. Appendix D incorporates collateralized insurance contracts into our model and shows that the results are robust to this extension. There, we consider the possibility that banks could purchase CDS insurance on their counterparties at date 0, once the nature of the shock is understood (which is the main difference from blanket insurance). We also require each CDS contract to be individually and fully collateralized. In this setting, while banks demand counterparty insurance at date 0, the supply of this insurance is restricted because of sellers' collateral constraints. In particular, sellers within the financial network choose not to pledge their collateral in an insurance contract in view of their own payoff uncertainty (in fact, they would rather demand insurance for their own cross exposures). Thus, the only insurance supply comes from sellers that are outside the network. When the collateral of these sellers is small relative to the size of the network, allowing for collateralized counterparty insurance does not change our qualitative results. This analysis is consistent with the behavior of the CDS markets during the recent Bear Sterns and Lehman debacles. As described by Duffie (2011), the demand for counterparty insurance in both episodes spiked, but this demand could not be met by insurance sellers.

B. The Role of the Network Structure

For simplicity, we assumed that the network structure takes the form of a circle. This is an example of a highly incomplete network, which tends to increase the likelihood of contagion (see, for example, Allen and Gale (2000)). To illustrate the point, consider a more complete network in which each bank has exposures (of $z/2$ each) to two other banks. In this case, a shock θ that leads to the insolvency of bank b^0 would lead to a smaller liquidity need for each of the backward neighbor banks [$\frac{\theta - l(p)}{2}$ instead of $\theta - l(p)$]. Consequently, it would take a larger shock, θ , to generate a domino effect.

In contrast, the incompleteness of the network does *not* play an important role for our main result (panics driven by complexity) as long as we control for banks' information appropriately. To see this, consider the two networks illustrated in Figure 7. The left panel concerns the circle network we have analyzed, with the additional assumption that banks know the identities of their two forward neighbors (as opposed to just their immediate neighbor). It is easy to see that all banks choose the precautionary action, $A_0^j = S$, if there

¹⁷ The unreliability of CDS insurance in the run-up to the crisis is also emphasized by Vause (2010), in the *BIS Quarterly Review* for December 2010, as follows: "Market participants responded to increased concern about counterparty risk by buying protection on CDS dealers But none of these trading responses represented a comprehensive solution to the problem. Buying protection on one dealer from another dealer is of limited value if there are systemic concerns about the robustness of counterparties in the market" (p. 61).




Figure 7. Role of network completeness. The two panels illustrate flight-to-quality episodes in two networks that differ in their completeness, but not in banks’ information about the network (measured as the number of other banks whose locations each bank knows).

is a domino effect of length at least three, which in turn happens if the initial losses satisfy $\theta > 3l(p)$.

The right panel illustrates an alternative network with three tiers, in which each bank in the bottom two tiers has exposures to two other banks (and the initially distressed bank is in the top tier). For this network, we assume that the banks know the identities of their immediate neighbors only. With this assumption, the banks have the same amount of information as in the first case in the sense that they know the locations of *exactly two other banks*. In this case, the two backward neighbors of the initially distressed bank each have a liquidity need of $\frac{\theta - l(p)}{2}$. If this is greater than their available liquidity, that is, $\frac{\theta - l(p)}{2} > l(p)$, then a number of the banks in the bottom row also have a positive liquidity need. Moreover, none of the banks in the bottom row can rule out an indirect hit because they only know the identities of their immediate neighbors. It follows that all of the banks in the bottom row choose the precautionary action if $\theta > 3l(p)$, which is the same condition as in the first case. Put differently, even though the second network is more complete, it takes a shock of the same size to generate a flight-to-quality episode in both networks.¹⁸

¹⁸ It is also instructive to show a more general version of this result. Consider the network on the right panel of Figure 7 with a large number of tiers (as opposed to three). Suppose banks know their neighbors up to distance k , so that they know the locations of $K \equiv 2 + \dots + k$ other banks. It is easy to see that a sufficient condition for a flight-to-quality episode is $\theta > (K + 1)l(p)$. This is the same condition that generates a flight-to-quality episode in the circle network when banks know their forward neighbors up to distance K . Thus, the key determinants of a flight-to-quality episode is the size of the shock, θ , the banks’ information, K (measured by the number of other bank locations a bank understands), and the available liquidity of each bank, $l(p)$. In particular, the completeness of the network is not an important determinant. Note, however, that this argument takes banks’ information, K , as a constant independent of the network structure. There are at least

Intuitively, while a more complete network features a shorter domino effect, it also generates a greater informational burden for banks. The more neighbors a bank has, the more balance sheets it needs to inspect to assess its financial health (and the financial health of its neighbors, and their neighbors, etc.). Consequently, controlling for the bank's information, it takes a smaller domino effect to trigger a flight-to-quality episode in a more complete network. Figure 7 illustrates that, under appropriate assumptions (related to the symmetry of the network), the two forces exactly cancel and the completeness of the network does not affect the incidence of panics.

It is also worth noting that in practice banking networks appear to be highly incomplete and specialized. Direct evidence for this point can be gleaned from the Bank of England's Financial Stability Report for June 2009. Chart 3.8 of this report plots the network of large exposures between U.K. banks. While the U.K. banking network does not exactly take the form of a circle (it looks more like a hub-and-spoke network), it shares the key characteristic of being highly incomplete. Similarly, the U.S. financial network is highly incomplete and concentrated, as evidenced by the recent regulatory attempts (in particular, the Dodd–Frank Act) to curb large and systemic banks' exposures to a single counterparty.¹⁹

two caveats to this approach. First, one could imagine that forming connections provides banks with information about their neighbors as a byproduct, so that banks in the more complete network naturally have a greater K . While this effect might be important for banks' information about their immediate neighbors, it is less relevant for their information about indirect links (neighbors' neighbors and so on). Consequently, when banks need to know the network at some depth to rule out a hit (i.e., when θ is large), the additional information provided by the more complete network becomes negligible. Second, one could also wonder whether the result (that network structure does not matter) would continue to hold if the information, K , was endogenously chosen by banks at some audit cost. This is in fact the case under appropriate assumptions. In particular, suppose the (nonpecuniary) audit costs are captured by the same function, $c(K)$, in both networks. Suppose also that $R < c(\lceil \frac{\theta}{K_{(P_{scrap})}} \rceil) - c(0)$, that is, the maximum benefit from obtaining sufficient information to rule out a hit is lower than the cost of acquiring that information. It then follows that banks in either network choose not to acquire information, which in turn leads to a flight-to-quality episode.

¹⁹ An important reason for the incompleteness of the network in practice is the heterogeneity in the benefits and costs of forming links. To see this, consider the example of a bank that enters an interest rate swap agreement in the OTC derivative markets. Depending on the movements of the yield curve, this contract might generate exposures to banks' counterparties. In this setting, a more complete network would correspond to the case in which the bank buys a small amount of swaps from a large number of counterparties, thereby spreading its exposures across the network. However, this is unlikely to be optimal for several reasons. First, some intermediaries might lack the necessary expertise to value this swap, especially if the terms of the contract are complex. If this is the case, the bank would have to enter the agreement with a few specialized intermediaries. Second, even the specialists' valuations often differ from one another, depending on their portfolio risks and beliefs. If this is the case, the bank would find it optimal to enter the agreement with the specialists that offer the best price, leading to a highly incomplete network. More generally, heterogeneity in financial institutions' valuations of derivative contracts would naturally lead to an incomplete network of exposures in the OTC derivative markets.

C. The Role of Knightian Uncertainty about the Network

The Knightian uncertainty, and the corresponding Maximin representation in (1), is not essential for our results. In particular, our qualitative results also apply in a standard expected utility framework as long as banks are risk averse (or as long as they are averse to bankruptcy in view of potential fixed costs). To see this, suppose instead that the bank chooses $A_0^j \in \{S, B\}$ to maximize $E^j[u(q_2^j(\tilde{\sigma}))]$, where $u(\cdot)$ denotes a concave utility function and $E^j[\cdot]$ denotes the bank's expectation over the random network $\tilde{\sigma}$. Suppose also that the bank assigns each network in $\mathcal{N}^j(\sigma)$ equal probability (i.e., it applies the principle of insufficient reason). Suppose $\theta > 2l(p)$, so that the version of the model with Knightian uncertainty features a flight-to-quality episode. We claim that, under appropriate assumptions, this version also generates flight to quality. To see this, consider a bank that is at distance $d \geq 2$, which can be in one of three distinct states. First, with probability $\frac{1}{n-2}$, the bank is exactly at distance $D(p)$. In this case, the bank obtains the utility level $u(0)$ if it chooses $A_0^j = B$, but a greater utility level $u((1 + D(p))l(p) - \theta)$ if it chooses the precautionary action $A_0^j = S$. Second, with probability $\frac{D(p)-2}{n-2}$, the bank's distance satisfies $2 \leq d \leq D(p) - 1$. In this case, the bank obtains the utility level $u(0)$ regardless of its choice. Third, with the remaining probability $1 - \frac{D(p)-1}{n-2}$, the bank's distance satisfies $d \geq D(p) + 1$. In this case, the bank obtains the utility level $u(R)$ if it chooses $A_0^j = B$, but a lower utility level $u(1)$ if it chooses $A_0^j = S$. Collecting the three cases, the bank chooses $A_0^j = S$ if

$$\frac{1}{n-2}(u((1 + D(p))l(p) - \theta) - u(0)) \geq \left(1 - \frac{D(p) - 1}{n - 2}\right)(u(R) - u(1)).$$

As this expression illustrates, the bank chooses the precautionary action as long as $u(0)$ is sufficiently low, for example if $u(0) \rightarrow -\infty$, which in turn leads to a flight-to-quality episode. Intuitively, the left side of the expression captures the precautionary benefit of avoiding bankruptcy, while the right side captures the opportunity cost of holding liquid assets. When the bank is sufficiently averse to risk (or bankruptcy), it optimally chooses the precautionary action $A_0^j = S$ even if it maximizes subjective expected utility.

In the main analysis, we consider Knightian uncertainty instead of expected utility for two reasons. First, it provides analytical tractability by enabling us to focus on the worst-case scenario, instead of specifying a distribution over $\mathcal{N}^j(\sigma)$ and taking expectations. Second, and more importantly, Knightian uncertainty seems more appropriate for our context than quantifiable risk. Given the complexity of the network of cross exposures in real financial markets, banks are unlikely to have a unique belief distribution over various possible networks. Microeconomic studies (both empirical and theoretical) have argued that economic agents are more averse to this type of uncertainty compared to quantifiable risks. The Maximin representation enables us to capture this feature in a tractable way. That said, the Maximin representation corresponds to an extreme form of ambiguity aversion that

provides quantitative amplification for our results. We could have instead used the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005), which would lead to similar qualitative conclusions but milder quantitative effects.

VI. Conclusion

In this paper, we provide a model that illustrates how fire sales can arise even when financial markets are deep and the shock is small relative to the wealth in the financial network. The key ingredient for this outcome is complexity, which we have captured as banks' uncertainty about the network of cross exposures. This feature, which is a dormant factor in normal times, generates counterparty risk when banks are unable to figure out their exposures to an indirect hit.

We also show that there is a powerful feedback between fire sales and complexity. More severe fire sales amplify the potential domino effects and increase banks' counterparty risk. This triggers a precautionary reaction from potential asset buyers, which pull back and exacerbate the fire sale. In extreme scenarios these potential buyers can turn into sellers, leading to a complete collapse in secondary markets.

We only partially explore policy questions, but it is apparent that our environment creates many policy opportunities. In particular, the complexity externality supports government actions during crises that are aimed at mitigating domino effects (e.g., bailing out distressed banks or asset purchases), as well as those that are aimed at reducing network uncertainty (e.g., stress testing, and widespread guarantees to banking liabilities or assets). In addition, the complexity externality also supports preemptive measures that are aimed at simplifying (and increasing the transparency of) the financial network, for example, moving OTC transactions to exchanges.

A question that emerges in our environment is whether banks can aggregate their (local) information about the financial network. In our model, banks cannot credibly share their information if we assume that distressed banks suffer losses resulting from revealing that they are distressed (which is likely to be the case in reality). This is because banks that are close to the original distressed bank have an incentive to misreport their distance, which prevents the aggregation of information. More broadly, one could imagine many other reasons why information production and sharing during a crisis is inefficient, which emphasizes the importance of policies that provide information (e.g., stress testing, collecting data on OTC transactions, living wills, etc.).

As a parting thought, we note that the particular insolvency motive we consider raises the question of what would happen if the distressed institutions chose to gamble for resurrection by not selling their assets, which would improve their outcome in good states at the cost of greater bankruptcy risk. Our model suggests that gambling for resurrection may be a mixed blessing in aggregate. Gambling by potential buyers, that is, institutions that are far from the domino effect but that do not know this, would limit fire sales and the downward spiral of prices. On the other hand, gambling by institutions near

the domino effect would exacerbate contagion and trigger the complexity mechanism. This issue also points to important policy trade-offs for the decision as to which institutions to guarantee during a systemic event.

Initial submission: April 2, 2011; Final version received: May 15, 2013
Editor: Campbell Harvey

Appendix A: Banks' Debt Rollover Decision

In the main text, we simplify the model by assuming that all short-term debt claims must be settled at date 1. In this appendix, we consider the extension of the model with banks' rollover actions, and we show that the equilibrium is unchanged as all banks choose to withdraw their debt claims immediately.

To see this, consider an extension in which each bank b^j has an additional action at date 1, $A_1^j = W(\tilde{z})$, for some $\tilde{z} \in [0, z]$. A bank that chooses $A_1^j = W(\tilde{z})$ withdraws \tilde{z} dollars of its debt claims on its forward neighbor bank at date 1, and rolls over the remaining $z - \tilde{z}$ dollars of its debt claims to date 2. In this setting, consider a distressed bank with a positive liquidity need at date 1 (e.g., the original distressed bank b^0). This bank could try to obtain the required liquidity by withdrawing its debt claims at date 1 (i.e., by choosing $A_1^j = W(\tilde{z})$ for some $\tilde{z} > 0$) and/or by taking the precautionary action at date 0 (i.e., by choosing $A_0^j = S$). Taking the precautionary action is strictly costly for the bank because it sacrifices equity value at date 2. However, withdrawing debt claims is not costly. In fact, either the forward neighbor bank is insolvent, in which case withdrawing is strictly better than rolling over (recall that each bank is small and takes the debt payment of the forward neighbor bank as given), or the forward bank is solvent, in which case withdrawing and rolling over generate the same amount of equity value. Hence, the bank always prefers ex post withdrawal to the ex ante precautionary actions. In other words, the liquidity pecking order is such that a bank that will need liquidity at date 1 first chooses $A_1^j = W(\tilde{z})$, and then (if there is need) resorts to ex ante precautionary measures.

Next consider the original distressed bank, b^0 , that needs at least θ dollars of liquidity. This bank withdraws a positive amount of its debt claims from its forward neighbor, that is, $A_1^0 = W(\tilde{z})$ for some $\tilde{z} > 0$. This puts the neighbor bank also in need of \tilde{z} dollars of liquidity. Consequently, the neighbor bank also withdraws \tilde{z} units of its debt claims on the forward neighbor. As in Allen and Gale (2000), this triggers further withdrawals until, in equilibrium, $A_1^j = W(\tilde{z})$ for all j . Hence, the original distressed bank tries, but cannot obtain, any net liquidity through cross-debt withdrawals. In particular, this bank still needs at least θ dollars of liquidity after cross-debt withdrawals. This further implies that, in equilibrium, the bank withdraws all of its debt claims, that is, $\tilde{z} = z$. Thus, no bank rolls over its debt and all debt claims are settled at date 1. It follows that the equilibria analyzed in the main text continue to be the equilibria in this setting with a more general action space at date 1.

Appendix B: Equilibrium with Certainty

This appendix presents the analysis and proofs omitted from Section II.

Characterizing banks' optimal actions: Consider a bank with distance d . This bank's optimal action can be characterized by comparing its liquidity need in (4) and the available liquidity in (5). There are three cases to consider. First, if the bank's liquidity need is zero, then it is not distressed. Since this bank does not need dollars at date 1, it chooses the aggressive action, $A_0^j = B$, to maximize its equity value. Second, if the bank's liquidity need lies in the interval, $(0, l(p)]$, then its available liquidity is sufficient to meet its liquidity need. This bank chooses the precautionary action, $A_0^j = S$, to avert insolvency at date 1 (which maximizes its equity value at date 2). Third, if the bank's liquidity need is greater than $l(p)$, then its available liquidity is not sufficient to meet its liquidity need. This bank is indifferent between choosing $A_0^j = S$ or $A_0^j = B$, because it will be insolvent regardless of the action. Nonetheless, choosing the precautionary action increases the liquidation outcome because it enables the bank to liquidate with time: more specifically, the bank's assets yield $l(p)$ dollars with the precautionary action and zero dollars with the aggressive action. Thus, the precautionary action increases the payoff to debt holders. Given that equity holders are indifferent, we restrict attention to equilibria in which the bank (with liquidity need $> l(p)$) chooses the precautionary action, $A_0^j = S$.

Combining the three cases, note that the bank chooses the precautionary action, $A_0^j = S$, if and only if its liquidity need is strictly positive. Moreover, the bank is insolvent at date 1 if and only if its liquidity need is strictly greater than $l(p)$. We next use this characterization to solve for the partial equilibrium, that is, banks' actions and payments for a given price p .

Proof of Proposition 1

Under the claim in the proposition, the original distressed bank (with distance 0) receives full payment from its debt claims on its forward neighbor, that is, $Q_1[n - 1] = z$. Hence, the liquidity need of the bank with distance 0 is $\theta > 0$. According to the earlier characterization, this bank chooses the precautionary action, $A_0[0] = S$. If $\theta \leq l(p)$, then this bank avoids insolvency and the size of the domino effect is $D(p) = 0$, which is consistent with (6).

Suppose instead $\theta > l(p)$. In this case, the bank with distance 0 is insolvent and pays

$$Q_1[0] = z + l(p) - \theta < z, \quad (\text{B1})$$

where $Q_1[0] \geq 0$ in view of the second condition in (3).²⁰ Note that the bank receives z dollars from its claims on its forward neighbor, has $l(p)$ units of liquidity at date 1, and has to make a payment of θ dollars. Thus, its backward

²⁰ Note that $Q_1[0] = 0$ when the second condition in (3) is violated. That is, the original distressed bank pays zero on its debt claims because it is unable to make the outside payment. To accommodate this case, equation (B1) could be modified to $Q_1[0] = \max(0, z + l(p) - \theta)$. The rest of the analysis would be identical at the expense of additional notation.

neighbor bank (which has distance 1) receives $Q_1[0] < z$ from its debt claims, and it has liquidity need of

$$z - Q_1[0] = \theta - l(p). \tag{B2}$$

Here, the second expression comes from using (B1) to substitute for $Q_1[0]$. Since we are considering the case $\theta > l(p)$, the neighbor bank also has a positive liquidity need, and thus it chooses $A_0[1] = S$. If $\theta \leq 2l(p)$, then the neighbor bank's available liquidity, $l(p)$, is greater than its liquidity need. In this case, this bank is able to avoid insolvency and there is a domino effect of size $D(p) = 1$. Otherwise, the neighbor bank is also insolvent, and it pays

$$Q_1[1] = l(p) + Q_1[0].$$

From this point onwards, a pattern emerges. The payment by an insolvent bank with distance $d - 1$ is

$$Q_1[d - 1] = l(p) + Q_1[d - 2] = l(p)(d - 1) + Q_1[0].$$

Here, the first equality shows that banks' payments are linearly increasing in their distance, and the second equality uses this property to solve for the payment of the bank with distance $d - 1$ in closed form. Using this expression along with equation (B2), the bank with distance d has the liquidity need

$$z - Q_1[d - 1] = \theta - l(p)d.$$

That is, banks' liquidity needs are linearly decreasing in their distance, d . If $\theta > l(p)d$, then the bank with distance d has a positive liquidity need, and thus chooses the precautionary action, $A_0[d] = S$. If $\theta \leq l(p)(d + 1)$, this bank is able to avoid insolvency. Otherwise, it is also insolvent despite taking the precautionary action.

Next note that $D(p)$ defined in equation (6) is the first nonnegative integer such that $\theta \leq l(p)(D(p) + 1)$. Consequently, all banks with distance $d \leq D(p) - 1$ are insolvent since their liquidity needs are greater than their available liquidity, $l(p)$. These banks choose $A_0[d] = S$ to improve their liquidation outcome. In contrast, a bank with distance $D(p)$ is solvent since it can meet its losses by choosing the precautionary action, $A_0[D(p)] = S$. Since this bank is solvent, all banks with distance $d \geq D(p) + 1$ are also solvent as they do not incur losses in cross-debt claims. These banks choose the aggressive action, $A_0[d] = B$, to optimize their equity value. It follows that there is a domino effect of size $D(p)$ and a flight to quality of size $F = D(p) + 1$. The first condition in (3) also implies that $D(p) < n$ and $F < n$, completing the proof of the proposition. Q.E.D.

Proof of Proposition 2

We complete the sketch proof provided in the paragraph following Proposition 2. Suppose $p \in [p_{\text{scrap}}, 1]$ and consider banks' asset supply and demand. There are $D(p) + 1$ banks that choose $A_0^j = S$. The supply of assets from these banks is given by $(1 - y)(D(p) + 1)$. The remaining $n - D(p) - 1$ banks choose

$A_0^j = B$. The demand for assets from these banks is given by $\frac{y(n-D(p)-1)}{p}$. We claim that the demand exceeds the supply regardless of the price, that is

$$\frac{y(n-D(p)-1)}{p} > (1-y)(D(p)+1) \text{ for each } p \in [p_{\text{scrap}}, 1]. \quad (\text{B3})$$

By the secondary market-clearing condition (2), this claim ensures that the equilibrium price is $p = 1$, proving part (i). Given price $p = 1$, D and F are characterized by Proposition 1, proving part (ii). Finally, the aggregate amount of new asset purchases is equal to banks' demand for assets net of the supply of legacy assets. Taking the difference of the left-hand side and right-hand side of the inequality in (B3) and using $p = 1$, we have

$$\mathcal{Y} = ny - (D(1) + 1) = ny - \lceil \theta \rceil,$$

proving part (iii).

The remaining step is to show the claim in (B3). Recall that $D(p) + 1 = \lceil \frac{\theta}{l(p)} \rceil$. Using this expression, the claim in (B3) can be written as

$$ny > \left\lceil \frac{\theta}{l(p)} \right\rceil (y + p(1-y)) = \left\lceil \frac{\theta}{l(p)} \right\rceil l(p),$$

where the equality follows from the definition of $l(p)$ in equation (5). To show this inequality, note that

$$ny > \lceil \theta \rceil \geq \left\lceil \frac{\theta}{l(p)} \right\rceil l(p),$$

where the first inequality follows from condition (3) and the second inequality follows since $l(p) \leq 1$. This implies the claim in (B3), completing the proof. Q.E.D.

Appendix C: Equilibrium with Complexity

This appendix presents the analyses and proofs omitted from Section III. We first present the characterization of the banks' optimal actions. We then present proofs of Propositions 3 and 4.

Consider banks' optimal actions, taking the cross-debt payments, $\{q_1^j\}_j$, as given. Note that a *sufficient statistic* for bank b^j with distance d to choose action $A_0^j \in \{S, B\}$ is the amount it will receive in equilibrium from its forward neighbor. In particular, to decide the level of its precautionary measure, this bank only needs to know its liquidity need in (4), which only depends on the debt payment of its forward neighbor. Formally, if the bank chooses A_0^j at date 0 and its forward neighbor pays x at date 1, then this bank's debt payment and equity value can be written as a function $(q_1[A_0^j, x], q_2[A_0^j, x])$. However, the bank chooses A_0^j while facing uncertainty about the financial network, and consequently about $x = q_1^j(\sigma)$. More specifically, the bank knows that x lies in

some interval

$$\left[x^{\text{worst}} = \min_{\tilde{\sigma} \in \mathcal{N}^j(\sigma)} q_1^j(\tilde{\sigma}), \quad x^{\text{best}} = \max_{\tilde{\sigma} \in \mathcal{N}^j(\sigma)} q_1^j(\tilde{\sigma}) \right],$$

but it is uncertain about the exact location of x in this interval. Note also that $q_1[A_0^j, x]$ and $q_2[A_0^j, x]$ are weakly increasing in x for any choice of action. That is, the bank's debt and equity payments are increasing in the amount it receives from its forward neighbor regardless of the ex ante precautionary measure it takes. In view of the Minimax optimization (see problem (1)), it follows that the bank will choose its action as if it will receive with certainty the lowest possible payment, x^{worst} . This completes the characterization of banks' optimal actions.

Proof of Proposition 3

Case (i): $\theta \leq 2l(p)$. We prove the statement by showing that banks' payoffs (and thus solvencies) and actions are the same as in the certainty benchmark (which is characterized in Proposition 1).

To this end, it suffices to show that the payments and actions of the certainty benchmark also constitute a partial equilibrium in this case. To show this, first consider the actions of banks with distance $d \leq 1$. Recall that these banks' optimal actions are characterized exactly as in the certainty benchmark. Moreover, in the conjectured equilibrium, they receive the same payment from their forward neighbors, $x = Q_1[d - 1]$, as in the certainty benchmark. Consequently, they optimally choose the same actions. Next consider the actions of banks with distance $d \geq 2$. Given $\theta \leq 2l(p)$ and the size of the domino effect $D(p) \leq 1$, these banks choose the aggressive action, $A_0^j = B$, in the certainty benchmark. In particular, the bank at distance $d = 2$ chooses the aggressive action since it does not make any losses from cross claims. With uncertainty, recall that banks with distance $d \geq 2$ act as if they are at distance 2. Thus, these banks optimally choose the aggressive action also with uncertainty. It follows that the partial equilibrium without uncertainty continues to be a partial equilibrium with uncertainty.

This analysis also verifies for this case that banks' actions and payments can be written as a function of their distance. Moreover, the function $Q_1[d]$ is weakly increasing because it is the same as in the certainty benchmark.

Case (ii): $\theta > 2l(p)$. To prove this claim, first consider the banks with distance $d \leq 1$. It can be seen that these banks' optimal actions and payments are the same as in the certainty benchmark. Since $\theta > 2l(p)$, the size of the domino effect satisfies $D(p) \geq 2$. Consequently, banks with distance $d \leq 1$ are insolvent and they choose $A_0^j = S$, proving the claim for these banks.

Next consider the banks with distance $d \geq 2$. Recall that these banks act as if they are at distance $\tilde{d} = 2$. Given the characterization for banks with distance $d \leq 1$, the bank at distance $\tilde{d} = 2$ receives the payment $Q_1[1]$, which is the same as in the certainty benchmark. Consequently, these banks choose the action that the bank at distance $\tilde{d} = 2$ would choose in the certainty benchmark. Since $D(p) \geq 2$, all of these banks optimally choose the precautionary action, $A_0^j = S$. It follows that the banks with distance $d \leq D(p) - 1$ are insolvent and

their debt payments and equity values are the same as in the certainty economy. The transition bank with distance $D(p)$ is solvent and its debt payment and equity value is also the same as in the certainty economy. The banks with distance $d \geq D(p) + 1$ are also solvent and they pay $Q_1[d] = z$ on their debt. However, the equity values of these banks are different from the certainty economy. In particular, the equity value of a bank with distance $d \geq D(p) + 1$ is given by

$$Q_2[d] = y + (1 - y)p < R.$$

This discussion proves the claim also for banks with distance $d \geq 2$, and completes the proof of the proposition.

This analysis also verifies for this case that the banks' actions and payments can be written as a function of their distance. Moreover, the function $Q_1[d]$ is weakly increasing because it is the same as in the certainty benchmark. Q.E.D.

Proof of Proposition 4

There are three cases to consider. The first case concerns a shock, θ , that is weakly smaller than the available liquidity of two banks even when the price of legacy assets is at its lowest level. In this case, part (i) of Proposition 3 applies regardless of the price. Consequently, the banks' payoffs and actions are the same as the certainty benchmark. In particular, all banks with distance $d \geq 2$ choose the aggressive action, $A_0^j = B$. In view of condition (3), the asset demand from these banks exceeds the asset supply from distressed banks. This leads to an equilibrium price $p = 1$ and a domino effect of size $D(1)$. Furthermore, aggregate purchase of new assets is the same as in Proposition 2.

The second case concerns a liquidity shock, θ , which is greater than the available liquidity of two banks even when the price of legacy assets is at its highest level. In this case, part (ii) of Proposition 3 applies regardless of the price. Consequently, there is a flight to quality of size n . In particular, all banks choose the precautionary action, $A_0^j = S$, which has two effects. First, since all banks are sellers in the secondary market (and there are no buyers), the market-clearing condition (2) implies that $p = p_{\text{scrap}}$. Second, since all banks choose to keep their dollars in cash, no new assets are purchased, that is, $\mathcal{Y} = 0$.

The third case concerns a liquidity shock, θ , which is weakly smaller than the available liquidity of two banks when the price is at its highest level, but not when the price is at its lowest level. In this case, there are multiple equilibria. To see this, first suppose legacy assets trade at their fair price, $p = 1$. With this price, the available liquidity, $l(1)$, is sufficiently large that part (i) of Proposition 3 applies. In particular, banks with distance $d \geq 2$ are potential buyers of the asset. This ensures that the fair price, $p = 1$, corresponds to an equilibrium. Suppose, instead, that legacy assets' price is at the fire sale level, $p = p_{\text{scrap}}$. With this price, the available liquidity, $l(p_{\text{scrap}})$, is sufficiently small that part (ii) of Proposition 3 applies. In particular, all banks (including banks with distance $d \geq 2$) are sellers in the secondary market. This ensures that the fire sale price, $p = p_{\text{scrap}}$, also corresponds to an equilibrium. Q.E.D.

Appendix D: Collateralized Credit Default Swap (CDS) Contracts

In our setting, partial domino effects lead to aggregate effects because they increase banks' *idiosyncratic* counterparty risk from cross exposures. A natural question is to what extent this risk could be insured at date 0 (once banks learn about the shock). In practice, banks could obtain some insurance by purchasing CDSs on their counterparties. However, as discussed in the main text, the CDS contracts that pay in systemic events need to be collateralized to protect insurance buyers against a potential default of the insurance seller. This appendix allows for collateralized CDS contracts at date 0 and shows that our results are robust to this extension. The key insight is that, while banks demand CDS insurance on their counterparties, the supply of insurance is also restricted because of sellers' collateral constraints.

Consider the setting of Proposition 3 with network uncertainty and fixed asset price (for simplicity). Consider also parameters such that $D(p) \geq 2$, so that banks are trying to maximize their available liquidity at date 1 (in their worst-case scenario). In this setting, banks have a demand for insurance contracts that pay when they are distressed at date 1. To capture this aspect, suppose banks can invest at date 0 in insurance contracts on the insolvencies of their forward neighbor banks. In particular, for each bank j , there is a contract, I^j , that pays one dollar if bank j is insolvent at date 1 (i.e., if it pays $q_1^j < 1$). Suppose also that insurance contracts, $\{I^j\}_j$, must be individually and fully collateralized. In particular, the insurance seller must pledge one unit of cash as collateral for each unit of insurance contract she sells at date 0. Each contract, I^j , is traded at date 0 in a competitive market at price $f^j \in (0, 1)$, which will be endogenously determined.

The collateral constraint implies that banks within the network choose not to sell insurance contracts. To see this, note that selling the contract, I^j , requires the bank to pledge $1 - f^j$ dollars of cash (in addition to f^j dollars, which she raises from the sale). In particular, selling insurance reduces banks' available liquidity at date 1 (even though it may increase their return at date 2). Given that banks are trying to maximize their available liquidity, they choose not to sell insurance. Put differently, network uncertainty not only increases banks' demand for insurance, but also naturally decreases their supply of insurance.

It follows that insurance contracts must be sold by an agent outside the financial network. Suppose the outside agent has y^{out} dollars at date 0 and consumes only at the end of date 1. Suppose the outside agent does not know the financial network. In addition, suppose also that the outside agent does not know the identity of the original distressed bank, b^0 .²¹ Importantly, the outside agent knows that the size of the domino effect will be exactly D (i.e., not all of the banking system can go under). This is the main feature that facilitates insurance.

Let x^j denote the amount of contract I^j sold by the outside agent. We conjecture an equilibrium for the insurance market in which $f^j \equiv f \in (0, 1)$ for each

²¹ This assumption is only made for simplicity. The results do not change if we assume that the outside agent knows b^0 (i.e., she knows as much as the inside banks).

j and $x^j \equiv x$ for each j . That is, all banks' insurance contracts trade at the same price and the outside agent sells an equal number of contracts.

To characterize this equilibrium, first consider the supply of insurance by the outside agent. This agent's portfolio choice problem can be written as

$$\begin{aligned} \max_{\tilde{x} \geq 0} y^{\text{out}} + \tilde{x}fn - \tilde{x}D, \\ \text{s.t. } n\tilde{x}(1 - f) \leq y^{\text{out}}. \end{aligned} \tag{D1}$$

The first line is the outside agent's expected profit: for each contract she sells, she collects fn dollars in premiums and she expects to pay D dollars. Note that, even though the outside agent does not know the network, it knows that exactly D banks will fail. The second line of (D1) is the outside agent's budget constraint. For each contract she sells, she raises f dollars. However, she needs to put an additional, $1 - f$, dollars as collateral. The total amount of collateral she pledges cannot exceed her available collateral, y^{out} .

Problem (D1) implies that, as long as $fn > D$, which we will verify in equilibrium, the outside bank sells as much insurance as possible. That is

$$x = \frac{y^{\text{out}}}{n(1 - f)}. \tag{D2}$$

Next consider the demand for insurance by banks. To maximize their available liquidity at date 1, banks, $\{b^j\}_{j=1}^{n-1}$, spend all of their date 0 resources to buy insurance on their respective forward neighbor banks. This is because they buy insurance at price $f < 1$ (which is fully collateralized), which gives them one dollar at date 1 in their worst-case scenario (when their forward neighbor is insolvent). Thus, these banks' demand for insurance is given by

$$x = \frac{l(p)}{f}. \tag{D3}$$

With these insurance purchases, their available liquidity at date 1 (when their forward neighbor is insolvent) becomes

$$l(p, f) = x = \frac{l(p)}{f}. \tag{D4}$$

Next consider the original distressed bank, b^0 . This bank cannot increase its available liquidity by buying insurance on its forward neighbor, because its forward neighbor will always be solvent. On the other hand, this bank is indifferent between any level of insurance (because it will be insolvent regardless of its action). To keep the analysis and the notation simple, consider equilibria in which this bank's insurance demand is also given by (D3), which leads to an available liquidity of zero dollars at date 1.

Given this characterization of insurance purchases and available liquidities, the size of the domino effect can be calculated as before. In particular, the

analogue of equation (6) in this setting is given by

$$D(p, f) = \left\lceil \frac{\theta}{l(p, f)} \right\rceil.$$

Note that, when f is lower, more liquidity is available to banks in distress, which leads to a smaller domino effect.

The equilibrium price of insurance is characterized by equating the supply of insurance in (D2) with the demand for insurance in (D3). This leads to the closed-form solution

$$f = \frac{l(p)}{\frac{y^{\text{out}}}{n} + l(p)} \text{ and } x = \frac{y^{\text{out}}}{n} + l(p).$$

Note that insurance is expensive when the aggregate collateral of the insurance sellers, y^{out} , is small relative to the number of banks that demand insurance, n . When $\frac{y^{\text{out}}}{n}$ is sufficiently small, f is close one, which has two implications. First, the condition $fn > D$ (which leads to equation (D2)) is verified because $n > D$. Second, banks' available liquidity in (D4) is close to $l(p)$. Consequently, the equilibrium is qualitatively similar to the earlier setting without insurance.

This analysis illustrates that, as long as the collateral of insurance sellers outside the financial network is scarce, the CDS market does not overturn our results. The behavior of the CDS market during the recent Bear Sterns and Lehman debacles is broadly consistent with this analysis. Duffie (2011) states that the demand for insurance spiked in both episodes (measured by novation requests), and that this demand could not be met by insurance sellers (dealers). These observations suggest that collateralized insurance contracts might fail to fully eliminate the counterparty risk in systemic episodes, because the supply of this type of insurance is also limited.

REFERENCES

- Acemoglu, Daron, Vasco Carvalho, Asu Ozdaglar, and Alireza Tahbaz-Salehi, 2012, Network origins of aggregate fluctuations, *Econometrica* 80, 1977–2016.
- Acharya, Viral, Douglas Gale, and Tanju Yorulmazer, 2011, Rollover risk and market freezes, *Journal of Finance* 66, 1177–1209.
- Allen, Franklin, and Ana Babus, 2009, Networks in finance, in Paul Kleindorfer and Joram Wind, eds.: *The Network Challenge* (Wharton School Publishing, Philadelphia, PA).
- Allen, Franklin, Ana Babus, and Elena Carletti, 2010, Financial connections and systemic risk, NBER Working paper 16177.
- Allen, Franklin, and Douglas Gale, 1994, Limited market participation and volatility of asset prices, *American Economic Review* 84, 933–955.
- Allen, Franklin, and Douglas Gale, 2000, Financial contagion, *Journal of Political Economy* 108, 1–33.
- Brunnermeier, Markus, Andrew Crockett, Charles Goodhart, Avinash Persaud, and Hyun Shin, 2009, The fundamental principles of financial regulation, ICMB-CEPR Geneva Report on the World Economy 11.
- Brunnermeier, Markus, and Lasse Pedersen, 2005, Predatory trading, *Journal of Finance* 60, 1825–1863.

- Brunnermeier, Markus, and Lasse Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.
- Brunnermeier, Markus, and Yuliy Sannikov, 2012, A macroeconomic model with a financial sector, Working paper, Princeton University.
- Caballero, Ricardo, and Arvind Krishnamurthy, 2008, Collective risk management in a flight to quality episode, *Journal of Finance* 63, 2195–2230.
- Caballero, Ricardo, and Pablo Kurlat, 2008, Flight to quality and bailouts: Policy remarks and a literature review, Working paper, MIT.
- Caballero, Ricardo, and Alp Simsek, 2009, Complexity and financial panics, NBER Working paper 14997.
- Cifuentes, Rodrigo, Gianluigi Ferrucci, and Hyun Shin, 2005, Liquidity risk and contagion, *Journal of European Economic Association* 3, 556–566.
- Dang, Tri Vi, Gary Gorton, and Bengt Holmstrom, 2010, Financial crises and the optimality of debt for liquidity provision, Working paper, Yale University.
- Dasgupta, Amil, 2004, Financial contagion through capital connections: A model of the origin and spread of bank panics, *Journal of European Economic Association* 2, 1049–1084.
- Diamond, Douglas, and Raghuram Rajan, 2011, Fear of fire sales and the credit freeze, *Quarterly Journal of Economics* 126, 557–591.
- Duffie, Darrell, 2011, *How Big Banks Fail and What to Do About It* (Princeton University Press, Princeton, NJ).
- Easley, David, and Maureen O'Hara, 2009, Ambiguity and nonparticipation: The role of regulation, *Review of Financial Studies* 22, 1817–1843.
- Eisenberg, Larry and Thomas Noe, 2001, Systemic risk in financial systems, *Management Science* 47, 236–249.
- Fender, Ingo, and Peter Hordahl, 2007, Overview: Markets hit by renewed credit woes, *BIS Quarterly Review* December, 1–16.
- Freixas, Xavier, Bruno 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.
- Garleanu, Nicolae, and Lasse Pedersen, 2011, Margin-based asset pricing and deviations from the law of one price, *Review of Financial Studies* 24, 1980–2022.
- Geanakoplos, John, 2003, Liquidity, default and crashes: Endogenous contracts in general equilibrium, Cowles Foundation Discussion Paper 1316.
- Geanakoplos, John, 2009, The leverage cycle, in Acemoglu Daron, Kenneth Rogoff, and Michael Woodford, eds.: *NBER Macroeconomics Annual* (National Bureau of Economic Research, Cambridge, MA).
- Gilboa, Ithzak, and David Schmeidler, 1989, Maxmin expected utility with a non-unique prior, *Journal of Mathematical Economics* 18, 141–153.
- Gromb, Denis, and Dimitri Vayanos, 2002, Equilibrium and welfare in markets with financially constrained arbitrageurs, *Journal of Financial Economics* 66, 361–407.
- Haldane, Andrew, 2009, Rethinking the financial network, Working paper, Bank of England.
- Hansen, Lars, and Thomas Sargent, 2010, Fragile beliefs and the price of uncertainty, *Quantitative Economics* 1, 129–162.
- Holmstrom, Bengt, and Jean Tirole, 1998, Private and public supply of liquidity, *Journal of Political Economy* 106, 1–40.
- Kiyotaki, Nobuhiro, and John Moore, 1997a, Credit chains, Working paper, Princeton University.
- Kiyotaki, Nobuhiro, and John Moore, 1997b, Credit cycles, *Journal of Political Economy* 105, 211–248.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2005, A smooth model of decision making under ambiguity, *Econometrica* 73, 1849–1892.
- Lagunoff, Roger, and Stacey 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.
- Lorenzoni, Guido, 2008, Inefficient credit booms, *Review of Economic Studies* 75, 809–833.

- Rochet, Jean-Charles, and Jean Tirole, 1996, Interbank lending and systemic risk, *Journal of Money, Credit, and Banking* 28, 733–762.
- Rotemberg, Julio, 2008, Liquidity needs in economies with interconnected financial obligations, NBER Working paper 14222.
- Routledge, Bryan, and Stanley Zin, 2009, Model uncertainty and liquidity, *Review of Economic Dynamics* 12, 543–566.
- Shleifer, Andrei, and Robert Vishny, 1992, Liquidation values and debt capacity: A market equilibrium approach, *Journal of Finance* 47, 1343–1366.
- Shleifer, Andrei, and Robert Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35–55.
- Shleifer, Andrei, and Robert Vishny, 2011, Fire sales in finance and macroeconomics, *Journal of Economic Perspectives* 25, 29–48.
- Stein, Jeremy, 2012, Monetary policy as financial-stability regulation, *Quarterly Journal of Economics* 127, 57–95.
- Upper, Christian, 2007, Using counterfactual simulations to assess the danger of contagion in interbank markets, BIS Working paper 234.
- Vause, Nicholas, 2010, Counterparty risk and contract volumes in the credit default swap market, *BIS Quarterly Review* December, 59–69.
- Zawadowski, Adam, 2013, Entangled financial systems, *Review of Financial Studies* 26, 1291–1323.

Copyright of Journal of Finance is the property of Wiley-Blackwell and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.