Wealth Inequality, Network Topology and Financial Crisis

Thomas Hauner*

March 28, 2017

Abstract

This paper asks if two, otherwise identical, economies were distinguished only by their distributions of wealth, are they equally stable in response to a random shock? A theoretical financial network model is proposed to understand the relationship between wealth inequality and financial crises. In a financial network, financial assets link individual asset and liability holders to form a web of economic connections. The total connectivity of an individual is described by their *degree*, and the overall distribution of connections in the network is imposed through a *degree distribution*—equivalent to the wealth distribution as incoming connections represent assets and outgoing connections liabilities. A network's topology varies with the level of wealth inequality and total wealth and together, simulations show, they determine network contagion in the event of a random negative income shock to some individual. Random network simulations, whereby each financial connection is randomly placed, reveal that increasing wealth inequality makes a wealthy network less stable—as measured by the share of individuals failing financially or the decline in financial asset values. These results suggest a unique architectural role for accumulated assets and their distribution in macro-financial stability.

Keywords: Wealth inequality, financial crisis, instability, financial network, degree distribution.

JEL-Classification: D31, D85, G01, L14

*The Graduate Center, City University of New York (CUNY).
Contact: thauner@gradcenter.cuny.edu
Thank you to Wim Vijverberg, Suresh Naidu, George Vachadze, Branko Milanovic, and seminar participants at the Graduate Center and Tobin Project for helpful comments. Also, thank you to the Tobin Project Graduate Student Forum for financial support. All errors are my own.

1 Introduction

The share of wealth held by top percentiles in the US provocatively peaks before both the Great Crash in 1929 and the global financial crisis of 2007–2008.[1] Is economic inequality a destabilizing economic force? Or is the pronounced correlation a symptom of deeper economic perturbations? The objective of this paper is to conceptualize the relationship between wealth inequality and macroeconomic instability as manifested through the financial sector. One goal is to demonstrate how the wealth distribution can alter the configuration of the financial economy into more or less stable arrangements. The approach can be summarized thusly: consider two identical economies that are observed at a point in time and distinguished only by their distributions of wealth. Which is more unstable in the event of a negative income shock? To answer, an interpersonal financial network model is constructed using elements of graph theory. The model is then repeatedly simulated, generating predictions about the endogenous role of wealth distributions on financial stability. It suggests one direct channel from top wealth inequality to the vulnerability of a financial economy in the event of a shock.

Axel Leijonhufvud once described a network economy as a "web of contracts and understandings" between individuals. Wealth, as a collection of financial assets, by definition creates financial links in a network economy. The model assumes one type of financial asset exists, one individual's claim on some future cash flow that is generated by another individual's labor income. The presence of a financial asset naturally links asset owners to liability holders. The total number of financial assets an individual owns represents their *in-degree* and the distribution of assets in the network economy is imposed through an *in-degree distribution*—equivalent to the wealth distribution since it is assumed real assets are homogeneous, and thus labor income is as well. As the distribution of wealth changes the distribution of links in the network also changes, thereby altering the topology of the interpersonal financial network. Though the network is static, with no individual optimization problems, contagion is a dynamic process. Contagion occurs when a random negative income shock decreases one individual's net worth to the point of financial failure, which prompts failure costs that wipe out collateral wealth. Importantly, an individual's net worth is assumed to collateralize their financial liabilities, much like an asset-backed security. The network structure implies one

[1] See Kopczuk & Saez (2004). Piketty & Saez (2014), Saez & Zucman (2014), and Kopczuk (2015).

individual's net worth is linked to, and dependent on, the net worth of others. Therefore decreases in net worth spread.

Network simulations, in which the arrangement of financial links is random, show that the model economy is more unstable in the event of a negative shock when it (a) exhibits high wealth inequality, and (b) is sufficiently wealthy in aggregate. Additionally, an inverted U-shaped relationship is found between aggregate network wealth and instability. Financial instability, in the model, is determined by the share of individuals whose net worth drops below a predetermined threshold or the total decline in financial asset values.

Our network model embeds several features of Hyman Minsky's Financial Instability Hypothesis. While not an explicit network model, it is a framework that generates endogenous instability in a financial economy of connected banks and firms rather than individuals.[2] The first is that balance sheets of individuals are interrelated, where one's asset is always another's liability. Second, assets and liabilities represent commitments to future cash flows, where the flows across network links are what Minsky called a "complex system of money in/money out transactions,"[3]. Third, a collapse in asset values stifles future cash flows and catalyzes a crisis, or as Kregel (2014) argues, only a "slight disturbance" in money flows is necessary to cause instability and "widespread financial distress." And finally, a growing financial economy increases the scale of contagion.

Financial network models are frequently used to model financial crises amongst banks. The model in this paper borrows from the framework of Elliott et al. (2014a), however, network nodes represent individuals rather than banks or countries, and financial links do not exist outside of the network. Their emphasis is also on the levels of financialization in the network, at both the intensive and extensive margins, rather than the skewness of financial asset distributions as in this paper. Allen & Gale (2000) were one of the first to show in a simple bank network model that the configuration of financial links mattered for contagion—complete networks were more stable than incomplete ones. More recently, Acemoglu et al. (2015) stress network structure as the determining factor in contagion, but they largely look at the magnitude and frequency of negative shocks to the network in order to analyze its stability, not network topology. And Glasserman & Young (2015)

[2]See Minsky (1975) and Minsky (1986b) for longer expositions, or Minsky (1992) for a brief summary.
[3]See (Minsky, 1986a, p. 69).

abandon topology measures altogether in favor of bank-specific sufficient statistics to evaluate bank network contagion risks. They conclude that factors beyond pure spillovers, such as confidence in counterparties and bankruptcy costs (included in our model), are responsible for substantial economic losses from contagion.[4] The only known network models of individual wealth inequality reside in a small statistical mechanics literature, demonstrating a power law degree distribution of wealth often results from network formation dynamics.[5] None consider contagion or network instability as this paper does.

While the finance network literature has largely ignored inequality, the *income* inequality literature has considered inequality's relation to financial crises—with mixed results. In a qualitative survey of 84 crises across 21 countries over the past century, Morelli & Atkinson (2015) examine both the levels of and changes in income inequality preceding a crisis episode. They conclude that either's impact is ambiguous. Rajan (2011) argues increasing US income inequality was but one cause of the crisis because it prompted policies that ultimately relaxed credit to unsustainable levels. Testing the Rajan hypothesis, Bordo & Meissner (2012) regress changes in real credit growth on lagged changes in top income shares. They find no effect amongst a panel of 14 countries between 1870 and 2010 and thus conclude no link between inequality and crisis exists.[6] A dynamic stochastic general equilibrium model by Kumhof & Ranciere (2010) is structured, like this paper's model, around assets linking households. In their case the top 5% own assets derived from the borrowing of the bottom 95%. The authors calibrate the model to show that increasing income inequality, rising household debt of the bottom 95%, and increasing financial assets of the top 5% cause higher leverage and thus a higher probability of crisis.

In the inequality literature cited above, the most common mechanism linking rising inequality to instability is rising household debt. Mason & Jayadev (2014) show, however, that a set of so-called "Fisher dynamics" (i.e. interest rate change, inflation, and income growth) account for most, if not all, of the increase in US household leverage since 1980—the same structural lever modeled by Kumhof & Ranciere (2010). In other words, increasing household debt-income ratios do not

[4] Both Acemoglu et al. (2015) and Glasserman & Young (2015) are derived from Eisenberg & Noe (2001), a network model of equilibrium clearing payments among banks used to measure contagion.

[5] See Lee & Kim (2007), Kim et al. (2008), and Coelho et al. (2005)

[6] Gu & Huang (2014), criticizing the econometric methods of Bordo & Meissner (2012), argue that income inequality, in Anglo-Saxon countries, *does* determine credit growth—and therefore leads to financial crisis.

necessarily imply newly issued debt, and new debt is the critical lever of the inequality-household debt-instability story.

Instead, this paper argues that an economy's financial network configuration, as determined by the wealth distribution, is the critical determinant. Of course, the model is a gross simplification of a financial capitalist economy—assuming a static network, with one type of financial asset serviced by (uniform) labor income cash flows and individual net worth acting as collateral. But by abstracting away layers of financial intermediaries, it becomes possible to expose the latent financial relationships between individual creditors and debtors and to understand how the interpersonal distribution of wealth in the economy may impact its overall stability. Though the setup also ignores network formation dynamics, or consumption and saving decisions by individuals, it provides a tractable model that can be simulated and whose results are generalizable.

An intuitive metaphor for understanding how network attributes wealth inequality and aggregate wealth may work in tandem to determine financial stability is to consider a Jenga tower, the block-building game. If each block represents a financial asset or link, then a short tower is relatively stable regardless of the distribution of the blocks. As the Jenga tower grows, however, the distribution of blocks becomes critical to its stability.

The rest of the paper is organized as follows: Section 2 derives the theoretical financial network model, presents its mechanics, and introduces concepts of instability. Financial network parameter estimates are shared in Section 3, to motivate model parameterization. Section 4 describes the method to simulate random static networks and also presents results, including the finding that increasing wealth inequality is destabilizing in wealthy networks. Section 5 concludes.

2 Financial Network Model

In this section we introduce the wealth inequality network model, using Elliott et al. (2014a) as a foundation. It notably disregards financial intermediaries and instead relies on the latent financial links between asset and liability holders to form an interpersonal financial network economy. This enables a more tractable model between the economy's wealth distribution, how it translates to network topology, and overall financial (in)stability. We present an extended example at the end to help elucidate concepts from the model.

2.1 Setup

Consider a static financial network composed of nodes $i = 1, \ldots, n \in N$. Each node represents a wealth owning individual or household. (We exclude firms, banks, and other types of organizations to simplify our model and to argue that variations in the distribution of wealth between individuals have network consequences.[7]) Links, or edges, connect two nodes and represent a financial claim between them. A financial asset is a claim on future cash flows. All network links, or financial assets, are represented by an $n \times n$ *adjacency matrix* matrix $\mathbf{G}$, where element $G_{ij} = 1$ if node i has some financial claim on node j and 0 otherwise. Claims are directional, implying the orientation of cash flows. Matrix $\mathbf{G}$ is thus composed of creditors (rows) making financial claims on debtors (columns). Though individuals are along both dimensions of the matrix, financial claims need not be reciprocated—and $\mathbf{G}$ need not be symmetric. The network can be summarized as an unweighted directed graph $\mathcal{G}(N, G)$ whose edges indicate the existence, and paths, of financial flows between individuals.

Assume there exists only one type of financial asset held by individuals and households, a type of asset-backed security. Each security is a claim on future labor income cash flows generated by the liability holder, with their net worth serving as collateral.[8] A node i owns d_i financial assets, where $d_i = \sum_j G_{ij}$ is the node's *in-degree*. It also represents the total number of individuals i holds claims against (a row sum in $\mathbf{G}$). A financial asset-owning node may also back the value of an asset themselves, a function of their own valuation. Let d_j^{out} represent the total number of financial liabilities node j is collateralizing, where $d_j^{out} = \sum_i G_{ij}$ (a column sum in $\mathbf{G}$) is the *out-degree*. Financial assets are distributed according to some probability distribution $f(d_i)$, the *degree distribution*.[9] Only some fraction $c \in (0, 1)$ of each individual's overall net worth is collateralized and can be claimed by, and owed to, asset holders in the network.

[7]To be sure, many individuals rely on opaque institutions and organizations to hide private wealth. See Zucman (2014) and Zucman (2015) for a detailed analysis on hidden private wealth.

[8]Node net worth is discussed in detail in Section 2.3.

[9]Note that $\sum d_i = \sum d_j^{out}$ so that total assets equal total liabilities and the economy's balance sheet balances.

Let matrix $\mathbf{C}$, the *cross-holdings matrix*, describe the relative ownership claims on each node in the network, with elements

$$C_{ij} = \begin{cases} \frac{c}{d_j^{out}} & \text{if } d_j^{out} > 0 \\ 0 & \text{else.} \end{cases} \quad \forall\ i \neq j. \tag{1}$$

The unweighted adjacency matrix $\mathbf{G}$ is now a weighted matrix $\mathbf{C}$ of financial claims between nodes. The total number of asset holders d_j^{out} holding assets backed by individual's j's wealth are each entitled to an equal portion of future cash flows. Cash flows not claimed by other individuals $(1 - c)$ are saved. (Savings do not accumulate as the model is static.) The savings of each node are summarized in a diagonal matrix $\hat{\mathbf{C}}$, with element $\hat{C}_{jj} = 1 - \sum_i C_{ij}$. It is possible to rewrite the total sum of claims made on individual j as $\sum_i C_{ij} = 1 - \hat{C}_{jj}$.

To illustrate, consider the network in Figure 1a, where $n = 4$ and $c = 0.5$. The corresponding adjacency and cross-holdings matrices are in Figure 1b. Notice, from $\mathbf{G}$'s bottom row, that node 4 has financial assets which are claims on the cash flows of nodes 1, 2 and 3 ($d_4 = 3$), but has no cash flow obligations itself ($d_4^{out} = 0$). Because $c = 0.5$, half of nodes 1, 2, and 3's future incomes flow to node 4.

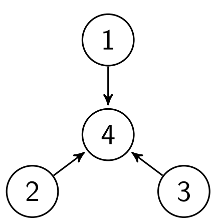

(a) Graph of network

$$\mathbf{G} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ .5 & .5 & .5 & 0 \end{bmatrix}$$

(b) Corresponding adjacency and cross-holdings matrices.

Figure 1: Example of a four-node network

2.2 Gross value

There also exist $k = \{1, \ldots, m\} \in M$ real, or physical, assets (e.g. productive assets like land or human capital).[10] Let matrix $\mathbf{D}$, analogous to the cross-holdings matrix $\mathbf{C}$, describe the pattern of

[10] Consider one possible microfoundation for our network model thus far. Suppose our static network is an endowment economy, whereby all nodes are endowed with a single type of financial asset—like our asset-backed security. The endowments are randomly distributed between nodes according to some probability distribution $f(d_i)$, where d_i is the total number of financial assets node i owns. Which nodes then back each of the d_i securities node i owns is

real asset claims. Its elements D_{ik} denote individual i's share of real asset k . The *gross value* of individual i's total assets V_i is the sum of their real asset claims (each at their respective prevailing market price, p_k) and financial asset claims (backed by the liability holder's own gross value).

$$V_i = \sum_k D_{ik} p_k + \sum_j C_{ij} V_j, \tag{2}$$

or in matrix notation, $\mathbf{V} = \mathbf{Dp} + \mathbf{CV}$. Solving for the gross value of each individual in the network yields the vector of gross values $\mathbf{V}$

$$\mathbf{V} = (\mathbf{I} - \mathbf{C})^{-1} \mathbf{Dp}. \tag{3}$$

Note, however, that the gross value of individual i's total assets V_i double counts real asset claims D_{ik}. They appear not only in the first term of (2) but also in the second term as a component of other individuals' own valuations V_j. Therefore we derive a measurement of node net worth in the next section.

To specifically study how the distribution of financial assets $f(d_i)$ impacts the network's overall stability, the model is simplified by assuming there exists one type of real asset, human capital, with $m = n$ different units. Homogeneous, it also cannot be owned by anyone else, hence $\mathbf{D} = \mathbf{I}_n$—though others may have claim to the future cash flows generated by it.[11] Human capital prices are homogeneous and also normalized to one, such that $p_k = 1 \ \forall \ k$.

randomly determined.

A richer form of heterogeneity could allow for variation amongst nodes rather than endowments. Let ρ represent a symmetrically distributed stochastic discount factor where $\rho \in \{\rho_l, \rho_\mu, \rho_h\}$. (In Krusell & Smith (1998), a dynamic general equilibrium model using stochastic discount factors generates a Pareto wealth distribution in the tails that closely fits the empirical estimates of Wolff (1994).) If a node is assigned the lower discount factor ρ_l, then the node must borrow to consume single good y as they have a preference for consumption. In this circumstance $d_l < d_l^{out}$ and the individual is a net debtor. If a node receives the higher discount factor ρ_h it is a lender with a preference for accumulating assets. In this event $d_h > d_h^{out}$ and the individual is a net creditor. Should the node receive the mean discount factor ρ_μ, then $d_\mu = d_\mu^{out}$.

Yet another possibility is to consider an economy of entrepreneurs. Each node is endowed with some productive asset and an intermediary good, drawn from a distribution. The intermediary good may be consumed, but nodes prefer to consume a final consumption good that requires the interaction of at least two intermediary goods. Credit, fixed in aggregate, is extended between entrepreneurs to produce the final consumption good, which may be used to repay liabilities.

[11] Allowing $\mathbf{D}$ to represent human capital takes into consideration a common critique of Piketty (2014), best articulated by Blume & Durlauf (2015), that aggregating financial and physical assets at prevailing market prices crucially ignores the important contemporary role human capital plays in generating cash flows.

2.3 Net worth

Node net worth is defined as total assets (real and financial) less liabilities.[12] In other words, claims on one's own wealth (outflows) are subtracted from the sum of real assets and ownership claims on other individuals' wealth (inflows):

$$v_i = \sum_k D_{ik} p_k + \sum_{j \neq i} C_{ij} V_j - \left(\sum_{j \neq i} C_{ji} \right) V_i. \tag{4}$$

Note that the first two terms are simply individual i's gross value. In matrix form:

$$\mathbf{v} = \mathbf{Dp} + \mathbf{CV} - (\mathbf{I} - \hat{\mathbf{C}})\mathbf{V} = \mathbf{Dp} + [\mathbf{C} - (\mathbf{I} - \hat{\mathbf{C}})]\mathbf{V}, \tag{5}$$

where $\mathbf{I} - \hat{\mathbf{C}}$ is a diagonal matrix representing weighted total obligations in the network and $\mathbf{C}$ represents weighted total claims. Substituting the gross value from (3) for $\mathbf{V}$ in (5) and rearranging leads to the following definition of net worth.

$$\begin{aligned}
\mathbf{v} &= \mathbf{Dp} + [\mathbf{C} - (\mathbf{I} - \hat{\mathbf{C}})]\mathbf{V} \\
&= \mathbf{Dp} + [\mathbf{C} - (\mathbf{I} - \hat{\mathbf{C}})][(\mathbf{I} - \mathbf{C})^{-\mathbf{1}}\mathbf{Dp}] \\
&= ((\mathbf{I} - \mathbf{C}) + \mathbf{C} - \mathbf{I} + \hat{\mathbf{C}})(\mathbf{I} - \mathbf{C})^{-\mathbf{1}}\mathbf{Dp} \\
\mathbf{v} &= \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-\mathbf{1}}\mathbf{Dp} \\
\mathbf{v} &= \mathbf{ADp}
\end{aligned} \tag{6}$$

Net worth is derived from the overall claims between all nodes in the network (matrix $\mathbf{A}$) made on the underlying real assets (matrix $\mathbf{D}$ at price $\mathbf{p}$) of the economy. Since each real asset represents a node's human capital, net worth is a function of the cumulative claims on future output generated by another's human capital.

The benefit of matrix $\mathbf{A} = \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-\mathbf{1}}$, called the *dependency matrix*, is that it summarizes the total claims between all nodes, i.e. the sum of direct and indirect dependencies between individuals

[12] See, for example, Davies & Shorrocks (1999) and Davies et al. (2007). In Elliott et al. (2014a) this is called a node's *market value*, since their model's nodes represent firms or banks.

in the network.[13] It is possible for element A_{ij} to be nonzero even if the corresponding element in the cross-holdings matrix, C_{ij}, is zero—an indication of indirect claims by i on j via other nodes in the network but no direct claims. The dependency matrix $\mathbf{A}$ is not unlike Leontief's input-output matrix, Elliott et al. (2014a) posit, in its ability to summarize the interconnections of a network economy. It is instructive to examine the differences between direct holdings (from cross-holdings matrix $\mathbf{C}$) and total direct and indirect holdings (from dependency matrix $\mathbf{A}$) in the examples in Section 2.6.

The dependency matrix $\mathbf{A}$ also simplifies the accounting considerably. Claims on individual real assets, rather than both financial assets and liabilities, become a sufficient statistic to determine an individual's overall net worth when calculating the impacts of a shock as they reverberate through the network. In fact, all wealth is derived from human capital.

2.4 Wealth inequality

The wealth distribution of the network can be decomposed into its real and financial components. By assuming real assets, in the form of human capital, are fixed, equal for all individuals, and have the same market price, financial assets entirely determine the wealth distribution, defined by the degree distribution of financial assets $f(d_i)$. That is, wealthier individuals will have more positive financial claims and links to other individuals in the network than less wealthy individuals. A deterministic degree distribution, for example, captures perfect equality of financial wealth. Let a Pareto distribution describe the degree distribution of an unequal society where the probability of someone having d_i financial assets is given by $p(d_i) = ad_i^{-\gamma}$, with $\gamma > 0$.[14]

Aggregate financial wealth of the entire network equals the total number of financial claims $\sum d_i$. Because the network is static and the number of individuals n remains fixed, increasing the number of assets in the network increases total financial wealth. This is akin to the economy growing through increased credit, or financialization at the extensive margin.

Figure 2 illustrates how a random network's structure changes with increasing financial wealth

[13] Matrix $\mathbf{A}$ has all nonnegative elements and is also column-stochastic, thus each of its columns sums to 1 ($\sum_i A_{ij} = 1$).

[14] Our network model simulation results are robust to allowing wealth inequality to be determined by the out-degree distribution $f(d_i^{out})$ rather than the in-degree distribution $f(d_i)$.

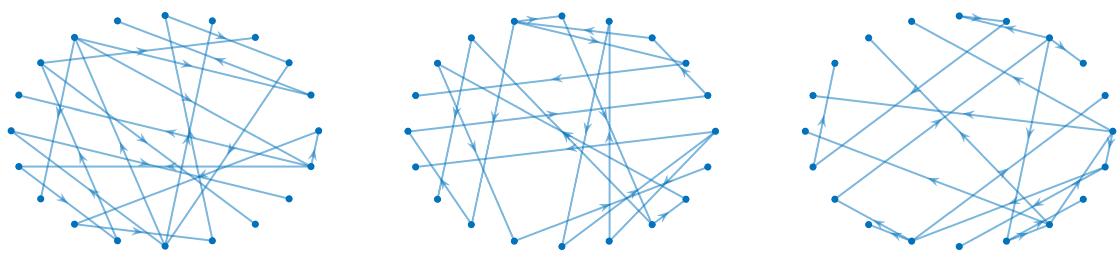

(a) $\gamma = 2.35$ **(b)** $\gamma = 1.65$ **(c)** $\gamma = 1.025$

Figure 2: Random network graphs ($n = 20$)

inequality via the Pareto degree distribution. Arrows indicate the direction of cash flows. With $n = 20$ and expected in-degree $E[d_i] = 1$, each network is generated randomly for a specified γ. The highest Pareto parameter ($\gamma = 2.35$) corresponds to the lowest inequality among the three graphs. Its financial claims are more evenly spread out compared to the most unequal random network graph ($\gamma = 1.025$).[15]

Given that we can calculate the net worth of an individual v_i using (6), why not directly model the wealth distribution with $f(v_i)$? Using $f(v_i)$, rather than $f(d_i)$, to model wealth inequality obscures the critical role that interconnectedness plays in the financial network. It is precisely the interlocking structure of the network that determines whether or not a shock leads to contagion and instability. In order to have a tractable link structure in our adjacency matrix, the random network's inequality must be derived from the degree distribution $f(d_i)$. Finally, the degree distribution of the network characterizes the same magnitude of wealth inequality given by the distribution of individual net worths, without loss of generality.

2.5 Shocks, financial failure, and instability

Though the model is static, contagion is dynamic. Let the time subscript t specify periods in relation to the initial shock in period $t = 0$. Recall that initial real asset prices are set to 1 so that $\mathbf{p}$ is a vector of ones. A random exogenous income shock at time $t = 0$ drops one individual's real

[15] Each graph is generated thusly: draw a random Pareto distribution of financial claims d_i, truncated at the top to ensure $E[d_i] = 1$ across distributions; randomly link financial claims d_i to other nodes to create adjacency matrix $\mathbf{G}$; plot directed graph $\mathbf{G}$.

asset price by the amount λ, such that $p_i = 1$ becomes $\tilde{p}_i = \lambda p_i = \lambda \; \forall \; t \geqslant 1$, where $\lambda \in [0, 1)$. The decline in individual i's market price for human capital will negatively impact their labor income and thus cash outflows. The magnitude of the negative real asset price shock $\tilde{p}_i$ is decreasing in λ. No other individuals in the network experience a real asset price shock, and thus the vector of real asset prices after the initial shock $\tilde{\mathbf{p}}$ contains a value λ in the i^{th} row and 1 everywhere else.[16] A uniform risk of shock $(\frac{1}{n})$ exists, therefore no risk premia are priced into financial assets.

The negative real asset price shock could represent the loss of a job or earning capacity. An individual experiences *financial failure* if, as a result of this income shock, the individual's wealth $v_{i,t}$ should fall below some threshold $\underline{v}_i$. Let the failure threshold $\underline{v}_i = \theta v_i > 0$, with $\theta \in (0, 1)$ remaining constant throughout the dynamic contagion process. Parameter θ describes individual financial fragility. A high θ implies a more easily breached threshold and likelier financial failure in the event of a shock, whereas a low value means more robust personal finances. The threshold $\underline{v}_i$ is positive because financial duress and accompanying cash flow strains need not imply negative net worth in our model, only a financial setback such that creditors are not repaid and penalties imposed.[17]

Individual financial failure triggers bankruptcy costs β_i. They are not to be taken literally (net worth remains positive), but instead as representative of increased financial burdens faced when an individual's net worth is depressed by some relative amount. Such burdens could include direct costs like attorney and accounting fees as well as indirect costs such as lost income, increased future borrowing costs, loss of collateral or counterparty confidence. Let $\mathbf{b}_{t-1}$ represent a vector of failure costs with element $b_{i,t-1} = \beta_i(\tilde{\mathbf{p}}) I_{v_{i,t-1} < \underline{v}_i}$, or $\beta_i(\tilde{\mathbf{p}}) I_{v_{i,t} < \underline{v}_i}$, where I is an indicator function taking a value of 1 if $v_{i,t} < \underline{v}_i$ and 0 otherwise. By definition, $\beta_i = 0 \; \forall \; i$ at $t = 1$ because no individuals have failed yet. The first iteration of calculating new node valuations occurs at $t = 1$, so (6) is rewritten to incorporate failure costs.[18]

$$\mathbf{v_t} = \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-\mathbf{1}}(\mathbf{D}\tilde{\mathbf{p}} - \mathbf{b_{t-1}}) = \mathbf{A}(\mathbf{D}\tilde{\mathbf{p}} - \mathbf{b_{t-1}}) \text{ for } t = 1, \ldots, T \tag{7}$$

[16] The notation $\tilde{\mathbf{p}}_i$ indicates that individual i experiences the negative real asset price shock.

[17] If $\underline{v}_i < 0$, it would imply individual human capital market price is negative, an unrealistic scenario.

[18] An algorithm in Appendix Section A.1 describes the process of iteratively calculating node valuations in the event of a negative shock in order to determine the number of total financial failures resulting from the initial shock.

The dependency matrix $\mathbf{A}$ not only describes the share of an individual's wealth owed to claimants, but also the share of failure costs absorbed by some individual's financial failure. When an individual fails financially their remaining net worth (collateral) is wiped out due to failure costs—we assume $\beta_{i,t} = v_{i,t}$ in our parameterization in Section 4.2. This is the contagion mechanism, where failure costs spread according to the dependency matrix $\mathbf{A}$.

Consider the following example of contagion dynamics (illustrated in Figure 3). Suppose some individual j is the first financial failure as a consequence of a real asset price shock $\tilde{p}_j = \lambda$ at time $t = 0$, such that $v_{j,1} < \underline{v}_j = \theta v_j$ (the first re-evaluation at $t = 1$). This prompts failure costs $\beta_{j,1}$ that deplete j's collateral wealth and are partially absorbed by, for example, individual i's dependency on j as represented by a nonzero value for element A_{ij} in the dependency matrix. Such codependence implies i's value decreases by the amount $A_{ij}\beta_{j,1}$ in period $t = 2$. Should i's value $v_{i,2}$ fall below $\underline{v}_i$, it would incur its own failure cost $\beta_{i,2}$ and as a consequence alter the values, in period $t = 3$, of all individuals i is financially connected to (directly or indirectly) through the dependency matrix $\mathbf{A}$.

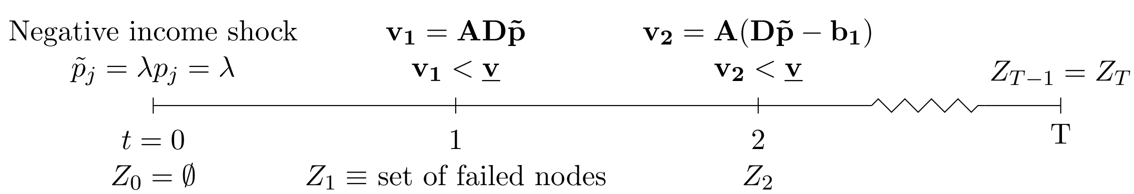

Figure 3: Timeline of Network Contagion

A static financial network gives way to a dynamic process of cascading failures. Network instability is defined by the share of the network that fails financially. The instability is initiated by a decrease in one individual's earning capacity and wealth, hindering their ability to service financial debts and thus provide cash flows for the financial claims creditors have on their output.[19] This cessation of cash outflows to creditors decreases each creditor's wealth, setting off progressive failures as any decline in a creditor's wealth below their own failure threshold would initiate additional failure costs. Any shock to individual net worth may ignite contagion. The model emphasizes the

[19]This appropriates Minsky's position on financial instability: "the behavior and particularly the stability of the economy change as the relation of payment commitments to the funds available for payments changes and the complexity of financial arrangements evolve." (Minsky, 1986a, p. 197)

role of network topology on instability by shocking one individual rather than the entire network, which could cause instability simply because of the scope of the shock and not necessarily network structure.

An algorithm to identify the set of failed nodes Z_t is outlined in Section A.1 of the Appendix. Each iteration of the algorithm solves for Z_t, the subset of nodes who fail as a direct result of the preceding $t-1$ group's failures. Contagion stops when no new individuals in the network fail, or $Z_T = Z_{T-1}$.

The level of network instability is defined as the share of individuals in the network who have failed financially $S = \frac{|Z_T|}{n}$. One possible interpretation of a financial crisis is a sufficient share of the network failing financially, though we are agnostic about a specific threshold. Network failure in the model is driven by drops in the value of, initially real but then financial, assets.

Congruent with empirical definitions of financial crisis that specify the magnitude asset values must drop, the share of nodes failing financially in the model S is equivalent to the percentage decline in financial asset values of the entire network. In other words, when 10 percent of nodes fail then total network wealth declines by 10 percent.

2.6 Example networks

Consider a simple network with $n = 3$ nodes and increasing numbers of financial assets. The example will be illustrative of the network and matrix structures, not contagion effects.[20] Throughout, we assume $\mathbf{D} = \mathbf{I}_3$ and $p_k = 1 \ \forall \ k$.

First, consider an unconnected network—no edges linking any nodes exist (Figure 4a). In a network with no financial claims, each individual keeps all future cash flows and their net worth depends only on their human capital—which is homogeneous. When a shock occurs, only the wealth of the individual experiencing the shock declines, but every other node is isolated. No contagion can occur.

Next, suppose two financial assets are introduced into the network (Figure 5). The total share

[20] Because this is the smallest possible network that can display a variety of link structures, a shock to any connected node may or may not immediately cause failure for all nodes.

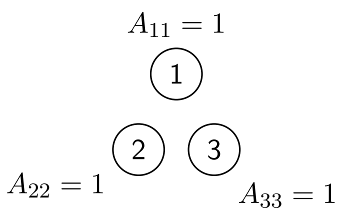

(a) Graph of network

$$\mathbf{C} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \mathbf{A} = \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-\mathbf{1}} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

(b) Corresponding cross-holdings and dependency matrices

Figure 4: Three-node network with no financial assets

of a node's net worth that may be claimed by other nodes, c, is 0.5. All elements in the diagonal savings matrix $\hat{\mathbf{C}}$ will be 1, unless financial claims are made on a node's value and it equals 0.5. The two financial assets represent two claims: node 1 has a claim on node 2's future cash flows and node 2 has a claim on node 3's ($d_1 = d_2 = 1 = d_2^{out} = d_3^{out}$, while $d_3 = 0 = d_1^{out}$). According to (1), $C_{12} = 0.5$. Node 1, therefore, has claim to half of node 2's cashflows while node 2 retains the other half. The same relationship holds between nodes 2 and 3, where $C_{23} = 0.5$. Importantly, nodes 1 and 3 are indirectly connected through node 2, though no direct link exists. (Note the dashed edge in Figure 5a.) Hence $A_{13} = 0.25 > C_{13} = 0$, because node 2 claims half of node 3's net worth, and node 1 claims half of 2's. Node 1 also has the highest net worth ($\sum A_{1j} = 1.75$) of which 0.75 is derived from the other two nodes. Node 2 has a net worth of 0.75, of which 0.25 is derived from node 3, and node 3 has no financial assets and thus a net worth of only 0.5 (equal to its own savings). A shock to node 1 would have no effect on the other nodes since no other nodes have financial claims on node 1 or are dependent on node 1's net worth. Only if nodes 2 or 3 were shocked could multiple nodes fail (nodes 1 and 2) since others are dependent on them.

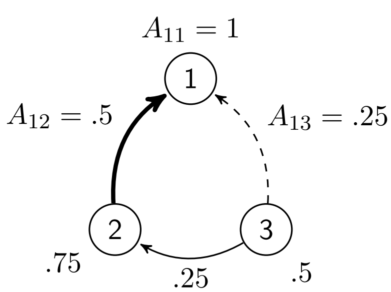

(a) Graph of network

$$\mathbf{C} = \begin{bmatrix} 0 & .5 & 0 \\ 0 & 0 & .5 \\ 0 & 0 & 0 \end{bmatrix}, \mathbf{A} = \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-\mathbf{1}} = \begin{bmatrix} 1 & .5 & .25 \\ 0 & .5 & .25 \\ 0 & 0 & .5 \end{bmatrix}$$

(b) Corresponding cross-holdings and dependency matrices

Figure 5: Three-node network with two financial assets

Next, we introduce another asset into the network for a total of three financial assets (Figure

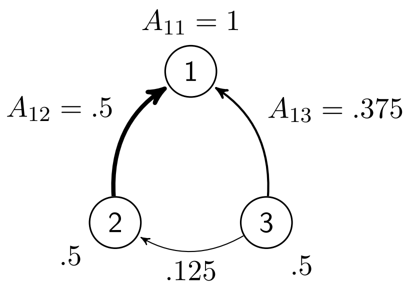

(a) Weighted network graph

$$\mathbf{C} = \begin{bmatrix} 0 & .5 & .25 \\ 0 & 0 & .25 \\ 0 & 0 & 0 \end{bmatrix}, \ \mathbf{A} = \begin{bmatrix} 1 & .5 & .375 \\ 0 & .5 & .125 \\ 0 & 0 & .5 \end{bmatrix}$$

(b) Corresponding cross-holdings and dependency matrices

Figure 6: THREE-NODE NETWORK WITH THREE FINANCIAL ASSETS

6). Node 1 gains an explicit financial claim on node 3. The in-degree of each node is now $d_1 = 2, d_2 = 1, d_3 = 0$. Of the 0.5 share of node 3's value that is securitized within the network, half goes to node 2 and the other half to node 1. But node 1 still receives indirect cash flows from node 3 via node 2. Thus its total cash inflows from node 3 are greater than its direct cash flows, or $A_{13} = 0.375 > C_{13} = 0.25$. Contagion would depend on which individual is initially shocked. For example, if $\lambda = 0$ and node 1 were shocked (such that $\tilde{p}_1 = 0$), then only node 1 would fail financially. No other nodes depend on its value so its failure would not disrupt the net worth of others. If, on the other hand, node 3 were shocked then because its value backs the financial assets held by the other nodes it would cause all three nodes to fail.

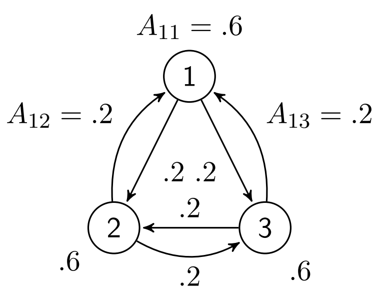

(a) Weighted network graph

$$\mathbf{C} = \begin{bmatrix} 0 & .25 & .25 \\ .25 & 0 & .25 \\ .25 & .25 & 0 \end{bmatrix}, \ \mathbf{A} = \begin{bmatrix} .6 & .2 & .2 \\ .2 & .6 & .2 \\ .2 & .2 & .6 \end{bmatrix}$$

(b) Corresponding cross-holdings and dependency matrices

Figure 7: THREE-NODE NETWORK WITH MAXIMUM $(n-1)$ FINANCIAL ASSETS

Finally, suppose all nodes are linked such that $d_i = d_i^{out} = n - 1 \ \forall \ i$ (Figure 7). The network has absorbed the maximum possible number of financial assets $(n^2 - n)$ and represents a *complete graph*—a special case of a *regular graph* where all nodes have equal degree. Each node has equal net worth: 0.6 from oneself and 0.2 from each of the other two nodes. Since everyone is connected in both directions, any shock will precipitate contagion.

3 Empirics on Financial Networks

To motivate the choice of a Pareto distribution to model inequality of financial assets (and thus financial connections), we summarize some empirical findings from the financial network literature on the connectivity of financial institutions through interbank lending as well as the distribution of those connections.[21] Then we present estimates fitting various datasets of individual wealth to Pareto (power-law) distributions along with their goodness of fit and tests against alternative distributions.

3.1 Interbanking networks

In a seminal work, Furfine (1999) developed an algorithm to parse transactions from federal funds market data for bilateral overnight lending.[22] Summarizing interbank lending market concentration during the first quarter of 1998, Furfine finds that the top 1% of financial institutions in the federal funds market account for two thirds of all assets. They also represent 86 percent of federal funds sold and 97 percent of federal funds bought. These levels of financial market concentration are within the range of parameter estimates tested in our simulations in the next section.

Empirical estimates of various financial network structural parameters from Blasques et al. (2015) are based on data from Dutch interbank markets between 2008 and 2011.[23] Amongst the top 50 lending banks, the authors estimate a mean in-/out- degree of 1.04, with standard deviations of 1.6 and 1.84, respectively. (Banks lend to or borrow from an average of 1.04 different banks.) At the same time, they find very positively skewed in-/out- degree distributions, supporting the Pareto distribution imposed on our network.

Bech & Atalay (2010) describe the topology of the federal funds market in the US between 1997

[21]We must rely on existing research as only aggregate lending data are publicly available. Fedwire Funds Service, a large value transfer service operate by the Federal Reserve—though not unique to federal funds lending—provides bank-level data of the US federal funds market.

[22]All subsequent papers cited in this section rely on the Furfine (1999) algorithm, or adaptations of it, to generate their interbank lending data from the broader Fedwire data. An important caveat of the resulting Furfine interbank lending data are their dependence on transactions occurring through Federal Reserve balance sheets, but not the banks' own lending.

[23]Unlike the Fedwire data in Furfine (1999), the authors use TARGET2 interbank lending data (the Eurosystem equivalent of Fedwire) which specifies individual borrowing and lending institutions for indicated bilateral credit payments. The Dutch interbank lending data has also been cross-validated against Italian and Spanish interbank lending data to minimize type I errors.

and 2006—also using Fedwire data and the Furfine (1999) algorithm. In 2006, banks had an average in-/out- degree of 3.3 ± 0.1 for overnight interbank lending.[24] Among many other parameters describing the topology of the federal funds market, they estimate the out-degree distribution for banks on a representative day in their sample period, concluding that a power-law distribution provides the best fit with a parameter estimate of 1.76 ± 0.02. Their results lend support to our model's degree distribution parameterization, described in Section 4.2.

The aforementioned papers only consider unsecured overnight interbank lending. Bargigli et al. (2015) study both secured and unsecured lending for varying maturities, reflecting our own model more closely—which posits financial assets are secured by an individual borrower's labor income and hence a longer maturity.[25] The authors estimate the the in-/out- degree distributions of the Italian Interbank Network (IIN) between 2008 and 2012, and for 2012 they report power-law parameters on the interval [1.8, 3.5]. A similar parameterization is applied in our model's Pareto degree distribution of individual financial assets. The authors' expected degrees of networks with long-term maturities are also within our range of mean degree values.

While omitting financial intermediaries, our interpersonal financial network framework emphasizes the latent interconnectedness of parties in a financial economy. Estimates on existing networks are therefore helpful guides for reasonable parameterizations.

3.2 Financial distributions

The Pareto distribution, or power law, is typically used to estimate top shares.[26] Thus our model more accurately describes a network of top financial asset holders where we assume financial assets are Pareto distributed. According to the Survey of Consumer Finances (SCF), between 1989 and 2007 US households in the top 1% of households by net worth typically owned one third of all wealth, around 29 percent of all assets, and also nearly one third of all financial assets. The top 10% held nearly two thirds of all wealth and assets, and over 70 percent of all financial assets. The

[24] The authors define a directed link as going from lender to borrower. Thus their definition of a bank's out-degree corresponds to our own definition of an individual's in-degree (cash flows directed in towards the asset holder).

[25] Bargigli et al. (2015) argue that the overnight market is a poor approximation of other interbank lending and thus an inaccurate barometer of systemic risk.

[26] See Kennickell (2009) for estimates of the nonparametric wealth distribution in the US using Survey of Consumer Finances data and Vermeulen (2014) for a detailed discussion on estimating top tails in wealth distributions.

bottom 50%, however, never held more than 3 percent of financial assets or 6.7 percent of all assets (which almost entirely consisted of real estate). We argue that since top wealth holders describe the majority of financial assets, their network topology is a sufficient determinant of overall financial instability.

Given that the power-law relationship $p(x) = \Pr(X = x) = Cx^{-\gamma}$ implies $\ln p(x) = constant + \gamma \ln x$, the approximate linear relationship on a log-log plot suggests its absolute slope is a reasonable estimate of the parameter γ. Since Pareto (1896), power-law distributions have been traditionally estimated thusly: construct a histogram representing the frequency distribution of the variable x; plot on a log-log scale; finally, if approximately linear, estimate its slope to find the scaling parameter.[27]

For numerous reasons outlined in Clauset et al. (2009), the above estimation method is problematic. Instead, the authors propose a maximum likelihood estimation method whereby the scaling parameter γ is estimated conditional on a correct estimate of the lower bound value for power-law behavior x_{min}—as chosen by Kolmogorov-Smirnov statistics. Following the methodology of Clauset et al. (2009) and applying it to the 1989 and 2010 Survey of Consumer Finances, we find a wide range of plausible power-law fittings for US household data on total net worth, financial assets, and total debt.[28] We repeat the exercise for comparable variables using three international datasets from the Luxembourg Wealth Study (LWS): the UK in 2007, and Australia and Italy in 2010.

Results vary by country (Table 1). The US data are the least representative of a Pareto, or power-law, distribution. Though parameter estimates are easily fitted to the data, hypothesis testing rejects a statistically significant goodness of fit between generated data and fitted data.[29] The Pareto distribution fits US financial asset data from 1989 best, though only 60 percent of comparisons between generated and fitted data fail to reject the null that they come from the same Pareto distribution. In all other sets of US data we reject the null the majority of the time. However, we also reject any alternative distributions (the exponential and lognormal, both with

[27] In Pareto (1896), $\hat{\gamma}$ was approximately 1.5—and conjectured to be fixed.

[28] Estimation programs are available online at http://tuvalu.santafe.edu/∼aaronc/powerlaws/.

[29] Generated data come from 2,500 randomly generated Pareto distributions simulated from our fitted parameter estimates.

and without cutoff values) as good fits of the US data.[30] Fitted Pareto parameters range from 1.450 (US net worth in 2010), indicating high inequality, to 2.208 (US financial assets in 1989), indicating much lower inequality.

Data for the UK, Australia and Italy consistently fit a Pareto distribution, across all household variables. In at least 87 percent of the comparisons between generated and fitted Pareto distributions, we cannot reject a difference between the two. Alternative distributions are also unanimously rejected as possible models. Though the Pareto is a uniformly good fit of the LWS data, the scaling parameter estimates are much higher than for the US data, with a minimum of only 2.224 (AUS financial assets in 2010) and a maximum of 3.571 (AUS liabilities in 2010). One reason why may be that over-sampling of high-earning households occurs in the SCF survey population, but not the other national surveys.

Because a Pareto distribution estimates top wealth inequality in the tail of the distribution, the interpersonal financial network model is representative of top financial asset holders and their influence on financial stability. Along with the empirical literature on interbank networks, our estimates of Pareto parameters for 15 different wealth series suggest that our range of calibrated γ values [1.025, 2.375] for the simulation in the proceeding section are reasonable.

4 Network Simulations

4.1 Setup

In a static random network the number of nodes is fixed and links are established following some probabilistic rule. Let d_i be drawn independently from the Pareto distribution $p(d_i) = ad_i^{-\gamma}$, where γ is the Pareto, or power-law, parameter and a is a normalizing constant. For example, suppose a random draw from the degree distribution yields an in-degree of 10 for individual i. Ten financial assets are owned by i, each backed by the net worths of 10 different individuals. As a creditor, i is represented by a row in the adjacency matrix $\mathbf{G}$. Those 10 financial claims are randomly assigned to debtors, represented along columns in $\mathbf{G}$, where $G_{ii} = 0$. In short, the Pareto draw tells us the

[30] Using the R package `poweRlaw` , we also test against alternative poisson distributions for the US data. Using a Vuong test, also outlined in Clauset et al. (2009), we prefer a Pareto distribution against a poisson in all cases.

row sum of G_i, which is randomly allotted to columns along row i. Also, an increase in aggregate wealth $\sum d_i$ directly increases $d = E[d_i]$ because the network size is fixed at $n = 100$

One characteristic of the Pareto distribution is that its scaling parameter γ decreases in the distribution's skewness. Not only is γ a natural inequality measure, but it is also directly related to top percentile shares: if a random variable is Pareto distributed, then the share going to the top q percent of the population is equal to $S(q) = (\frac{100}{q})^{\frac{1-\gamma}{\gamma}}$. The Gini index can also be derived from the Pareto shape parameter with $GINI = (2\gamma - 1)^{-1}$ when $\gamma > \frac{1}{2}$. Each relationship illustrates that wealth inequality is decreasing in γ.

For each simulation, a random network is generated, one individual is randomly shocked, and then (according to the algorithm in Section A.1 of the Appendix) the total percentage of nodes in the network that have failed financially S is evaluated—our measure of instability. Each simulation specifies a unique parameterization and is repeated 1,000 times. The share of failing nodes S reported is the average across iterations. Note that each iteration generates a unique graph $\mathcal{G}(N, G)$ with a network structure that conforms to an exogenously imposed financial wealth distribution and level of total wealth.[31] The below procedure describes the process in full.

Step 1 Generate a static, directed random network $\mathbf{G}$ with parameter d_i represented by a truncated Pareto probability distribution. (The distribution is truncated to isolate the effect of γ for a given d. At each level of γ a maximum in-degree is set so that d remains constant.)

Step 2 Derive the cross-holdings matrix $\mathbf{C}$ from $\mathbf{G}$ using (1).

Step 3 Calculate individuals' starting values $v_i \; \forall \; i \in n$, given an initial real asset price of $p_k = 1$, and determine failure threshold values $\underline{v}_i = \theta v_i$ for some $\theta \in (0, 1)$.

Step 4 Randomly choose an individual j to experience a negative income shock and decrease its real asset price to $\tilde{p}_j = \lambda p_j$.

Step 5 Assume all other real asset prices remain at 1 and calculate the number of nodes failing according to the algorithm in Section A.1 of the Appendix.

The set of all nodes Z_T who have failed financially, calculated at the algorithm's terminal step,

[31] Additional simulations (not reported) consider $n = 500$ and $n = 1,000$ and yield indistinguishable results. For computational ease, all simulation results are generated with $n = 100$.

yields the share of nodes in the network who have failed S. Results are reported graphically, with S plotted against the wealth inequality parameter γ for varying levels of aggregate wealth d.

Given the assumptions built into the model (i.e. $n = 100$, $\mathbf{D}$=$\mathbf{I}$, $\mathbf{p}$=ι) the share of failing nodes S is also equivalent to the percentage decrease in total financial asset values. The baseline results in Figure 8 can therefore be interpreted as asset value percentage declines, an alternative measure of financial instability.

4.2 Calibration

The share of an individual's net worth that can be securitized c characterizes the percentage of future income flows claimed by creditors in the model. An analogous, and available, macroeconomic variable measuring the burden of liabilities is the debt service ratio (DSR), the share of an individual's income used to repay debt. Aggregate estimates from Drehmann & Juselius (2012) necessitate several assumptions concerning average credit maturity, lending rates, and total outstanding credit. Across a panel of both advanced and developing economies, the aggregate debt-service ratio for households ranges from 5.1 percent in Italy in 2010 to 20.3 percent for Denmark.[32]

Because nodes in our model represent individuals or households who also produce and have presumably made financing decisions, we also consider the DSR of private non-financial firms and corporations. In 2010, Italy has a private non-financial firm aggregate DSR of 12.9 percent and Denmark's equals 29.5 percent. (For non-financial corporations the rates are even higher in 2010: 40.6 percent in Italy and 55.5 percent in Denmark.)

The Federal Reserve produces two similar aggregate DSR estimates for the US: household debt service payments and household financial obligations, both as shares of personal disposable income.[33] Financial obligations include rent payments on tenant-occupied property, auto lease payments, homeowners' insurance, and property tax payments. Thus its ratio is larger, peaking at 18.1 percent in the fourth quarter of 2007 while the DSR was only 13.1 percent in the same period.

[32]Data are available online at http://www.bis.org/statistics/dsr.htm

[33]Data are available from FRED online.
Household debt service payments series: https://research.stlouisfed.org/fred2/series/TDSP
Household financial obligations series: https://research.stlouisfed.org/fred2/series/FODSP

The BIS data for private non-financial firms (corporations) in the US in 2010 is 15.8 percent (39.4 percent).

Heterogeneity of debt burdens may skew aggregate estimates, thus the distributions are examined. In 1989 and 2010 in the US, for example, top wealth holders have a greater DSR than middle portions of the wealth distribution but lower than the household average. (See Figures A.2.1 and A.2.2 in Section A.2 of the Appendix.) Generally, BIS household aggregate estimates are lower than averages calculated from household survey data for overlapping years.[34]

Setting $c \in [0.05, 0.5]$ captures the full range of DSR estimates. In the baseline model $c = 0.3$, modeling an economy with reasonable household cash flow obligations. While higher c values are more akin to firms than individuals, they are also more congruent with the units of analysis in the network literature (Section 3).

In the event of a financial failure, such that $v_i < \underline{v}_i = \theta v_i$, an individual incurs bankruptcy costs or some increased economic burden as a consequence of their depressed net worth. We follow Elliott et al. (2014a) and let θ take on a range of values in [0.8, 0.98]. This provides a wide enough spectrum such that individuals are either very robust to valuation changes or very sensitive.

Since the advent of the US Bankruptcy Act in 1978, the majority of consumer bankruptcy cases are filed under Chapter 7 protection, where assets (above some exemption threshold) are liquidated to pay off creditors of secure debt but the debtor's future income streams are untouched. For example, in 2014 approximately two thirds of all consumer bankruptcy petitions filed in US courts were under Chapter 7.[35] The model assumes that, as in Chapter 7, financially failing individuals liquidate their remaining asset position to cover failure costs. And because failure costs equal the value of the individual's wealth after failure in period t, or $\beta_i = v_{i,t}$, a failed individual's remaining assets (or collateral wealth) are liquidated and wealth drops to zero.

Recall that, an income shock lowers an individual's real asset price, so that $\tilde{p}_i = \lambda p_i = \lambda$. A negative shock may theoretically decrease an individual's labor-earning capacity by varying amounts, depending on an individual's level of savings, the number of wage earners in a household,

[34]The distribution of household DSRs is calculated using Survey of Consumer Finance (SCF) data for the US and Luxembourg Income Study (LIS) data for France and other countries (not shown).

[35]Data are available online at http://www.uscourts.gov/statistics-reports/bapcpa-report-2014.

support systems of friends and family and other financial coping mechanisms. The human capital price decline could be very large if, for example, it was caused by some physical injury preventing a wage earner from earning any labor income through their human capital. In such an instance λ would be small. On the other hand, the income shock may be very small if earning capacity is not greatly inhibited, and so λ is large. The range of λ values tested is in $[0, 0.9]$. So long as $\lambda < \theta$ a failure, and thus contagion, can occur.

Concerning historical US wealth inequality measurements, Wolff (1992) finds the top 1% of individuals own approximately as little as 19 percent of household wealth (excluding retirement wealth) in 1976 and as much as 38 percent in 1922. These translate to Pareto parameter values of 1.56 and 1.27, assuming top wealth shares are described by a power law.[36] In 1962, the first iteration of the Federal Reserve's household survey, the Survey of Consumer Finances (SCF)[37], found a Gini coefficient of 0.72 in wealth with a corresponding top 1% wealth share of 32 percent. In its second iteration in 1983, the SCF found a Gini coefficient for wealth of 0.74 (top 1% wealth share of 31 percent). Using more recent SCF waves, Kennickell (2009) decomposes the wealth distribution. In 1989 the top 1% owned 28.3 percent of financial assets and in 2007 it owned 31.5 percent. Assuming a power law describes top wealth shares for the US in those years, the equivalent Pareto parameters are 1.38 in 1989 and 1.33 in 2007.

Values for the Pareto parameter γ are in the interval [1.025, 2.375], which corresponds to a range of Gini coefficients from 0.9524 to 0.2667. The corresponding range of top 1% shares is from 89.4 percent to 6.95 percent. The parameter space is credible and within the range of empirical estimates of wealth, asset, and liability inequalities estimated in Section 3, described above, and in the literature.[38]

Changes in γ also change the mean $d = E[d_i]$ of the in-degree distribution $f(d_i)$. Therefore the Pareto distribution must be truncated in order to hold d constant as γ varies. It becomes possible to isolate the distribution effect from the aggregate wealth effect. With $n = 100$, possible

[36] Solve for γ in $S(0.01) = 100^{\frac{1-\gamma}{\gamma}}$.

[37] The earliest Federal Reserve Board wealth survey was called the Survey of Financial Characteristics of Consumers.

[38] See Vermeulen (2014), Table 8, for Pareto parameter estimates which merge Forbes billionaire data with national surveys, such as the SCF. Vermeulen's Pareto parameter estimates range from 1.02 in the US (very unequal) to 3.55 in the Netherlands. In his broad survey of power laws in economics, Gabaix (2009) finds 1.5 to be the median estimate found for top wealth.

d values are restricted to the interval [1, 2]. For example, suppose $\gamma = 2.375$ (minimal inequality). The maximum possible d_i is 99 (it is not feasible to have $d_i \geqslant n$). When $\max\{d_i\} = 99$ and $\gamma = 2.375$, then $d = 2$ and represents an upper bound on expected in-degree values under our Pareto distribution. For each level of γ we adjust the maximum d_i accordingly.

The baseline model calibration is the fullowing: $c = 0.3, \theta = 0.92, \beta_{i,t} = \underline{v}_{i,t}, \lambda = 0, \gamma = [1.025, 2.375]$ and $d = [1, 2]$. The full range of parameterizations is summarized in Table 2.

4.3 Results

Two results from the baseline simulation (Figure 8) are emphasized. First, wealth inequality positively increases network instability for moderate to wealthy networks; and second, aggregate wealth has an inverted U-shaped relationship with instability—initially increasing, but then decreasing it. As the network becomes more unequal (γ decreases), the share of individuals in the economy failing financially S increases, but only when the network is sufficiently wealthy. Wealth inequality, in other words, is destabilizing only when the economy attains a minimum level of wealth. In our baseline model this approximately occurs when $d = 1.4$. At or above this level of financial wealth, increasing wealth inequality causes greater financial contagion, greater financial asset losses, and therefore a greater likelihood of financial crisis. The positive contribution of inequality on instability is relatively linear and is most pronounced when our network's wealth has an expected in-degree of 1.6, doubling the amount of instability.

Unlike wealth inequality, the effect of increasing aggregate wealth on stability is notably non-monotonic. Initial increases in aggregate wealth (from $d = 1.0$ to 1.2) increase the share of financial failures but are immune to any inequality effects. At moderate levels of aggregate wealth ($d = 1.4$) instability is higher still, but now inequality begins to have a destabilizing impact as it goes up (γ decreases). The strongest effect of wealth inequality occurs in a moderately wealthy network ($d = 1.6$), when moving from very low wealth inequality to very high inequality roughly doubles the size of the contagion—from around 25 to over 50 percent of the network failing. Finally, at the highest levels of aggregate wealth ($d \geqslant 1.8$), inequality remains positively and significantly related to contagion, however the level, or the share of the network failing financially, is smaller than at

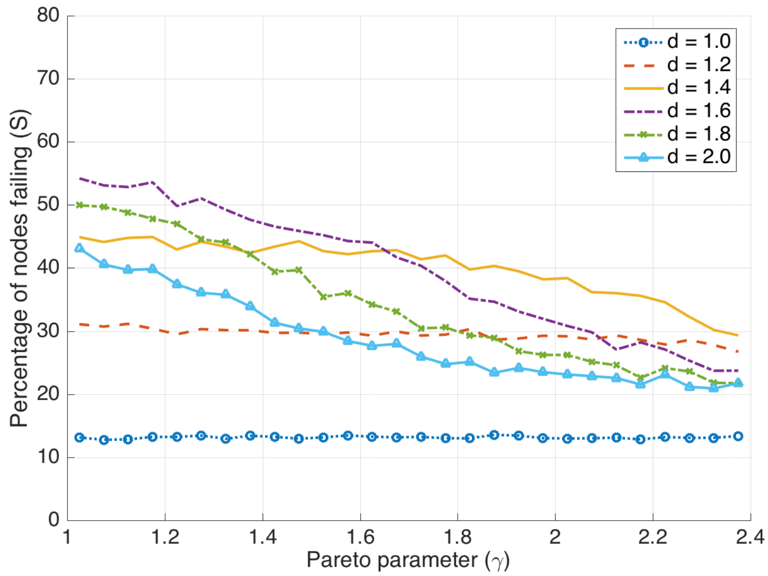

Figure 8: Baseline model
Notes: Aggregate wealth is increasing in expected in-degree d. Calibrated with $c = 0.3$, $\theta = 0.92$, and $\lambda = 0$. As γ increases wealth inequality decreases. The domain of $\gamma = [1.025, 2.375]$ corresponds to Gini coefficients of [0.952, 0.267] and top 1% wealth shares of [0.894, 0.070]. Percentage of financial failures is average of 1,000 iterations.

moderate levels of wealth—nearly 20 percentage points less at some levels of inequality.

The network economy is therefore most unstable, or vulnerable to negative shocks, when it is both wealthier (higher d) and unequal (low γ). The interaction between an economy's level of wealth inequality and total aggregate wealth reflects the "robust-yet-fragile" nonmonotonicity found in other network models.[39]

Importantly, instability occurs independent of the shocked node. Figure 9 depicts two scenarios of identically calibrated networks. In the left panel, the poorest node in the network receives the negative income shock and in the right the richest node is shocked. The level of contagion between them, while significant and nearly identical to our baseline model in both, is noticeable different as is the likelihood of financial crisis. When the poorest node ($\min\{v_i\}$) receives the income shock, a greater share of the network fails for both a given level of inequality and aggregate wealth than when the richest node ($\max\{v_i\}$) is shocked. This makes sense because poorer nodes have more liabilities, and thus their failure costs spread to a greater number of nodes than when the richest node is shocked. The stronger effect from shocking the poorest node is more muted at lower

[39] Gai & Kapadia (2010), Nier et al. (2007), and Elliott et al. (2014a).

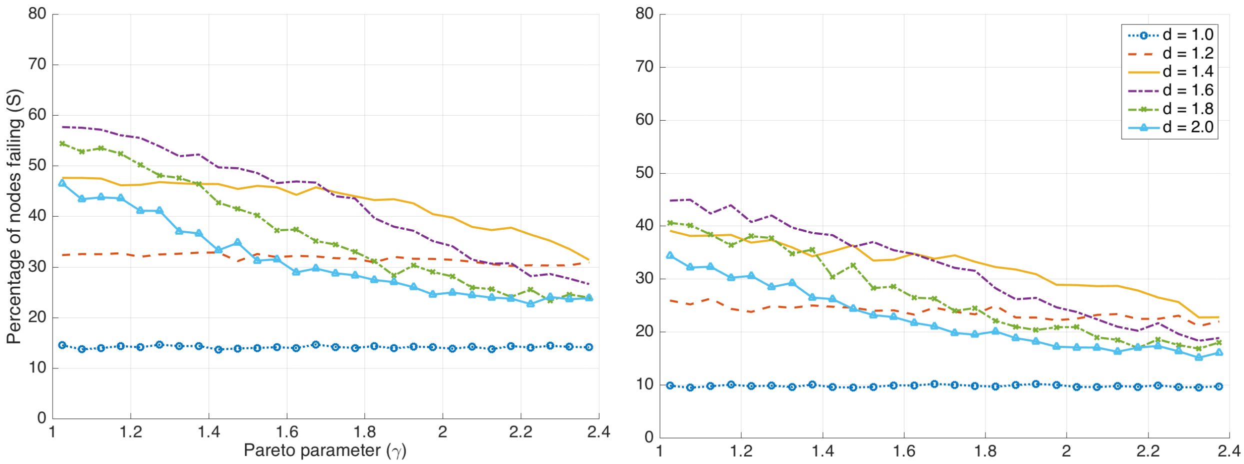

(a) Poorest node ($\min\{v_i\}$) shocked. **(b)** Richest node ($\max\{v_i\}$) shocked.

Figure 9: TARGETED SHOCKS

NOTES:Aggregate wealth is increasing in expected in-degree d. Calibrated with $c = 0.3$, $\theta = 0.92$, and $\lambda = 0$. As γ increases wealth inequality decreases. The domain of $\gamma = [1.025, 2.375]$ corresponds to Gini coefficients of [0.952, 0.267] and top 1% wealth shares of [0.894, 0.070]. Percentage of financial failures is average of 1,000 simulations.

aggregate wealth levels. When the richest node is shocked, networks are more robust by over ten percentage points for the wealthiest networks ($d \geqslant 1.6$) and approximately five percentage points for the least wealthy networks ($d = 1$).

The overall pattern of our baseline model, observed when a random node is shocked, however, persists: increasing wealth inequality (decreasing γ) causes a greater share of individuals to fail in networks of at least moderate wealth while increasing the aggregate wealth (increasing d) of the network is initially destabilizing but then stabilizing.

4.3.1 Regular graphs

To emphasize the importance of both aggregate wealth and the financial wealth distribution on network stability, random regular graphs are simulated for comparison. Regular graphs have equal in-degrees and thus represent perfect financial asset equality in the model. The only parameters changing are c, the percentage of future cash flows owed by an individual to other claimants, and d, the in-degree of all individuals. No longer restricted by the degree distribution parameter γ, d can take on a broader set of values. Results are presented in Figure 10.

As d increases the aggregate wealth of the network increases, though the levels are not necessarily

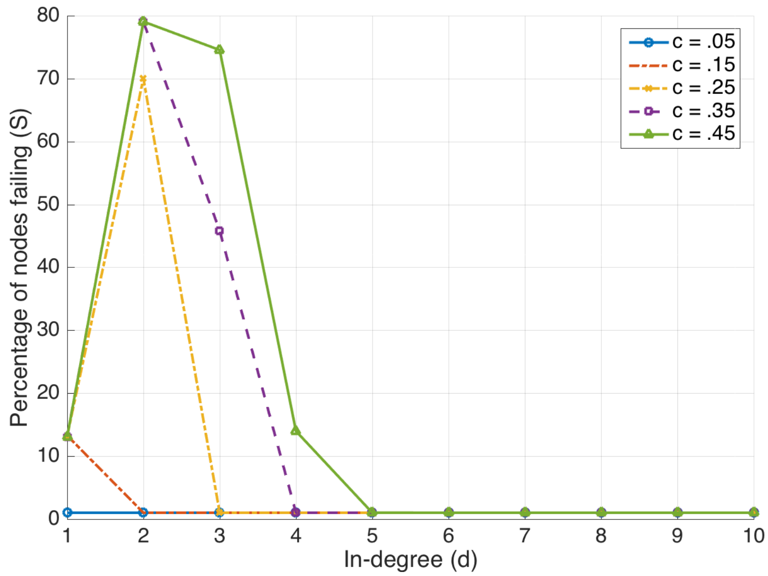

Figure 10: Regular (equal) network

Notes: Regular network contains fixed in-degree d_i for each node, hence there exists perfect wealth equality. Aggregate wealth is increasing in expected in-degree d. Calibrated with $\theta = 0.92$, and $\lambda = 0$. Percentage of financial failures is average of 1,000 iterations.

comparable to the baseline model. When $c > 0.15$, there exists a stark pattern: the share of nodes failing increases sharply when aggregate wealth is low ($d = 1$), but quickly drops again as aggregate wealth increases beyond some level. (The particular level depends on c.) At $d = 5$, instability disappears. Like the models in Figures 8 and 9, the regular network displays increasing instability as aggregate wealth increases from low to moderate levels, but decreasing instability as wealth increases further. Decreasing c, or financialization at the intensive margin, however, also significantly lowers instability.[40]

Simulation results for the full range of parameterizations described in Table 2 are presented in the Appendix, Section A.3. The model is particularly sensitive to the c parameter. This makes intuitive sense because it captures financialization at the intensive margin. The higher c's value is, the more dependent asset holders are on incoming cash flows, and the greater the risk in the event of some financial failure. The θ parameter, the measure of an individual's personal robustness under financial stress (or the economy's ability to absorb depleted cash flows on asset claims) is also critical as it itself determines the failure thresholds. The parameter λ, inversely proportional to the size of the income shock $\tilde{p}_i = \lambda p_i$, is nearly indiscriminate in its effect on contagion (Figure

[40]The step-function-like behavior of the regular network results are due to the fact that individuals must have integer values of $d_i = d$. A rounding function in the program simply rounds up to the next integer.

A.3.3). So long as $\lambda < \theta$, an initial negative income shock will always cause at least one financial failure which will catalyze contagion within the network.

5 Conclusion

Keynes once described the relationship between debtors and creditors as forming "the ultimate foundation of capitalism."[41] This paper's central goal was to examine the relationship between inequality and financial crisis by reducing the financial economy into a network of creditors and debtors who are linked through financial assets. While the income inequality literature has posited that inequality's association with debt may relate to financial crises, studying the distribution of the asset side of the balance sheet illuminates a possible topographical relationship between wealth inequality and crises.

The financial network model presented is a radically simplified interpretation of a financial economy, one that eliminates intermediaries and instead relies on the latent financial pathways that link individual asset and liability holders. Implicit financial links observed between individuals are made explicit in a directed network graph. A link indicates the presence of a financial asset and the direction of cash flows between individuals. It follows that changing the distribution of financial assets changes the arrangement of links in the network. This model of wealth inequality and financial instability, building from the framework of Elliott et al. (2014a), suggests that changes to the network topology have two main effects: first, increasing top wealth inequality, conditional on a network's overall wealth, increases instability; and second, aggregate network wealth should have an increasing and then decreasing effect on instability—measured by the share of nodes in the network that is determined to have failed financially. The model's assumptions of homogeneous real assets and prices allow for another interpretation of financial instability: decreases in total network financial asset values.

One implication of the model is that future increases in wealth inequality in the US and other financially advanced economies (as predicted in Piketty (2014)) may increase macroeconomic instability. The consequences for moral hazard, systemic risk, and too-big-to-fail, among other regulatory

[41](Keynes, 1920, p. 236).

concerns, could also be great. Another broader implication is the incitement to reduce inequality for cogent economic—not simply moral—reasons. Rising inequality will always have broad welfare effects, but macroeconomic health may also be at stake.

References

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

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

Bargigli, L., Di Iasio, G., Infante, L., Lillo, F., & Pierobon, F. (2015). The multiplex structure of interbank networks. *Quantitative Finance*, *15*(4), 673–691.

Bech, M. L., & Atalay, E. (2010). The topology of the federal funds market. *Physica A: Statistical Mechanics and its Applications*, *389*(22), 5223–5246.

Blasques, F., Bräuning, F., & van Lelyveld, I. (2015). A dynamic network model of the unsecured interbank lending market. Working Paper 491, BIS.

Blume, L. E., & Durlauf, S. N. (2015). Capital in the twenty-first century: A review essay. *Journal of Political Economy*, *123*(4), 749–777.

Bordo, M. D., & Meissner, C. M. (2012). Does inequality lead to a financial crisis? *Journal of International Money and Finance*, *31*(8), 2147–2161.

Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review*, *51*(4), 661–703.

Coelho, R., Neda, Z., Ramasco, J., & Santos, M. A. (2005). A family-network model for wealth distribution in societies. *Physica A: Statistical Mechanics and its Applications*, *353*, 515–528.

Davies, J. B., & Shorrocks, A. (1999). The distribution of wealth. In A. B. Atkinson, & F. Bourguignon (Eds.) *Handbook of Income Distribution*, vol. 1, chap. 11, (pp. 605–675). Elsevier.

Davies, J. B., Shorrocks, A., Sandstrom, S., & Wolff, E. N. (2007). The world distribution of household wealth.

Drehmann, M., & Juselius, M. (2012). Do debt service costs affect macroeconomic and financial stability? *BIS Quarterly Review*, *September*, 21–34.

Eisenberg, L., & Noe, T. H. (2001). Systemic risk in financial systems. *Management Science*, *47*(2), 236–249.

Elliott, M., Golub, B., & Jackson, M. O. (2014a). Financial networks and contagion. *The American Economic Review*, *104*(10), 3115–53.

Elliott, M., Golub, B., & Jackson, M. O. (2014b). Online appendix: Financial networks and contagion.

Furfine, C. H. (1999). The microstructure of the federal funds market. *Financial Markets, Institutions & Instruments*, *8*(5), 24–44.

Gabaix, X. (2009). Power laws in economics and finance. *Annual Review of Economics*, *1*, 255–293.

Gai, P., & Kapadia, S. (2010). Contagion in financial networks. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, *466*(2120), 2401–2423.

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

Gu, X., & Huang, B. (2014). Does inequality lead to a financial crisis? revisited. *Review of Development Economics*, *18*(3), 502–516.

Kennickell, A. B. (2009). *Ponds and Streams: Wealth and Income in the US, 1989 to 2007*. Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board.

Keynes, J. M. (1920). *The Economic Consequences of the Peace*. New York: Harcourt, Brace, and Howe.

Kim, S., Kim, G. I., & Lee, G. (2008). Wealth networks with local redistribution. *Physica A: Statistical Mechanics and its Applications*, *387*(19), 4973–4981.

Kopczuk, W. (2015). What do we know about the evolution of top wealth shares in the united states? *Journal of Economic Perspectives*, *29*(1), 47–66.

Kopczuk, W., & Saez, E. (2004). Top wealth shares in the united states, 1916-2000: Evidence from estate tax returns. *National Tax Journal*, *57*(2), 445–87.

Kregel, J. (2014). Regulating the financial system in a minskian perspective. In L. C. Bresser-Pereira, J. Kregel, & L. Burlamaqui (Eds.) *Financial Stability and Growth: Perspectives on Financial Regulation and New Developmentalism*, chap. 9, (pp. 127–142). Routledge.

Krusell, P., & Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, *106*(5), 867–896.

Kumhof, M., & Ranciere, R. (2010). Inequality, leverage and crises. IMF Working Papers 10/268, International Monetary Fund.

Kumhof, M., Rancière, R., & Winant, P. (2015). Inequality, leverage and crises. *The American Economic Review*, *105*(3), 1217–45.

Lee, G., & Kim, G. I. (2007). Degree and wealth distribution in a network induced by wealth. *Physica A: Statistical Mechanics and its Applications*, *383*(2), 677–686.

Mason, J., & Jayadev, A. (2014). "fisher dynamics" in us household debt, 1929–2011. *American Economic Journal: Macroeconomics*, *6*(3), 214–234.

Minsky, H. P. (1975). *John Maynard Keynes*. Columbia University Press New York.

Minsky, H. P. (1986a). *Stabilizing an Unstable Economy*. Yale University Press.

Minsky, H. P. (1986b). Stabilizing an unstable economy. *Hyman P. Minsky Archive*, (Paper 144).

Minsky, H. P. (1992). The financial instability hypothesis. Working Paper 74, Jerome Levy Economics Institute of Bard College.

Morelli, S., & Atkinson, A. B. (2015). Inequality and crises revisited. *Economia Politica*, (pp. 1–21).

Nier, E., Yang, J., Yorulmazer, T., & Alentorn, A. (2007). Network models and financial stability. *Journal of Economic Dynamics and Control*, *31*(6), 2033–2060.

Pareto, V. (1896). *Cours d'économie politique: professé à l'Université de Lausanne*. F. Rouge.

Piketty, T. (2014). *Capital in the 21st Century*. Harvard University Press.

Piketty, T., & Saez, E. (2014). Inequality in the long run. *Science*, *344*(6186), 838–843.

Rajan, R. G. (2011). *Fault Lines: How Hidden Fractures Still Threaten the World Economy*. Princeton University Press.

Saez, E., & Zucman, G. (2014). Wealth inequality in the united states since 1913: Evidence from capitalized income tax data. Tech. Rep. NBER WP 20625, National Bureau of Economic Research.

Stiglitz, J. E. (2012). Macroeconomic fluctuations, inequality, and human development. *Journal of Human Development and Capabilities*, *13*(1), 31–58.

Vermeulen, P. (2014). How fat is the top tail of the wealth distribution? Working Paper Series No 1692, European Central Bank.

Wolff, E. N. (1992). Changing inequality of wealth. *The American Economic Review*, *82*(2), 552–558.

Wolff, E. N. (1994). Trends in household wealth in the united states, 1962–83 and 1983–89. *Review of Income and Wealth*, *40*(2), 143–174.

Zucman, G. (2014). Taxing across borders: Tracking personal wealth and corporate profits. *The Journal of Economic Perspectives*, *28*(4), 121–148.

Zucman, G. (2015). *The Hidden Wealth of Nations*. University of Chicago Press.

Tables

Table 1: Empirical Pareto Estimates

		US 1989	US 2010	UK 2007	Australia 2010	Italy 2010
net worth	$\hat{\gamma}$	1.475	1.450	2.809	2.729	2.904
	$\hat{x}_{min}$	146,468	206,670	940,162	978,558	495,000
Hypothesis testing	PL	reject	reject	fail (98)	fail (92)	fail (98)
	Alt.	reject	reject	reject	reject	reject
financial assets	$\hat{\gamma}$	2.208	1.493	3.254	2.224	2.382
	$\hat{x}_{min}$	5,102,103	184,330	788,000	495,660	59,777
	PL	fail (60)	reject	fail (98)	fail (87)	fail (98)
	Alt.	reject	reject	reject	reject	reject
liabilities	$\hat{\gamma}$	1.988	2.036	3.086	3.571	3.393
	$\hat{x}_{min}$	158,376	217,700	147,000	554,457	109,900
	PL	fail (16)	fail (6)	fail (93)	fail (98)	fail (94)
	Alt.	reject	reject	reject	reject	reject

Sources: US: Survey of Consumer Finances (SCF); UK, Australia, Italy: Luxembourg Wealth Study (LWS)

Notes: Australian, Italian, UK and US data are all in local currency units. SCF (US only) financial asset data are the total market value of financial investments and products, deposit accounts, cash and other financial assets owned by household members, including pension assets as well as life insurance. LWS (GBR, AUS, ITA) financial asset data exclude pension assets and other long-term savings. Net worth data are total assets minus total liabilities, except Italy 2010, where disposable net worth is measured. Hypothesis testing: (PL) null hypothesis of fitted power-law distribution and generated power-law distribution (using estimated parameters) being the same, by Kolmogorov-Smirnov statistic; and (Alt.) null hypothesis of fitted alternative distribution and generated alternative distribution (using estimated parameters) being the same, by Kolmogorov-Smirnov statistic. Alternative distributions tested are an exponential distribution and log normal distribution, both with and without cutoff values ($\hat{x}_{min}$). If we fail to reject a null, the percentage of 2,500 simulated fittings of generated and fitted data which fail to reject null is reported in parentheses.

Table 2: Parameter calibration for static random network simulations

Variable	Values	Source(s)
c	$[0.05, 0.5]$	Author's estimates (Section A.2), Drehmann & Juselius (2012), BIS, FRB St. Louis
θ	$[0.8, 0.98]$	Elliott et al. (2014a)
β_i	$\underline{v}_i$	UScourts.gov (Federal Caseload Statistics)
λ	[0, 0.9]	
γ	[1.025, 2.375]	Author's estimates (Section 3.2), Elliott et al. (2014b)
d	[1, 2]	Blasques et al. (2015), Elliott et al. (2014b)

Appendix

A.1 Failure Algorithm

This algorithm is used the determine the ordering of individuals who fail financially in the event of an initial income shock. It finds what Elliott et al. (2014a) refer to as the *best-case* equilibrium, i.e. there exist the fewest number of failures and highest values $v_{i,t}$ possible.

The initial financial shock occurs at period $t = 0$, changing real asset price values to $\tilde{\mathbf{p}}$. Let Z_t represent the set of financially failed individuals at period t, where $Z_0 = \emptyset$. Then for periods $t \geqslant 1$:

Step 1 Let $\mathbf{b}_{t-1}$ be a vector of failure costs with element $b_{i,t-1} = \beta_i$ if $i \in Z_{t-1}$ and 0 otherwise. By definition, $\beta_i = 0\ \forall\ i$ at $t = 1$.

Step 2 Let Z_t be the set of all j where $v_{j,t} < 0$ and:

$$\mathbf{v_t} = \mathbf{A}(\mathbf{D}\tilde{\mathbf{p}} - \mathbf{b_{t-1}}) - \underline{\mathbf{v}}.$$

Step 3 Stop iterations if $Z_t = Z_{t-1}$, otherwise return to Step 1.

The resulting set Z_T, at terminal period T, is the corresponding set of individuals who have failed financially. An important feature is that the individuals added each period $(Z_t - Z_{t-1})$ are those individuals whose financial failures were catalyzed by the preceding set of cumulative failures. For example, Z_1 is the first group of individuals to fail and Z_2 includes the group of individuals who fail in the second period as a direct result of the individuals failing during period $t = 1$.

A.2 Distributions of Household Debt Service Ratio (DSR)

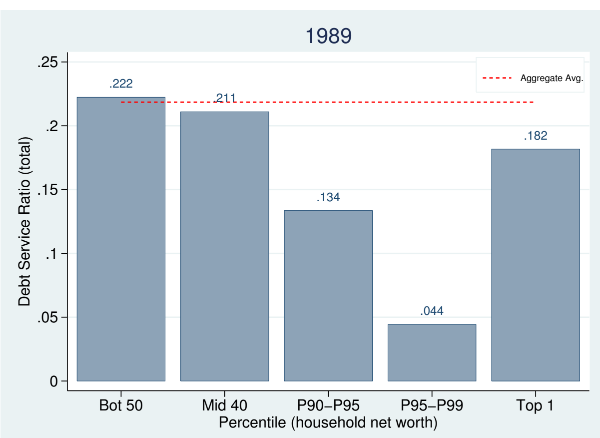

SOURCE: Survey of Consumer Finances (SCF)

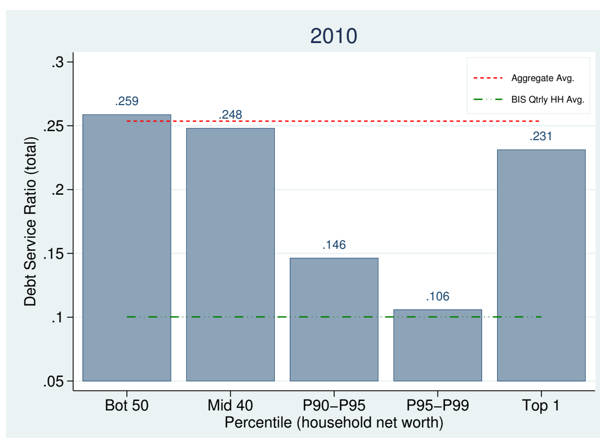

SOURCE: Survey of Consumer Finances (SCF)

Figure A.2.1: US: 1989, 2010

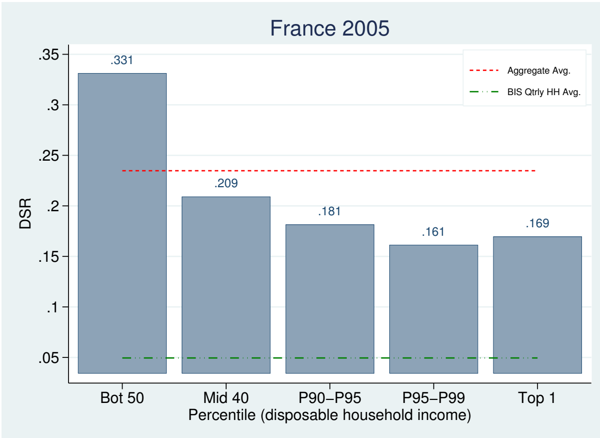

SOURCE: Luxembourg Income Study (LIS)

Figure A.2.2: FRANCE: 2005

A.3 Additional Parameterizations

A.3.1 Changes in parameter c

The parameter c determines the share of each node's value that can be securitized and claimed by other nodes. It measures the share of a node's cash flows that are sent to creditor nodes, an approximation of the level of financialization in the network at the intensive margin. Simulation results of the static random network for values of $c = \{0.1, 0.2, 0.3, 0.4, 0.6, 0.8\}$, and $\theta = 0.92$ and $\lambda = 0$, are shown in Figure A.3.1.

When $c = 0.1$, less than seven percent of the nodes fail under all levels of wealth inequality and aggregate wealth. Only as the share of individual wealth that can be claimed increases ($c \geqslant 0.2$) does the positive effect of wealth inequality on S begin to assert itself at moderate levels of wealth ($d \in [1.4, 1.6]$). As financialization at the intensive margin, c, continues increasing instability at the highest levels of aggregate wealth keeps increasing until we reach a maximum amount of contagion at approximately $c = 0.6$. (See Figure A.3.1.) Hereafter, instability declines as c increases. Thus the inverted U-shaped effect of financialization (at the extensive margin) on instability observed from increasing network wealth (d) appears to also take shape when financialization at the intensive margin (c) is also increased—though the downward sloping portion occurs at values of c that are well beyond any reasonably estimated debt servicing burden, commercial or private. These results broadly echo those of Drehmann & Juselius (2012) who show debt service burdens positively predict economic downturns.

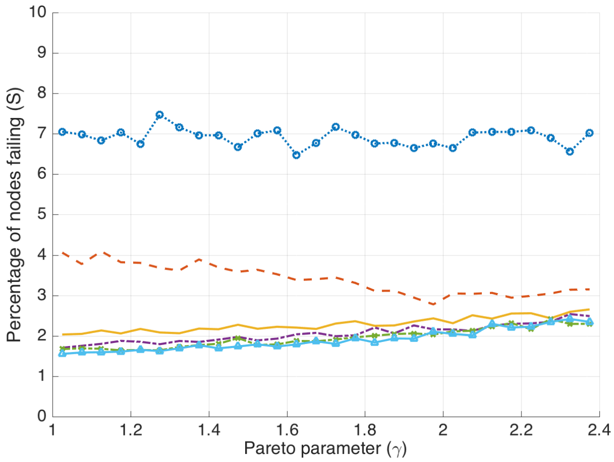

(a) $c = 0.1$

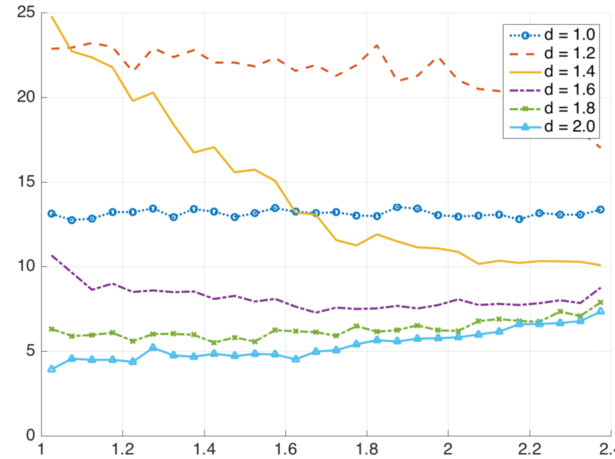

(b) $c = 0.2$

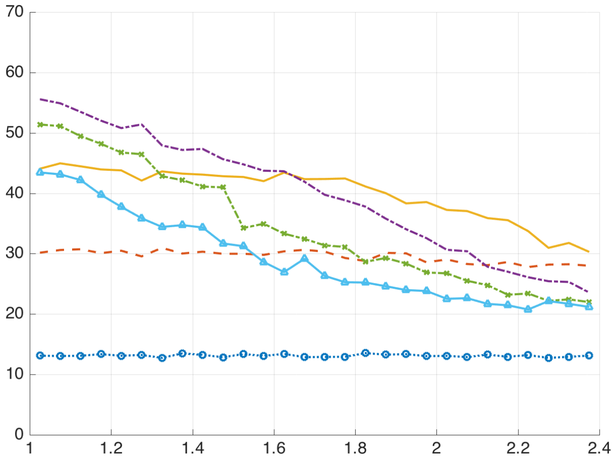

(c) $c = 0.3$

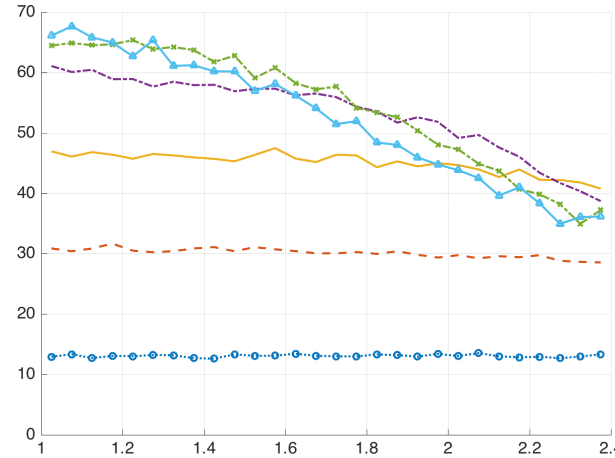

(d) $c = 0.4$

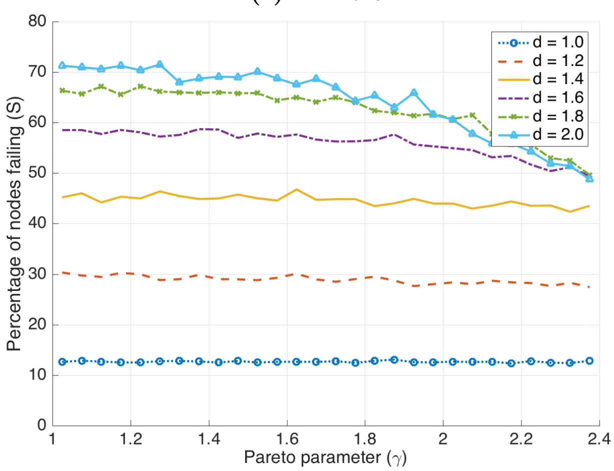

(e) $c = 0.6$

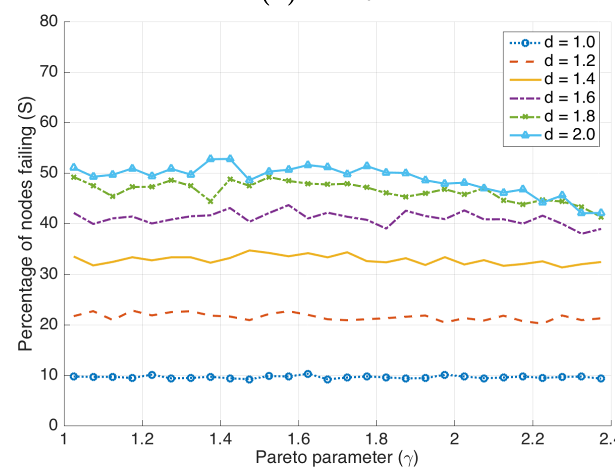

(f) $c = 0.8$

Figure A.3.1: Changes in parameter c

Notes: Pareto distributed in-degree d. Aggregate wealth is increasing in expected in-degree d. $\theta = 0.92$ and $\lambda = 0$. As γ increases wealth inequality decreases. The domain of $\gamma = [1.025, 2.375]$ corresponds to Gini coefficients of [0.952, 0.267] and top 1% wealth shares of [0.894, 0.070]. Percentage of financial failures is average of 1,000 simulations.

A.3.2 Changes in parameter θ

The parameter θ determines the financial robustness of an individual node in the event of an income shock. Since financial failure is predicated on $v_i < \underline{v}_i$ and $\underline{v}_i = \theta v_i$, the smaller θ is the more financially robust an individual is. An individual's financial fragility is increasing in θ. Simulation results of the static random network for values of $\theta = \{0.8, 0.84, 0.88, 0.92, 0.94, 0.98\}$, and $c = 0.3$ and $\lambda = 0$, are shown in Figure A.3.2.

As θ increases an individual is more likely to breach $\underline{v}_i$ in the event that they personally experience an income shock or absorb failures indirectly through the dependency matrix $\mathbf{A}$. When θ is smallest (0.8), individuals are especially robust to any shock and the share of failing nodes is very low ($S < 2\%$). See Figure A.3.2. When θ increases (0.88) individual financial vulnerability increases, but contagion is still very low and unaffected by inequality. When θ is high ($\geqslant 0.92$), only a slight disturbance can tip an individual into financial failure and contagion spreads easily. The impact of wealth inequality on contagion is also strongly felt, but, again, is dependent on the network's aggregate wealth level.

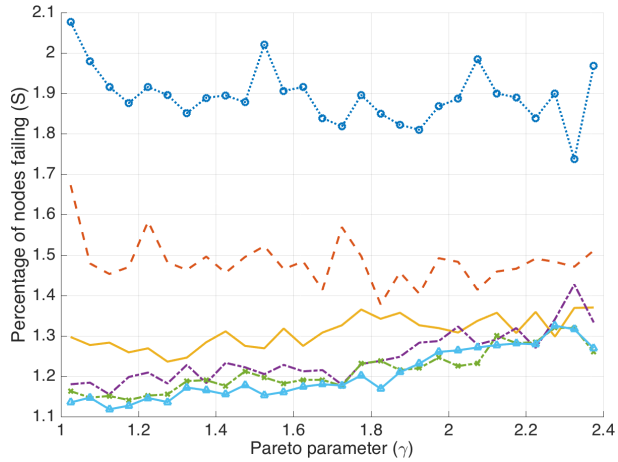

(a) $\theta = 0.8$

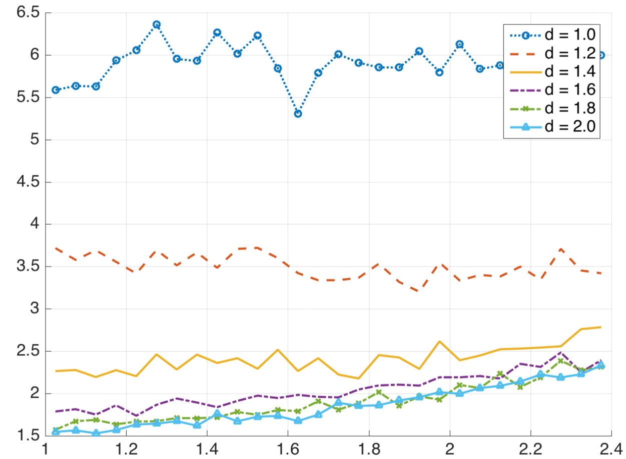

(b) $\theta = 0.84$

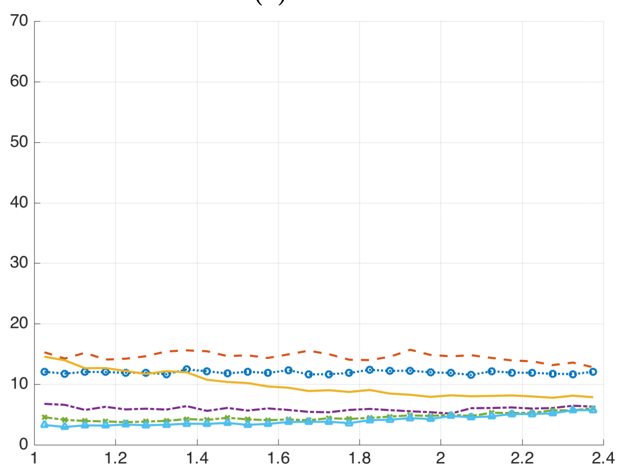

(c) $\theta = 0.88$

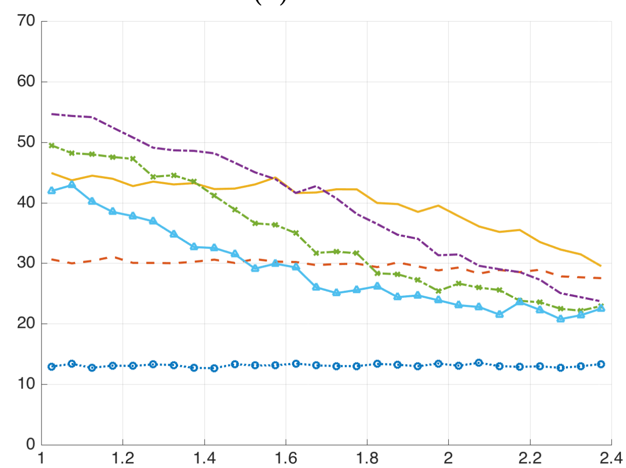

(d) $\theta = 0.92$

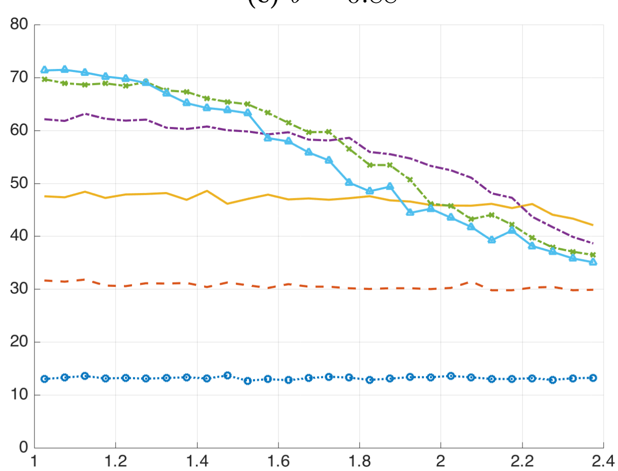

(e) $\theta = 0.94$

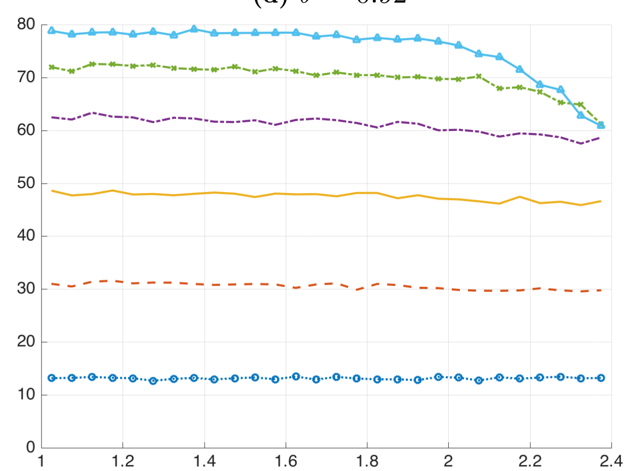

(f) $\theta = 0.98$

Figure A.3.2: CHANGES IN PARAMETER θ

NOTES: Pareto distributed in-degree d. Aggregate wealth is increasing in expected in-degree d. $c = 0.3$ and $\lambda = 0$. As γ increases wealth inequality decreases. The domain of $\gamma = [1.025, 2.375]$ corresponds to Gini coefficients of [0.952, 0.267] and top 1% wealth shares of [0.894, 0.070]. Percentage of financial failures is average of 1,000 simulations.

A.3.3 Changes in parameter λ

The parameter λ determines the magnitude of the random income shock imposed on a single node. An income shock decreases the market price of the node's real asset to $\tilde{p}_k = \lambda p_k$, where $p_k = 1$ and $\lambda \in [0, 1)$. Therefore as the magnitude of the income shock is decreasing in λ. Simulation results of the static random network for values of $\lambda = \{0, 0.25, 0.5, 0.75, 0.9\}$, and $c = 0.3$ and $\theta = 0.92$, are shown in Figure A.3.3.

Only when λ is very close to θ in value (0.9) is there any significant decrease in contagion. If $\lambda < \theta$, no matter the size of the shock the overall pattern of our simulation results holds: increasing inequality causes an increase in the percentage of nodes failing, conditional on a certain level of aggregate wealth; and increasing the aggregate wealth of the network, first increases then decreases network stability.[42]

[42]We test one counterfactual simulation in which a random individual receives a positive income shock, setting $\lambda = 2$. Because contagion is a property of net worth decreasing below some threshold value, we expect increases in net worth to have no effect on contagion. As our model would predict, the network is perfectly stable and no financial failures occur at any level of aggregate wealth. Contagion is conditional upon some negative shock.

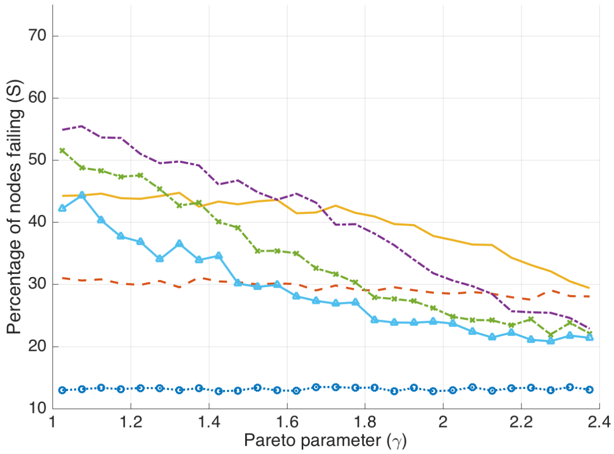

(a) $\lambda = 0$

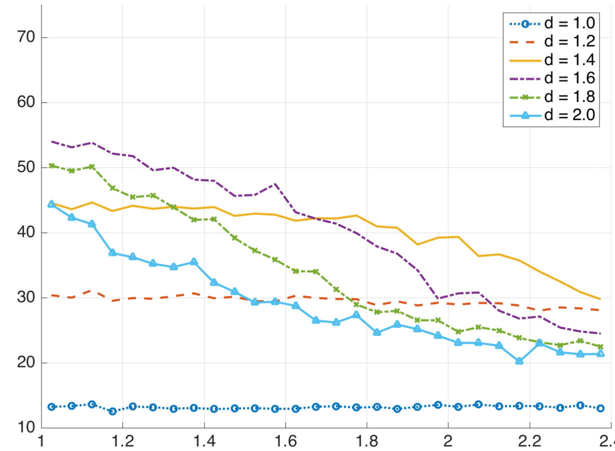

(b) $\lambda = 0.25$

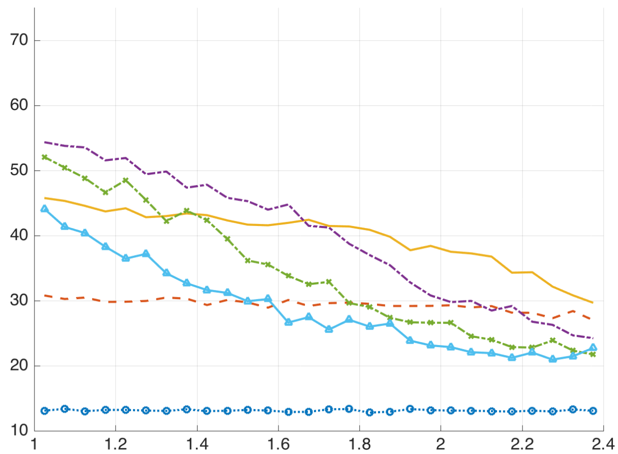

(c) $\lambda = 0.5$

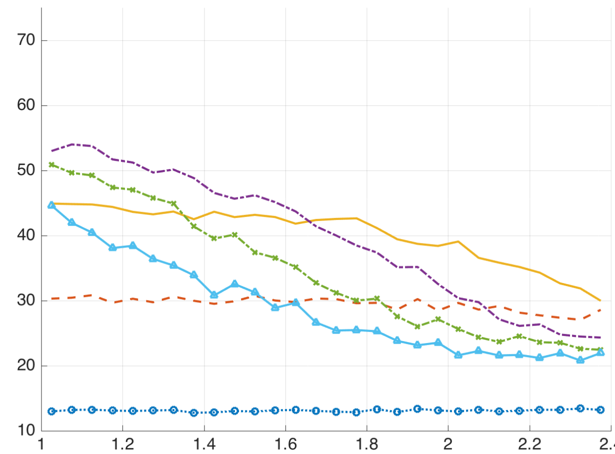

(d) $\lambda = 0.75$

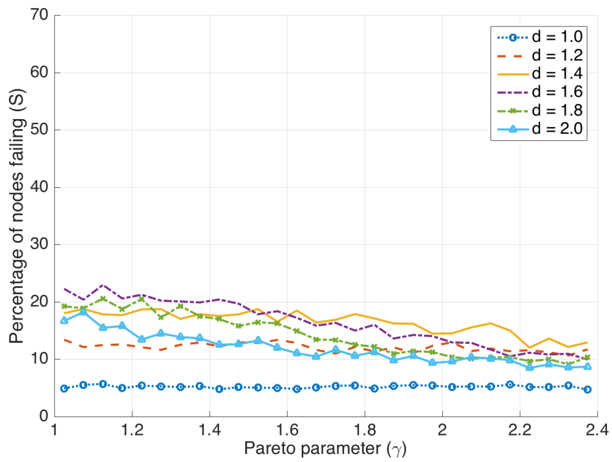

(e) $\lambda = 0.9$

Figure A.3.3: Changes in parameter λ.

Notes: Pareto distributed in-degree d. Aggregate wealth is increasing in expected in-degree d. $c = 0.3$ and $\theta = 0.92$. As γ increases wealth inequality decreases. The domain of $\gamma = [1.025, 2.375]$ corresponds to Gini coefficients of [0.952, 0.267] and top 1% wealth shares of [0.894, 0.070]. Percentage of financial failures is average of 1,000 simulations.