CHAPTER

Models of Financial Stability and Their Application in Stress Tests*

6

Christoph Aymanns*, J. Doyne Farmer†,‡,§,¶,1, Alissa M. Kleinnijenhuis†,‡,‖, Thom Wetzer†,‖,**

*Swiss Institute of Banking and Finance, University of St. Gallen, St. Gallen, Switzerland
†Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford, United Kingdom
‡Mathematical Institute, University of Oxford, Oxford, United Kingdom
§Department of Computer Science, University of Oxford, Oxford, United Kingdom
¶Santa-Fe Institute, Santa Fe, NM, United States
‖Oxford-Man Institute of Quantitative Finance, University of Oxford, Oxford, United Kingdom
**Faculty of Law, University of Oxford, Oxford, United Kingdom
1Corresponding author: e-mail address: doyne.farmer@inet.ox.ac.uk

CONTENTS

*The authors thank Tobias Adrian, Fabio Caccioli, Agostino Capponi, Darrell Duffie, Luca Enriques, Cars Hommes, Sujit Kapadia, Blake LeBaron, Alan D. Morrison, Paul Nahai-Williamson, James Paulin, Peyton Young, Garbrand Wiersema, an anonymous reviewer, and the participants of the Workshop for the Handbook of Computational Economics for their valuable comments and suggestions. The usual disclaimers apply.

Handbook of Computational Economics, Volume 4, ISSN 1574-0021, https://doi.org/10.1016/bs.hescom.2018.04.001

1 INTRODUCTION

The financial system is a classic example of a complex system. It consists of many diverse actors, including banks, mutual funds, hedge funds, insurance companies, pension funds and shadow banks. All of them interact with each other, as well as interacting directly with the real economy (which is undeniably a complex system in and of itself). The financial crisis of 2008 provided a perfect example of an emergent phenomenon, which is the hallmark of a complex system.

While the causes of the crisis remain controversial, a standard view goes like this: A financial market innovation called mortgage-backed securities made lenders feel more secure, causing them to extend more credit to households and purchase large quantities of securities on credit. Liberalized lending fueled a housing bubble; when it crashed, the fact that the portfolios of most major financial institutions had significant holdings of mortgage backed securities caused large losses. This in turn caused a credit freeze, cutting off funding for important activities in the real economy. This generated a global recession that cost the world an amount that has been estimated to be as high as fifty trillion dollars, the order of half a year of global GDP. Here

we will present an alternative hypothesis, suggesting the possibility that the housing bubble was only the spark that lit the fire, and a deeper underlying cause might have been the build up of systemic risk over time. As we will discuss below, a substantial part of this systemic risk may have been due to backward looking, procyclical risk management of leveraged financial institutions.[1] In either case, the crisis provides a clear example of an emergent phenomenon.

The crisis has made everyone aware of the complex nature of the interactions and feedback loops in the economy, and has driven an explosive amount of research attempting to better understand the financial system from a systemic point of view. It has also underlined the policy relevance of the complex systems approach. Systemic risk occurs when the decisions of individuals, which might be prudent if considered in isolation, combine to create risks at the level of the whole system that may be qualitatively different from the simple combination of their individual risks. By its very nature systemic risk is an emergent phenomenon that comes about due to the nonlinear interaction of individual agents. To understand systemic risk we need to understand the collective dynamics of the system that gives rise to it.

The financial system is sufficiently complicated that it is not yet possible to model it realistically. Existing models only attempt a stylized view, trying to elucidate the underlying mechanisms driving financial stability. There are currently two basic approaches. The mainstream approach has been to focus on situations where it is possible to compute an equilibrium. This generally requires making very strong simplifications, e.g. studying only a few actors and interactions at a time. The equilibrium approach has been useful to clarify some of the key mechanisms driving financial instabilities and financial contagion, but it comes at the expense of simplifications that limit the realism of the conclusions. There is also a concern that, particularly during a crisis, the assumptions of rationality and equilibrium are too strong.

The alternative approach abandons equilibrium and rationality and replaces them with behavioral assumptions.[2] This approach often relies on simulation, which has the advantage that it is easier to study more complicated situations, e.g. with more actors and more realistic institutional constraints. It also makes it possible to study multiple channels of interaction; even though research in this direction is still in its early stages, it is clear that this plays an important role.

The use of behavioral assumptions as an alternative to utility maximization is controversial. Unlike utility, behavioral assumptions have the advantage of being directly observable, and in many cases the degree to which they are followed can be

[1] One example of procyclical risk management practices based on backward looking risk estimates are Value-at-Risk constraints. Such constraints were imposed on the trading book of banks under Basel II. However, many other leveraged institutions, not subject to Basel II, also used Value-at-Risk constraints in their internal risk management.

[2] Rationality vs. behavioral assumptions can be regarded as two poles in a continuum. A useful intermediate alternative is to place more emphasis on agents that learn. The computational approach is often forced to abandon rationality because in more realistic settings perfect rationality may be computationally intractable, but numerical approximations with optimizing behavior may be feasible.

confirmed empirically. The disadvantage of this approach is that behavior may be context dependent, and as a result, such models typically fail the Lucas critique. We will show examples here where models based on behavioral assumptions are nonetheless very useful because they make it possible to directly investigate the consequences of a given set of behaviors. We will show examples where it leads to simple models that make clear predictions, at the same time that it can potentially be extended to complex real world situations.

This review will focus primarily on the simulation approach, though we will attempt to discuss key influences and interactions with the more traditional equilibrium approach. Our view is that the two approaches are complements rather than substitutes. The most appropriate approach depends on the context and the goals of the modeling exercise. We predict that the simulation approach will become increasingly important with time, for several reasons. One is that this approach can be easier to bring to the data, and data is becoming more readily available. Many central banks are beginning to collect comprehensive data sets that make it possible to monitor the key parts of the financial system. This makes it easier to test the realism of behavioral assumptions, making such models less ad hoc. With such models it is potentially feasible to match the models to the data in a literal, one-to-one manner. This has not yet been done, but it is on the horizon, and if successful such models may become valuable tools for assessing and monitoring financial stability, and for policy testing. In addition, computational power is always improving. This is a new area of pursuit and the computational techniques and software are rapidly improving.

The actors in the financial system are highly interconnected, and as a consequence network dynamics plays a key role in determining financial stability. The distress of one institution can propagate to other institutions, a process that is often called *contagion*, based on the analogy to disease. There are multiple channels of contagion, including counterparty risk, funding risk, and common assets holdings. *Counterparty risk* is caused by the web of bilateral contracts, which make one institution's assets another's liabilities. When a borrower is unable to pay, the lender's balance sheet is affected, and the resulting financial distress may in turn be transmitted to other parties, causing them to come under stress or default. *Funding risk* occurs when a lender comes under stress, which may create problems for parties that routinely borrow from this lender because loans that they would normally expect to receive fail to be extended. Institutions are also connected in many indirect ways, e.g. by common asset holdings, also called *overlapping portfolios*. If an institution comes under stress and sells assets, this depresses prices, which can cause further selling, etc. There are of course other channels of contagion, such as common information, that can affect expectations and interact with the more mechanical channels described above.

These channels of contagion cause nonlinear interactions that can create positive feedback loops that amplify external shocks or even generate purely endogenous dynamics, such as booms and busts. Nonlinear feedback loops can also be amplified by behavioral and institutional constraints and by bounded rationality (often in the context of incomplete information and learning).

Behavioral and institutional constraints force agents to take actions that they would prefer to avoid in the absence of the constraint. Such behavioral constraints can be imposed by a regulator but they can also result from bilateral contracts between private institutions. In principle, regulatory constraints, such as capital or liquidity coverage ratios, are designed to increase financial stability. In many cases however, these constraints are designed to increase the resilience of an individual financial institution to idiosyncratic shocks rather than the resilience of the system as a whole. Take the example of a leverage constraint. If a financial institution has high leverage, a small shock may be enough to push it into insolvency. Hence, from a regulatory perspective, a cap on leverage seems like a good idea. However, as we will discuss below, a leverage constraint may have the adverse side effect that it forces distressed institutions to sell into falling asset markets, causing prices to fall further and amplifying a crisis. Of course, leverage constraints are needed, but the point is that their effects can go far beyond the failure of individual institutions, and the way in which they are enforced can make a big difference. Similar positive feedback can result from other behavioral constraints as well.

This brings up the distinction between *microprudential regulation*, which is designed to benefit individual institutions without considering the effect on the system as a whole, vs. *macroprudential regulation*, which is designed to take systemic effects into account. These can come into dramatic conflict. For example, we will discuss the base of Basel II, which provided perfectly sensible rules for risk management from a microprudential point of view, but which likely caused substantial systemic risk from a macroprudential point of view, and indeed may have been a major driver of the crisis of 2008. It is ironic that prudent behavior of an individual can cause such significant problems for society as a whole.

Rational agents with complete information might be able to navigate the risks inherent to the financial system. Indeed, optimal behavior might well mitigate the positive feedback resulting from interconnectedness and behavioral constraints. However, we believe that optimal behavior in the financial system is rare. Instead, agents are restricted by bounded rationality. Their limited understanding of the system in which they operate forces agents to rely on simple rules as well as biased methods to learn about the state of the system and form expectations about its future states (Farmer, 2002; Lo, 2005). Suboptimal decisions and biased expectations can exacerbate the destabilizing effects of interconnectedness and behavioral constraints but can also lead to financial instability on their own.

The remainder of this paper is organized as follows: In Section 2 we briefly contrast and compare traditional equilibrium models with agent-based models. In Section 3 we introduce the dynamical systems perspective on the financial system that will underlie many of the models of financial stability that we discuss in subsequent sections. In Sections 4 and 5 we discuss models of systemic risk resulting from leverage constraints and models of financial contagion due to interconnectedness, respectively. Sections 6 to 9 consider various different stress tests. In particular, Section 6 gives a brief conceptual overview of stress tests; Section 7 introduces and critically evaluates standard, micro-prudential stress tests; Section 8 discusses exam-

ples of macroprudential stress tests and how to bring them to data; finally Section 9 outlines a vision for the next generation of system-wide stress tests.

2 TWO APPROACHES TO MODELING SYSTEMIC RISK

As mentioned in the introduction, traditionally finance has focused on modeling systemic risk in highly stylized models that are analytically tractable. These efforts have improved our understanding of a wide range of phenomena related to systemic risk ranging from bank runs (Diamond and Dybvig, 1983; Morris and Shin, 2001), credit cycles (Kiyotaki and Moore, 1997; Brunnermeier and Sannikov, 2014), balance sheet (Allen and Gale, 2000), and information contagion (Acharya and Yorulmazer, 2008) over fire sales (Shleifer and Vishny, 1992), to the feedback between market and funding liquidity (Brunnermeier and Pedersen, 2009). A comprehensive review that does justice to this literature is beyond the scope of this paper. However, we would like to make a few observations in regard to the traditional modeling approach and contrast it with the agent-based approach.

Traditional models place great emphasis on the incentives and information structure of agents in a financial market. Given those, agents behave strategically, taking into account their beliefs about the state of the world, and other agents' strategies. The objects of interest are then the game theoretic equilibria of this interaction. This allows for studying the effects of properties such as asymmetric information, uncertainty or moral hazard on the stability of the financial system. While these models provide valuable qualitative insights, they are typically only tractable in very stylized settings. Models are usually restricted to a small number or a continuum of agents, a few time periods and a drastically simplified institutional and market set up. This can make it difficult to draw quantitative conclusions from such models.

Agent-based models typically place less emphasis on incentives and information, and instead focus on how the dynamic interactions of behaviorally simple agents can lead to complex aggregate phenomena, such as financial crises, and how outcomes are shaped by the structure of this interaction and the heterogeneity of agents. From this perspective, the key drivers of systemic risk are the amplification of dynamic instabilities and contagion processes in financial markets. Complicated strategic interactions and incentives are often ignored in favor of simple, empirically motivated behavioral rules and a more realistic institutional and market set up. Since these models can easily be simulated numerically, they can in principle be scaled to a large number of agents and, if appropriately calibrated, can yield quantitative insights.

Two common criticisms leveled against heterogeneous agent-based models are the lack of strategic interactions and the reliance on computer simulations. The first criticism is fair and, in many cases, highlights an important shortcoming of this approach. Hard wired behavioral rules need to be carefully calibrated against real data, and even when they are, they can fail in new situations where the behavior of agents may change. For computer simulations to be credible, their parameters need to be calibrated and the sensitivity of outcomes to those parameters needs to be understood.

The latter in particular is more challenging in computational models than in tractable analytical models.

In our view, what is not fair is to regard computer simulations as inherently inferior to analytic results. Analytic models have the benefit of the relative ease with which they can be used to understand the concepts driving structural cause-and-effect relationships. But many aspects of the economic world are not simple, and in most realistic situations computer simulations are the only possibility. Good practice is to make code freely available and well documented, so that results are easily reproducible.

Traditional and heterogeneous agent-based models are complements rather than substitutes. Some heterogeneous agent-based models already use myopic optimization, and in the future the line between the two may become increasingly blurred.[3] As methods such as computational game theory or multi-agent reinforcement learning mature, it may become possible to increasingly introduce strategic interactions into computational heterogeneous agent-based models. Furthermore, as computational resources and large volumes of data on the financial system become more accessible, parameter exploration and calibration should become increasingly feasible. Therefore, we are optimistic that, provided technology progresses as expected,[4] in the future heterogeneous agent-based models will be able to overcome some of the shortcomings discussed above. And as we demonstrate here, they have already led to important new results in this field, that were not obtainable via analytic methods.

3 A VIEW OF THE FINANCIAL SYSTEM

At a high level, it is useful to think of the financial system as a dynamical system that consists of a collection of institutions that interact via centralized and bilateral markets. An institution can be represented by its balance sheet, i.e. its assets and liabilities, together with a set of decision rules that it deploys to control the state of its balance sheet in order to achieve a certain goal. Within this framework, a market can be thought of as a mechanism that takes actions from institutions as inputs and changes the state of their balance sheets based on its internal dynamics. Anyone wishing to construct an agent model of the financial system therefore has to answer three fundamental questions: (i) what comprises the institutions' balance sheets, (ii) what determines their actions conditional on the state of the world, and (iii) how do markets respond to these actions? In the following, we will sketch the balance sheet of a generic leveraged investor, which will serve as the fundamental building block of the models of financial stability that we will discuss in this review. We will also

[3] In fact, this is already the case in the literature on financial and economic networks, see for example Goyal (2018).

[4] It seems unlikely that scientists' ability to analytically solve models will improve as quickly as numerical techniques and heterogeneous agent-based simulations, which benefit from rapid improvements in hardware and software.

briefly touch on (ii) and (iii) when discussing the important channels through which leveraged investors interact. In the subsequent sections we will then discuss concrete models of financial stability that fall within this general framework.

3.1 BALANCE SHEET COMPOSITION

When developing a model of a financial system, it is useful to distinguish between two types of agents which we refer to as "active" and "passive" agents. Active agents are the objects of interest and their internal state and interactions are carefully modeled. Passive agents represent parts of the financial system that interact with active agents but are not the focus of the model, and are therefore typically represented by simple stochastic processes. For the remainder of this review, consider a financial system that consists of a set $\mathcal{B}$ of active leveraged investors and a set of passive agents which will remain unspecified for now. We are particularly interested in systemic risk that is driven by borrowing, and thus we focus on agents that use leverage (defined as purchasing assets with borrowed funds). However, the setup below is sufficiently general to accommodate unleveraged investors as a special case with leverage equal to one.

Leveraged investors need not be homogeneous and may differ, among other aspects, in their balance sheet composition, strategies or counterparties. In practice, a leveraged investor might be a bank or a leveraged hedge fund and other active investors might include unleveraged mutual funds. Passive agents could be depositors, noise traders, fund investors that generate investment flows or banks that lend to hedge funds. The choice of active vs. passive investors of course varies from model to model.

The balance sheet of an investor $i \in \mathcal{B}$ is composed of assets A_i, liabilities L_i, and equity E_i, such that $A_i = L_i + E_i$. The investor's leverage is simply the ratio of assets to equity $\lambda_i = A_i/E_i$. It is useful to decompose the investor's assets into three classes: bilateral contracts A_i^B between investors, such as loans or derivative exposures; traded securities A_i^S, such as stocks; and external assets A_i^R, whose value is assumed exogenous. Throughout this review, we assume that the value A_i^S of traded securities is marked to market.[5] That is, the value of a traded security on the investor's balance sheet will be determined by its current market price. Of course we must have that $A_i = A_i^B + A_i^S + A_i^R$. Each asset class can be further decomposed into individual loan contracts, stock holdings and so on.

The investor's liabilities can be decomposed in a similar fashion. For now, let us decompose the investor's liabilities simply into bilateral contracts L_i^B between investors, such as loans or derivative exposures, and external liabilities L_i^R which can be assumed exogenous. In the case of a bank, these external liabilities might be deposits. Again we must have that $L_i = L_i^B + L_i^R$, and bilateral liabilities can be

[5]The term *marked to market* means that the value of assets is recomputed in every period based on current market prices. This is in contrast to valuing assets based on an estimate of their fundamental value.

further decomposed into individual bilateral contracts. Bilateral assets and liabilities might be secured, such as repurchase agreements, or unsecured such as interbank loans. Naturally, bilateral liabilities are just the flip side of bilateral assets such that summing over all investors we must have $\sum_i A_i^B = \sum_i L_i^B$.

3.2 BALANCE SHEET DYNAMICS

Of all the factors that affect the dynamics of the investors' balance sheets, three are of particular importance for financial stability: leverage, liquidity, and interconnectedness. Below, we discuss each factor in turn.

Leverage: Leverage amplifies returns, both positive and negative. Therefore, investors typically face a leverage constraint to limit the investors' risk.[6] However, at the level of the financial system, binding leverage constraints can lead to substantial instabilities. On short time scales, a leveraged investor may be forced to sell into falling markets when she exceeds her leverage constraint. Her sales will in turn depress prices further as we explain in the next paragraph on market liquidity. Leverage constraints can thus lead to an unstable feedback loop between falling prices and forced sales. On longer time scales dynamic leverage constraints that depend on backward looking risk estimates can lead to entirely endogenous volatility – so called leverage cycles.[7]

Liquidity: Broadly speaking, one can distinguish between two types of liquidity: market liquidity and funding liquidity.

Market liquidity can be understood as the inverse to price impact. When market liquidity is high, the market can absorb large sell orders without large changes in the price. If markets were perfectly liquid it would always be possible to sell assets without affecting prices and most forms of systemic risk would not exist.[8] Leverage is dangerous both because it directly increases risk, amplifying gains and losses proportionally, but also because the market impact of liquidating a portfolio to achieve a certain leverage increases with leverage. This point was stressed by Caccioli et al. (2012a), who showed how, due to her own market impact, an investor with a large leveraged position can easily drive herself bankrupt by liquidating her position. They showed that this can be a serious problem even under normal market conditions, and recommended taking market impact into account when valuing portfolios in order to reduce this problem. The problem can become even worse if investors are forced to

[6] This constraint may be imposed by a regulator, a counterparty or internal risk management.

[7] Beyond leverage, investors may also face other constraints. Regulators have imposed restrictions on the liquidity of assets that some investors may hold (with a preference for more liquid assets) and the stability of their funding (with a preference for more long term funding). The effect of these constraints on systemic risk is much less studied than the effect of leverage constraints. A priori however, one would expect these constraints to improve stability. This is because of the absence of feedback loops similar to the leverage-price feedback loop that drives forced sales.

[8] Prices would of course still change to reflect the arrival of new information.

sell too quickly, inducing *fire sales* in which a market is overloaded with sell orders, causing a dramatic decrease in liquidity for sellers.[9] Fire sales can be induced when investors hit leverage constraints, forcing them to sell, which in turn causes leverage constraints to be more strongly broken, inducing more selling.

Funding liquidity refers to the ease with which investors can borrow to fund their balance sheets. When funding liquidity is high, investors can easily roll over their existing liabilities by borrowing again, or even expand their balance sheets. In times of crises, funding liquidity can drop dramatically. If investors rely on short term liabilities they may be forced to liquidate a large part of their assets to pay back their liabilities. This forced sale can trigger fire sales by other investors.

Interconnectedness: Investors are connected via their balance sheets and so are not isolated agents. Connections can result from direct exposures due to bilateral loan contracts, or from indirect exposures due to investments into the same assets. Interconnectedness together with feedback loops resulting from binding leverage constraints and endogenous liquidity can lead to financial contagion. In analogy to epidemiology, financial contagion refers to the process by which "distress" may spread from one investor to another, where distress can be broadly understood as an investor becoming uncomfortably close to insolvency or illiquidity. Typically financial contagion arises when, via some mechanism or channel, a distressed investor's actions negatively affect some subset of other investors. This subset of investors is said to be connected to the distressed investor. A simple example of such connections are the bilateral liabilities between investors. Taken together, the set of all such connections form a network over which financial contagion can spread. For an in-depth review of financial networks see Iori and Mantegna (2018).

The aim of the subsequent sections is to introduce the reader to a number of models that tackle the effect of leverage, liquidity and interconnectedness on financial stability in isolation. These models then form the building blocks of more comprehensive models discussed in Section 6. Below, in Section 4, we first focus on the potentially destabilizing effects of leverage as they form the basis of fire sale models discussed later, and because they are thought to have contributed to the build up of risk prior to the great financial crisis. In Section 5 we then proceed to models of financial contagion as they form the scientific bedrock of the stress testing models that will be discussed in Section 6 and beyond. While liquidity is of great importance, we will only discuss it implicitly in Sections 4 and 5, rather than dedicating a separate section to it. This is because, unfortunately, there are currently only few dedicated models on this topic, see Bookstaber and Paddrik (2015) for an example. We will not be able to provide a complete overview of the agent-based modeling literature devoted to various aspects of financial stability. Important topics that we will not be able to discuss include the role of heterogeneous expectations or time scales in the

[9]There is always market impact from buying or selling. The term "fire sales" technically means selling under stress, but often means simply a case where the sale of assets is forced (even when markets remain orderly). See the discussion in Section 5.3.

dynamics of financial markets, see for example Hommes (2006), LeBaron (2006) for early surveys, and Dieci and He (2018) for a recent overview.

4 LEVERAGE AND ENDOGENOUS DYNAMICS IN A FINANCIAL SYSTEM

4.1 LEVERAGE AND BALANCE SHEET MECHANICS

Many financial institutions borrow and invest the borrowed funds into risky assets. Three simple properties of leverage are worth noting at this point. First, ceteris paribus, leverage determines the size of the investor's balance sheet. Second, leverage boosts asset returns. Third, leverage increases when the investor incurs losses, again ceteris paribus. In the following, we discuss each property in turn. For a fixed amount of equity, an investor can only increase the size of its balance sheet by increasing its leverage. Further, it is easy to show that, if r_t is the asset return, the equity return is $u_t = \lambda r_t$, where, as above, λ is the investor's leverage. In good times, leverage thus allows an investor to boost its return. In bad times however, even small negative asset returns can drive the investor into bankruptcy provided leverage is sufficiently high. Given the potential risks associated with high leverage, an investor typically faces a leverage limit which may be imposed by a regulator, as is the case for banks, or by creditors via a haircut[10] on collateralized debt. Finally, why does leverage increase when the investor incurs losses? Suppose the investor holds S units of a risky asset at price p such that $A = Sp$. Holding the investor's liabilities fixed, it is easy to see that $\lambda > 1$ implies $\partial\lambda/\partial p < 0$. In other words, whenever an investor is leveraged ($\lambda > 1$), a decrease (increase) in asset prices leads to an increase (decrease) in its leverage.

In what follows, we discuss how these three properties of leverage, in combination with reasonable assumptions about investor behavior, can lead to financial instability. We begin by discussing how leverage constraints can force investors to sell into falling markets even if they would prefer to buy in the absence of leverage constraints. We then show how a leverage constraint based on a backward looking estimator of market risk can lead to endogenous volatility and leverage cycles.

4.2 LEVERAGE CONSTRAINTS AND MARGIN CALLS

Consider again the simple investor discussed above. Suppose the investor faces a leverage constraint $\overline{\lambda}$ and has leverage $\lambda_{t-1} < \overline{\lambda}$.[11] The investor has to decide on an action at time $t-1$ to ensure that it does not violate its leverage constraint at

[10] A *haircut* is the difference between the face value of a loan and the market value of the assets used as loan collateral. Typically, the haircut is such that the dollar amount of collateral is higher than the face value of the loan. This ensures that the lender can recover his investment in case of default of the borrower even if the collateral has lost some of its value.

[11] As mentioned above, a leverage constraint can be the result of regulation or contractual obligations.

time t. Suppose the investor expects the price of the risky asset to drop sufficiently from one period to the next, such that its leverage is pushed beyond its limit, i.e. $\lambda_t > \overline{\lambda}$. In this situation the investor has two options to decrease its leverage: raise equity or reduce its assets (or some combination of the two). Raising equity can be time consuming or even impossible during a financial crisis. Therefore, if the leverage constraint has to be satisfied quickly or if new equity is not available and no assets are maturing in the next period, the investor has to sell at least $\Delta A_{t-1} = \max\{0, (\mathbb{E}_{t-1}[\lambda_t] - \overline{\lambda})\mathbb{E}_{t-1}[E_t]\}$ of its assets to satisfy its leverage constraint, where $\mathbb{E}_t[\cdot]$ is the conditional expectation at time t. In the following we will set $\mathbb{E}_t[\lambda_{t+1}] = \lambda_t$ and $\mathbb{E}_t[E_{t+1}] = E_t$. This can be done for simplicity or because a contract forces the investor to make adjustments based on current rather than expected values. In this case we have simply $\Delta A_t = \max\{0, (\lambda_t - \overline{\lambda})E_t\}$. If λ_t exceeds the leverage limit due to a drop in prices, the investor will sell into falling markets which may lead to a feedback loop between leverage and falling prices as outlined in the previous section.

This simple mechanism has been discussed by a number of authors, see for example Gennotte and Leland (1990), Geanakoplos (2010), Thurner et al. (2012), Shleifer and Vishny (1997), Gromb and Vayanos (2002), Fostel and Geanakoplos (2008). Gorton and Metrick (2012) study the effect of haircuts on repo markets during the financial crisis empirically. Thurner et al. (2012) incorporate this mechanism in a heterogeneous agent model of leverage-constrained value investors. In the remainder of this section we will introduce their model and discuss some of the quantitative results they obtain for the effect of leverage constraints on asset returns.

Consider our set $\mathcal{B}$ of leveraged investors introduced in Section 3. Suppose that investors have no bilateral assets or liabilities and only invest into a single traded security, i.e. $A_i = A_i^S$. Furthermore, assume that the investor has access to a credit line from an unmodeled bank such that $L_i = L_i^R$. For brevity and to guide intuition, we will refer to these leveraged investors as funds for the remainder of this section. In addition to the funds, there is a representative noise trader and a representative "fund investor" that allocates capital to the funds. There is an asset of supply N with fundamental value V that is traded by the funds and the noise trader at discrete points in time $t \in \mathbb{N}$. Every period a fund i takes a long position $A_{it} = \lambda_{it} E_{it}$ provided its equity satisfies $E_{it} \geq 0$. The fund's leverage is given by the heuristic

$$\lambda_{it} = \min\{\beta_i m_t, \overline{\lambda}\},$$

where $m_t = \max\{0, V - p_t\}$ is the mispricing signal and β_i is the fund's aggressiveness. In other words, the fund goes long in the asset if the asset is underpriced relative to its fundamental value V. The noise trader's long position follows a transformed AR(1) process with normally distributed innovations. The price of the asset is determined by market clearing. Every period, the fund investor adjusts its capital allocation to the funds, withdrawing capital from poorly performing funds and investing into successful funds relative to an exogenous benchmark return.

Before considering the dynamics of the full model, let us briefly discuss the limit where the funds are small, i.e. $E_{it} \to 0$. In this case, in the absence of any significant

effect of the funds, log price returns will be approximately iid normal due to the action of the noise trader. This serves as a benchmark. The authors then calibrate the parameters of the model such that funds are significant in size and prices may deviate substantially from fundamentals. This corresponds to a regime where arbitrage is limited as in Shleifer and Vishny (1997). The authors also assume that funds differ substantially in their aggressiveness β_i but share the same leverage constraint $\overline{\lambda}$ and initial equity E_{i0}.

In this setting the funds' leverage and wealth dynamics can lead to a number of interesting phenomena. When the noise trader's demand drives the price below the asset's fundamental value, funds will enter the market in proportion to their aggressiveness β_i. Due to the built-in tendency of the price to revert to its fundamental value due to the action of the noise traders, these trades will be profitable for the funds on average and even more profitable for more aggressive funds. Hence, the equity of aggressive funds grows quicker due to a combination of profits and capital reallocation of the fund investor. Importantly, as the equity of funds grows, their influence on prices increases and the volatility of the price decreases, due to the fact that they buy into falling markets and sell into rising markets.

Aggressive funds are also more likely to leverage to their maximum. Consider an aggressive fund i that has chosen $\lambda_{it-1} = \overline{\lambda}$. Now suppose the price drops such that $\lambda_{it} > \overline{\lambda}$. In response the fund sells parts of its assets as outlined above. Thurner et al. (2012) refer to this forced selling as a margin call, as they interpret the leverage constraint as arising from a haircut on a collateralized loan. Recall that the amount the fund will sell is $\Delta A_{it} = \max\{0, (\lambda_{it} - \overline{\lambda})E_{it}\}$, i.e. it is proportional to the fund's equity. As the aggressive fund is likely also the most wealthy fund, its selling can be expected to lead to a significant drop in prices. This drop may push other, less aggressive funds past their leverage limits. A margin spiral ensues in which more and more funds are forced to sell into falling markets. In an extreme outcome, most funds will exit or will have lost most of their equity in the price crash. As a result, their impact on prices is limited and the price is dominated by the noise trader. Thus following a margin spiral, price volatility increases due to two forces. First, it spikes due to the immediate impact of the price collapse. But then, it remains at an elevated level due to lack of value investors that push the price toward its fundamental value. These dynamics, which are illustrated in Fig. 1, reproduce some important features of financial time series in a reasonably quantitative way, in particular fat tails in the distribution of returns and clustered volatility (cf. Cont, 2001), as well as a realistic volatility dynamics profile before and after shocks (Poledna et al., 2014). These are difficult to reproduce in standard models.

One would expect these dynamics to be less drastic if funds took precautions against margin calls and stayed some $\epsilon > 0$ below their maximum leverage allowing them to more smoothly adjust to price shocks. However, it is important to note that a single "renegade" fund that pushes its leverage limit while all other funds remain well below it can be sufficient to cause a margin spiral.

It should be noted that the deleveraging schedule ΔA_{it} that a fund follows can depend on how the leverage constraint is implemented. In Thurner et al. (2012), the

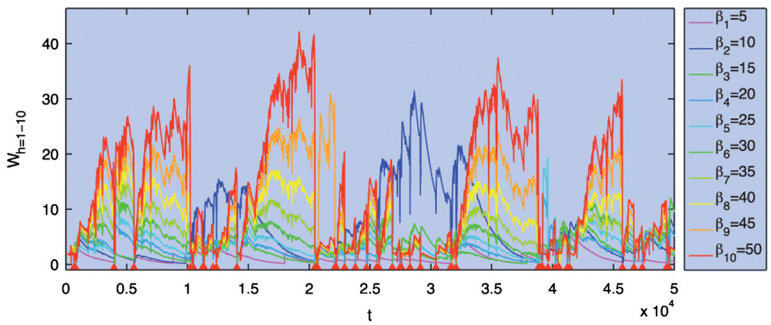

FIGURE 1

Time series of fund wealth dynamics from Thurner et al. (2012). Each time series corresponds to the wealth dynamics of a fund with different aggressiveness ranging from $\beta = 5$ to $\beta = 50$. Aggressive funds grow in size, become highly leveraged and susceptible to margin spirals and subsequent rapid collapse. The resulting asset return time series displays several realistic features including fat tails and clustered volatility.

leverage constraint results from a haircut applied to a collateralized loan, i.e. the fund obtains a short term loan from a bank, purchases the asset with the loan and its equity and then posts the asset as collateral for the loan. The haircut is equivalent to leverage and determines how much of its assets the fund can finance via borrowing. When the value of the asset drops, the bank will make a margin call as outlined above and the fund will have to sell assets immediately. However, a leverage constraint can, for example, also be imposed by a regulator. In this case, the fund may be allowed to violate the leverage constraint for a few time steps while smoothly adjusting to satisfy the constraint in later periods. Such an implementation will increase the stability of the system. Finally, the schedule $\Delta A_{it} = \max\{0, (\lambda_{it} - \overline{\lambda})E_{it}\}$ assumes the price remains unchanged from the current to the next period. A more sophisticated fund might take its own price impact into account when determining the deleveraging schedule.

4.3 PROCYCLICAL LEVERAGE AND LEVERAGE CYCLES

In the model presented in the previous section, funds actively increase their leverage when the price falls until they reach a leverage limit. Of course, a variety of other leverage management policies are possible. In an effort to study leverage management policies, Adrian and Shin (2010) analyze how changes in leverage $\Delta\lambda_t$ relate to changes in total assets ΔA_t (at mark-to-market prices) during the period 1963–2006 for three types of investors: households, commercial banks and security broker dealers (such as Goldman Sachs). Below we focus on households on one extreme and broker dealers on the other.

For households and broker-dealers the authors find a distinct correlation between leverage and asset changes, see Fig. 2. For households, changes in leverage are negatively correlated to changes in assets: $\mathrm{Corr}(\Delta\lambda_t, \Delta A_t) < 0$. For broker dealers they find a positive correlation $\mathrm{Corr}(\Delta\lambda_t, \Delta A_t) > 0$. This points toward at least two distinct leverage management policies.

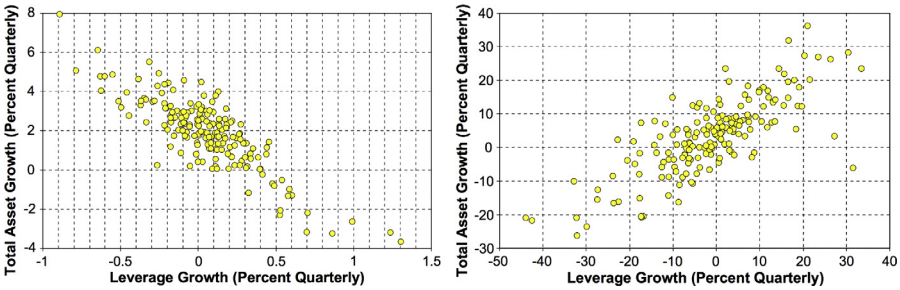

(A) Households (B) Broker dealers.

FIGURE 2

Change in total assets vs change in leverage from Adrian and Shin (2010).

Households appear to be passive investors since leverage decreases when assets appreciate, ceteris paribus. Broker-dealers however, appear to follow a state-contingent target leverage which they try to reach through balance sheet adjustments. To see this, suppose an investor has a leverage target which is high in good times and low in bad times. Let us say that good times are identified by increasing asset prices while bad times are identified by falling asset prices (there are other ways of identifying the state of the world as we will discuss below). In this case, in response to an increase (decrease) in the price of the asset, the investor will increase (decrease) its target leverage and adjust its balance sheet accordingly. Importantly, the leverage adjustment often occurs via debt and asset adjustment rather than equity adjustment. Adrian and Shin (2010) call this a *procyclical* leverage policy. With such a leverage policy we expect $\text{Corr}(\Delta\lambda_t, \Delta A_t) > 0$. Hence, it appears that broker-dealers follow a procyclical leverage policy.

A procyclical leverage policy could arise if the broker-dealers face a time varying leverage constraint and choose to leverage maximally. In fact, Adrian and Shin (2010), Danıelsson et al. (2004), and others show that a time varying leverage constraint arises when the investor faces a Value-at-Risk (VaR) constraint as was required under the Basel II regulatory framework. As we will show below, the effect of a VaR constraint is that the investor faces a leverage constraint that is inversely proportional to market risk. Thus, when market risk is high (low), the leverage constraint is low (high). In this setting the level of risk identifies the state of the world: in good times risk is low, while in bad times risk is high.

In summary, two leverage management policies are borne out by the data: passive leverage and procyclical target leverage. The type of leverage management policy used by the investor can have significant implications for financial stability. Indeed, at least anecdotally, the time series of broker-dealer leverage,[12] perceived risk (as

[12]Broker-dealer leverage is defined as the ratio of the series "Total Assets" (Fed time series identifier Z1/OTHER/FL664090663.Q) to "Equity capital" (Fed time series Z1/OTHER/FL665080003.Q), available at https://www.federalreserve.gov/datadownload/.

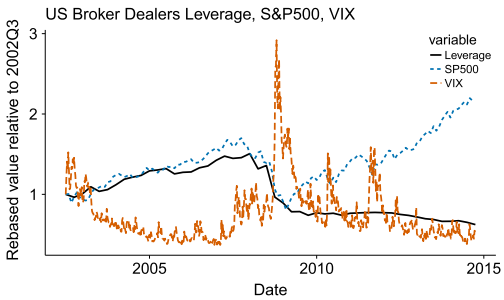

FIGURE 3

Time series of broker-dealer leverage, perceived risk (as measured by the VIX) and asset prices (as measured by the S&P500).

measured by the VIX) and asset prices (as measured by the S&P500) in Fig. 3 suggests a relationship between these three variables that is potentially induced by the dealers' procyclical leverage policy. In the following, we will introduce a model developed by Aymanns and Farmer (2015) that links leverage, perceived risk and asset prices in order to illustrate the effect of procyclical leverage and VaR constraints on the dynamics of asset prices.

Consider again our set $\mathcal{B}$ of leveraged investors (banks for short) and a representative noise trader. As above, we assume that there are no bilateral assets or liabilities. There is a risk free asset (cash) and a set $\mathcal{A}$ of risky assets that are traded by banks and the noise trader at discrete points in time $t \in \mathbb{N}$. At the beginning of every period, the banks and the noise trader determine their demand for the assets. For this, each bank i picks a vector $\mathbf{w}_{it}$ of portfolio weights and is assigned a target leverage $\overline{\lambda}_{it}$. The noise trader is not leveraged and therefore only picks a vector $\mathbf{v}_t$ of portfolio weights. Once the agent's demand functions have been fixed, the markets for the risky assets clear which fixes prices. Given the new prices, banks choose their next period's balance sheet adjustment (buying or selling of assets) in order to hit their target leverage. We refer the reader to Aymanns and Farmer (2015) for a detailed description of the model.

As mentioned above, banks are subject to a Value-at-Risk constraint.[13] Here, a bank's VaR is the loss in market value of its portfolio over one period that is exceeded with probability $1 - a$, where a is the associated confidence level. The VaR constraint then requires that bank holds equity to cover these losses, i.e. $E_{it} \geq \mathrm{VaR}_{it}(a)$. We approximate the Value-at-Risk by $\mathrm{VaR}_{it} = \sigma_{it} A_{it}/\alpha$, where σ_{it} is the estimated portfolio variance of bank i and α is a parameter. This relation becomes exact for normal asset returns and an appropriately chosen α. Rearranging the VaR constraint yields the bank's leverage constraint $\overline{\lambda}_{it} = \alpha/\sigma_{it}$. We assume that the bank chooses to be maximally leveraged, e.g. for profit motives. The leverage

[13] Within the context of this model the Value-at-Risk constraint should be understood as a placeholder for a procyclical leverage policy. We choose Value-at-Risk here for modeling convenience.

constraint is therefore equivalent to the target leverage we discussed above. To evaluate their VaR, banks compute their portfolio variance as an exponentially weighted moving average of past log returns.

Let us briefly discuss the implications of this set up. As mentioned at the outset of this section, banks follow a procyclical leverage policy. In particular, the banks' VaR constraint, together with its choice to be maximally leveraged at all times, imply a target leverage that is inversely proportional to the banks' *perceived risk* as measured by an exponentially weighted moving average of past squared returns. Why is such a leverage policy procyclical? Suppose a random drop in an asset's price causes an increase in the level of perceived risk of bank i. As a result the bank's target leverage will decrease (while its actual leverage simultaneously increases) and it will have to sell some of its assets, similar to the funds in the previous section.[14] The banks selling may lead to a further drop in prices and a further increase in perceived risk. In other words, the bank's leverage policy together with its perception of risk can lead to an unstable feedback loop. It is in this sense that the leverage policy is procyclical.

Banks in this model have a very simple, yet realistic, method of computing perceived (or expected) risk. Similar backward looking methods are well established in practice, see for example Andersen et al. (2006). It is important to note that perceived risk σ_{it} and realized volatility over the next time step can be very different. Since banks have only bounded rationality and follow a simple backward looking rule in this model, their expectations about volatility are not necessarily correct on average and tend to lag behind realizations.

Let us now consider the dynamics of the model in more detail. In Fig. 4 we show two simulation paths (with the same random seed) of the price of a single risky asset for two leverage policy rules. In the top panel, banks behave like the households in Adrian and Shin (2010) – they are passive and do not adjust their leverage to changes in asset prices or perceived risk. In the bottom panel, banks follow the procyclical leverage policy outlined above. The difference between the two price paths is striking. In the case of passive banks, the price follows what appears to be a simple mean reverting random walk. However, when banks follow the procyclical leverage policy, the price trajectory shows stochastic, irregular cycles with a period of roughly 100 time steps. These complex, endogenous dynamics are the result of the unstable feedback loop outlined above.

Aymanns and Farmer (2015) refer to these cycles as *leverage cycles*. Leverage cycles are an example of *endogenous* volatility – volatility that arises not because of the arrival of exogenous information but due to the endogenous dynamics of the agents in the financial system. To better understand these dynamics, consider the state of the system just after a crash has occurred, e.g. at time $t \approx 80$. Following the crash, banks' perceived risk is high, their leverage is low and prices are stable. Over time, perceived risk declines and banks increase their leverage. As they increase their leverage, they buy more of the risky assets and push up their prices. At some point

[14] Note that this selling will be spread across all assets according to the bank's portfolio weight matrix.

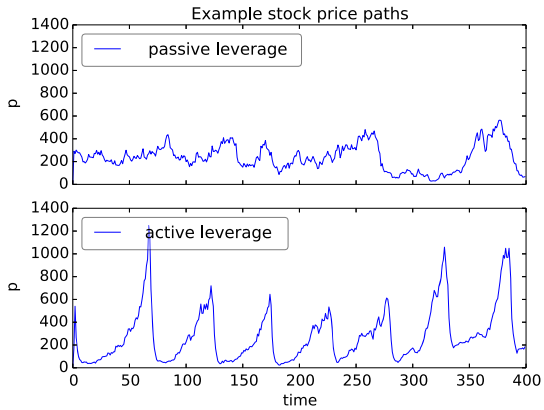

FIGURE 4

Example paths for the price of the risky asset from Aymanns and Farmer (2015). In the top panel, banks behave like households in Adrian and Shin (2010), i.e. they do not adjust their balance sheets following a change in leverage. Prices roughly follow a random walk – volatility is driven by the exogenous noise fed into the system. In the bottom panel, banks actively manage their leverage attempting to achieve a risk dependent target leverage, similar to broker-dealers in Adrian and Shin (2010). Prices now show endogenous, stochastic, irregular cycles.

leverage is sufficiently high and perceived risk sufficiently low, so that a relatively small drop of the price of an asset leads to large downward correction in leverage. A crash follows and prices fall until the noise trader's action stops the crash and the cycle begins anew. Naturally, these dynamics depend on the choice of parameters. In particular, when the banks are small relative to the noise trader, banks' trading has no significant impact on asset prices and leverage cycles do not occur. For a detailed discussion of the sensitivity of the results to parameters see Aymanns and Farmer (2015), and for a more realistic model that is better calibrated to the data, see Aymanns et al. (2016).

These results show that simple behavioral rules, grounded in empirical evidence of bank behavior (Adrian and Shin, 2010; Andersen et al., 2006), can lead to remarkable and unexpected dynamics which bear some resemblance to the run up to and crash following the 2008 financial crises. The results originate from the agents' bounded rationality and their reliance on past returns to estimate their Value-at-Risk. These features would be absent in a traditional economic models in which agents are fully rational. Indeed, rational models rarely display the dynamic instabilities that Aymanns and Farmer (2015) observe. If we believe that real economic actors are rarely fully rational, we should take note of this result. Of course, the agents in this model are really quite dumb. For example, they do not adjust to the strong cyclical pattern in the time series. However, they also live in an economy that is significantly simpler than the real world. Thus, their level of rationality in relation to the complexity of the world they inhabit might not be too far off from real economic agents' level of rationality.

The model discussed above can also yield insights for policy makers on how bank risk management might be modified in order to mitigate the effects of the leverage cycle. Aymanns et al. (2016) present a reduced form version of the model outlined above in order to investigate the implications of alternative leverage policies on financial stability. They show that, depending on the size of the banking sector and the properties of the exogenous volatility process, either a constant leverage policy or a Value-at-Risk based leverage policy is optimal from the perspective of a social planner. This finding lends support to the use of macroprudential leverage ratios as discussed in ESRB (2015). The authors also show that the timescale for the bubbles and crashes observed in the model is around 10–15 years, roughly corresponding to the run-up to the 2008 crash. Another important insight from Aymanns et al. (2016) is that the time scale on which investors need to achieve their leverage constraint plays a crucial role in the stability of the financial system: slower adjustment toward the constraint (corresponding to more slackness) increases stability.

The effect of leverage targeting on asset price dynamics has also been studied by other others in the multi-asset case. For example, Capponi and Larsson (2015) show that the deleveraging of banks may amplify asset return shocks and lead to large fluctuations in realized returns which in turn can cause spillover effects between different assets.

5 CONTAGION IN FINANCIAL NETWORKS

5.1 FINANCIAL LINKAGES AND CHANNELS OF CONTAGION

A channel of contagion is a mechanism by which distress can spread from one financial institution to another.[15] Often the channel of contagion is such that distress can only spread from one institution to a subset of all institutions in the system. These susceptible institutions are said to be linked to the stressed institution. The set of all links then forms a financial network associated with the channel of contagion.[16] Depending on the channel, links in this network may arise directly from bilateral contracts between banks, such as loans, or indirectly via the markets in which the banks operate. In the literature, one typically distinguishes between three key channels of contagion: counterparty loss, overlapping portfolios, and funding liquidity contagion.[17] Counterparty loss and overlapping portfolio contagion affect the value of the assets on the investors' balance sheet while funding liquidity contagion affects the availability of funding for the investors' balance sheets. In the following we will first introduce the investor's balance sheet relevant for this section. We will then

[15] For the first micro-level evidence of the transmission of shocks through the financial network, see Morrison et al. (2016).

[16] See Iori and Mantegna (2018) for a review of financial networks. For a review of models of financial contagion see Young and Glasserman (2016).

[17] Information contagion (cf. Acharya and Yorulmazer, 2008) is another channel of contagion but won't be discussed in this section.

give a brief overview of the channels of contagion before discussing each in more detail.

Balance sheet: Throughout this chapter we will consider a set $\mathcal{B}$ of leveraged investors (banks for short) whose assets can be decomposed into three classes: bilateral interbank contracts A_i^B, traded securities A_i^S that are marked to market and external, unmodeled assets A_i^R. Furthermore bank liabilities can be decomposed into bilateral interbank contracts L_i^B, and external, unmodeled liabilities L_i^R such that $L_i = L_i^B + L_i^R$. Note that bilateral interbank contracts need not be loans, they can also be derivative contracts for example. For simplicity however, we will think of bilateral interbank contracts as loans for the remainder of this section.

Counterparty loss: Suppose bank i has lent an amount C to bank j such that $A_i^B = L_j^B = C$. Now suppose the value of bank j's external assets A_j^R drops due to an exogenous shock. As a result the probability of default of bank j is likely to increase, which will affect the value of the claim A_i^B that bank i holds on bank j. If bank i's interbank assets are marked to market, a change in bank j's probability of default will affect the market value of A_i^B. In the worst case, if bank j defaults, bank i will only recover some fraction $r \leq 1$ of its initial claim A_i^B. If the loss of bank i exceeds its equity, i.e. $(1-r)A_i^B > E_i$, bank i will default as well.[18] Now, how can this lead to financial contagion? To elaborate on the above stylized example, suppose that bank i in turn borrowed an amount C from another bank k such that $A_k^B = L_i^B = C$.[19] In this scenario, it can be plausibly argued that an increase in the probability of default of j increases the probability of default of i which in turn increases the probability of default of k. If all banks mark their books to market, an initial shock to j can therefore end up affecting the value of the claim that bank k holds on bank i. Again, in the extreme scenario, the default of bank j may cause bank i to default which may cause bank k to default. This is the essence of counterparty loss contagion. Naturally, in a real financial system the structure of interbank liabilities will be much more complex than in the stylized example outlined above. However, the conceptual insights carry over: the financial network associated with the counterparty loss contagion channel is the network induced by the set of interbank liabilities.

Overlapping portfolios: The overlapping portfolio channel is slightly more subtle. Suppose bank i and bank j have both invested an amount C in the same security l such that $A_{il}^S = A_{jl}^S = C$, where we have introduced the additional index to reference the security.[20] Now, suppose the value of bank j's external assets A_j^R drops due to some exogenous shock. How will bank j respond to this loss? In the extreme

[18] In reality, this scenario is excluded due to regulatory large exposure limits which require that $A_i^B < E_i$.
[19] We assume that the contract between i and j as well as i and k has the same notional purely for expositional simplicity and all conceptual insights carry over for heterogeneous notionals.
[20] Again, we assume that both banks invest the same amount purely of expositional simplicity.

case, when the exogenous shock causes bank j's bankruptcy ($E_i < 0$), the bank will liquidate its entire investment in the security in a fire sale. However, even if the bank does not go bankrupt, it may wish to liquidate some of its investment. This can occur for example when the bank faces a leverage constraint as discussed in Section 4. Bank j's selling is likely to have price impact. As a result, the market value of A^S_{il} will fall. If bank i also faces a leverage constraint, or even goes bankrupt following the fall in prices, it will liquidate part of its securities portfolio in response. How will this lead to contagion? Suppose that bank i also has invested an amount C into another security m and that another bank k has also invested into the same security, such that $A^S_{im} = A^S_{km} = C$. If bank i liquidates across its entire portfolio, it will sell some of security m following a fall in the price of security l. The resulting price impact will then affect the balance sheet of bank k which was not connected to bank j via an interbank contract or a shared security. This is the essence of overlapping portfolio contagion. Banks are linked by the securities that they co-own and the fact that they liquidate with market impact across their entire portfolios. Empirical evidence from the 2007 Quant meltdown for this contagion channel has been provided in Khandani and Lo (2011).

Funding liquidity contagion often occurs when a lender is stressed, and so often occurs in conjunction with overlapping portfolio contagion and counterparty loss contagion. To see this, let us reconsider the scenario we discussed for counterparty loss contagion. Suppose bank i has lent an amount C to bank j such that $A^B_i = L^B_j = C$. As before, suppose the value of bank j's external assets A^R_j drops due to some exogenous shock and as a result, the probability of default of bank j increases. Now, suppose that every T periods bank i can decide whether to roll over its loan to bank j. Further assume that bank i is bank j's only source of interbank funding and L^R_j is fixed. Given bank j's increased default probability, bank i may choose not to roll over the loan at the next opportunity. Ignoring interest payments, if bank i does not roll over the loan, bank j will have to deliver an amount C to bank i. In the simplest case, bank j may choose not to roll over its own loans to other banks which in turn may decide against rolling over their loans. This is the essence of funding liquidity contagion. As for counterparty loss contagion, the associated financial network is induced by the set of interbank loans. Empirical evidence on the fragility of funding markets during the past financial crisis has been provided for example by Afonso et al. (2011) and Iyer and Peydro (2011). In a further complication, bank j may also choose to liquidate part of its securities portfolio in order to pay back its loan. Funding liquidity contagion can therefore lead to fire sales and overlapping portfolio contagion and vice versa. This interdependence of contagion channels makes the funding liquidity and overlapping portfolio contagion processes the most challenging from a modeling perspective.

In the remainder of this section, we will discuss models for counterparty loss, overlapping portfolio and funding liquidity contagion, as well as models for the interaction of all three contagion channels.

5.2 COUNTERPARTY LOSS CONTAGION

Let P denote the matrix of nominal interbank liabilities such that banks hold interbank assets $A_i^B = \sum_j P_{ij}^T$, where T denotes the matrix transpose. In addition, banks hold external assets A_i^R which can be liquidated at no cost. Banks have interbank liabilities $L_i^B = \sum_j P_{ij}$ only. Assume all interbank liabilities mature at the same time and have the same seniority. We further assume that all banks are solvent initially. There is only one time period. At the end of that period all liabilities mature, external assets are liquidated and banks pay back their loans if possible. Now suppose banks are subject to a shock $s_i \geq 0$ to the value of their external assets such that $\hat{A}_i^R = A_i^R - s_i$. Given an exogenous shock, we can ask a number of questions. First, which loan payments are feasible given the exogenous shock? Second, which banks will default on their liabilities? And finally, how do the answers to the first two questions depend on the structure of the interbank liabilities P? There is a large literature that studies counterparty loss contagion in a set up similar to the above, including Eisenberg and Noe (2001), Gai and Kapadia (2010), May and Arinaminpathy (2010), Elliott et al. (2014), Acemoglu et al. (2015), Battiston et al. (2012), Amini et al. (2016), and Capponi et al. (2015). In the following, we will briefly introduce the seminal contribution by Eisenberg and Noe (2001), who provide a solution to the first two questions. We will then consider a number of extensions of Eisenberg and Noe (2001) and alternative approaches to addressing the above questions.

Define the relative nominal interbank liabilities matrix as $\Pi_{ij} = P_{ij}/L_i^B$ for $L_i^B > 0$ and $\Pi_{ij} = 0$ otherwise. The relative liabilities matrix corresponds to the adjacency matrix of the weighted, directed network $\mathcal{G}$ of interbank liabilities. Let $\mathbf{p} = (p_1, \ldots, p_N)$ denote the vector of total payments made by the banks when their liabilities mature, where $N = |\mathcal{B}|$. Naturally, a bank pays at most what it owes in total, i.e. $p_i \leq L_i^B$. However, it may default and pay less if the value of its external assets plus the payments it receives from its debtors is less than what it owes. The individual payments that bank i makes are given by $\Pi_{ij} p_i$ since by assumption all liabilities have equal seniority. The vector of payments, also known as the clearing vector, that satisfies these constraints is the solution to the following fixed point equation

$$p_i = \min\{L_i^B, \hat{A}_i^R + \sum_j \Pi_{ij}^T p_j\}. \tag{1}$$

Eisenberg and Noe (2001) show that such a fixed point always exists. In addition, if within each strongly connected component of $\mathcal{G}$ there exists at least one bank with $\hat{A}_i^R > 0$, Eisenberg and Noe (2001) show that the fixed point is unique.[21] In other words, there exists a unique way in which losses incurred due to the adverse shock $\{s_i\}$ are distributed in the financial system via the interbank liabilities matrix. The

[21] In a strongly connected component of a directed graph there exists a directed path from each node in the component to each other node in the component. The strongly connected component is the maximal set of nodes for which this condition holds.

clearing vector and the set of defaulting banks can be found easily numerically by iterating the fixed point map in Eq. (1). As the map is iterated, more and more banks may default, resulting in a default cascade propagating through the financial network.

It is important to note that in this setup losses are only redistributed and the system is conservative – contagion acts as a distribution mechanism but does not, in the aggregate, lead to any further losses to bank shareholders beyond the initial shock. To see this, define the equity of bank i prior to the exogenous shock as $E_i = A_i^B + A_i^R - L_i^B$ and after the exogenous shock as $\hat{E}_i = \hat{A}_i^B(\mathbf{p}) + A_i^R - s_i - \hat{L}_i^B(\mathbf{p})$. Note that post-shock both bank i's assets and liabilities depend on the clearing vector $\mathbf{p}$. Taking the difference and summing over all banks we obtain $\sum_i E_i - \hat{E}_i = \sum_i A_i^R - (A_i^R - s_i) = \sum_i s_i$ since $\sum_i A_i^B = \sum_i L_i^B$ and $\sum_i \hat{A}_i^B(\mathbf{p}) = \sum_i \hat{L}_i^B(\mathbf{p})$. Also note that, while bank shareholder losses are not amplified, losses to the total value of bank assets are amplified due to indirect losses, i.e. losses not stemming from the initial exogenous shock but due to revaluation of interbank loans. This can be seen by taking the difference between pre- and post-shock total assets in the system. The total pre-shock assets of bank i are $A_i = A_i^B + A_i^R$ and its total post-shock assets are $\hat{A}_i = \hat{A}_i^B(\mathbf{p}) + A_i^R - s_i$, then $\sum_i A_i - \hat{A}_i = \sum_i A_i^B - \hat{A}_i^B(\mathbf{p}) + s_i \geq \sum_i s_i$. Some authors argue that this total asset loss can be useful measure of systemic impact of the exogenous shock, see Glasserman and Young (2015). Finally, note that the mechanism of finding a clearing vector ignores any potential frictions in the financial system and ensures that the maximal payment is made given the exogenous shocks. Several authors have argued that this is too optimistic and assume instead that once a default has occurred, some additional bankruptcy costs are incurred, see for example Rogers and Veraart (2013) and Cont et al. (2010).[22,23] In this case, aggregate bank shareholder losses may be larger than the aggregate exogenous shock. Further shortcomings of the Eisenberg and Noe model include the lack of heterogeneous seniorities or maturities and the lack of the possibility of strategic default.

The extent of the default cascade triggered by an exogenous shock depends on the structure of the financial network induced by the matrix of interbank liabilities P. One key property of the financial network is the average degree of a bank, i.e. the number of other banks it lends to. A well-known result is that, as banks' interbank lending A_i^B becomes more diversified over $\mathcal{B}$, i.e. the average degree increases, the expected number of defaulting banks first increases and then decreases, see Fig. 5. If banks lend only to a very small number of other banks, the network is not fully connected. Instead, it consists of several small and disjoint components. A default in a particular component cannot spread to other components, hence limiting the size of the default cascade. As banks become more diversified, the network will become fully connected and default cascades can spread across the entire network. As banks

[22] Such bankruptcy cost might for example capture the cost of forced liquidation of the banks' external assets.

[23] Papers that do not assume bankruptcy costs are essentially treating the system as it were conservative in equity: losses to one party are gains to the other, but there is no deadweight loss that ravages welfare. Hence, they fail to capture the negative externalities imposed by the banking system on society.

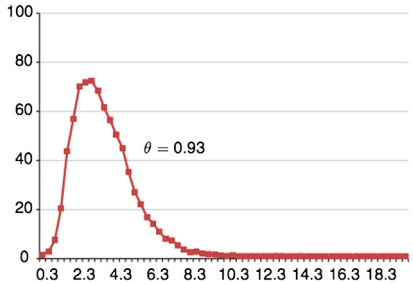

FIGURE 5

Expected number of defaults as a function of diversification in Elliott et al. (2014).

diversify further, the size of the individual loans between banks declines to the point that the default of any one counterparty becomes negligible for a given bank. Thus default cascades become unlikely. However, if they do occur, they will be very large. This is often referred to as the "robust-yet fragile" property of financial networks and has been observed for specifications of the financial network and the default cascade mechanism, see for example Elliott et al. (2014), Gai and Kapadia (2010), Battiston et al. (2012), or Amini et al. (2016). However, not only the average of the network's degree distribution is important for the system's stability. Caccioli et al. (2012b) show that if the degree distribution is very heterogeneous, i.e. there are a few banks that lend to many banks while most only lend to a few, the system is more resilient to contagion triggered by the failure of a random bank, but more fragile with respect to contagion triggered by the failure of highly connected nodes. In addition, Capponi et al. (2015) show that the level of concentration of the liability matrix, as defined by a majorization order, can qualitatively change the system's loss profile.

The models and solution methods discussed above tend to be simple to remain tractable and usually reduce to finding a fixed point.[24] However, these equilibrium models often form useful starting points for heterogeneous agent models that try to incorporate additional dynamic effects and more realism into the counterparty loss contagion process. See for example Georg (2013) where the effect of a central bank on the extent of default cascades is studied.

Finally, note that it is widely believed that large default cascades are quite unlikely for reasonable assumptions about the distribution of the exogenous shock and nominal interbank liabilities matrix, see for example Glasserman and Young (2015). For larger cascades to occur, default costs or additional contagion channels are necessary. Nevertheless, the existence of a counterparty loss contagion channel is important in practice as it affects the decisions of agents, for example in the way they form lending

[24]Gai and Kapadia (2010) for example make similarly restrictive assumptions on the structure of bank balance sheets as Eisenberg and Noe (2001). In addition several technical assumptions on the structure of the matrix of liabilities are necessary to solve for the fixed point of non-defaulted banks via a branching process approximation.

relationships. In other words, while default cascades are unlikely to occur in reality, they form an "off-equilibrium" path that shapes reality, see Elliott et al. (2014).

5.3 OVERLAPPING PORTFOLIO CONTAGION

In the following we will formally discuss the mechanics of overlapping portfolio contagion. To this end, consider again our set of banks $\mathcal{B}$. There is an illiquid asset whose value is exogenous and a set of securities $\mathcal{S}$, with $M = |\mathcal{S}|$, traded by banks at discrete points in time $t \in \mathbb{N}$. Let $\mathbf{p}_t = (p_{1t}, \ldots, p_{Mt})$ denote the vector of prices of the securities and let the matrix $\mathbf{S}_t \in \mathbb{R}^{N \times M}$ denote the securities ownership of all banks at time t. Thus S_{ijt} is the position that bank i holds in security j at time t. The assets of bank i are then given by $A_{it} = \mathbf{S}_{it} \cdot \mathbf{p}_t + A_i^R$, where A_i^R is the bank's illiquid asset holding. Let E_{it} and $\lambda_{it} = A_{it}/E_{it}$ denote bank i's equity and leverage, respectively. There are no interbank assets or liabilities.

As mentioned above, overlapping portfolio contagion occurs when one bank is forced to sell and the resulting price impact forces other banks with similar asset holdings to sell. What might force banks to sell? In an extreme scenario, a bank might have to liquidate its portfolio if it becomes insolvent, i.e. $E_{it} < 0$. But even before becoming insolvent, a bank might be forced to liquidate part of its portfolio if it violates a leverage constraint $\overline{\lambda}$ as we have shown in Section 4.[25] Both of these were considered by Caccioli et al. (2014) and by Cont and Schaanning (2017). In fact Caccioli et al. (2014) showed that such pre-emptive liquidations only make the problem worse due to increasing the pressure on assets that are already stressed. (This is closely related to the problem that liquidation can in and of itself cause default as studied by Caccioli et al. (2012a).) Other papers that discuss the effects of overlapping portfolios include Duarte and Eisenbach (2015), Greenwood et al. (2015), Cont and Wagalath (2016, 2013). An important early contribution to this topic is Cifuentes et al. (2005).

Let us first discuss the simpler case where liquidation occurs only upon default. Suppose bank i is subject to an exogenous shock $s_i > 0$ that reduces the value of its illiquid assets to $\hat{A}_{it}^R = A_{it}^R - s_i$. If $s_i > E_{it}$, the bank becomes insolvent and liquidates its entire portfolio. Let $Q_{jt} = \sum_{i \in \mathcal{I}_t} S_{ijt}$ denote the total amount of security j that is liquidated by banks in the set $\mathcal{I}_t$ of banks that became insolvent at time t. The sale of the securities is assumed to have market impact such that $p_{jt+1} = p_{jt}(1 + f_j(Q_{jt}))$, where $f_j(\cdot)$ is the market impact function of security j. Caccioli et al. (2014) assume an exponential form $f_j(x) = \exp(-\alpha_j x) - 1$, where x is volume liquidated and $\alpha_j > 0$ is chosen to be inversely proportional to the total shares outstanding of security j. In the next period, banks reevaluate their equity at the new securities prices. The change in equity is equal to $\Delta E_{it+1} = \sum_j S_{ijt} p_{jt} f_j(Q_{jt}) - s_i$. Note that

[25] Other "constraints" might also lead to forced sales and overlapping portfolio contagion. For example, investor redemptions that depend on past performance, as in Thurner and Poledna (2013), can force a fund to liquidate, which may result in an overlapping portfolio contagion similar to the one induced by leverage constraints.

in this setting we hold S_{ijt} fixed unless a bank liquidates its entire portfolio. Thus, banks who share securities with the banks that were liquidating in the previous period will suffer losses due to market impact. These losses may be sufficiently large for additional banks to become insolvent. If this occurs, contagion will spread and more banks will liquidate their portfolios, leading to further losses. Over the course of this default cascade, banks may suffer losses that did not share any common securities with the initially insolvent banks.

The evolution of the default cascade can be easily computed numerically by following the procedure outlined above until no further banks default. Caccioli et al. (2014) also show that the default cascade can be approximated by a branching process, provided suitable assumptions are made about the network structure. For their computations, Caccioli et al. (2014) assume that a given bank i invests into each security with a fixed probability μ_B/M, where μ_B is the expected number of securities that a bank holds. The bank distributes a fixed investment over all securities it holds. When μ_B/M is high, the portfolios of banks will be highly overlapping, i.e. banks will share many securities in their portfolios. Similar to the results for counterparty loss contagion, the authors find that as banks become more diversified, that is μ_B increases while M is held fixed, the probability of default (blue circles) first increases and then decreases, see Fig. 6. The intuition for this result is again similar to the counterparty loss contagion case. If banks are not diversified, their portfolios are not overlapping and price impact from portfolio liquidation of one bank affects only a few banks. As banks become more diversified, their portfolios become more overlapping and price impact spreads throughout the set of banks leading to large default cascades. Eventually, when banks become sufficiently diversified, the losses resulting from a price change in an individual security become negligible and large default cascades become unlikely. However, when they do occur, they encompass the entire set of banks. Thus, here again the financial network displays the robust-yet fragile property. Interestingly, the authors also show that for a fixed level of diversification, there exists a critical bank leverage λ_{it} at which default cascades emerge. The intuition for this result is that, when leverage is low, banks are stable and large shocks are required for default to occur, as leverage grows banks become more susceptible to shocks and defaults occur more easily.

As mentioned above, banks are likely to liquidate a part of their portfolio even before bankruptcy, if an exogenous shock pushes them above their leverage constraint. This is the setting studied in Cont and Schaanning (2017) and Caccioli et al. (2014). In this case, the shocks for which banks start to liquidate as well as the amount liquidated are both smaller than in the setting discussed above. If banks breach their leverage constraint due to an exogenous shock s_i to the value of their illiquid assets, Cont and Schaanning (2017) require that banks liquidate a fraction Γ_i of their entire portfolio such that $((1-\Gamma_i)\mathbf{S}_{it} \cdot \mathbf{p}_t + \hat{A}^R_{it})/E_{it} = \overline{\lambda}$. The corresponding liquidated monetary amount for a security j is then $Q_{jt} = \sum_{i \in \mathcal{B}} \Gamma_i S_{ijt} p_{jt}$. Again, the sale of the securities is assumed to have market impact such that $p_{jt+1} = p_{jt}(1 + f_j(Q_{jt}))$. In contrast to Caccioli et al. (2014), the authors assume that the market impact function $f_j(x)$ is linear in x, where x is the total monetary amount sold rather than the

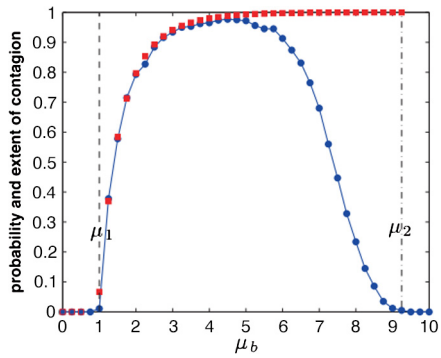

FIGURE 6

Blue circles: probability of contagion. Red squares: conditional on a contagion, the fraction of banks that fail (i.e. the extent of contagion). Taken from Caccioli et al. (2014).

number of shares. Similar market impact functions are used by Greenwood et al. (2015) and Duarte and Eisenbach (2015).

The shape and parameterization of the market impact function is crucial for the practical usage of models of overlapping portfolio contagion. There is a large body of market microstructure literature addressing this question. This literature begins with Kyle (1985), who derived a linear impact function under strong assumptions. More recent theoretical and empirical work indicates that under normal circumstances market impact is better approximated by a square root function (Bouchaud et al., 2008).[26]

The use of the term "fire sales" in this literature is confusing. Overlapping portfolio contagion can occur even if markets are functioning normally. A good example was the quant meltdown in August 2007. There are also likely to be circumstances in which normal market functioning breaks down due to overload of sellers, generating genuine fire sales; in this circumstance one expects market impact to behave anomalously and the square root approximation to be violated. There is little empirical evidence for this – genuine fire sales most likely occur only for extreme situations such as the crash of 1987 or the flash crash. In common usage the term "fire sale" often refers to any situation where selling of an asset depresses prices, even when the market is orderly.

Cont and Schaanning (2017) calibrate their model to realistic portfolio holdings and market impact parameters and obtain quantitative estimates of the extent of losses due to overlapping portfolio contagion. This provides a good starting point for more sophisticated financial system stress tests that will be discussed in the following sections. The models outlined above can be improved in many ways. Cont and Wagalath

[26]The market impact function takes the form $k\sigma\sqrt{\Delta V/V}$, where σ is volatility, ΔV is the size of the trade, V is market trading volume, and k is a constant of order one, whose value depends on the market.

(2016) study the effect of overlapping portfolios and fire sales on the correlations of securities in a continuous time setting, where securities prices follow a stochastic process rather than being assumed fixed up to the price impact from fire sales.

5.4 FUNDING LIQUIDITY CONTAGION

Funding liquidity contagion has been much less studied than overlapping portfolio or counterparty loss contagion. In the following we will briefly outline some of the considerations that should enter a model of funding liquidity contagion.

In modeling funding liquidity contagion, it is useful to partition an investor's funding into long term funding as well as short term secured and unsecured funding. Only short term funding should be susceptible to funding liquidity contagion as long term funding cannot be withdrawn on the relevant time scales. The availability of secured and unsecured short term funding may be restricted via two channels: a deleveraging channel and a default anticipation channel. The deleveraging channel applies equally to secured and unsecured funding: when a lender needs to deleverage, she can refuse to roll over short term loans, which may in turn force the borrower to deleverage, resulting in a cascade. This channel can be modeled using the same tools applied to overlapping portfolio and counterparty loss contagion. A paper that studies this channel is Gai et al. (2011). The default anticipation channel is different for secured and unsecured funding. In the case of secured funding, a lender might withdraw funding if the quality of the collateral decreases so that the original loan amount is no longer adequately collateralized. In the case of unsecured funding, a lender that questions the credit quality of one its borrowers might anticipate the withdrawal of funding of other lenders to that borrower and therefore withdraw her funding. This mechanism is similar to a bank run and therefore should be modeled as a coordination game, see Diamond and Dybvig (1983) and Morris and Shin (2001). This poses a challenge for heterogeneous agent models and might explain the relative scarcity of the literature on this topic. One notable exception that tries to combine both mechanisms is Anand et al. (2015).

5.5 INTERACTION OF CONTAGION CHANNELS

So far we have focused on counterparty loss and overlapping portfolio contagion in isolation. Of course, focusing on one channel in isolation only provides a partial view of the system and thus ignores important interaction effects. Indeed, it has been shown by a number of authors that the interaction of contagion channels can substantially amplify the effect of each individual channel (e.g. Poledna et al., 2015; Caccioli et al., 2015; Kok and Montagna, 2013; Arinaminpathy et al., 2012). Although constructing models with multiple contagion channels is difficult, some progress has been made.

Cifuentes et al. (2005) and Caccioli et al. (2015) study the interaction of counterparty loss and overlapping portfolio contagion by combining variants of the contagion processes outlined above into a comprehensive simulation model. In particular, using data from the Austrian interbank system, Caccioli et al. (2015) show that the expected size of a default cascade, conditional on a cascade occurring, can increase by orders

of magnitude if overlapping portfolio contagion occurs alongside counterparty loss contagion, rather than in isolation.

In an equilibrium model Brunnermeier and Pedersen (2009) show that market liquidity and funding liquidity can be tightly linked. In particular, consider a market in which intermediaries trade a risky asset and use it as collateral for their secured short term funding. A decline in the price of the risky asset can lead to an increase in the haircut applied on the collateral. An increase in the haircut can be interpreted as a decrease in funding liquidity and can force intermediaries to sell some of their assets. This in turn can lead to a decrease in market liquidity of the asset. Aymanns et al. (2017) show that a similar link between market and funding liquidity can also result from the local structure of liquidity in over-the-counter markets (OTC). The authors show that, when the markets for secured debt and the associated collateral are both OTC, the withdrawal of an intermediary from the OTC markets can cause a liquidity contagion through the networks formed by the two OTC markets. Similar to Caccioli et al. (2015), the authors show that under certain conditions the interaction of two contagion channels – funding and collateral – can drastically amplify the resulting cascade.

Finally, Kok and Montagna (2013) construct a model that attempts to combine counterparty loss, overlapping portfolio and funding liquidity contagion. Such comprehensive stress testing models are the subject of the remainder of this chapter and will be discussed in detail in the following sections.

6 FROM MODELS TO POLICY: STRESS TESTS

6.1 WHAT ARE STRESS TESTS?

The insights from the models discussed so far are increasingly used in the tools designed to assess and monitor financial stability. After the crisis of 2008, maintaining financial stability has become a core objective of most central banks.[27] One example of such a tool, which has become increasingly prominent over the past years, has been the stress test.[28] Stress tests assess the resilience of (parts of) the financial system to crises (Siddique and Hasan, 2012; Scheule and Roesh, 2008; Quagliariello, 2009; Moretti et al., 2008). The central bank designs a hypothetical but plausible adverse scenario, such as a general economic shock (e.g. a negative shock to house prices or GDP) or a financial shock (e.g. a reduction in market liquidity, increased market volatility, or the collapse of a financial institution). Using simulations, the central bank then evaluates how this shock – in the event this scenario were to take

[27] For example, the mission statement of the US Federal Reserve (FED): 'The Federal Reserve promotes the stability of the financial system and seeks to minimize and contain systemic risks through active monitoring and engagement in the U.S. and abroad'; https://www.bankofengland.co.uk.

[28] Timothy Geither, who played a key role in fighting that crisis as President of the New York Fed and U.S. Secretary of the Treasury, has named his memories after the tool he helped introduce, see Geithner (2014).

place – would affect the resilience of the institution or financial system it tests. Say, for example, that the central bank submits a bank to a stress test. In this case, it would provide the bankers with a hypothetical adverse scenario, and ask them to determine the effect this scenario would have on the bank's balance sheet. If a bank's capital drops below a given threshold, it must raise additional capital. Stress tests evaluate resilience to shocks and link that evaluation to a specific policy consequence intended to enhance that resilience (e.g. raising capital). The process also provides valuable information to regulators and market participants, and helps both to better identify and evaluate risks in the financial system.

6.2 A BRIEF HISTORY OF STRESS TESTS

Stress tests are a relatively novel part of the regulatory toolkit. The potential utility of stress tests had been extensively discussed in the years preceding the financial crisis, and were already used by the International Monetary Fund to evaluate the robustness of countries' financial systems. Banks already designed and conducted stress tests for internal risk management under the Market Risk Amendment of the Basel I Capital Accord, but it was only during the financial crisis that regulators introduced them on a large scale and took a more proactive role in their design and conduct (Armour et al., 2016).

In February 2009 the U.S. Treasury Department introduced the Supervisory Capital Assessment Program (SCAP). This effort was led by Timothy Geither, at a time when uncertainty about the capitalization of banks was still paramount (Schuermann, 2014; Geithner, 2014). Under the auspices of this program the Federal Reserve Board introduced a stress test and required the 19 largest banks in the U.S. to apply it. The immediate motivation was to determine how much capital a bank would need to ensure its viability even under adverse scenarios, and relatedly, whether capital injections from the U.S. tax payer were needed. A secondary motivation was to reduce uncertainty about the financial health of these banks to calm markets and restore confidence in U.S. financial markets (Anderson, 2016; Tarullo, 2016).

In later years, SCAP was replaced by the Comprehensive Capital Analysis and Review (CCAR) and the Dodd–Frank Act Stress Test (DFAST), which have been run on an annual basis since 2011 and 2013, respectively (FED, 2017b, 2017a). These early stress tests gave investors, regulators and the public at large insight into previously opaque balance sheets of banks. They have been credited with restoring trust in the financial sector and thereby contributing to the return of normalcy in the financial markets (Bernanke, 2013).

Across the Atlantic European authorities followed suit and introduced a stress test of their own (EBA, 2017a). This resulted in the first EU stress tests in 2009, overseen by the Committee of European Banking Supervisors (CEBS) (Acharya et al., 2014). Due to concerns about their credibility, the CEBS stress test was replaced in 2011 by stress tests conducted by the European Banking Authority (EBA) (see Ong and Pazarbasioglu, 2014). These have been maintained ever since (EBA, 2017b).

In 2014 the Bank of England also introduced stress tests in line with the American example (BoE, 2014). Around that time, stress tests became a widely used regulatory

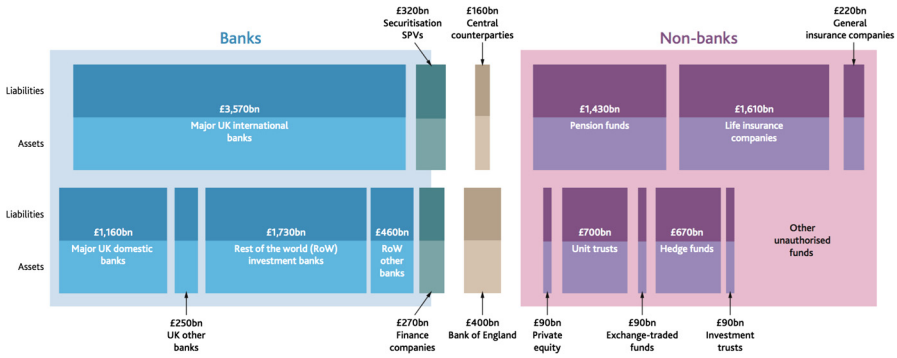

FIGURE 7

Map of the UK financial system. Source: Burrows et al. (2015).

tool in other countries too (Boss et al., 2007). Now stress tests are regarded as a cornerstone of the post-crisis regulatory and supervisory regime. Daniel Tarullo, who served on the board of the U.S. Federal Reserve from 2009 to 2017 and was responsible for the implementation of stress tests in the U.S., has hailed stress tests as 'the single most important advance in prudential regulation since the crisis' (Tarullo, 2014).

Stress tests are not a uniform tool. They can take a variety of forms, which can be helpfully classified along two dimensions. The first dimension concerns their *object*, or the types of agent that the stress test covers; does the stress test only cover banks, or non-banks as well? In the early days of stress testing, only banks were considered, but now there is an increasing trend toward including non-banks. Given the composition of the financial system in most advanced economies, and the importance of non-banks in these financial systems, it is increasingly acknowledged that excluding non-banks from stress tests would leave regulators with a partial picture of financial stability risks in their jurisdiction. In the United Kingdom, for example, almost half of the assets in the financial system are held by non-banks (Burrows et al., 2015), as is illustrated by a stylized map of the UK financial system depicted in Fig. 7.

The second dimension concerns the *scope* of the stress test. Generally speaking, stress tests can be used to evaluate the resilience of individual institutions (*microprudential* stress tests), but could also assess the resilience of a larger group of financial institutions or even of the financial system as a whole (*macroprudential* stress tests) (Cetina et al., 2015; Bookstaber et al., 2014a). Methodologically speaking, the key difference is that *macroprudential* stress tests take the feedback loops and interactions between (heterogeneous) financial institutions – as described in Section 4 and Section 5 of this chapter – into account, whereas the *microprudential* stress tests do not.

Perhaps more than any other financial stability tool, stress tests rely explicitly on the models introduced so far. The following sections will cover micro- and macropru-

dential stress tests in depth. In each instance, we will first review some representative stress tests and subsequently conclude with an evaluation of their strengths and weaknesses.

7 MICROPRUDENTIAL STRESS TESTS

7.1 MICROPRUDENTIAL STRESS TESTS OF BANKS

As noted, microprudential stress tests evaluate the resilience of an individual institution, in this case a bank. Regulators subject the bank to an adverse scenario and evaluate whether a bank has sufficiently high capital buffers[29] (and, in some cases, liquid assets[30]) to withstand it.[31] If this is not the case, regulators can require the bank to raise additional capital (or liquidity) to enhance its buffers. The idea is that this will make the bank more resilient, and by implication increase the resilience of the financial system as a whole.

Given this general approach, microprudential stress tests for banks tend to follow three steps. First, the regulator designs the adverse scenario the bank is subjected to. As noted, this scenario usually involves an economic and/or financial shock. In some cases, the scenario consists of multiple (exogenous) shocks operating at the same time, sometimes with specified ripple effects affecting other variables, which together create a 'crisis narrative' for the bank. The hypothetical scenario a bank is subjected to must be adverse, plausible and coherent. That is, it cannot consist of a set of shocks that, taken together, violate the relationships among variables historically observed or deemed conceivable. Typically, the exogenous shocks affect a set of macro-variables (such as equity prices, house prices, unemployment rate or GDP) as well as financial variables (such as interest rates and credit spreads).

Second, the effect of this scenario on the bank's balance sheet is determined. This determination primarily relates to the effect of the scenario on the bank's capital (and liquidity) buffer, usually expressed as a ratio of capital (liquidity) buffers to assets,[32]

[29] The simplest measure of a capital buffer is that of a bank's net assets – the value of its assets minus its liabilities. This represents a buffer that protects the bank against bankruptcy when its assets decline in value. In most models described earlier in this chapter, this buffer corresponds to a bank's equity. When describing whether a bank has a sufficiently high buffer, the term 'capital adequacy' is commonly used. For a more comprehensive overview, see Armour et al. (2016), Chapter 14.

[30] A liquidity buffer is intended to ensure that, when liquidity risks of the type discussed in Section 3.2 materialize, a bank has sufficient liquid assets to meet demands for cash withdrawals. Although microprudential liquidity stress tests for banks have been developed, they are currently not yet widely used for regulatory purposes. Hence, we will focus on microprudential capital stress tests here.

[31] Note that this capital buffer is an example of a regulatory leverage constraint as introduced in Section 3.2.

[32] When capital ratios are computed as capital over total (unweighted) assets, this amounts to the inverse of the leverage ratio, as defined in Section 4. Regulators typically use a more complex measure for the capital buffer to account for the fact that some assets are riskier than others. Suppose a bank holds two assets with the same value, but one (asset Y) is riskier than the other (asset X). When regulators take the riskiness of these assets into account, to meet regulatory requirements the bank would have to hold a higher capital

and profits. This calculation is based on an evaluation of how the shocks change the values of the assets and liabilities on the bank's balance sheet, as well as on the bank's expected income. Value changes on the balance sheet materialize either through a re-evaluation of the market value (if the asset or liability is marked-to-market), or through a credit shock re-evaluation. These effects are captured by market risk models and credit risk models (such as those described in Siddique and Hasan, 2012; Scheule and Roesh, 2008; Quagliariello, 2009; Moretti et al., 2008). Credit losses for specific assets or asset classes are commonly computed by multiplying the probability of default (PD), the exposure at default (EaD), and the loss given default (LGD). Estimating these variables is therefore key to the credit risk component of stress testing (Foglia, 2009). Value changes to the expected income stream result largely from shocks that affect income on particular assets or asset classes, such as interest rate shocks. This determination matters in the context of the stress test because such income can, in the form of retained earnings, feedback into capital buffers.[33] Usually, these microprudential stress test models therefore equate the post-stress regulatory capital buffer to the sum of post-stress retained earning plus regulatory capital[34] over the post-stress (risk-weighted[35]) assets.

Third, once the bank's post-stress capital buffer has been determined, regulators compare it to a hurdle rate. This hurdle rate is usually set at such a level that, when passing it, the bank would withstand the hypothetical scenario without being at risk of bankruptcy. Consequently, if the bank does not meet this hurdle rate, it fails the stress test and is said to be 'undercapitalized' (that is, its capital buffer is insufficient). When that happens, the regulator commonly has the authority to require the bank to raise extra capital to increase its buffer, so as to leave it better prepared for adverse scenarios. Microprudential stress tests are thus used as a tool to recapitalize undercapitalized banks, thereby reducing their leverage and increasing their resilience.

7.2 MICROPRUDENTIAL STRESS TEST OF NON-BANKS

Given the importance of non-bank financial institutions to the financial system,[36] it was only a matter of time before the scope of microprudential stress tests would be extended beyond banks. The rationale for doing so is similar to the one that applies to

buffer for asset Y than for asset X, corresponding to their relative riskiness. This process is referred to as 'risk-weighting', and the resulting capital buffer is commonly expressed relative to 'risk-weighted assets' (RWAs).

[33] This is true unless part of this income is being paid to shareholders as dividends, which stress tests commonly assume not to be the case.

[34] If the scenario results in a loss to the bank's equity and lowers its income, the capital buffer drops (ceteris paribus).

[35] In most cases the model also updates the assets' risk-weights to reflect that the adverse scenario has altered the riskiness of the asset (class). For an overview of the methodologies commonly used by banks, see Capgemini (2014). The 'standard' approach as proposed by regulators is set out in BIS (2015).

[36] See e.g. FSB (2015), ECB (2015), Burrows et al. (2015), Pozsar et al. (2010), Pozsar and Singh (2011), Mehrling et al. (2013), Pozsar (2013).

banks: regulators want to understand the resilience of non-bank financial institutions, and where they find fragility they want to be able to amend it. So far, at least three types of non-bank financial institutions are subjected to stress tests: insurers, pension funds and central clearing parties (CCPs).

Like the microprudential stress tests for banks, those for non-bank financial institutions are primarily used to assess capital adequacy in times of distress. However, the methodology used in that assessment differs between the various institutional types, because each type faces a different set (and type) of risks. For example, the balance sheet composition differs for each institution, which in turn means each institution is exposed to different tail risks which should be reflected in the scenario design used in the stress test. Moreover, because of the differences in balance sheet composition, losses materialize in different ways and should be determined using methodologies suitable to each institutional type. Finally, the benchmark each type of institution has to meet in order to 'pass' the stress test differs too, because the regulatory requirements vary among different institutional types.[37]

In sum, the heterogeneity of non-bank financial institutions requires bespoke microprudential stress tests. They all, however, follow the same pattern: they start by setting a hypothetical adverse scenario, evaluate the effect of that scenario on the institution's balance sheet, and compare the post-stress balance sheet to regulatory requirements (hurdles). Table 1 sets out some of the most salient differences between microprudential stress tests for various types of institutions. In what follows, we provide a high-level overview of regulatory stress tests for insurers, pension funds, and central clearing parties.

Insurers and Pension Funds: Insurance stress tests are becoming increasingly common, and have been conducted by the Bank of England (BoE, 2015), the IMF (under its FSAP Program) (Jobst and Andreas, 2014), the U.S. Federal Reserve (Accenture, 2015; Robb, 2015), and the European Insurance and Occupational Pensions Authority (EIOPA) (EIOPA, 2016). Similarly, pension fund stress tests have been conducted by the International Organisation of Pension Fund Supervisors (IOPFS) (Ionescu and Yermon, 2014) and, in the EU, by EIOPA (EIOPA, 2017).

EIOPA's 2016 stress test of life insurers tested each insurer's capital and liquidity adequacy[38] (EIOPA, 2016). When evaluating capital adequacy, the benchmark was that an insurer's assets should exceed its liabilities.[39] Liquidity adequacy was assessed by performing a cash-flow analysis to investigate whether the timing of insurer's incoming cash-flow (from its assets) matched the insurer's expected cash outflow (resulting from its insurance liabilities).

[37] This reflects the different loss absorption mechanisms that have been designed for each type of non-bank financial institution.

[38] For insurance companies, this refers to the ability to meet insurance obligations.

[39] In other words, its Assets-over-Liability ratio (AoL) should exceed a hundred percent.

Table 1 Key distinguishing characteristics of microprudential stress tests for banks, insurers, pension funds, and central clearing parties

	Banks	Insurers	Pension Funds	CCPs
Primary Objective	Capital Adequacy	Capital and Liquidity Adequacy		Capital Adequacy
Measure of Capital Adequacy	Capital Ratio	Assets over Liability Ratio (AoL)	Coverage Ratio	Default Waterfall
Loss Assessment	How asset losses affect capital buffers	How asset and liability re-evaluations affect capital, and whether* liabilities can be met		If, and how, clearing member default losses are absorbed by the default waterfall

* *Other than in the case of banks, insurance and pension fund liabilities re-evaluate significantly with certain types of adverse shocks, such as interest rate shocks. In such scenarios, the discounted future value of promised insurance schemes and pension fund schemes adjusts. Pension funds must also assess whether there is no significant maturity mismatch between promised pension fund payments (liabilities) and income from assets to meet them.*

EIOPA's 2015 stress test of occupational pension funds assessed whether pension payment promises[40] could be met in the face of adverse market conditions. The hypothetical adverse scenario was tailored to risks specific to a pension fund. For example, the effect of increased life expectancy (which lengthens the time a pension fund must pay out a pension, and thus increases the cumulative amount of pension payments a pension fund must make) on the pension fund's ability to meet its pension obligations was tested.

Central Clearing Parties: Central clearing parties (CCPs) have been created to mitigate counterparty risk, for example in (simple, or 'over-the-counter' (OTC)) derivatives transactions. By doing so, they also reduce the likelihood that counterparty risk causes a cascade of losses, and generates contagion (as has been discussed in Section 5). In this way, well-functioning CCPs can mitigate systemic risk in financial systems.

CCPs operate by stepping in between two contractual counterparties, and becoming, as is commonly noted, 'the buyer to every seller, and the seller to every buyer'.[41] Once a contract is 'cleared' through a CCP, its counterparties are referred to as 'clearing members'. As long as no clearing member defaults, the assets and liabilities of the CCP balance out, so that the CCP faces no market risk. That changes when clearing members default, in which case the CCP is exposed to losses. To absorb such

[40] Specifically, those related to defined benefit and hybrid pension schemes.

[41] Because many counterparties that are also exposed to each other engage in contractual relationships via the CCP, the CCP can also net out exposures, thereby reducing the complexity of exposures that counterparties must manage and reducing bilateral exposures (Cont and Kokholm, 2014).

losses, the CCP has an elaborate process in place that distributes these losses among all its clearing members as well as its own equity, which is referred to as a 'default waterfall' (see Murphy, 2013; Capponi et al., 2017). This default waterfall consists of various contributions of the clearing members (e.g. initial and variation margin, default fund contributions) and some of the CCPs own capital buffer (equity). When losses materialize, these are absorbed by each of the layers in turn, until the CCP's own capital buffer is exhausted and it defaults.[42]

CCPs have become increasingly important after the financial crisis, as regulators require counterparties to frequently-used contracts to clear these contracts through CCPs (ESMA, 2017; Ey, 2013). Individual CCPs have also grown substantially and process very high volumes of trades, leading some to argue that their failure would be catastrophic for the financial system in which they operate (and perhaps beyond) (ESMA, 2015; Murphy, 2013). That is why regulators around the world increasingly carry out microprudential stress tests for CCPs, including the Commodity Futures and Trading Commission (CFTC) in the U.S. (CFTC, 2016), the British and German regulatory authorities (Erbenova, 2015) (this will include a U.S. regulator in 2017) (Robb, 2015), and the European Securities and Markets Authority (ESMA) (ESMA, 2015).

Microprudential stress tests for CCPs focus on whether a CCP's capital buffer (its default waterfall) can absorb losses in a crisis event, to avoid that the CCP defaults. ESMA's CCP stress tests illustrates how such a stress test can be designed (ESMA, 2015).

In ESMA's microprudential CCP stress test from 2015, the adverse scenario included the default of the CCP's two largest clearing members (that is, those two clearing members with the largest contribution to the CCP's default fund),[43] while the CCP was simultaneously hit by a severe adverse market shift.[44] Because clearing members often trade in multiple CCPs, the two defaulted clearing members for each CCP were assumed to default in all CCPs where they cleared – which is referred to as cross default contagion.

To assess the extent to which the CCP's capital buffer has been depleted as a consequence of this adverse scenario, the test calculates the losses to each step of the CCP's default waterfall. Losses beyond the absorption capacity of the default

[42] A comprehensive overview of the operation of CCPs and the risks they create is beyond the scope of this chapter. Examples of excellent overviews include: Cont (2015), Murphy (2013), Duffie et al. (2015), Duffie and Zhu (2011).

[43] This scenario tests whether CCPs meet the minimum requirement set under the EU's EMIR regulation (see Art. 42(2)).

[44] It is important that both the default of two counterparties and a severe adverse market shift hit simultaneously. If two clearing members default, but there is no market shift, the total margin posted should be sufficient to absorb all losses of the defaulted clearing members. If only the market conditions change, the CCP is not at risk of default because its variation margin ensures it is not exposed to changes in market conditions.

waterfall are calculated as well. Taken together, these two criteria are used to judge whether a CCP is sufficiently capitalized.[45]

7.3 STRENGTHS AND WEAKNESSES OF CURRENT MICROPRUDENTIAL STRESS TESTS

Microprudential stress tests are valuable from at least three perspectives. First, they give market participants more insight into the opaque balance sheets of the financial institutions being evaluated (Bookstaber et al., 2014a). Opacity coupled with asymmetric information can, especially in times of financial distress, lead to a loss of confidence (Diamond and Dybvig, 1983; Brunnermeier, 2008). If the type and quality of a financial institution's assets and liabilities are unclear, outsiders may conceivably fear the worst and, for example, pull back their funding.[46] Such responses feed speculative runs which can turn into self-fulfilling prophecies and, ultimately, (further) destabilize the financial system at the worst possible time (He and Xiong, 2012; Diamond and Dybvig, 1983; Martin et al., 2014; Copeland et al., 2014). Credibly executed microprudential stress tests provide insight into an institution's balance sheet, can signal confidence about the institution's ability to withstand severe stress, and create a separating equilibrium that allows solid banks to avoid runs (Ong and Pazarbasioglu, 2014; Bernanke, 2013).[47]

Second, microprudential stress tests help financial institutions to improve their own risk-management. By forcing them to assess their resilience to a variety of novel scenarios, stress tests require banks to take a holistic look at their own risk-management practices (Bookstaber et al., 2014a). As a consequence, more banks are now also engaged in serious internal stress tests (Wackerbeck et al., 2016).

Third, microprudential stress tests have proven to be an effective mechanism to recapitalize banks (Armour et al., 2016). In the EU, the stress tests have forced banks to raise their capital by 260 billion euros from 2011 to 2016 (Arnold and Jenkins, 2016), and in the US the risk-weighted regulatory ratio of the banks that took part in the stress test went up from 5.6 percent at the end of 2008 to 11.3 at the end of 2012 (Bernanke, 2013). Against a backdrop of frequent questions about the adequacy of banks' capital buffers,[48] in part due to the gaming of risk weights (Behn et al., 2016;

[45] Although the 2015 ESMA CCP stress test was microprudential in nature, an attempt was made to make it somewhat macroprudential by capturing a (first-order) contagion effect. In case the adverse scenario affected the default fund contributions of the non-defaulting clearing members, the resulting hit to the equity of these clearing members was included. If that hit, in turn, caused a loss to the non-defaulting clearing member that exceeded a threshold percentage of its equity, that clearing member was said to be vulnerable. If that hit wiped out all of the previously non-defaulting clearing member's equity, it was said to default as well.

[46] The general economic principle at play is that of asymmetric information causing market failures, see Akerlof (1970).

[47] Weaker banks, however, may be exposed by the stress test. But regulators would learn this information first, giving them an opportunity to intervene before the information reaches the market.

[48] See, for example, Admati and Hellwig (2014).

Fender and Lewrick, 2015; Groendahl, 2015), many regulators have welcomed the role that stress tests have played to enhance the resilience of banks. Even if microprudential stress tests are not, strictly speaking, designed to assess and evaluate systemic risk, their role in raising capital adequacy standards can have the effect of enhancing resilience (Greenwood et al., 2015).

Despite their strengths in specific areas, the current microprudential stress tests have been criticized on at least four grounds. First, and most importantly from the perspective of this chapter, microprudential stress tests ignore the fact that economies are complex systems (as noted in Section 1) and therefore are ill-suited to capture systemic risk. As discussed in Sections 4 and 5 of this chapter, systemic risk materializes due to interconnections between heterogeneous agents (for example due to overlapping portfolios and funding liquidity contagion). By considering institutions in isolation, microprudential stress tests (largely[49]) ignore the interconnections and interaction between financial institutions that serve to propagate and amplify distress caused by the initial shock resulting from the adverse scenario. Empirical research suggests that this approach substantially underestimates the losses from adverse scenarios (Bookstaber et al., 2014b, also see Section 3). Bernanke (2015), for example, notes that the majority of the losses in the last financial crisis can be traced back to such interactions as opposed to the initial shock emerging from credit losses in subprime mortgage loans.

Second, microprudential stress tests tend to impose an unrealistically large initial shock. Because regulators are aware of the fact that a microprudential modeling strategy does not capture the higher order losses on the balance sheets of individual financial institutions, they use a more severe initial scenario that causes direct losses to compensate for that. To generate a sufficiently large initial shock, the scenario tends to depart quite strongly from reality. Often, the initial scenario posits a substantial increase in the unemployment rate as well as a sharp drop in GDP.[50] In reality, however, it is uncommon for these conditions to *precede* a financial crisis, so the stress test might be testing for the wrong type of scenario.[51] Imposing an unrealistic shock – and excluding higher-order effects – can also affect the outcome of the stress test in unexpected ways. In particular, while stress tests with large initial shocks might get the overall losses right, they might fail to accurately capture the distribution of losses across institutions, which ultimately determines which banks survive and which do not. For an investigation of this issue, see for example Cont and Schaanning (2017).

Third, the value of the information produced by microprudential stress tests is increasingly being questioned. The outcomes of stress tests have converged (Glasserman et al., 2015), perhaps because banks seem increasingly able to 'train to the test'. This has left some to wonder what the information produced by the stress tests is

[49] In some cases a proxy for such contagious effects is included in the microprudential stress test.

[50] See, for example, FED (2016), BoE (2016), ESRB (2016).

[51] Instead, exogenous shocks such as declining house prices or stock markets precede financial crises. These are commonly also part of the initial scenario.

actually worth (Hirtle et al., 2016), and others to conclude that the value of such information has declined over time (Candelon and Sy, 2015). Such concerns have been further fueled by the apparent willingness of some regulators to allow banks to pass the test on the basis of dubious assumptions.[52]

Finally, the stress tests are commonly calibrated to the losses incurred during the last financial crisis, raising questions about their relevance in relation to current, let alone future, scenarios – not least because the financial system constantly changes.

8 MACROPRUDENTIAL STRESS TESTS

Because the financial system is a complex system (see Section 1), the whole is different from the sum of its parts (Anderson et al., 1972; Farmer, 2012; Battiston et al., 2016). In other words, measures focused on the health of individual institutions (as microprudential stress tests would prescribe) will not necessarily guarantee the health of the financial system as a whole. In fact, such measures might destabilize the system. To understand the system as a whole – and, by implication, systemic risk – stress tests have to account for feedback loops and non-linearities.

The inability of microprudential stress tests to appropriately account for systemic risk has prompted the development of a specific type of stress tests focused on this goal; the macroprudential stress test. Macroprudential stress tests aim to assess the resilience of a whole sector, or even the whole financial system, rather than that of one particular institution. To do so, they extend the microprudential stress test by including contagion effects between interconnected financial institutions that can arise following the initial adverse scenario. This means that the regulators must not only assess the effect of the initial shocks on the individual balance sheets, but must capture how the balance sheets are interlinked (see Section 5). They should also address what consequences such interlinkages have for the potential of financial distress to propagate throughout the system. The contagion models discussed in Sections 4 and 5 can help inform regulators on how to model these higher order spill-over effects.

This section discusses two macroprudential models for banks, and one that combines banks and non-banks. The first two models, the Bank of England's 'Risk Assessment Model of Systemic Institutions' (RAMSI) and the Bank of Canada's 'MacroFinancial Risk Assessment Framework' (MFRAF), have been used in stress tests. The last model, U.S. Office of Financial Research's (OFR's) 'Agent-Based Model for Financial Vulnerabilities' (ABMFV), has not.[53]

[52] Deutsche Bank, which has seen its share price fall significantly in 2016 on fears that it could face a US fine of up to USD 14bn, was given special treatment by the European Central Bank in the 2016 EBA stress tests, so that it could use the result of the stress test as evidence of its healthy finances (Noonan et al., 2016).

[53] We focus on comparing these three models. However, other relevant macroprudential stress tests have recently been developed. Baranova et al. (2017), for example, study market liquidity in a corporate bond

The ABMFV and the RAMSI are examples of cases where heterogeneous agent models have been applied to macroprudential stress tests.[54] The MFRAF is an example of another, somewhat more traditional approach.

After introducing these three models, their differences and similarities are outlined. The section ends with a discussion of the strengths and weaknesses of these macroprudential stress tests.

8.1 THREE MACROPRUDENTIAL STRESS TESTS

8.1.1 RAMSI Stress Test of the Bank of England

The Bank of England has pioneered the development and use of a macroprudential banking stress test, called the RAMSI model.[55] The model evaluates how adverse shocks transmit through the balance sheets of banks and can cause further contagion effects (Burrows et al., 2012). It is based on earlier research that has been conducted by Bank of England researchers and others (Aikman et al., 2009; Kapadia et al., 2013; Alessandri et al., 2009).

The RAMSI stress test begins as a microprudential stress test. Subsequently, possible feedback effects within the banking system are considered. If the initial shocks have caused a bank to fall below its regulatory capital ratio, or have caused the bank to be shut off of all unsecured funding[56] markets, the bank respectively suffers an insolvency or illiquidity default. Subsequently, the default causes two interbank contagion effects: common asset holding contagion and interbank contagion. The combined effect of the marked-to-market losses and the credit losses can cause other banks to default through insolvency or illiquidity by being shut out of the funding market. If this happens, the loop is repeated. If this does not happen, each bank's net operating expenses are invested in assets such that the bank targets its regulatory risk-weighted target ratio. The credit losses persist, but the marked-to-market losses are assumed

market by modeling broker-dealers, hedge funds and (unlevered) asset managers. They capture common asset holding contagion. Like Baranova et al. (2017), Feroli et al. (2014) highlight that subdued leverage of financial intermediaries is no sufficient ground to rule out stability concerns. Instead, unlevered investors (such as unlevered funds) may be the locus of potential financial instability. Dees and Henry (2017) offer a host of macroprudential (stress testing) tools. The multi-layered network model (and ABM) of Kok and Montagna (2013) (discussed in Section 5.5) can also be considered to be macroprudential stress testing model, as it is a data-driven stress simulation of the European Union (EU) banking system. The model of Kok and Montagna (2013) is similar in style to the ABMFV discussed here.

[54]Indeed, these models combine the contagion models discussed in Section 5.

[55]The model is currently being phased-out. We discuss this model to showcase its strengths and weaknesses. These are further treated in Section 8.3.

[56]This causes funding liquidity contagion. The bank is shut off of all unsecured funding based on a rating system. Based on the shocked balance sheets and profit and losses (PL), the credit score for the bank is computed, which the authors assume affects the funding cost of the bank and its ability to access the long-term and short-term funding market. This credit score takes into account liquidity and solvency characteristics of the bank's balance sheet, but also system-wide market distress. If its credit score is above a certain threshold, the bank is shut out of the unsecured funding markets altogether (both long-term and short-term) and is assumed to default.

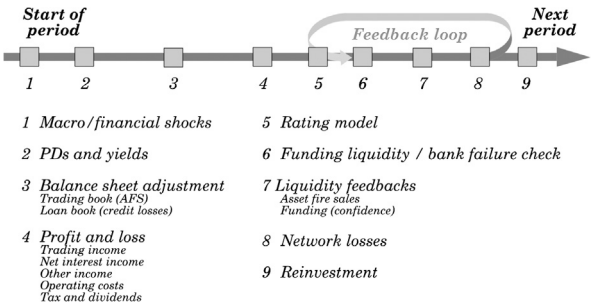

FIGURE 8

Description of the RAMSI stress test of the Bank of England. Source: Aikman et al. (2009).

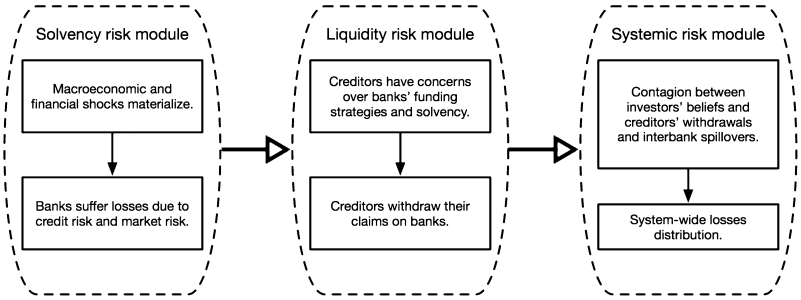

FIGURE 9

Description of the MFRAF stress test of the Canada. Source: Anand et al. (2014).

to disappear as each asset price returns to its fundamental value. Then, the next time step starts, and the process can be repeated, starting with a balance sheet that includes the credit losses incurred in the previous time step.

Thus, the RAMSI stress test turns a microprudential foundation into a macroprudential model by including interbank contagion effects via common asset holdings, interbank losses and funding liquidity contagion. Fig. 8 summarizes what happens at each step of the RAMSI model.

8.1.2 MFRAF Stress Test of the Bank of Canada

Contrary to the RAMSI model, the Bank of Canada's MacroFinancial Risk Assessment Framework (MFRAF) is at its core not a heterogeneous agent model, but a global games model, such as those described in Morris and Shin (2001). In the way it sets up funding runs (i.e. as a global coordination game) it is similar to the seminal model of Diamond and Dybvig (1983) (discussed in Section 5). It captures three sources of risk that banks face (Anand et al., 2014; BoC, 2014, 2012): solvency, liquidity, and spill-over risk (see Fig. 9).

The MFRAF stress test has been applied to the Financial Sector Stability Assessment (FSAP) of the Canadian financial sector conducted by the International

$t = 0$	$t = 1$ (round 1)	$t = 1$ (round 2)	$t = 2$
1. Debt issuance	1. Interim shock	1. Belief updated	1. Investment matures
2. Investments	2. Private signals	2. New pooling price	2. Final shock
	3. Debt withdrawals	3. New private signals	3. Debts honored
		4. Debt withdrawals	

FIGURE 10

Time steps in theoretical model of MFRAF stress test of the Bank of Canada. Source: Anand et al. (2015).

Monetary Fund (IMF) in 2014 (IMF, 2014). The 2014 FSAP stress test, which considers the direct effects of adverse shocks on the solvency of banks, is microprudential. When extending it to capture system-wide effects (i.e. liquidity effects and spill-over effects) using MFRAF, overall losses to the capital of the Canadian banks rose with 20 percent. This again shows that microprudential stress tests significantly underestimate system-wide losses. We will now discuss the theoretical underpinnings of the MFRAF stress tests, which builds on research at the Bank of Canada and elsewhere (Anand et al., 2015; Gauthier et al., 2012, 2014).

The theoretical model that underpins the MFRAF stress test is described in Anand et al. (2015) and will be discussed here.[57] The model captures how solvency risks, funding liquidity risks, and market risks of banks are intertwined. In essence, this works as follows: a coordination failure between a bank's creditors and adverse selection in the secondary market for the bank's assets interact, leading to a vicious cycle that can drive otherwise solvent banks to illiquidity. Investors' pessimism over the quality of a bank's assets reduces the bank's access to liquidity, which exacerbates the incidence of runs by creditors. This, in turn, makes investors more pessimistic, driving down other banks' access to liquidity. The model does not capture interbank contagion upon default, although this is captured in MFRAF (IMF, 2014).

The key components of the model according to the evolution of the model over time is summarized in Fig. 10.

8.1.3 ABM for Financial Vulnerabilities

The final system-wide stress testing model that will be discussed, the Agent-Based Model (ABM) for Financial Vulnerabilities (Bookstaber et al., 2014b),[58] captures similar contagion mechanisms as MFRAF, but it does so using a different methodology. The model is designed to investigate the vulnerability of the financial system to asset- and funding-based firesales that can lead to common asset holding contagion.

[57] The degree to which the theoretical model of Anand et al. (2015) is in unaltered form translated into the MFRAF stress tests is not made explicit in the IMF (2014) documentation of the MFRAF stress test.

[58] A further discussion of some agent-based models of the financial crisis and stress testing can be found in Bookstaber and Kirman (2018).

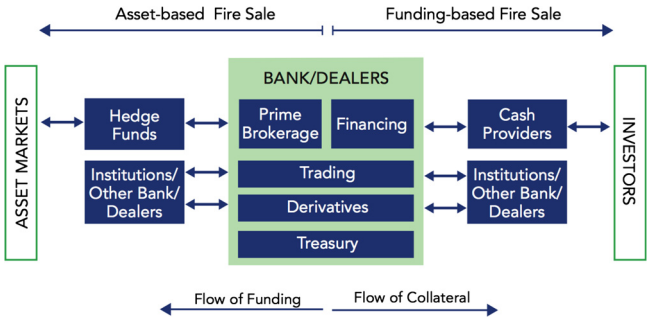

FIGURE 11

Map of the financial system and its flows, as considered in the ABM for Financial Vulnerabilities. Source: OFR (2014).

The financial system is modeled as a combination of banks that act as intermediaries between the cash provider (a representative agent for various types of funds) and the ultimate investors (i.e. the hedge funds). Hedge funds can receive funding from banks for long positions in return for collateral. Banks, in turn, receive funding from the cash provider in return for collateral. Funding and collateral therefore flow in opposite directions, as is illustrated in Fig. 11.

The role of the cash provider c in the model is to provide secured funding to banks.[59] Although the cash provider is not actively modeled, it can take two actions. First, it can set the haircut (this can force the hedge fund to engage in fire sales), and second it can pull funding from the banks (this may lead the bank to contribute to pre-default contagion or default).

Hedge funds have a balance sheet that consists of cash and tradable assets on the asset side, and secured loans and equity (and possibly short positions) on the liability side. A hedge fund funds its long positions in assets using funding from banks in the form of repurchase contracts (often referred to as repos).[60] When funding themselves this way, hedge funds receive cash in return for collateral they pledge to the bank. Although the hedge fund does not face a regulatory leverage constraint, it

[59]The cash provider is a representative agent that represents financial institutions that typically provide funding to banks, such as asset managers, pension funds, insurance companies, and security lenders, but most importantly, money market funds.

[60]In a repo, one party sells an asset to another party at one price at the start of the transaction and commits to repurchase the fungible assets from the second party at a different price at a future date. If the seller defaults during the life of the repo, the buyer (as the new owner) can sell the asset to a third party to offset his losses. The asset therefore acts as collateral and mitigates the credit risk that the buyer has on the seller. Although assets are sold outright at the start of a repo, the commitment of the seller to buy back the fungible assets in the future means that the buyer has only temporary use of those assets, while the seller has only temporary use of the cash proceeds of the sale. Thus, although repo is structured legally as a sale and repurchase of securities, it behaves economically like a collateralized loan or secured deposit. For an overview, see https://www.icmagroup.org/Regulatory-Policy-and-Market-Practice/repo-and-collateral-markets/icma-ercc-publications/frequently-asked-questions-on-repo/1-what-is-a-repo/.

faces an implicit leverage constraint based on the haircut it receives on its collateral. The haircut determines how much equity a hedge fund needs for a given amount of repo funding. If the haircuts on all types of collateral (i.e. on all types of assets that can be pledged as collateral) is the same, and assuming that the bank passes on the haircut it receives from the cash provider, the maximum leverage $\bar{\lambda}_{jt}$ of the hedge fund j at time t is given by $\bar{\lambda}_{jt} = \frac{1}{h_{cjt}}$. If the leverage of the hedge fund exceeds the maximum leverage,[61] the hedge fund is forced to de-lever. It will do so by fire selling assets. This can cause marked-to-market losses for other banks or hedge funds who hold the same assets.

The banks act as an intermediary between buyers and sellers of securities and between lenders and borrowers of funding.[62] On the whole, the bank can contribute to financial distress pre-default and post-default in various ways. Pre-default, the bank may have to fire sell assets or to pull funding from the hedge fund (which consequently may also have to engage in firesales) in order to raise cash, de-lever, or pay back funding to the cash provider (if the cash provider pulled its funding). In addition, by passing on an increased haircut to the hedge fund, it can trigger a hedge fund to engage in firesales. Post-default, the bank contributes to exposure losses and further firesale losses.

8.2 COMPARING AND EVALUATING MACROPRUDENTIAL STRESS TESTS: FIVE BUILDING BLOCKS

To comprehensively design, study and evaluate macroprudential stress tests, we introduce a general framework consisting of five building blocks that allow us to break down each stress test in discrete components: (1) types of financial institutions (agents), (2) financial contracts, (3) markets, (4) constraints, and (5) behavior.[63] This framework also offers an analytically coherent way to combine the various heterogeneous agent models discussed in Section 4 and Section 5 in order to capture their interactions (see Section 5.5). With such a framework one can capture *critical features* necessary to be able to capture systemic risk. This section covers these five building blocks and compares the three macroprudential stress tests discussed

[61] A hedge fund's leverage can exceed the maximum leverage due to asset price depreciations (as a consequence of firesales, for example) or increases in the haircut (due to the cash provider's downward assessment of the bank's solvency and/or liquidity). If the hedge fund is forced to de-lever, it will attempt to go back to a 'buffer leverage' level, which is below the maximum leverage value.

[62] In its role, it facilitates maturity, liquidity, and risk transformations. The banks have various desks that play a role in these processes: the prime broker, the finance desk, the trading desk, the derivatives desk, and the treasury. The various equations associated with the functioning of the bank dealer and its various subdesks can be found in Bookstaber et al. (2014b).

[63] With these five building blocks, many relevant features of a financial system can be captured by initializing bespoke implementations for each building block. Once financial institutions and financial contracts are defined, a multi-layered network can be initialized. When, subsequently, markets, constraints, and behavior are chosen, the dynamics of the system can be studied.

Table 2 Comparison between the three macroprudential stress tests (RAMSI, MFRAF, ABMFV) regarding the (system-wide stress test) building blocks: (1) financial institutions; (2) financial contracts; (3) markets; (4) constraints; and (5) behavior. Note that rc, cc, mc stand for regulatory, contractual and market-based constraints respectively. Remark that MFRAF captures unsecured interbank loans, counterparty loss contagion and a leverage constraint, the theoretical model of Anand et al. (2015) does not. We list the behavior that impacts the state of the system

		RAMSI	MFRAV	ABMFV
(1)	Financial institutions:			
	Banks	✓	✓	✓
	Creditors (exogenous)	✓	✓	✓
	Hedge funds	–	–	✓
(2a)	Financial contracts:			
	Traded securities	✓	✓	✓
	Unsecured interbank loans	✓	✓	–
	Unsecured term deposits	–	✓	–
	Secured interbank loans (repos)	–	–	✓
(2b)	Channels of contagion:			
	Overlapping portfolios	✓	✓	✓
	Counterparty loss	✓	✓	–
	Funding liquidity	✓	✓	✓
	Margin spirals	–	–	✓
(3)	Modeled markets:			
	Traded securities	✓	✓	✓
(4)	Constraints:			
	Leverage constraints (rc)	✓	✓	✓
	Liability payment obligations (cc)	✓	✓	✓
	Margin call obligations (cc)	–	–	✓
	Funding run (mc)	✓	✓	✓
(5)	Behavior:			
	Pre-default			
	– no action (banks)	✓	✓	–
	– action (banks, hedge funds)	–	–	• control leverage • meet contractual obligations • maximize profits
	– exogenous action (creditor run)	✓	✓	✓
	Post-default			
	– Default procedure	• fire sales • exposure losses	• fire sales (of collateral)	• fire sales (implicit)

above[64] as we go along (these findings are summarized in Table 2). We will see that these stress tests implement each building block with varying degrees of fidelity to the real world.

8.2.1 Financial Institutions

Financial institutions are at the heart of any financial stability analysis and form a key component of macroprudential stress tests. In most models they are represented by balance sheets filled out with a collection of financial contracts that are unique to that institution. Moreover, each institution comes with its own set of constraints and behavioral rules. By endowing an institution with its unique collection of financial contracts, combination of constraints, and behavioral rules, various types of heterogeneous financial institutions (e.g. banks, insurance companies, hedge funds, unlevered funds, central clearing parties) can be characterized. This allows for the inclusion of the many types of financial institutions that need to be studied to capture the dynamics of a financial system under stress.

None of the macroprudential models discussed in this chapter capture all relevant financial institutions, which limits their claim to be a truly 'system-wide' macroprudential stress test. Specifically, the RAMSI and MFRAF model only capture the banking system, and though the ABMFV also considers non-banks it only covers a subset (hedge funds and cash providers).[65]

8.2.2 Financial Contracts: Interlinkages and Associated Contagion Channels

Contracts sit on the balance sheet of each institution, but because contracts are between institutions, they also stipulate the interconnections between institutions. Taking institutions as the nodes in the network the contracts define the edges of the network. (Common asset holdings also define connections, though a more accurate approach is to treat these as bipartite networks.) Contagion dynamics, such as those described in Section 5, operate over these financial contracts to jump from institution to institution. It is therefore important to ensure that the models representing these contracts capture the features that create the interconnections between institutions (e.g. contractual counterparties) and enable contagion (e.g. valuation method, contractual obligations).

The three macroprudential stress tests capture these three contractual characteristics for a subset of contracts (leaving out some relevant contractual types), but do study how the contagion dynamics operating over them can interact. Specifically, models capture the interaction between contagion channels discussed in Section 5:

[64] See Section 8.1.1, Section 8.1.2, and Section 8.1.3.

[65] Each model also considers exogenous creditors. The balance sheets of exogenous agents are not explicitly modeled. As such exogenous agents cannot default. When exogenous creditors withdraw a loan, the cash exists the system.

common asset holding contagion, counterparty loss contagion, and funding liquidity contagion. The ABMFV also captures 'collateral contagion'.[66]

8.2.3 Markets

In most models (as in reality), markets are the places where asset prices are determined, as well as the place where new contracts are agreed upon and existing ones modified or terminated. It is their role in the price formation process and the provision of liquidity that makes the modeling of markets particularly relevant to macroprudential stress tests. Markets are diverse in their institutional characteristics; they can be bilateral (such as the interbank loan market), exchange-based (like the stock market), intermediated (like a dealer-based market for, say, corporate bonds), or centrally cleared (i.e. by a CCP). Typically, there is a specific market for each financial contract on the balance sheet of an institution.

However, although each of our three macroprudential models consider multiple types of contractual linkages, they only model one market: the market for common asset holdings.[67] Moreover, although all three models consider bilateral funding, they do not include a market for these contracts. Therefore, when an (un)secured loan is not rolled over, institutions have no opportunity to seek funding elsewhere. That potentially causes these models to overestimate financial distress.

Because financial stability critically depends on price formation and the ability of institutions to forge contractual links (or break them), it is important to model the markets that exist for each type of contract (and do so with sufficient realism).[68] An understudied challenge is thus to determine whether and how the dynamics in a given market contribute to financial (in)stability, and to reflect that in stress testing models. This is complicated, because ideally it would require an understanding of the supply and demand functions for each market.[69]

8.2.4 Constraints

Financial institutions typically face four types of constraints: regulatory constraints, contractual constraints, market-based constraints, and internal risk limits. Regulatory

[66] 'Collateral contagion' refers to the contagious spill-overs that can arise from margin calls associated to repo contracts (e.g. secured funding contracts). Institutions receive a margin call when the asset collateral value drops (or haircuts increase) so that it is not enough to cover the loan amount. If institutions are not able to meet the margin call they may be forced to engage in fire sales. Collateral contagion is especially relevant as it interacts with common asset holding contagion. Indeed, price falls due to fire sales can trigger collateral contagion.

[67] The modeling of price formation is approached differently in the three models. In the case of the RAMSI and MFRAF model a price impact function is used. The MFRAF model updates prices based on the investors' beliefs about the quality of the assets.

[68] For example, Baranova et al. (2017) show for the case of corporate bond markets that market liquidity (and common asset holding contagion) critically depends on the ability of intermediaries to make markets.

[69] To capture price formation (or counterparties for a bilateral contract), the model must produce well-balanced supply and demand as observed in normal times and allow for imbalances in times of distress.

constraints are constraints set by the regulator. Most regulatory constraints are specific to a type of institution; banks face different regulatory constraints than insurers, for example. The models capture a subset of the regulatory constraints that banks face[70] and do not capture the regulatory constraints that non-banks confront.[71]

Contractual constraints arise out of contractual obligations. Because, as noted before, each financial institution holds a unique collection of contracts, the contractual constraints of each institution are unique too. Each model covers repayment obligations, because they capture (un)secured funding contracts. The ABMFV also considers margin call obligations as part of the secured funding contracts. Because each of the macroprudential stress tests discussed above only captures a subset of the relevant contracts, the contractual constraints they capture are incomplete as well. Banks, for example, typically hold derivatives contracts (e.g. credit default swaps) that can give liquidity shocks that may foster pre- or post-default contagion.

Market-based constraints (commonly referred to as 'market discipline') are those that are enforced by market participants. Sometimes, market participants set higher standards than regulators do; a bank might, for example, be cut off from funding markets because its leverage is judged to be too high, even though it still meets the regulatory leverage requirements. In this case, the market constraint could be formalized as a leverage constraint that is stricter than the regulatory leverage constraint. The most relevant market-based constraint, which entails that creditors run if the liquidity and/or solvency characteristics of a bank are sufficiently negative, is captured by all three models.[72]

Finally, internal risk limits are set by the financial institutions themselves, as part of their risk-management practices. An example could be a value-at-risk (VaR) constraint on a portfolio.[73]

Taken together, these constraints (and their various interactions) can drive an institution's behavior, especially under stress. First, institutions may act in a precautionary manner to avoid breaching constraints in order to avoid defaults. These actions, which are often prudent for each institution separately, may contribute to pre-default contagion (e.g. firesales in order to meet payment obligations). Second, institutions may fail to avoid breaching a constraint and default, which then leads to post-default con-

[70]The RAMSI model and the ABMFV consider one regulatory constraint for banks, an (unweighted) leverage ratio and a risk-weighted leverage ratio respectively. The theoretical model of Anand et al. (2015) that underpins the MFRAF stress test does not consider a regulatory leverage constraint. (But the MFRAF (IMF, 2014) stress test presumably does.) Banks default when they no longer meet their minimum (risk-weighted) leverage constraint. The models do not capture other regulatory constraints of banks that may affect financial stability, such as liquidity constraints (e.g. the liquidity coverage ratio) for banks.

[71]A few important (solvency) constraints that non-banks face have been covered in Section 7.2. As has been discussed in that section, insurers face a Solvency II constraint, pension funds face a coverage ratio, and CCPs must fulfill default fund requirements.

[72]The models consider the creditors to be exogenous to the system. A more realistic approach would be to make these creditors endogenous to the system. That way, cash does not leave the system but ends up in an institution's pockets.

[73]None of the macroprudential models discussed here consider internal risk limits.

tagion (e.g. due to exposure losses). Given their vital role in driving interactions under stressed conditions, it is important to consider whether the constraints included in a given stress test model represent those most relevant to the description of the system or sector that is being studied. More specifically, for any given institution the nature of its contribution to contagion will be critically determined by the set of constraints it faces. In sum, a failure to consider the relevant constraints makes it unlikely that the stress test model will correctly identify which channels of contagion operate and which institution are affected (Cetina et al., 2015).

8.2.5 Behavior

Behavior drives the dynamics of the financial system and the evolution of the multi-layered network representation thereof. It therefore critically affects the inherent stability of the financial system and can be an important driver of contagion. Behavior of institutions is typically not known and must thus be reasonably estimated.

Institutions can *affect the state of the system* when they default (*i.e. post-default*) or when they are still alive (*i.e. pre-default*). When institutions are alive they act for two reasons: *to fulfill objectives* (e.g. seek profits) and *to avoid default*.[74] When institutions default either through insolvency (i.e. breaching *regulatory constraints*) or illiquidity (i.e. when an institution does not meet its *contractual obligations*) they also affect the system. Through these pre- and post-default actions institutions can contribute to contagion.[75]

The three macroprudential stress testing models capture the critical drivers of financial stability dynamics to various degrees. The ABMFV most realistically simulates a financial market and its (contagious) dynamics. It captures that institutions can contribute to 'pre-default contagion' when they aim to avoid default,[76] but can also contribute to 'post-default contagion' once they have defaulted.[77] In addition, the ABMFV captures normal-time behavior, presumably to ensure that contagion is not overestimated (e.g. some may be willing to buy when others are forced to sell). The MFRAF and the Aikman et al. (2009), Alessandri et al. (2009) versions of the RAMSI model[78] assume that institutions are largely *passive*: they do not act until they default (only when they do default, institutions affect the system). Barring any

[74]Note that many financial stability models (see Section 5) abstract away from profit-seeking behavior. This may be a reasonable abstraction because in times of distress behavior is typically mostly driven by the objective to avoid default. However, by doing so, these models might overestimate contagion. As in crisis times, the institutions that are not under pressure (e.g. do not experience binding constraints) can stabilize the market.

[75]Or act as stabilizers.

[76]For example, institutions, which must meet the contractual obligation to repay a loan, may engage in fire sales to do so.

[77]To capture the contagion consequences that may ensue following a default, the relevant aspects of a default procedure must be modeled. For example, models must not only capture one contagion effect (e.g. exposure loss contagion), but various relevant contagion effects (e.g. including common asset holding contagion, etc.).

[78]The Kapadia et al. (2013) version of the RAMSI model does capture pre-default contagion.

defaults, these models thus only capture dynamics to a limited extent. By not capturing pre-default contagion, these may significantly underestimate losses (see e.g. Bardoscia et al., 2017). Table 2 summarizes the implementation of behavior in the three macroprudential stress testing models.

8.3 THE CALIBRATION CHALLENGE

Calibration is a process to ensure that the estimated parameters of a model match existing data (Turrell, 2016). The design of stress tests can make calibration easier or harder. Calibration is made easier when models are designed so as to either avoid free parameters entirely (by initializing all components to data), or to set up the model so that its parameters can be measured independently on input data rather than based on target data (the data that one wants to fit).[79] In general, it is therefore useful to design the stress test model so as to closely fit the market infrastructure, because this allows regulators to collect data on each component and then put it together – for example by using the five building blocks used above. A stress test that relies heavily on latent parameters[80] will require more assumptions will therefore introduce more uncertainty.

When using the five building blocks (institutions, contracts, markets, constraints and behavior) it becomes clear that the first four can (to a large extent) be data-driven.[81] Balance sheet data is already collected by regulators, although not always on a contractual level.[82] This latter step is important, given the importance of contractual constraints in driving specific contagion dynamics (as discussed in Section 8.2). Regulators increasingly recognize this, and have started collecting contract-level data for contractual types considered especially important to financial stability (e.g. Abad et al., 2016).[83] Because data gaps still persist, the (multi-layered) networks that make up system-wide financial models cannot be completely calibrated to data. In such cases, network reconstruction techniques can generate 'realistic' networks based on the known information (e.g. Anand et al., 2018).

Markets are complicated, because the market mechanism has to be modeled correctly. So far, most stress tests abstract away from market infrastructure and instead rely on price impact functions to move prices. The problem is that these functions are

[79] In other words, (loosely speaking) the more a model can be a one-to-one fit with the available data, the better.

[80] Latent parameters are unobservable and can – at best – only be calibrated by fitting model outputs to data.

[81] Provided, of course, that such data is indeed collected. On this front, more progress is desirable (see Section 9).

[82] Examples of data sources that capture contract level data are: the securities holding database and the trade reporting of derivatives (under EMIR). See https://sdw.ecb.europa.eu/browseExplanation.do?node=9691594 and https://www.esma.europa.eu/policy-rules/post-trading/trade-reporting, respectively.

[83] Regulators should collect data on three dimensions of contracts: counterparties, valuation methods (and inputs), and contractual obligations (and inputs). The first is needed stipulate interconnections, the second is needed to understand contagion dynamics arising through valuation and liquidity shocks.

driven by 'market depth', which is a latent variable and can only be approximated with data about the daily volume of trades and the volatility in the asset class (e.g. Cont and Schaanning, 2017).

The most relevant constraints that drive dynamics, regulatory and contractual constraints could in principle be calibrated to data,[84] but market constraints have to be inferred and can easily change over time. Internal risk limits can change too and are often proprietary.[85]

The building block for which calibration is most complicated is the last one, behavior. Although behavioral assumptions can be informed by supervisory data and surveys (Bookstaber, 2017), and perhaps even inferred using machine learning techniques, it is bound to change when a new type of crisis hits.[86] That is why it is important that even if the first four building blocks are (relatively closely) calibrated, the resulting stress tests are explicitly conditional on the behavioral assumption chosen.[87] It is then possible to change that assumption and run parameter sweeps to get a sense of the effect size of that behavioral assumption. In other words, stress tests might not predict exactly what will happen in a given scenario, but can explore directionally what might happen under a set of what-if scenarios.[88] This can be useful to (1) assess the level of systemic risk; (2) identify potential vulnerabilities in the system; and (3) evaluate the effectiveness of policies designed to mitigate systemic risk.

Calibrating the three macroprudential stress tests discussed above is challenging. The ABMFV may be the easiest to calibrate,[89] as it captures a realistic market infrastructure (consisting of the first four building blocks). Only behavioral parameters (related to block five) are uncertain, as is the usual case in a social system. It models realistic behavior, but it does not investigate how different behavioral parameters give different dynamics. The RAMSI model relies heavily on a careful calibration of the initial shock and the effect this has on the balance sheet.[90] That emphasis on calibrating the initial shock may distract from the most relevant function of any macroprudential model; its evaluation of a financial system's capacity for

[84] Regulatory constraints are typically publicly known. Most relevant contractual constraints can be known by a regulator that collects contract-level data, and includes the third dimension of contracts: contractual obligations.

[85] It is unclear whether, and how, these constraints will be enforced in times of crisis. It may be, for example, that a financial institution loosens its internal risk limits to avoid fire sales, or that contractual counterparties agree to suspend their obligations because strict enforcement would be costly for both.

[86] Based on a range of observations, representative behavior for certain types of institutions, and for various circumstances, can be inferred. Many models in the ABM literature capture somewhat realistic behavior (e.g. Kok and Montagna, 2013 discussed in Section 5.5).

[87] This is not a problem particular to macroprudential stress tests. Anyone who models a social system – rather than a physical system (where the dynamics are governed by physical laws) faces this problem.

[88] As far as we are aware, none of the financial stability models and macroprudential stress tests focus on prediction.

[89] Although, in practice, this model has not yet been calibrated.

[90] As noted, this is the microprudential component of the model.

shock amplification and endogenous dynamics.[91] Finally, Anand et al. (2015) is a more traditional model that seems much harder to calibrate to data. Its structure is more rigid, and therefore harder to map onto a given market structure. For example, the model's dynamics revolve around a two-period funding contract (and an abstract asset market), but it is not clear how the model can accommodate the variety of contracts that exist in the financial system. On the other hand, Anand et al. (2015) has been used in a data-driven MFRAF stress test.

8.4 STRENGTHS AND WEAKNESSES OF THE CURRENT MACROPRUDENTIAL STRESS TESTS

Macroprudential stress tests are strongly complementary to microprudential stress tests, because they allow regulators to assess the resilience of the financial system as a whole (or a larger subset of it) rather than that of individual financial institutions. The current macroprudential stress tests have three related strengths.

First, they provide insights into the interlinkages between financial institutions, mapping out how financial shocks transmit through individual balance sheets and affect other institutions. The data-driven methodology to establish the model setup (as well as the subsequent calibration) provide a promising avenue for future stress tests, but also for further data-driven research into the structure of the financial system (Aikman et al., 2009).

Second, they capture the interactions between various financial institutions and contagion channels that can drive distress, and therefore capture (some of) the feedback effects that characterize the complex nature of the financial system (see Section 4.1 and Section 5). Especially the ABM for Financial Vulnerabilities makes an important contribution by including heterogeneous financial institutions, which is key to allow for emergent phenomena (Bookstaber, 2017).

Third, in addition to capturing solvency risk, or separately investigating solvency and liquidity risk, the current macroprudential stress tests capture funding liquidity risk and the interactions between solvency and liquidity (the interaction between contagion channels has been discussed in Section 5.5). The RAMSI model, for example, not only considers defaults through insolvency, but also through illiquidity, and takes their interaction into account. In case of the MFRAF, a particular strength is that market risk and funding liquidity are endogenously determined. Market risk is based on the degree of adverse selection. Because of asymmetric information, investors offer banks a pooling price for their assets. The pooling price (and hence the market liquidity) lowers if investors become more pessimistic and the quality of the assets is

[91] In addition, the RAMSI model is a one-size fits all model. The aim seemed to be to flexibly use this model for multiple analyses and macroprudential policy purposes. But this was inhibited by rigid outfit of inflexible code and entangled parts in which it was dressed. In such a model it is harder to disentangle the effects of various components of the model, and such a model is easily prone to becoming a 'black box'. That is why we propose that, as a matter of model design, it is preferable to create a *modular* and *transparent* plug-and-play model, where (building block) components can be *flexibly* added and removed.

lower. Funding liquidity risk is determined as a function of the bank's credit and market losses (based on general market confidence, and thus as a function of information contagion), its funding composition and maturity profile, and concerns that creditors may have over its future solvency.

Despite these strengths, there is substantial scope for improvement. First, most macroprudential stress tests only cover banks and their creditors, and therefore fail to capture interactions with non-banks that make up a substantial part of the financial system. Non-banks have played an important role in amplifying distress to the banking sector during the 2007–2009 financial crisis (Bernanke, 2015). Therefore, failing to capture non-banks does not just exclude many institutions from the analysis, but also leaves regulators less well-equipped to understand the resilience of the subset of financial institutions they do study. The ABM for Financial Vulnerabilities is an exception, since it does include multiple types of financial institutions, but contrary to the RAMSI and the MFRAF models it is not used as a regulatory stress test.

Second, and relatedly, most macroprudential stress tests capture only a few types of interconnections, even though it is clear that the multiplicity of channels and interconnections between financial institutions plays a critical role in spreading distress (Brunnermeier, 2008) (see also Section 5.5). Notable examples of such contractual linkages include securitized products and credit default swaps.

Third, most current macroprudential stress tests only capture post-default contagion. However, in financial crises pre-default contagion is rampant, often resulting from actions that are prudent from a firm-specific risk-management perspective, but destabilizing from a system-wide perspective. A bank, for example, might engage in precautionary de-leveraging to avoid insolvency (i.e. breaking a leverage constraint), which can add to further negative price spirals. Not capturing such dynamics implies that the total size of contagion, as well as the timing of contagion, is misunderstood.

These three areas of improvement essentially come down to the same point: the current macroprudential stress tests insufficiently capture the diversity of agents and interactions that make up the financial system, and therefore do not do justice to the complex nature of the financial system (or, for that matter, to the insights of the heterogeneous agent model literature, see Sections 4 and 5). One of the important challenges is to devise a modeling strategy that can capture these various effects, and the ABM for Financial Vulnerabilities offers a promising start; the model could easily be extended to capture more types of financial institutions (e.g. central clearing parties, pension funds), financial contracts (e.g. derivative contracts, securitized products), and constraints that drive behavior under stressed circumstances (Cetina et al., 2015; Farmer et al., 2018).

Finally, macroprudential stress tests must be more data-driven[92] and more carefully calibrated to be credible. Thus, suitably designed system-wide stress tests are enabled to become more credible as regulators collect better (contract-level) data.

[92]This depends on data availability.

9 THE FUTURE OF SYSTEM-WIDE STRESS TESTS

So far, we have spoken largely about *what* is. When thinking about what *should be*, we start by setting an overarching objective: to study systemic risk in the financial system. Such risk would not exist if firms operated in isolation, so adopting a system-wide perspective that takes account of the heterogeneity of the agents that inhabit it, as well as their interconnectedness and interactions (see Section 4 and Section 5), is critical. This view is gaining popularity among central bankers. Alex Brazier, head of financial stability at the Bank of England, recently made a statement that aligns with our observation (made in Section 1) that the economy is a complex system.[93] Brazier warned that a salient principle for macroprudential policy is to realize that 'the system is not the sum of its parts'. Instead, he emphasized, 'feedback loops within the system mean that the entities in the system can be individually resilient, but still collectively overwhelmed by the stress scenario'. Brazier related this statement explicitly to the stress tests, suggesting that these tools should be developed so that they can take that system-wide view. We agree (Farmer et al., 2018), but also observe that current macroprudential stress tests are not yet 'system-wide'. What should a genuine system-wide stress test be able to do?

System-wide stress tests[94] serve at least three important goals: to monitor financial stability, identify vulnerabilities in the financial system, and evaluate policies designed to mitigate systemic risk. The first, monitoring financial stability, involves developing metrics that would allow regulators to track whether systemic risks are building up over time, and to have early-warning indicators to ensure that they can intervene in a timely manner.

The second, identifying vulnerabilities in the financial system, implies that stress tests should enable regulators to become aware of structural deficiencies in the financial system that render it vulnerable to systemic risk. Another way of phrasing the same point would be to say that it should identify sources of systemic risk, the factors that contribute to such risk, and the relative importance of those factors. For example, regulators should be able to analyze the network structure of the financial system (Cont et al., 2010; e Santos et al., 2010; Battiston et al., 2012; Caccioli et al., 2014; Acemoglu et al., 2015), evaluate asset-holding patterns and concentration risk, identify systemically important nodes (Battiston et al., 2012), and examine the maturity structure and leverage of a financial institution's balance sheet (Puhr et al., 2003; Hirtle and Lehnert, 2014).

The third objective of a system-wide stress test would be that it can evaluate policies designed to mitigate systemic risk. In part, this objective touches on the concerns related to *microprudential* policies.[95] Such policies, meant to enhance the

[93] A substantial body of research using network theory to study financial systems finds emergent properties at the system-level which arise out of interactions between agents, see e.g. Battiston et al. (2007).

[94] Note that our conception of a 'stress test' is broader than the one commonly used; when we describe a 'system-wide stress test', we are not merely referring to the regulatory tool, but also to the underlying models that enable it.

[95] In other words, risk regulation itself can cause systemic risk.

resilience of individual institutions, can increase the fragility of the system in times of crisis when these requirements have procyclical effects (Danıelsson et al., 2004; Aymanns and Farmer, 2015).[96] At the same time, the objective to be able to evaluate the system-wide effects of proposed policies recognizes the significant design challenges associated with the development of *macroprudential* policies. To evaluate their efficacy *ex ante* is a significant challenge, and one that by definition requires a system-wide evaluation of their impact (see Armour et al., 2016). The interaction of multiple risk management policies, each of which would be beneficial on its own, may combine to produce effects that are undesirable. A system-wide stress testing model should be able to evaluate this, even if not in point-estimate terms.[97]

An example that highlights the potential for policies to pro-actively dramatically reduce systemic risk is provided by the work of Poledna and Thurner (2016). They use the debt-rank methodology of Battiston et al. (2012) to quantify the marginal systemic risk contribution of a given transaction, in this case a potential new loan. They then tax individual transactions according to that transaction's marginal contribution to systemic risk. In an agent-based simulation of the economy they find that this tax causes the agents to alter their transactions to re-organize the network and drastically decrease systemic risk at little cost. They demonstrate that this is far more effective than a Tobin (transaction) tax, which is both ineffectual and has substantial and potentially detrimental side effects. More generally, agent-based models have the advantage for policy evaluation that it is easy to change policies and explore their effects, though of course here one must work to properly take into account the Lucas critique[98] (Turrell, 2016; Farmer et al., 2018).

Before system-wide stress tests can credibly serve these important goals, the frontiers of financial stability models have to be pushed. One of the frontiers of financial stability modeling is to better understand the effect of interacting channels of contagion, and more generally, of multi-layered networks (like the ones used in Kok and Montagna, 2013). Setting up a system-wide stress test using multi-layered networks is useful because it allows for the representation of different types of relationships between various agents. That in turn allows for the interaction between different contractual types, corresponding to different layers of the model, enabling a richer set of contagion and amplification mechanisms (see Poledna et al., 2015, discussed in Section 5.5). Using a fine grained and comprehensive dataset Poledna et al. (2015) quantify the daily contribution to systemic risk from four layers of the Mexican banking system from 2007 to 2013. They find that focusing on a single layer underestimates the total systemic risk by up to 90%. A lingering question is how the

[96] As Andrew Crockett of the BIS observed as early as in 2000: 'actions that may seem desirable from the perspective of individual institutions may result in unwelcome system outcomes' (Crockett, 2000).

[97] Instead, the model may produce stylized facts, which help policymakers evaluate whether the policy is directionally efficacious (Haldane and Turrell, 2018). That way, it could serve as a laboratory for policy experiments.

[98] For example, conditional on a particular calibration for a proposed macroprudential policy.

interaction between the layers is modeled. Often, systemic risk or contagion estimates in each layer are simply added up (albeit jointly considered), so that richness of the interaction effects between contractual types is ignored[99] – despite their importance to the overall contagion dynamics.

Another frontier is more realistic modeling of agent behavior. In most macroprudential stress tests agents are naive and often are simply static, so that they do not take precautionary action and often take no action at all, even when catastrophic events occur. There are good reasons why agents could be modeled using fixed heuristics (Bookstaber, 2017), whether geared toward leverage targeting or to avoiding default (Kok and Montagna, 2013; Bookstaber et al., 2014b). But operating on fixed heuristics is also a limiting factor. Lo (2017), for example, has noted that some of the most interesting and salient behavioral phenomena – which translate into the dynamics of financial markets – result from the updating of behaviors by agents in response to changing circumstances. Simple learning protocols such as reinforcement learning and switching between heterogeneous expectations (e.g. Brock and Hommes, 1997; Brock and Hommes, 1998; Hommes et al., 2017), that allow agents to display goal-seeking, optimizing behavior while learning from their past interactions, have been shown to be effective in explaining behavioral experiments.

Calibration and validation remain key challenges for heterogeneous agent modeling. Methodological advances are required to provide better solutions to this problem and to convince policymakers that system-wide stress testing models are reliable. A key aspect of this is creating fine grained data sets.[100] Heterogeneous ABMs typically model the behavior of agents at a detailed level, and with appropriate microdata, they can also be calibrated and validated at this level.[101] This potentially offers a huge advantage over aggregated models that can only be calibrated and validated at an aggregate level. It is also essential that fine-grained data in anonymized form be made available to academics. Such models need to be designed to be more modular and flexible, so that it is easy to test alternative hypotheses and understand the key factors that drive observed behavior, and so that they can be easily adapted to new situations.

[99] For example, common asset holding interacts with repo contracts – when collateral price falls this can lead to margin calls.

[100] An encouraging development in this respect is that central banks and other regulators have started to collect high-quality and fine-grained data. Perhaps the best example is the 'trade repository data set'. Art. 9 of the European Markets Infrastructure Regulation (EMIR) requires counterparties resident in the EU (including central clearing counterparties) to report the details of new and outstanding derivatives transactions to trade repositories on a daily basis. Sufficient information for each contract is gathered to determine the counterparties, the valuation and contractual obligations of a contract. To do so, around 85 variables are reported for each transaction. Such comprehensive reporting under EMIR implies huge data volumes. For a description of this data set, see Abad et al. (2016). On the basis of this data, it is possible to initialize derivatives networks and study contagion dynamics operating on that network.

[101] A good example is the housing ABM developed by Baptista et al. (2016).

10 CONCLUSION

Computational agent-based models provide a useful complement to more traditional equilibrium based methods. They have already been shown to be essential for understanding the dynamics of systemic risk and for investigating the network properties of the financial system. Their role is likely to become even more important in the future, as increasingly comprehensive fine-grained data becomes available, making it possible to carefully calibrate such models so that they can yield more quantitative conclusions. Due to the inherent complexity of the financial system, and in particular its nonlinear feedback loops, analytic methods are unlikely to be sufficient.

We expect that computational and simulation methods will soon begin to go beyond hard wired behavioral rules and move increasingly toward myopic optimization. Models of boundedly rational heterogeneous agents, who learn and adapt their behavior in response to observed market realizations and newly adopted policies, withstand the Lucas critique. Behavioral economists have documented more and more situations in which people are not fully rational, emphasizing the obvious point that realistic behavior lies somewhere between full rationality and zero intelligence. Computational models offer the possibility of implementing realistic levels of strategic behavior, while allowing one to model the complex institutional structure of the financial system. We think that computational models will play an expanding role for understanding financial stability and systemic risk.

REFERENCES

Abad, J., Aldasoro, I., Aymanns, C., D'Errico, M., Rousová, L.F., Hoffmann, P., Langfield, S., Neychev, M., Roukny, T., 2016. Shedding light on dark markets: first insights from the new EU-wide OTC derivatives dataset. ESRB Occasional Paper Series 10, 1–32.

Accenture, 2015. Federal Reserve Involvement in Insurance Industry Capital Standards. Technical report. Accenture.

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

Acharya, V., Engle, R., Pierret, D., 2014. Testing macroprudential stress tests: the risk of regulatory risk weights. Journal of Monetary Economics 65, 36–53.

Acharya, V.V., Yorulmazer, T., 2008. Information contagion and bank herding. Journal of Money, Credit, and Banking 40 (1), 215–231.

Admati, A., Hellwig, M., 2014. The Bankers' New Clothes: What's Wrong with Banking and What to Do About It. Princeton University Press.

Adrian, T., Shin, H.S., 2010. Liquidity and leverage. Journal of Financial Intermediation 19 (3), 418–437.

Afonso, G., Kovner, A., Schoar, A., 2011. Stressed, not frozen: the federal funds market in the financial crisis. The Journal of Finance 66 (4), 1109–1139.

Aikman, D., Alessandri, P., Eklund, B., Gai, P., Kapadia, S., Martin, E., Mora, N., Sterne, G., Willison, M., 2009. Funding Liquidity Risk in a Quantitative Model of Systemic Stability. Bank of England Working Paper No. 372.

Akerlof, G.A., 1970. The market for "lemons": quality uncertainty and the market mechanism. The Quarterly Journal of Economics, 488–500.

Alessandri, P., Gai, P., Kapadia, S., Mora, N., Puhr, C., et al., 2009. Towards a framework for quantifying systemic stability. International Journal of Central Banking 5 (3), 47–81.

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

Amini, H., Cont, R., Minca, A., 2016. Resilience to contagion in financial networks. Mathematical Finance 26, 329–365.

Anand, K., Bédard-Pagé, G., Traclet, V., 2014. Stress testing the Canadian banking system: a system-wide approach. Financial System Review 61.

Anand, K., Gauthier, C., Souissi, M., 2015. Quantifying Contagion Risk in Funding Markets: A Model-Based Stress-Testing Approach. Technical report.

Anand, K., van Lelyveld, I., Banai, Á., Friedrich, S., Garratt, R., Hałaj, G., Fique, J., Hansen, I., Jaramillo, S.M., Lee, H., et al., 2018. The missing links: a global study on uncovering financial network structures from partial data. Journal of Financial Stability 35, 107–119.

Andersen, T.G., Bollerslev, T., Christoffersen, P.F., Diebold, F.X., 2006. Volatility and correlation forecasting. In: Handbook of Economic Forecasting, vol. 1, pp. 777–878.

Anderson, P.W., et al., 1972. More is different. Science 177 (4047), 393–396.

Anderson, R.W., 2016. Stress testing and macroprudential regulation: a transatlantic assessment. In: Stress Testing and Macroprudential Regulation, p. 1.

Arinaminpathy, N., Kapadia, S., May, R.M., 2012. Size and complexity in model financial systems. Proceedings of the National Academy of Sciences 109 (45), 18338–18343.

Armour, J., Awrey, D., Davies, P., Enriques, L., Gordon, J.N., Mayer, C., Payne, J., 2016. Principles of Financial Regulation. Oxford University Press.

Arnold, M., Jenkins, P., 2016. EBA boss says recapitalization of banks has been 'successful'. https://www.ft.com/content/5496cc86-5417-11e6-befd-2fc0c26b3c60.

Aymanns, C., Caccioli, F., Farmer, J.D., Tan, V.W., 2016. Taming the Basel leverage cycle. Journal of Financial Stability 27, 263–277.

Aymanns, C., Farmer, J.D., 2015. The dynamics of the leverage cycle. Journal of Economic Dynamics and Control 50, 155–179.

Aymanns, C., Georg, C.-P., Bundesbank, D., Golub, B., 2017. Illiquidity spirals in coupled over-the-counter markets. Posted on SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2881814.

Baptista, R., Farmer, J.D., Hinterschweiger, M., Low, K., Tang, D., Uluc, A., 2016. Macroprudential Policy in an Agent-Based Model of the UK Housing Market. Bank of England Working Paper No. 619.

Baranova, Y., Coen, J., Lowe, P., Noss, J., Silvestri, L., 2017. Simulating Stress Across the Financial System: The Resilience of Corporate Bond Markets and the Role of Investment Funds. Bank of England Financial Stability Paper 42.

Bardoscia, M., Barucca, P., Brinley, A., Hill, J., 2017. The Decline of Solvency Contagion Risk. Bank of England Staff Working Paper No. 662.

Battiston, S., Farmer, J.D., Flache, A., Garlaschelli, D., Haldane, A.G., Heesterbeek, H., Hommes, C., Jaeger, C., May, R., Scheffer, M., 2016. Complexity theory and financial regulation. Science 351 (6275), 818–819.

Battiston, S., Gatti, D.D., Gallegati, M., Greenwald, B., Stiglitz, J.E., 2007. Credit chains and bankruptcy propagation in production networks. Journal of Economic Dynamics and Control 31 (6), 2061–2084.

Battiston, S., Puliga, M., Kaushik, R., Tasca, P., Caldarelli, G., 2012. Debtrank: too central to fail? Financial networks, the Fed and systemic risk. Scientific Reports 2.

Behn, M., Haselmann, R.F., Vig, V., 2016. The Limits of Model-Based Regulation. ECB Working Paper No. 1928.

Bernanke, B., 2013. Stress testing banks: what have we learned? https://www.federalreserve.gov/newsevents/speech/bernanke20130408a.htm.

Bernanke, B., 2015. The Courage to Act: A Memoir of the Crisis and Its Aftermath. W.W. Norton and Company.

BIS, 2015. Revisions to the Standardised Approach for Credit Risk. Technical report. Bank for International Settlements, Basel Committee on Banking Supervision.

BoC, 2012. Understanding Systemic Risk in the Banking Sector: A Macrofinancial Risk Assessment Framework. Technical report. Bank of Canada.

BoC, 2014. Financial System Review. Technical report. Bank of Canada.

BoE, 2014. Stress Testing the UK Banking System: 2014 Results. Technical report. Bank of England.

BoE, 2015. General Insurance Stress Test 2015 – Scenario Specification, Guidelines and Instructions. Technical report. BoE.

BoE, 2016. Stress Testing the UK Banking System: Key Elements of the 2016 Stress Test. Technical report. Bank of England.

Bookstaber, R., 2017. The End of Theory: Financial Crises, the Failure of Economics, and the Sweep of Human Interaction. Princeton University Press.

Bookstaber, R., Cetina, J., Feldberg, G., Flood, M., Glasserman, P., 2014a. Stress tests to promote financial stability: assessing progress and looking to the future. Journal of Risk Management in Financial Institutions 7 (1), 16–25.

Bookstaber, R., Kirman, A., 2018. Modeling a heterogeneous world. In: Handbook of Computational Economics, vol. 4. Elsevier, pp. 769–795 (this Handbook).

Bookstaber, R., Paddrik, M., Tivnan, B., 2014b. An agent-based model for financial vulnerability. Office of Financial Research Working Paper Series 14 (05).

Bookstaber, R.M., Paddrik, M.E., 2015. An Agent-Based Model for Crisis Liquidity Dynamics. Office of Financial Research (OFR) Working Paper 15-18.

Boss, M., Krenn, G., Puhr, C., Schwaiger, M.S., et al., 2007. Stress Testing the Exposure of Austrian Banks in Central and Eastern Europe. Oesterreichische Nationalbank (ONB) Financial Stability Report 13.

Bouchaud, J.-P., Farmer, J.D., Lillo, F., 2008. How markets slowly digest changes in supply and demand. arXiv:0809.0822.

Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. Econometrica, 1059–1095.

Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economic Dynamics & Control 22 (8), 1235–1274.

Brunnermeier, M.K., 2008. Deciphering the Liquidity and Credit Crunch 2007–08. Technical report. National Bureau of Economic Research.

Brunnermeier, M.K., Pedersen, L.H., 2009. Market liquidity and funding liquidity. The Review of Financial Studies 22 (6), 2201–2238.

Brunnermeier, M.K., Sannikov, Y., 2014. A macroeconomic model with a financial sector. The American Economic Review 104 (2), 379–421.

Burrows, O., Learmonth, D., McKeown, J., Williams, R., 2012. RAMSI: a top-down stress-testing model developed at the Bank of England. Bank of England Quarterly Bulletin 52 (3), 204–212.

Burrows, O., Low, K., Cumming, F., 2015. Mapping the UK financial system. Bank of England Quarterly Bulletin 55 (2), 114–129.

Caccioli, F., Bouchaud, J.-P., Farmer, J.D., 2012a. Impact-adjusted valuation and the criticality of leverage. Risk, 74–77.

Caccioli, F., Catanach, T.A., Farmer, J.D., 2012b. Heterogeneity, correlations and financial contagion. Advances in Complex Systems 15 (Suppl. 02), 1250058.

Caccioli, F., Farmer, J.D., Foti, N., Rockmore, D., 2015. Overlapping portfolios, contagion, and financial stability. Journal of Economic Dynamics and Control 51, 50–63.

Caccioli, F., Shrestha, M., Moore, C., Farmer, J.D., 2014. Stability analysis of financial contagion due to overlapping portfolios. Journal of Banking & Finance 46, 233–245.

Candelon, B., Sy, A.N., 2015. How Did Markets React to Stress Tests? International Monetary Fund.

Capgemini, 2014. Basel III: Comparison of Standardized and Advanced Approaches. Technical report. Capgemini.

Capponi, A., Chen, P.-C., Yao, D.D., 2015. Liability concentration and systemic losses in financial networks. Operations Research 64 (5), 1121–1134.

Capponi, A., Cheng, W., Sethuraman, J., 2017. Clearinghouse default waterfalls: risk-sharing, incentives, and systemic risk. Available at SSRN: https://ssrn.com/abstract=2930099 or http://dx.doi.org/10.2139/ssrn.2930099.

Capponi, A., Larsson, M., 2015. Price contagion through balance sheet linkages. The Review of Asset Pricing Studies 5 (2), 227–253.

Cetina, J., Lelyveld, I., Anand, K., 2015. Making Supervisory Stress Tests More Macroprudential: Considering Liquidity and Solvency Interactions and Systemic Risk. BCBS Working Paper.

CFTC, 2016. Supervisory Stress Test of Clearinghouses. Technical report. Commodity Futures and Trading Commission.

Cifuentes, R., Ferrucci, G., Shin, H.S., 2005. Liquidity risk and contagion. Journal of the European Economic Association 3 (2–3), 556–566.

Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. Quantitative Finance 1 (2), 223–236.

Cont, R., 2015. The end of the waterfall: default resources of central counterparties. Journal of Risk Management in Financial Institutions 8 (4), 365–389.

Cont, R., Kokholm, T., 2014. Central clearing of OTC derivatives: bilateral vs multilateral netting. Statistics & Risk Modeling 31 (1), 3–22.

Cont, R., Moussa, A., Santos, E.B., 2010. Network structure and systemic risk in banking systems. Available at SSRN: https://ssrn.com/abstract=1733528 or http://dx.doi.org/10.2139/ssrn.1733528.

Cont, R., Schaanning, E.F., 2017. Fire sales, indirect contagion and systemic stress-testing. Available at SSRN: https://ssrn.com/abstract=2541114 or http://dx.doi.org/10.2139/ssrn.2541114.

Cont, R., Wagalath, L., 2013. Running for the exit: distressed selling and endogenous correlation in financial markets. Mathematical Finance 23 (4), 718–741.

Cont, R., Wagalath, L., 2016. Fire sales forensics: measuring endogenous risk. Mathematical Finance 26 (4), 835–866. https://doi.org/10.1111/mafi.12071.

Copeland, A., Martin, A., Walker, M., 2014. Repo runs: evidence from the tri-party repo market. The Journal of Finance 69, 2343–2380.

Crockett, A.D., 2000. Marrying the Micro- and Macro-Prudential Dimensions of Financial Stability. Technical report. Bank of International Settlements.

Danıelsson, J., Shin, H.S., Zigrand, J.-P., 2004. The impact of risk regulation on price dynamics. Journal of Banking & Finance 28 (5), 1069–1087.

Dees, S., Henry, J., 2017. Stress-test analytics for macroprudential purposes: introducing stamp. In: Satellite Models, p. 13.

Diamond, D.W., Dybvig, P.H., 1983. Bank runs, deposit insurance, and liquidity. Journal of Political Economy 91 (3), 401–419.

Dieci, R., He, T., 2018. Heterogeneous agent models in finance. In: Hommes, C., LeBaron, B. (Eds.), Handbook of Computational Economics, vol. 4. Elsevier, pp. 257–328 (this Handbook).

Duarte, F., Eisenbach, T.M., et al., 2015. Fire-Sale Spillovers and Systemic Risk. Federal Reserve Bank of New York Staff Report No. 645. October 2013, revised February 2015.

Duffie, D., Scheicher, M., Vuillemey, G., 2015. Central clearing and collateral demand. Journal of Financial Economics 116 (2), 237–256.

Duffie, D., Zhu, H., 2011. Does a central clearing counterparty reduce counterparty risk? The Review of Asset Pricing Studies 1 (1), 74–95.

e Santos, E.B., Cont, R., et al., 2010. The Brazilian Interbank Network Structure and Systemic Risk. Technical report.

EBA, 2017a. CEBS stress testing results. http://www.eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2009.

EBA, 2017b. EU-wide stress testing. http://www.eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing.

ECB, 2015. Report on Financial Structures. Technical report. European Central Bank.

EIOPA, 2016. Insurance Stress Test 2016 Technical Specifications. Technical report. European Insurance and Occupational Pensions Authority.

EIOPA, 2017. Occupational pensions stress test. https://eiopa.europa.eu/financial-stability-crisis-prevention/financial-stability/occupational-pensions-stress-test.

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

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

Erbenova, M., 2015. Germany Financial Sector Assessment Program. Technical report. International Monetary Fund.

ESMA, 2015. EU-Wide CCP Stress Test 2015. Technical report. European Securities and Markets Authority.

ESMA, 2017. OTC derivatives and clearing obligation. https://www.esma.europa.eu/regulation/post-trading/otc-derivatives-and-clearing-obligation.

ESRB, 2015. The ESRB handbook on operationalising macroprudential policy in the banking sector, macroprudential leverage ratios. https://www.esrb.europa.eu/pub/pdf/other/140303_esrb_handbook_mp.en.pdf.

ESRB, 2016. Adverse Macro-Financial Scenario for the EBA 2016 EU-Wide Bank Stress Testing Exercise. Technical report. European Systemic Risk Board.

EY, 2013. Dodd–Frank's Title VII – OTC derivatives reform. http://www.ey.com/Publication/vwLUAssets/Key_questions_board_members_should_ask_about_Title_VII/$FILE/Americas_FAAS_Dodd_Frank_derivatives_reform.pdf.

Farmer, D., Kleinnijenhuis, A., Nahai-Williamson, P., Tanin, R., Wetzer, T., 2018. A Nesting Model for System-Wide Stress Tests. Working paper.

Farmer, J.D., 2002. Market force, ecology and evolution. Journal of Industrial and Corporate Change 11 (5), 895–953.

Farmer, J.D., 2012. Economics needs to treat the economy as a complex system. In: Papers for the INET Conference. Rethinking Economics and Politics, vol. 14.

FED, 2016. Dodd–Frank Act Stress Test 2016: Supervisory Stress Test Methodology and Results. Technical report. Board of Governors of the Federal Reserve.

FED, 2017a. Stress tests and capital planning: comprehensive capital analysis and review. https://www.federalreserve.gov/supervisionreg/ccar.htm.

FED, 2017b. Stress tests and capital planning: Dodd–Frank act stress tests. https://www.federalreserve.gov/supervisionreg/ccar.htm.

Fender, I., Lewrick, U., 2015. Calibrating the leverage ratio. BIS Quarterly Review. https://EconPapers.repec.org/RePEc:bis:bisqtr:1512f.

Feroli, M., Kashyap, A.K., Schoenholtz, K.L., Shin, H.S., 2014. Market tantrums and monetary policy. Chicago Booth Research Paper No. 14-09. Available at SSRN: https://ssrn.com/abstract=2409092 or http://dx.doi.org/10.2139/ssrn.2409092.

Foglia, A., 2009. Stress testing credit risk: a survey of authorities' approaches. International Journal of Central Banking 5 (3), 9–45.

Fostel, A., Geanakoplos, J., 2008. Leverage cycles and the anxious economy. The American Economic Review 98 (4), 1211–1244.

FSB, 2015. Global Shadow Banking Monitoring Report. Technical report.

Gai, P., Haldane, A., Kapadia, S., 2011. Complexity, concentration and contagion. Journal of Monetary Economics 58 (5), 453–470.

Gai, P., Kapadia, S., 2010. Contagion in financial networks. Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, rspa20090410.

Gauthier, C., Lehar, A., Souissi, M., 2012. Macroprudential capital requirements and systemic risk. Journal of Financial Intermediation 21 (4), 594–618.

Gauthier, C., Souissi, M., Liu, X., et al., 2014. Introducing funding liquidity risk in a macro stress-testing framework. International Journal of Central Banking 10 (4), 105–141.

Geanakoplos, J., 2010. The leverage cycle. NBER Macroeconomics Annual 24 (1), 1–66.

Geithner, T.F., 2014. Stress test: reflections on the financial crisis. Business Economics 49 (3), 201–203.

Gennotte, G., Leland, H., 1990. Market liquidity, hedging, and crashes. The American Economic Review, 999–1021.

Georg, C.-P., 2013. The effect of the interbank network structure on contagion and common shocks. Journal of Banking & Finance 37 (7), 2216–2228.

Glasserman, P., Tangirala, G., et al., 2015. Are the Federal Reserve's Stress Test Results Predictable? Technical report.

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

Gorton, G., Metrick, A., 2012. Securitized banking and the run on repo. Journal of Financial Economics 104 (3), 425–451.

Goyal, S., 2018. Heterogeneity and networks. In: Hommes, C., LeBaron, B. (Eds.), Handbook of Computational Economics, vol. 4. Elsevier, pp. 687–712 (this Handbook).

Greenwood, R., Landier, A., Thesmar, D., 2015. Vulnerable banks. Journal of Financial Economics 115 (3), 471–485.

Groendahl, B., 2015. Leverage ratio for banks can rise as high as 5%, BIS says. https://www.bloomberg.com/news/articles/2015-12-06/leverage-ratio-for-banks-can-be-raised-as-high-as-5-bis-says.

Gromb, D., Vayanos, D., 2002. Equilibrium and welfare in markets with financially constrained arbitrageurs. Journal of Financial Economics 66 (2), 361–407.

Haldane, A.G., Turrell, A.E., 2018. An interdisciplinary model for macroeconomics. Oxford Review of Economic Policy 34 (1–2), 219–251.

He, Z., Xiong, W., 2012. Dynamic debt runs. The Review of Financial Studies 25 (6), 1799–1843.

Hirtle, B., Kovner, A., Zeller, S., 2016. Are stress tests still informative. http://libertystreeteconomics.newyorkfed.org/2016/04/are-stress-tests-still-informative.html.

Hirtle, B., Lehnert, A., 2014. Supervisory Stress Tests. FRB of New York Staff Report 696.

Hommes, C.H., Jump, R., Levine, P., 2017. Internal rationality, heterogeneity, and complexity in the new Keynesian model. Available at SSRN: https://ssrn.com/abstract=3086851 or http://dx.doi.org/10.2139/ssrn.3086851.

Hommes, C.H., 2006. Heterogeneous agent models in economics and finance. In: Handbook of Computational Economics, vol. 2, pp. 1109–1186.

IMF, 2014. Canada Financial Sector Stability Assessment. Technical report. International Monetary Fund.

Ionescu, L., Yermon, J., 2014. Stress Testing and Scenario Analysis of Pension Plans. Technical report. International Organisation of Pension Fund Supervisors.

Iori, G., Mantegna, R., 2018. Empirical analyses of networks in finance. In: Hommes, C., LeBaron, B. (Eds.), Handbook of Computational Economics, vol. 4. Elsevier, pp. 637–685 (this Handbook).

Iyer, R., Peydro, J.-L., 2011. Interbank contagion at work: evidence from a natural experiment. The Review of Financial Studies 24 (4), 1337–1377.

Jobst, A., Andreas, S.N.B.T., 2014. Macroprudential Solvency Stress Testing of the Insurance Sector. Technical report. IMF.

Kapadia, S., Drehmann, M., Elliott, J., Sterne, G., 2013. Liquidity risk, cash flow constraints, and systemic feedbacks. In: Quantifying Systemic Risk. University of Chicago Press, pp. 29–61.

Khandani, A.E., Lo, A.W., 2011. What happened to the quants in August 2007? Evidence from factors and transactions data. Journal of Financial Markets 14 (1), 1–46.

Kiyotaki, N., Moore, J., 1997. Credit cycles. Journal of Political Economy 105 (2), 211–248.

Kok, C., Montagna, M., 2013. Multi-Layered Interbank Model for Assessing Systemic Risk. Technical report. European Central Bank.

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica, 1315–1335.

LeBaron, B., 2006. Agent-based computational finance. In: Handbook of Computational Economics, vol. 2, pp. 1187–1233.

Lo, A.W., 2005. Reconciling efficient markets with behavioral finance: the adaptive markets hypothesis. Journal of Investment Consulting, 21–44.

Lo, A.W., 2017. Adaptive Markets: Financial Evolution at the Speed of Thought. Princeton University Press.

Martin, A., Skeie, D., Von Thadden, E.-L., 2014. Repo runs. The Review of Financial Studies 27, 957–989.

May, R.M., Arinaminpathy, N., 2010. Systemic risk: the dynamics of model banking systems. Journal of the Royal Society Interface 7 (46), 823–838.

Mehrling, P., Pozsar, Z., Sweeney, J., Neilson, D.H., 2013. Bagehot was a shadow banker: shadow banking, central banking, and the future of global finance. Available at SSRN: https://ssrn.com/abstract=2232016 or http://dx.doi.org/10.2139/ssrn.2232016.

Moretti, M., Stolz, S.M., Swinburne, M., 2008. Stress Testing at the IMF. Citeseer.

Morris, S., Shin, H.S., 2001. Global games: theory and applications.

Morrison, A.D., Vasios, M., Wilson, M.I., Zikes, F., 2016. Identifying Contagion in a Banking Network. Saïd Business School WP 2016-37. Available at SSRN: https://ssrn.com/abstract=2987848 or https://doi.org/10.2139/ssrn.2987848.

Murphy, D., 2013. OTC Derivatives: Bilateral Trading and Central Clearing: An Introduction to Regulatory Policy, Market Impact and Systemic Risk. Springer.

Noonan, L., Binham, C., Shotter, J., 2016. Deutsche Bank received special treatment in EU stress tests. https://www.ft.com/content/44768ea8-8c71-11e6-8aa5-f79f5696c731.

OFR, 2014. Office of Financial Research: 2014 Annual Report. Technical report. Office of Financial Research.

Ong, L.L., Pazarbasioglu, C., 2014. Credibility and crisis stress testing. International Journal of Financial Studies 2 (1), 15–81.

Poledna, S., Molina-Borboa, J.L., Martínez-Jaramillo, S., Van Der Leij, M., Thurner, S., 2015. The multilayer network nature of systemic risk and its implications for the costs of financial crises. Journal of Financial Stability 20, 70–81.

Poledna, S., Thurner, S., 2016. Elimination of systemic risk in financial networks by means of a systemic risk transaction tax. Quantitative Finance 16 (10), 1599–1613.

Poledna, S., Thurner, S., Farmer, J.D., Geanakoplos, J., 2014. Leverage-induced systemic risk under Basel II and other credit risk policies. Journal of Banking & Finance 42, 199–212.

Pozsar, Z., 2013. Shadow banking and the global financial ecosystem. VoxEU 6.

Pozsar, Z., Adrian, T., Ashcraft, A., Boesky, H., 2010. Shadow banking. New York 458, 3–9.

Pozsar, Z., Singh, M., 2011. The Nonbank–Bank Nexus and the Shadow Banking System. IMF Working Paper WP/11/289.

Puhr, M.C., Santos, M.A., Schmieder, M.C., Neftci, S.N., Neudorfer, M.B., Schmitz, M.S.W., Hesse, H., 2003. Next Generation System-Wide Liquidity Stress Testing. No. 12–13. International Monetary Fund.

Quagliariello, M., 2009. Stress-Testing the Banking System: Methodologies and Applications. Cambridge University Press.

Robb, G., 2015. Fed adopting stress test for insurance companies: Fisher. http://www.marketwatch.com/story/fed-adopting-stress-test-for-insurance-companies-fischer-2015-06-24.

Rogers, L.C., Veraart, L.A., 2013. Failure and rescue in an interbank network. Management Science 59 (4), 882–898.

Scheule, H., Roesh, D., 2008. Stress Testing for Financial Institutions. Risk Publications.

Schuermann, T., 2014. Stress testing banks. International Journal of Forecasting 30 (3), 717–728.

Shleifer, A., Vishny, R.W., 1992. Liquidation values and debt capacity: a market equilibrium approach. The Journal of Finance 47 (4), 1343–1366.

Shleifer, A., Vishny, R.W., 1997. The limits of arbitrage. The Journal of Finance 52 (1), 35–55.

Siddique, A., Hasan, I., 2012. Stress Testing: Approaches, Methods and Applications. Risk Books.

Tarullo, D.K., 2014. Stress Testing After Five Years: A Speech at the Federal Reserve Third Annual Stress Test Modeling Symposium, Boston, Massachusetts. Technical report. Board of Governors of the Federal Reserve System.

Tarullo, D.K., 2016. Next steps in the evolution of stress testing. In: Remarks at the Yale University School of Management Leaders Forum (September 26, 2016).

Thurner, S., Farmer, J.D., Geanakoplos, J., 2012. Leverage causes fat tails and clustered volatility. Quantitative Finance 12 (5), 695–707.

Thurner, S., Poledna, S., 2013. Debtrank-transparency: controlling systemic risk in financial networks. arXiv:1301.6115.

Turrell, A., 2016. Agent-based models: understanding the economy from the bottom up.

Wackerbeck, P., Saxena, A., Crijns, J., Karsten, C., 2016. Stress testing: from regulatory burden to strategic capability. https://www.strategyand.pwc.com/reports/stress-testing.

Young, P., Glasserman, P., 2016. Contagion in financial networks. Journal of Economic Literature 54 (3), 779–831.