Zoltán Sass

Systemic Risk and Contagion in Interbank Networks

Institute of Finance

Supervisors: Péter Csóka, Ph.D., Péter Kerényi, Ph.D.

Corvinus University of Budapest

Doctoral School of Economics and Business Informatics

Systemic Risk and Contagion in Interbank Networks

From Ex-Post Measurement to Forward-Looking Valuation

Doctoral dissertation

Zoltán Sass

Budapest, 2026

Contents

List of Figures

List of Tables

Chapter 1

Introduction and motivation

1.1 The importance of measuring and managing systemic risk

The global financial crisis of 2007–2009 exposed critical vulnerabilities in the interconnected structure of modern financial systems. The collapse of Lehman Brothers in September 2008 triggered a cascade of failures and near-failures that propagated through the global financial network, ultimately requiring unprecedented government interventions to prevent a complete systemic collapse (Reinhart and Rogoff, 2009). This crisis demonstrated that the failure of individual institutions, previously considered manageable through traditional microprudential regulation, could threaten the stability of the entire financial system when those institutions were sufficiently interconnected (Haldane and May, 2011).

Systemic risk—the risk that the failure of one or more financial institutions triggers widespread failures throughout the system—has since become a central concern for regulators, policymakers, and academics alike. The Financial Stability Board, established in the aftermath of the crisis, explicitly identified the measurement and management of systemic risk as a key priority (Financial Stability Board, 2009). Post-crisis prudential policy responded with several tools that have different objectives. Basel III[1] strengthened capital and liquidity standards and introduced macroprudential[2] elements, but it is important to distinguish structural systemic-importance tools from cyclical macroprudential buffers (Basel Committee on Banking Supervision, 2011). The Countercyclical Capital Buffer (CCyB) is a time-varying buffer designed

[1] Basel III is the comprehensive set of post-crisis reform measures developed by the Basel Committee on Banking Supervision to strengthen the regulation, supervision, and risk management of banks.

[2] While microprudential regulation focuses on the soundness of individual institutions in isolation, macroprudential policy targets the stability of the financial system as a whole, addressing risks that arise from the collective behavior and interconnectedness of institutions.

to build system-wide resilience against cyclical risk associated with excessive aggregate credit growth, whereas Global Systemically Important Bank (G-SIB) and Other Systemically Important Institution (O-SII) frameworks focus on institutions whose distress would have especially large system-wide consequences (Basel Committee on Banking Supervision, 2010, 2012, 2013; European Banking Authority, 2014).

Tool / framework	**Primary objective**	**Institution-specific?**	**Time orientation**	**Closest relation to dissertation**
CCyB	Build system-wide resilience against cyclical macro-financial risk, especially excessive aggregate credit growth	No	Ex ante, time-varying	Context only
G-SIB / O-SII buffers	Increase resilience of structurally important institutions whose distress would have large system-wide consequences	Yes	Ex ante, structural	Closest official comparator for Chapter 2, but broader and more implementable
Supervisory stress testing	Evaluate resilience under adverse scenarios; quantify capital depletion and vulnerabilities	Usually institution and system	Ex ante, scenario-based	Practical analogue for the scenario-based contagion and loss-amplification logic in Chapters 2–4
Resolution / solvency-restoration planning	Manage failure, recapitalization, loss absorption and orderly resolution	Yes, with system-wide implications	Failure-time and planning	Relevant to Chapter 2's scenario-level capital-injection object; the chapter then aggregates those injections ex ante across states

Table 1.1: Regulatory and supervisory map.

Alongside capital regulation, supervisory practice has increasingly relied on system-wide stress testing and network-based contagion analysis to evaluate vulnerabilities under adverse scenarios. Practical frameworks developed at international institutions provide guidance on mapping interbank and cross-sector exposures and on combining first-round losses with second-round amplification mechanisms (Bricco and Xu, 2019). Recent extensions of large-scale stress tests in the euro area have also started to incorporate contagion effects between banks and non-banks, reflecting the broader financial ecosystem in which banks operate (Grassi et al., 2025).

In the Hungarian institutional context, the Magyar Nemzeti Bank (MNB, the central bank of Hungary) describes its macroprudential toolkit through regular Macroprudential Reports and annually reviews the identification and buffer rates of other systemically important institutions.(Magyar Nemzeti Bank, 2025, 2026).

The challenge of systemic risk is fundamentally rooted in the network structure of financial obligations. Financial institutions are connected through a complex web of claims and liabilities: interbank loans, derivatives contracts, payment obligations, and common asset holdings (Allen and Babus, 2009). Empirical analyses of interbank payment and lending networks often document a sparse, core-periphery structure[3], which can shape both the likelihood and the pattern of contagion (Soramäki et al., 2007; Craig and von Peter, 2014). This interconnectedness serves important economic functions—it facilitates risk sharing, liquidity provision, and efficient capital allocation. However, it also creates channels through which shocks can propagate from one institution to another, potentially amplifying small initial disturbances into system-wide crises (Acemoglu et al., 2015b). In practice, detailed bilateral exposure data are often confidential or incomplete. As a result, central banks and researchers frequently rely on reconstruction techniques and simulation-based stress tests to bound potential contagion and to evaluate the sensitivity of results to assumptions on the network structure (Upper, 2011).

Understanding and quantifying systemic risk requires mathematical models that capture both the network structure of financial obligations and the mechanisms through which defaults propagate. The seminal work of Eisenberg and Noe (2001) established a foundational framework for analyzing clearing[4] in

[3]In a core-periphery network, a small, densely interconnected core of institutions transacts with a larger periphery of institutions that have few connections among themselves. This topology concentrates counterparty risk in the core and can shape both the likelihood and the direction of contagion.

[4]A clearing mechanism determines which liabilities can be repaid when some agents are insolvent and interbank obligations form directed cycles. In the Eisenberg–Noe model, equilibrium payments are characterized as a fixed point of a monotone mapping under two constraints: limited liability (no agent pays more than its available assets) and proportional repayment (an insolvent agent distributes its assets to creditors in proportion to the face val-

networks of financial obligations with circular dependencies. Their model determines equilibrium payments when multiple institutions are simultaneously insolvent, providing a rigorous basis for studying default contagion. Subsequent research has extended this framework in numerous directions, incorporating bankruptcy costs (Rogers and Veraart, 2013), multiple maturities (Kusnetsov and Veraart, 2019), and decentralized clearing (Csóka and Herings, 2018).

Beyond modeling default cascades, a crucial question concerns the allocation of systemic risk to individual participants. If systemic risk can be measured at the system level, how should this risk be attributed to individual institutions? Such allocation serves multiple purposes: it can guide the design of capital requirements, inform the pricing of deposit insurance, and provide signals for regulatory intervention (Tarashev et al., 2016). Various approaches have been proposed, including methods based on cooperative game theory[5] (Denault, 2001), marginal contribution measures such as CoVaR[6] (Adrian and Brunnermeier, 2016), and network-based indicators such as DebtRank[7] (Battiston et al., 2012b).

The literature on systemic risk has identified two primary channels through which distress can spread: direct contagion through contractual obligations and indirect contagion through common exposures or fire sales (Glasserman and Young, 2016). Direct contagion occurs when the default of one institution imposes losses on its creditors, potentially causing them to default as well. Indirect contagion arises when institutions are exposed to the same risk factors—either because they hold overlapping portfolios or because funding and market-liquidity feedback loops generate fire sales and mark-to-market losses (Cifuentes et al., 2005; Brunnermeier and Pedersen, 2009; Greenwood et al., 2015). Empirical and theoretical work emphasizes that endogenous deleveraging can increase return comovement precisely in stress periods, turning diversification into common exposure (Cont and Wagalath, 2016). Both channels interact in complex ways, and realistic models of systemic risk must account for their interplay (Caccioli et al., 2015).

ues of its debts). Under mild conditions, a unique clearing vector—the vector of equilibrium payments—exists and can be computed efficiently. (Eisenberg and Noe, 2001)

[5]In a transferable-utility (TU) cooperative game, a characteristic function v assigns a value to every coalition of players. Applied to systemic risk, coalitions correspond to sets of institutions and v captures a risk quantity (e.g. losses or required capital) for each coalition, enabling the distribution of the system's total risk among its members (Denault, 2001).

[6]CoVaR measures the Value-at-Risk of the financial system conditional on a particular institution being in distress, thereby capturing that institution's contribution to systemic tail risk (Adrian and Brunnermeier, 2016).

[7]DebtRank measures the fraction of total economic value lost in the network as a consequence of an initial shock to a given institution, thereby quantifying how much damage a single institution's distress can propagate through the system (Battiston et al., 2012b).

1.2 Regulatory and supervisory positioning

This dissertation develops theoretical and computational models of solvency contagion in interbank liability networks. Importantly, "solvency" is used in the clearing-model sense of being able to meet due nominal obligations after shocks and network payments. It is therefore closest to the parts of the prudential toolkit that ask how adverse scenarios propagate through balance-sheet links, how much capital would be needed to restore solvency in a realized state, and how such scenario-wise costs can be attributed across institutions.

The prudential world is broader than the modeling object used in this dissertation. Real financial systems involve direct counterparty losses, common asset exposures, fire sales, liquidity hoarding, collateral and margin dynamics, CCP default waterfalls, payment-system frictions and real-economy feedbacks. The dissertation deliberately zooms in on a narrower solvency-network object. The gain is analytical and computational clarity: each chapter isolates one mechanism under a clean liability-network framework. The cost is institutional incompleteness, so the results should be read as mechanism-focused and supervisory-relevant rather than as a complete regulatory implementation. It is not a regulatory rulebook, not a direct implementation of the G-SIB or O-SII frameworks, and not an empirical calibration of optimal prudential buffers.

The distinction is especially important for Chapter 2. At the state level, the chapter computes a minimal capital injection after the shock scenario has been realized. At the systemic-importance stage, however, those state-by-state injection vectors are aggregated across scenarios by a coherent risk measure[8]. The resulting object is therefore a static but hybrid-timing framework: scenario-wise ex-post solvency-restoration needs are transformed into an ex-ante risk-allocation and systemic-importance indicator. This is closer to a stylized stress-test capital-shortfall or intervention-cost attribution exercise than to a full structural designation methodology.

Chapters 3 and 4 have an even narrower regulatory interpretation. They are best read as stylized contagion modules or sensitivity analyses. Chapter 3 studies how dependence assumptions for external assets interact with direct default contagion, while Chapter 4 studies forward-looking economic valuation feedback. Neither chapter calibrates an official capital-buffer formula, and neither chapter should be interpreted as replacing supervisory stress tests.

[8] A coherent risk measure (Artzner et al., 1999) is a risk quantification function satisfying four axioms: monotonicity, subadditivity, positive homogeneity, and translation invariance. See Chapter 2 for a formal definition.

1.3 What this dissertation models—and what it omits

The formal object throughout the dissertation is a stylized network of fixed nominal bilateral obligations among banks, combined with exogenous external assets and chapter-specific default or valuation rules. The cleanest market interpretation is unsecured interbank lending or, more cautiously, reduced-form net unsecured bilateral exposures measured at a given horizon. Empirical interbank markets often distinguish unsecured lending segments from secured repo segments, which reinforces why the present fixed-nominal-claim abstraction should be anchored in unsecured exposures rather than collateralized transactions (Upper, 2011; Forte et al., 2022). The framework does not literally represent repo, interest-rate swaps, credit-default swaps, CCP-cleared exposures, collateralization, close-out netting, variation or initial margin, or payment-system settlement rules (Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions, 2012; Basel Committee on Banking Supervision and International Organization of Securities Commissions, 2015).

Channel / mechanism	Ch. 2	Ch. 3	Ch. 4	Status	Comment
Direct default contagion through unpaid nominal interbank obligations	Yes	Yes	Yes	In scope	Core mechanism throughout
Common external shocks / common factors	Yes	Yes	Yes	In scope	Studied explicitly in Chapter 3
Forward-looking mark-to-market / valuation contagion	No	No	Yes	In scope	Chapter 4 only
Fire sales / overlapping portfolios	No	No	No	Out of scope	Omitted price-mediated channel
Funding liquidity stress, rollover risk, hoarding or runs	No	No	No	Out of scope	Liquidity channel omitted
Margin calls and collateral dynamics	No	No	No	Out of scope	Important for repo and derivatives
Close-out netting, CSA and legal set-off	No	No	No	Out of scope	Important for OTC derivatives
CCP novation and default waterfall	No	No	No	Out of scope	CCP mechanics omitted
Payment-system settlement mechanics and intraday liquidity	No	No	No	Out of scope	Settlement layer omitted
Real-economy feedback from bank distress	Exogenous only	Exogenous only	Exogenous only	Mostly out of scope	External assets are exogenous

Table 1.2: Contagion-channel scope map.

Throughout the dissertation, the liability network is best interpreted as a stylized network of unsecured interbank liabilities or, more cautiously, reduced-form net unsecured bilateral exposures measured at a given horizon. The model does not represent collateralization, close-out netting, margining, CCP novation, or contract-specific settlement rules. Accordingly, results should not be transferred mechanically to repo, IRS, CDS, or payment-system exposures.

1.4 Overview of contributions

This dissertation develops three stylized, mechanism-focused models of solvency contagion in interbank liability networks. The chapters share a common balance-sheet backbone, but each isolates a different mechanism and a different practical lens. The value added is mechanism identification under controlled network and model assumptions instead of focusing on breadth of institutional coverage.

1.4.1 Chapter 2: Systemic risk measurement via capital injections

Chapter 2 proposes a framework for measuring systemic importance based on minimally required capital injections[9]. For any subset of banks in the system and realized shock scenario, the model computes the minimum ex-post capital injection needed to keep all the selected banks solvent. It then takes an ex-ante view by collecting those coalition-specific injections across scenarios, aggregating them with a coherent risk measure, and using the induced cooperative game to define systemic importance.

The best practical reading is therefore an ex-ante capital-shortfall or intervention-cost attribution indicator built from scenario-wise ex-post solvency-restoration costs. This places the chapter close in spirit to capital-allocation and intervention approaches to systemic risk (Feinstein et al., 2017; Amini et al., 2017; Ahn and Kim, 2019), while also making it distinct from official G-SIB, D-SIB and O-SII designation methodologies. Official frameworks are broader, indicator-based and designed for transparent implementation; Chapter 2 is narrower and network-clearing based.

1.4.2 Chapter 3: Correlated shocks in stylized networks

Chapter 3 examines how default probabilities in a stylized financial network depend on two interacting channels: direct default contagion through nominal obligations and common or correlated shocks to external assets. The focus is on understanding how the correlation structure of external shocks interacts with the network topology to determine systemic risk.

The chapter employs a parsimonious one-period model in which agents hold external assets (endowments)[10] subject to correlated shocks, modeled via a one-factor Gaussian copula[11]. The key innovation of this chapter is the analytical treatment of how default probabilities vary with the correlation parameter,

[9]A capital injection is an ex-post transfer of cash into selected institutions.

[10]Throughout this dissertation, external assets (also called endowments) refer to all assets of an institution excluding claims on other agents in the network, such as investment portfolios, loan books to the real economy, and other non-interbank holdings.

[11]In the Gaussian one-factor copula, each agent's returns on their endowments is written as $\mathcal{A}_i = \sigma(\sqrt{\rho}\, Z + \sqrt{1-\rho}\,\varepsilon_i) + \mu$, where Z is a standard-normal common factor, the ε_i are i.i.d. standard-normal idiosyncratic factors, and $\rho \in [0, 1]$ controls the pairwise correlation. This is the standard dependence structure used in credit-risk modelling.

holding network structure fixed. Two complementary technical approaches are developed: a Bayesian network[12] formulation that enables exact calculations under specific symmetric system structures, and a mean-field approximation[13] that characterizes behavior in the large-system limit.

A striking finding emerges: default probabilities are non-monotonic in the asset correlation[14]. Under typical reasonable choices of parameters, a characteristic "hump shape" appears—default probabilities first increase with correlation, reach a maximum at intermediate values (around 0.8–0.9), and then decrease as correlation approaches one (i.e. as endowments become comonotonic[15]. Under comonotonicity in a perfectly homogeneous system, either all agents default together or none does, so contagion disappears and default probabilities are minimized. To quantify how strongly this result relies on ex-ante homogeneity, the chapter also introduces a Monte Carlo robustness experiment with heterogeneous bank-specific distances to default DD_i held fixed within each simulated system. This extension shows when the symmetry-based $\rho = 1$ result survives and when contagion can reappear even under comonotonic endowments.

The practical interpretation is a controlled dependence-sensitivity analysis. The chapter highlights why stress-test modules should treat dependence assumptions carefully, but it does not model liquidity stress, collateral dynamics, fire sales, endogenous portfolio choice or real-economy feedback.

1.4.3 Chapter 4: Forward-looking network valuation

The third contribution (Chapter 4) relaxes the one-period assumption and develops a multi-period, forward-looking framework for network valuation. The key departure from earlier models is the treatment of interbank claims: rather than carrying claims at face value until default, this chapter marks claims to market. This periodic revaluation makes defaults possible before the terminal time where claims are due.[16]

[12]A Bayesian network is a directed acyclic graph in which each node represents a random variable and directed edges encode conditional dependencies. In the present context, the default indicator of each institution is a node, and the edges follow the pattern of interbank liabilities, so that the joint default distribution factorises into local conditional probabilities that depend only on each node's parents (direct borrowers).

[13]A mean-field approximation replaces the detailed interactions among all agents with interaction with an aggregate quantity (typically the average default rate), yielding tractable equations for large systems. It is asymptotically exact for dense, homogeneous networks as the number of agents grows (Sznitman, 1991).

[14]Later denoted by ρ, which is the pairwise correlation between external assets of any two agents, to be precise.

[15]Comonotonicity is the strongest form of positive dependence: all agents' endowments move monotonically with a single common random variable. In the Gaussian one-factor model this corresponds to $\rho = 1$ (no idiosyncratic component)

[16]Mark-to-market valuation of interbank claims means adjusting the book value of a claim on counterparty j by a valuation function $V_j(t) \in [0, 1]$, so that the claim is recorded

This forward-looking approach introduces a fundamental feedback loop. The value of a claim on bank j depends on j's probability of default, but this probability in turn depends on how j's own claims are valued—which depends on the solvency of j's counterparties, and so on. Resolving this circularity requires finding a solution in which valuations are self-consistent[17]: the default probabilities used to discount claims must equal the actual default probabilities implied by the resulting valuations and system dynamics.

The chapter develops a temporal-difference (TD) learning[18] approach to compute these self-consistent valuations. Neural networks are trained to predict to-maturity default probabilities as functions of current balance sheet states, with training proceeding backwards in time. The method sidesteps the curse of dimensionality that afflicts grid-based approaches, enabling analysis of networks with more than a handful of banks.

Key findings include the persistence of the hump-shaped relationship between default probability and external shock correlation, but with an important modification: under fully state-dependent haircuts, the minimal default probability may no longer occur at perfect correlation. This suggests that mark-to-market valuation introduces additional contagion channels that can shift the optimal correlation away from one.

The cleanest practical interpretation is an economic valuation model for unsecured or uncollateralized bilateral interbank claims, or for reduced-form credit-sensitive exposures. The chapter is closest to dynamic distress-contagion and network valuation work (Feinstein and Sojmark, 2022; Banerjee and Feinstein, 2022), and only indirectly related to supervisory policy.

1.4.4 Comparing the three contributions

The three chapters share a common foundation: all build on the same network clearing framework and all consider interbank obligations due at a single

at $V_j(t)$ times its face value. The complement $1 - V_j(t)$ is called the haircut and reflects the discount for counterparty credit risk. In this dissertation, the haircut equals the anticipated probability that j defaults by maturity.

[17]A valuation rule (by this chapter's definition) is weakly self-consistent if claims on already-defaulted counterparties are always valued at zero. It is strongly self–consistent if, in addition, the default probabilities embedded in the haircuts coincide with the actual conditional default probabilities generated by the full network dynamics under those very haircuts. Achieving strong self–consistency amounts to finding a fixed point in the space of valuation functions.

[18]TD-learning (temporal-difference learning) is a family of prediction algorithms where a conditional expectation is learned by using bootstrapped targets—–i.e., targets that reuse the model's next-step prediction instead of waiting for the final outcome of a full simulation. (Sutton and Barto, 2018) In the present context, the to-maturity default probability at time t is updated towards a one-step target formed from the post-clearing survival indicator at t and the predicted default probability at $t{+}1$. Freezing the $t{+}1$ predictor while fitting the t predictor is a standard stabilisation device.

maturity T. Despite this shared structure, the chapters differ progressively in their treatment of uncertainty, time, and the valuation of interbank claims.

In Chapter 2, the clearing problem and the scenario-wise injection are evaluated after shocks are realized, and the ex-ante step comes from aggregating scenario-wise injections. (see Figure 1.1).

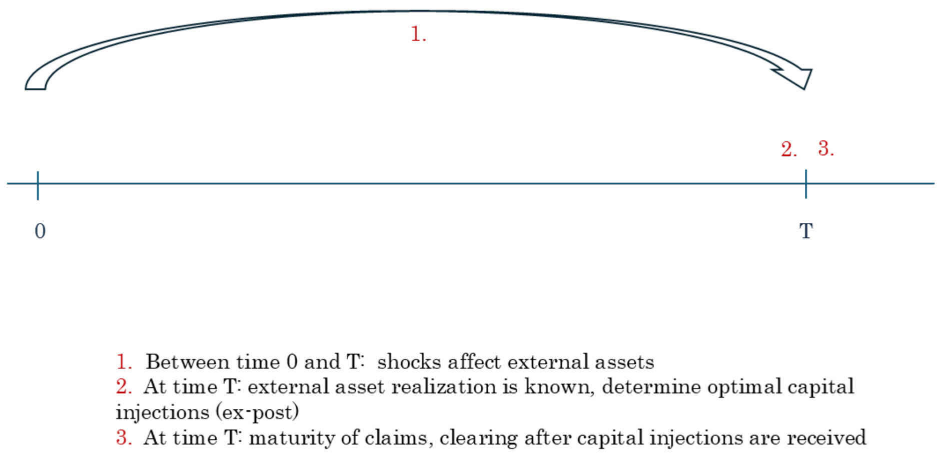

Figure 1.1: Illustration of Chapter 2's time grid.

Chapter 3 focuses on the stochastic external assets within the one-period framework. There is still only one clearing event at maturity, but the analysis in inherently ex-ante: default probabilities are computed over the distribution of shocks (see Figure 1.2). The key addition is the correlation structure of external assets and its interaction with network topology.

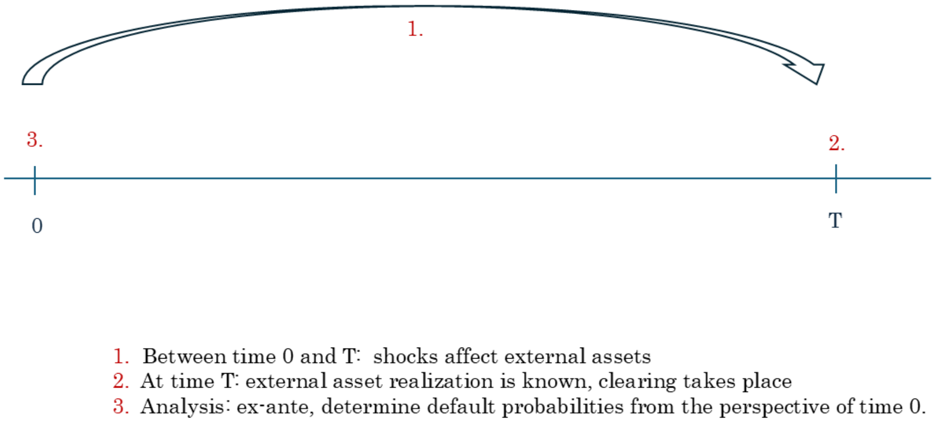

Figure 1.2: Illustration of Chapter 3's time grid.

Chapter 4 retains the single maturity but adds multiple interim revaluation dates at which interbank claims are marked to market. This creates a genuinely multi-period setting where anticipated future defaults can trigger current defaults through endogenous haircuts—a feedback channel absent from the first two chapters (see Figure 1.3). Computing the resulting self-consistent valuations requires machine-learning tools (TD learning) rather than purely analytical or simulation-based methods.

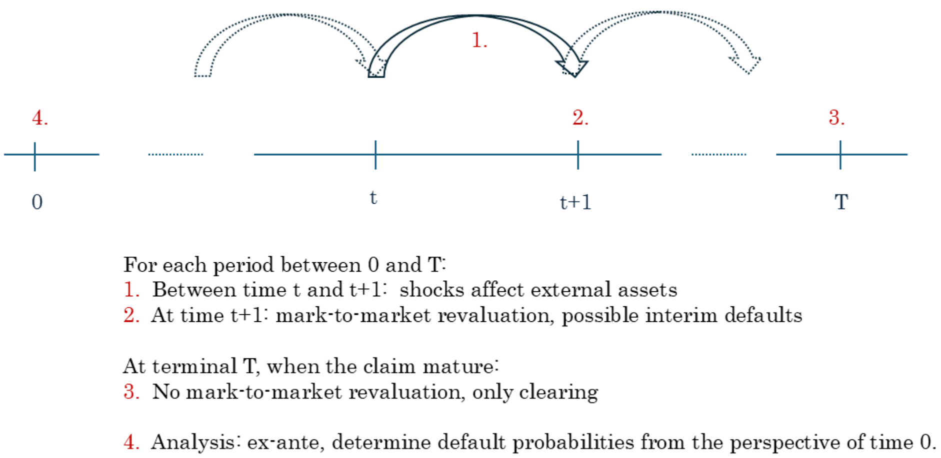

Figure 1.3: Illustration of Chapter 4's time grid.

The dissertation measures and allocates risk contributions under stylized assumptions. It does not determine whether current prudential buffers are too high or too low, nor does it attempt an optimal calibration of the regulatory framework. Such a calibration exercise would require additional empirical evidence on credit supply, macroeconomic costs, institutional substitution, supervisory objectives and legal implementation constraints.

1.5 Literature review

1.5.1 Financial networks and contagion

The study of financial networks draws on graph-theoretic and economic approaches to understand how the structure of financial connections affects systemic stability. Allen and Gale (2000) initiated the modern literature on financial contagion, showing that the network structure of interbank deposits determines whether liquidity shocks remain localized or spread throughout the system. Their finding that complete networks are more resilient than incomplete ones sparked extensive research on the relationship between connectivity and stability.

Acemoglu et al. (2015b) provide a comprehensive theoretical analysis of how network structure affects the propagation and amplification of shocks. They identify a phase transition—a sharp, qualitative change in system behavior at a critical threshold—in network behavior: up to a threshold, greater connectivity enhances stability by allowing losses to be shared across more counterparties, but beyond this threshold, connectivity becomes a source of fragility as it enables shocks to spread more widely. This non-monotonic relationship between connectivity and stability is a recurring theme in the literature.

In a related strand, contagion is studied as a cascade or threshold phenomenon on networks and highlights a "robust-yet-fragile" trade-off: networks that withstand small shocks can still be vulnerable to rare, large cascades (Gai and Kapadia, 2010). More general models emphasize that not only overall connectivity but also the distribution of exposures and the identity of key nodes matter, which can make targeted interventions more effective than uniform policies (Elliott et al., 2014). Analytically tractable results for large, heterogeneous networks provide resilience criteria and identify "contagious links" that have an outsized impact on systemic stability (Amini et al., 2016).

Empirical studies have documented the structure of interbank networks in various countries. Boss et al. (2004) analyze the Austrian interbank market, finding a scale-free degree distribution[19] and core-periphery structure. Simi-

[19] A scale-free network has a degree distribution that follows a power law, meaning a few nodes have very many connections while most have few.

lar patterns have been documented for other banking systems, including the Hungarian interbank market before and after the 2008 crisis (Berlinger et al., 2011, 2017).

Beyond interbank lending data, payment-system networks also exhibit pronounced heterogeneity and core-periphery structures, which can affect both the speed and the reach of liquidity shocks (Soramäki et al., 2007; Craig and von Peter, 2014). Because full bilateral exposures are rarely observable, a common empirical approach is to combine partial data with reconstruction techniques and to run counterfactual default simulations to bound contagion risk (Furfine, 2003; Upper, 2011).

1.5.2 Clearing mechanisms and fixed points

The Eisenberg–Noe model (Eisenberg and Noe, 2001) provides the mathematical foundation for analyzing clearing in networks with circular obligations. The model determines a clearing vector—the payments made by each institution—as the fixed point of a mapping that respects limited liability and proportional repayment. Under mild conditions, a unique clearing vector exists and can be computed efficiently.

Subsequent research has extended this framework in numerous directions. Rogers and Veraart (2013) introduce bankruptcy costs, which create discontinuities in payoff functions and can lead to multiple equilibria. Glasserman and Young (2015) analyze networks where institutions also have obligations to entities outside the network, providing conditions under which contagion is likely or unlikely. Csóka and Herings (2018) show that under relatively mild conditions, the clearing vector can be reached through decentralized bilateral negotiations, making it unnecessary to collect and process all the sensitive data of all the agents simultaneously and run a centralized clearing procedure.

The question of uniqueness versus multiplicity of equilibria has received considerable attention. Roukny et al. (2018) show that interconnectedness can create "default ambiguity"—situations where multiple clearing outcomes are possible—which adds an additional source of systemic risk. Schuldenzucker et al. (2020) demonstrate that the introduction of credit default swaps can create severe multiplicity problems, with the number of possible equilibria growing exponentially. Csóka and Herings (2024) study clearing under general bankruptcy laws and provide a network-structure–based sufficient condition (using a decomposition into strongly connected components) that guarantees uniqueness of the clearing payment matrix.

Beyond the static, one-period setting, several extensions allow liabilities and endowments to evolve over time. Dynamic clearing models separate short-term cash flows from longer-term capital positions and can distinguish illiquidity-driven defaults from insolvency (Banerjee et al., 2025). These dynamic frameworks connect naturally to forward-looking valuation problems because today's equity values depend on expectations about future clearing states.

1.5.3 Systemic risk measurement and allocation

Measuring systemic risk at the system level requires defining what constitutes a systemic event and how to quantify its severity. Various approaches exist in the literature. Adrian and Brunnermeier (2016) propose CoVaR, which measures the value-at-risk of the financial system conditional on an institution being in distress. Acharya et al. (2017) measures each institution's contribution to systemic risk as its systemic expected shortfall[20] (SES), that is, its propensity to be undercapitalized when the system as a whole is undercapitalized. Brownlees and Engle (2017) develop SRISK, which estimates the capital shortfall of an institution in a systemic crisis and provide econometric methods for estimating SRISK from market data.

A complementary line of research infers interconnectedness directly from market data. For example, connectedness measures based on Granger-causality networks or variance-decomposition spillovers summarize how shocks to one firm's returns propagate to others, providing a time-varying view of spillovers in the financial sector (Billio et al., 2012; Diebold and Yilmaz, 2014). While these measures are useful for monitoring, structural network models remain necessary when the objective is to evaluate counterfactual interventions (e.g., capital injections) and to separate balance-sheet linkages from common-factor effects.

The allocation of systemic risk to individual institutions connects to the broader theory of risk allocation in cooperative settings. Denault (2001) establishes the axiomatic foundations for coherent risk allocation, showing how game-theoretic solution concepts can be applied to distribute risk across portfolios or business units. The Shapley value[21], which measures marginal contributions, has been widely applied to systemic risk allocation (Drehmann and Tarashev, 2013; Tarashev et al., 2016).

[20] Expected Shortfall (also called CVaR) is a coherent risk measure that equals the average loss in the tail of the loss distribution beyond a given quantile. It is more sensitive to tail risk than Value-at-Risk and satisfies the subadditivity property required for coherent risk allocation.

[21] The Shapley value allocates a coalition's total worth to individual players by averaging each player's marginal contribution over all possible orders in which the coalition can be formed. For player i in an N-player game with characteristic function v, it equals $\phi_i(v) = \sum_{C \subseteq N \setminus \{i\}} \frac{|C|!\,(N-|C|-1)!}{N!} \big(v(C \cup \{i\}) - v(C)\big)$ (Shapley, 1953). In systemic risk applications, the Shapley value measures each institution's average marginal contribution to system-wide risk.

A body of work also formalizes systemic risk measures as sets of capital allocations that render the system acceptable from a regulator's perspective, bridging ex post bailout costs and ex ante capital requirements (Feinstein et al., 2017). Closely related, several papers treat the regulator's problem as choosing targeted equity infusions or interventions to restore solvency under scenario stress, often cast as optimization problems on top of a clearing mechanism (Amini et al., 2017; Ahn and Kim, 2019).

Network-based measures provide an alternative perspective. DebtRank (Battiston et al., 2012b) measures the fraction of value lost in the network following an initial shock to a given institution. Demange (2018) develops a threat index that captures each institution's contribution to systemic fragility. These network-centric measures complement market-based measures by explicitly accounting for the structure of obligations. Chapter 2 contributes to this literature by measuring systemic importance with the additional capital required to guarantee solvency. The methodology can also allocate the risk of the whole system to individual institutions capturing coalitional effects.

1.5.4 Forward-looking valuation and mark-to-market

Traditional clearing models value interbank claims at face value until default occurs. This historical-cost accounting approach ignores the possibility that counterparty credit risk should be reflected in current valuations. Barucca et al. (2020) develop a unified framework for network valuation that encompasses various approaches to marking interbank claims, from face value to full mark-to-market.

These modelling choices tie into the broader debate on fair-value (mark-to-market) accounting and its potential procyclicality. While mark-to-market reporting can improve transparency, it may also amplify stress by mechanically transmitting price declines and funding pressures into balance sheets (Allen and Carletti, 2008; Plantin et al., 2008; Laux and Leuz, 2009). Such amplification mechanisms are closely related to liquidity spirals, where declining asset prices tighten funding constraints and trigger further deleveraging (Brunnermeier and Pedersen, 2009).

Bardoscia et al. (2019) introduce forward-looking solvency contagion, where the deterioration of a counterparty's creditworthiness—even without actual default—can erode a bank's equity. This anticipatory channel can trigger defaults that would not occur under historical-cost accounting, providing a more conservative assessment of systemic risk.

The most sophisticated treatment appears in Feinstein and Sojmark (2022), who develop a fully endogenous framework where mark-to-market valuations and default probabilities are jointly determined in equilibrium. Their forward-backward iterative scheme solves for self-consistent valuations but faces computational challenges as the number of institutions grows. The present dissertation (Chapter 4) contributes to this literature by developing machine-learning methods that can handle larger networks. More broadly, the computational burden of solving large-scale network fixed points has motivated the use of approximation and learning methods, including temporal-difference learning (Sutton and Barto, 2018).

1.5.5 Correlation, common factors, and dependence modelling

Many systemic risk scenarios are driven by common macroeconomic factors that simultaneously affect many institutions. In the context of credit portfolios, a canonical representation is the one-factor model underlying the Basel internal-ratings-based formulas, where defaults become conditionally independent given a systematic factor (Vasicek, 2002; Gordy, 2003). For pricing and risk management of multi-name credit derivatives, dependence is often modelled using copulas, with the Gaussian copula being the workhorse specification (Li, 2000). At the same time, it is well known that linear correlation is an incomplete summary of dependence—particularly in the tails—and that different dependence structures can have large effects on joint default risk (Embrechts et al., 2002).

In network models, correlation interacts with balance-sheet linkages in a nontrivial way. Higher correlation increases the likelihood of large system-wide shocks, yet it may reduce the frequency of purely idiosyncratic triggers. This tension is reflected in the debate on whether diversification is stabilizing: some work argues that diversification can increase the probability of systemic crises through higher asset correlation (Wagner, 2010), and that institutions may have incentives to herd into similar exposures when they anticipate government support (Acharya and Yorulmazer, 2007). Simulation evidence in financial network models shows that the interaction of correlation and contagion can generate non-monotonic relationships between asset correlation and the probability of systemic crises (Frey and Hledik, 2014). Chapter 3 contributes to this literature by providing a Bayesian-network based characterization and mean-field approximations that help characterize and explain when and why such non-monotonicity emerges.

1.5.6 Supervisory and regulatory context

The methodological literature reviewed above is the main academic foundation of the dissertation. A compact supervisory map is nevertheless useful for interpreting the models. Post-crisis macroprudential architecture distinguishes structural systemic-importance frameworks from cyclical tools and scenario-based assessments. The BCBS G-SIB framework and the EBA O-SII guidelines use broad indicator sets—including size, interconnectedness, substitutability or importance, complexity and, for G-SIBs, cross-jurisdictional activity—to identify institutions whose distress would have large systemic consequences (Basel Committee on Banking Supervision, 2012, 2013; European Banking Authority, 2014). By contrast, the CCyB is designed to vary over time with cyclical system-wide risk (Basel Committee on Banking Supervision, 2010).

Stress-testing practice provides the closest supervisory analogue to the dissertation's scenario-based logic. Practical contagion frameworks map direct exposures and amplification mechanisms across institutions, and recent supervisory work increasingly studies contagion and amplification effects as complements to standard solvency stress tests (Bricco and Xu, 2019; Grassi et al., 2025). This dissertation speaks only to selected parts of that architecture: direct solvency contagion, correlated external shocks, scenario-wise capital-injection needs and forward-looking valuation feedback.

Finally, the market-scope limitations are essential. Repo markets, OTC derivatives, CCP-cleared exposures and payment systems involve collateral, margining, netting, settlement and default-waterfall mechanisms that are absent from the models. Those mechanisms are central in financial-market-infrastructure and derivatives regulation (Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions, 2012; Basel Committee on Banking Supervision and International Organization of Securities Commissions, 2015). The dissertation therefore isolates mechanisms in an unsecured-liability-network setting rather than attempting to reproduce the full institutional machinery of modern financial markets.

1.6 Organization of the dissertation

The remainder of this dissertation is organized as follows. Chapter 2 presents the capital injection framework for systemic risk measurement, including the cooperative game-theoretic foundations and a simulation example. Chapter 3 develops the analysis of correlated shocks using Bayesian networks and mean-field methods. Chapter 4 introduces the forward-looking valuation framework with TD-learning. Each chapter is self-contained but builds on the common foundation established in this introduction.

Chapter 2

Measuring systemic importance with minimally required capital injections

This chapter is based on Csóka and Sass (2023).

2.1 Introduction

Building on the foundations established in Chapter 1, this chapter addresses the measurement of systemic importance in financial networks through the lens of capital injections. Like the other chapters in this dissertation, the model considers obligations due at a single maturity; the distinguishing feature here is a static but hybrid-timing perspective. For each realized shock scenario and coalition (i.e., any subset of the banking agents)[22], the minimal capital injection is computed after the shocks have materialized, so at the scenario level it is an ex-post solvency-restoration object. Systemic importance, however, is not defined at that stage alone: the chapter collects the coalition-specific injections across scenarios, aggregates them with a coherent risk measure, and thereby induces an ex-ante systemic-risk game. While the introductory chapter surveyed the broad landscape of systemic risk and network-based models, here we develop a specific framework for allocating that risk to individual participants.

The chapter is not an attempt to reproduce current G-SIB, D-SIB or O-SII methodology. Official structural systemic-importance frameworks are broad, indicator-based and designed for transparent supervisory implementation, including dimensions such as size, interconnectedness, substitutability or importance, complexity and cross-border activity (Basel Committee on Banking

[22]Coalitions are not strategic groups; they are hypothetical subsets/groups of agents used for marginal-contribution accounting.

Supervision, 2012, 2013; European Banking Authority, 2014). The question studied here is narrower and deliberately model-based: first, for each coalition and realized shock scenario, how much capital must be injected ex post to keep the coalition solvent; second, across scenarios, how should those coalition-specific injections be aggregated by a risk measure and attributed across institutions. This places the chapter closer to capital-allocation, intervention-cost and stress-scenario approaches to systemic risk than to official designation exercises (Feinstein et al., 2017; Amini et al., 2017; Ahn and Kim, 2019; Jackson and Pernoud, 2024).

We wish to devote particular attention to cases where circular indebtedness makes it difficult to determine the course of contagion. In the study of such situations, Eisenberg and Noe (2001) is pioneering. Rogers and Veraart (2013) extend this model to allow for bankruptcy costs, while in Glasserman and Young (2015) obligations outside the network may also be present. Numerous studies build on the clearing algorithm of Eisenberg and Noe (2001), for example Amini et al. (2016) or Gai et al. (2011). The clearing mechanism, which originally governs the settlement of currently due payments, thus also determines, in these models, the course of shocks and the propagation of default. Most models build on a central clearing mechanism that assumes knowledge of the entire structure of the system. This assumption is relaxed by Csóka and Herings (2018), who analyze decentralized clearing.

A common feature in the study of systems with cycles of obligations is that the equilibrium state (the outcome of clearing) is a fixed point of some mapping representing the default rules. If this is not unique, additional systemic risk may arise due to the extra uncertainty, as in the models of Roukny et al. (2018) or Schuldenzucker et al. (2020). The outcome of clearing either consists of the equilibrium payments or, if clearing is a market-value revaluation of liabilities, the equilibrium equity values. In the latter case, a domino effect can be triggered not only by insolvency, i.e., an actually realized default event, but "closeness to default" can also propagate through the network. Such an ex-ante valuation within the network, based on equilibrium equity values of liabilities (which in its perspective builds on Fischer (2014)), appears in Battiston et al. (2012b) and, building on it, in Bardoscia et al. (2015) and Battiston et al. (2016). Within this perspective, Visentin et al. (2016) show how various standard models can be fitted into a common framework. Moreover, Veraart (2020) develop the model describing that equilibrium state, within which Veraart (2022) investigate the possibility of transforming the structure to reduce systemic risk. In the model of Barucca et al. (2020), it is possible to incorporate uncertainty regarding future asset values (future shocks as random variables). By linking the ex-ante and ex-post approaches, they show that the earlier, ex-post oriented models also arise as special parameterizations of this model.

The measurement of the systemic importance of individual agents, on the one hand, can help us understand the role agents play within the system. On the other hand, the allocation of systemic risk to agents can be used as a stylized risk-capital allocation method and as an input into supervisory risk-attribution exercises.[23]. The Shapley value (Shapley (1953)) is used to identify institutions that are important from a systemic risk perspective, among others, by Drehmann and Tarashev (2013), Bluhm et al. (2014), and Borsos et al. (2020). Brunnermeier and Cheridito (2019) capture the riskiness of the system as a whole by the cost of insurance against lower-tail losses, and then allocate it to individual participants. Demange (2018) likewise measures each participant's contribution to systemic risk. He defines an index that expresses the change in the payment capacity of the other participants in the event of the potential default of the given participant. Brunnermeier and Cheridito (2019) also address the measurement and allocation of systemic risk using the Euler (gradient) method.

It is natural to discuss the contribution of coalitions of participants to the risk of the system as a whole within the framework of cooperative game theory; this is grounded by Denault (2001). The usage of cooperative game theory does not imply there is strategic coordination, bargaining or coalition formation among banks[24]. It is only a tool used to determine the contribution of a group of banks to the risk of the whole system. If the risk of the system as a whole is smaller than the sum of the risks of the components, the diversification benefit must somehow be distributed among the components. According to Balog et al. (2011), applications of risk allocation include, for example, the design of capital requirements/risk limits and (individual) performance measurement. Related to the former, Walter (2002) points out the danger of extreme investment strategies arising from the indiscriminate application of VaR-based limit systems. The study by Csóka et al. (2009) examines this problem in a system without circular liabilities; in this paper we focus on the case with circular liabilities.

The structure of the chapter is as follows. In Section 2.2.1 we present the cooperative game-theoretic framework that enables the distribution of the risk of the system as a whole across individual participants. In Section 2.2.2 we describe the (stylized) financial network we model, the rules of its possible interactions, and we introduce the notion of capital injection. In Section 2.2.3 we connect the foregoing and define two possible versions of the systemic risk capital allocation game and the resulting systemic risk indicators, which measure the contribution of individual participants to systemic risk. In Section 2.3

[23]It should not be read as a direct calibration rule for prudential capital buffers.

[24]There is a hypothetical/implicit bargaining behind the fair allocation of the total capital injection, but this is not an explicit part of our model.

we illustrate the above through an example.

2.2 The systemic risk capital allocation game

2.2.1 Cooperative game theory and capital allocation games

Throughout this chapter, coalition is used in the cooperative game-theoretic sense: any nonempty subset of the set of banking agents. It does not imply any strategic cooperation among the members. Similarly, the term game is used in the sense of cooperative game theory—a cost-sharing or value-allocation problem—not in the sense of strategic or non-cooperative games.

If a financial firm consists of multiple subunits, then the sum of the risks of the subunits is greater than the risk of the financial firm; a diversification effect appears, which must somehow be allocated across the subunits[25]. The literature mostly uses cooperative game theory for this purpose. The capital to be reserved serving as a buffer is determined by the risk measure; hence the name of the problem: capital allocation/risk allocation game (for more details see Balog et al. (2017)). The range of possible applications is wide: strategic decision-making related to new business lines, product pricing, (individual) performance evaluation, the design of risk limits, and risk allocation in insurance companies.

Definition 2.1. Let $\mathcal{N}$ be a finite set of players, with $|\mathcal{N}| = N$, so that to every player one can assign a uniquely identifying index, a positive integer. An $(\mathcal{N}, v)$ *transferable-utility cooperative game* is specified by the *characteristic function* $v : 2^{\mathcal{N}} \to \mathbb{R}$, where $v(\emptyset) = 0$. Let $\mathcal{G}^{\mathcal{N}}$ denote the set of games with player set $\mathcal{N}$.

For a game $(\mathcal{N}, v) \in \mathcal{G}^{\mathcal{N}}$, let $x \in \mathbb{R}^N$ denote the vector of *payoffs*[26], where x_i is the payoff to player $i \in \mathcal{N}$. The payoff vector x provides the coalition $C \subseteq \mathcal{N}$ with payoff $x(C) = \sum_{i \in C} x_i$.

The payoff $x \in \mathbb{R}^N$ is *individually rational* if $x_i \geq v(\{i\})$ for every player $i \in \mathcal{N}$, and it is *coalitionally rational* if $x(C) \geq v(C)$ for every coalition $C \subseteq \mathcal{N}$. The *core of the cooperative game* is the set of efficient and coalitionally rational payoffs. Core payoffs are stable in the sense that no coalition wishes to block them.

[25] In our rather theoretical setup, acknowledging that in practice diversification effect might not be accounted for when calculating requirements of different subsidiaries.

[26] In cooperative game theory, *payoff* refers to the value allocated to a player by the solution concept, not to the cash flow from a financial instrument. Similarly, *efficient* means that the payoff vector exhausts the total value of the grand coalition, $x(\mathcal{N}) = v(\mathcal{N})$, which differs from the financial notions of Pareto efficiency or informational efficiency.

For capital allocation games we will define risk measures on discrete states of the world (scenarios). Let the set of realization vectors be $\mathcal{X} \in \mathbb{R}^s$, where S is the number of states of the world (scenarios). In what follows, to simplify notation, the symbol $s \in S$ may, depending on context, denote a given scenario from the set of scenarios or the positive integer index $s \leq S$ that is uniquely associated with it. Let the probability of state/outcome s be $p_s > 0$, with $\sum_{s=1}^{S} p_s = 1$. Henceforth we assume that the states of the world occur with equal probability, $p_1 = \cdots = p_S = 1/S$. The vector $X \in \mathbb{R}^S$ gives a portfolio's possible future profits over a given horizon. Thus positive values represent profits, while negative values represent losses.

Definition 2.2. The function $\rho : \mathbb{R}^S \to \mathbb{R}$ is a *coherent risk measure* (Artzner et al. (1999), rephrased by Csóka et al. (2009)) if it satisfies a few basic requirements (axioms). These are: monotonicity, subadditivity, positive homogeneity, and translation invariance.

Some popular risk measures do not, for example, satisfy the requirement of subadditivity, in which case it may occur that the risk of the union of two portfolios is greater than the sum of the risks of the portfolios.

Given a realization vector X, define the order statistic $X_{s:S}$ (i.e., the s-th smallest value among $X_1, \ldots, X_S$) after sorting the elements $X_1, \ldots, X_S$ into increasing order, that is, $X_{1:S}, \ldots, X_{S:S} = X_1, \ldots, X_S$ and $X_{1:S} \leq X_{2:S} \leq \cdots \leq X_{S:S}$. Then the definition of the coherent Expected Shortfall (see, e.g., Acerbi and Tasche (2002) and Ágoston (2010)) is as follows.

Definition 2.3. Let the outcomes have identical occurrence probabilities and let $k \in 1, \ldots, S$. Then the *k-expected shortfall* value of the realization vector X is

$$\mathrm{ES}_k(X) = -\delta \sum_{s=1}^{k} \frac{1}{k} X_{s:S}, \tag{2.1}$$

where δ is the price/discount factor of the risk-free instrument appearing in the translation invariance criterion, which we henceforth assume to be 1. For a more detailed description, see Csóka et al. (2009).

The k-expected shortfall is the discounted average loss of the worst k outcomes. Note that for $k = 1$ we obtain the discounted maximum loss.

From the elements defined so far we can build the risk environment as follows.

Definition 2.4. In a *risk environment* $(\mathcal{N}, S, \mathbf{X}, \rho)$, $\mathcal{N}$ is a finite set of portfolios (players), with $|\mathcal{N}| = N$; S is the number of states of the world (assumed to occur with equal probability); $\mathbf{X} \in \mathbb{R}^{S \times 2^N}$ is the matrix of realization vectors, whose columns are indexed by coalitions $C \subseteq \mathcal{N}$; and ρ is a coherent risk measure.

Thus a column of the matrix $\mathbf{X}$ contains the values realized by a given coalition across the different states of the world. We denote by $X(C) \in \mathbb{R}^S$ the fixed column of this realization matrix corresponding to coalition $C \subseteq \mathcal{N}$. A given risk environment defines the following cooperative game.

Definition 2.5. In the risk environment $(\mathcal{N}, S, \mathbf{X}, \rho)$, the *capital allocation game* $(\mathcal{N}, v)$ has characteristic function $v : 2^{\mathcal{N}} \to \mathbb{R}$ given by

$$v(C) = -\rho(X(C)) \quad \text{for all } C \subseteq \mathcal{N}. \tag{2.2}$$

For the maximum loss, when $\delta = 1$, we obtain that

$$v(C) = -\mathrm{ES}_1(X(C)) = -\big(-\min_{s \in S} X(C)_s\big) = \min_{s \in S} X(C)_s.$$

2.2.2 Financial networks and capital injection

In the foregoing we introduced a few concepts and results from the world of risk allocation games. In what follows we define a specific risk allocation game, to which we will refer as the systemic risk capital allocation game. First we present a stylized financial network in which systemic risk arises as a consequence of the interactions among agents. Next we show two ways to define the realization vector, and then, by applying coherent risk measures to these, we build the systemic risk capital allocation game, where the goal is to allocate systemic risk (for example, for capital allocation purposes) to the individual agents of the system.

In introducing the most important notations and definitions we build on Csóka (2017) and Csóka and Herings (2021). A financial network is determined by the set of agents, the value of the agents' initial endowments, and the magnitude of the agents' liabilities to other agents. We define these in turn as follows.

Let $\mathcal{N}$ denote the set of agents, which is a subset of the set of possible agents $\mathbb{N}$; formally, $N \in \mathcal{N}$, where $\mathcal{N}$ denotes the set of nonempty, finite subsets of $\mathbb{N}$. In what follows we will typically think of the agents as banks[27], with the exception of the distinguished agent indexed by 1, who represents the nonbank sector/real economy in aggregate. We assume that the nonbank sector has no liabilities towards the other agents.

The agents' *initial endowments (endowments, or nonbank asset holdings)* are given by the vector $z \in \mathbb{R}_+^N$, where z_i represents all tangible and intangible assets of agent i, excluding claims on other agents.

The agents' liabilities are given by the $L \in \mathbb{R}_+^{N \times N}$ *liability matrix*, in which the element L_{ij} shows how much the ith agent owes to j. Naturally, agents do

[27]Throughout this chapter, the terms agents and banks are used interchangeably for the banking agents ($i \geq 2$), while agent 1 specifically denotes the nonbank sector (the real economy).

not owe to themselves, so $L_{ii} = 0$ for every i. At the same time, it may occur due to various contracts that i owes j and j owes i as well, that is, we may have both $L_{ij} > 0$ and $L_{ji} > 0$. The vector of agents' total liabilities is $l \in \mathbb{R}^N_+$, where $l_i = \sum_{j=1}^N L_{ij}$.

Definition 2.6. We will later need the notion of the *relative liabilities matrix*.

$$\Pi_{ij} = \begin{cases} \dfrac{L_{ij}}{\sum_{k\in\mathcal{N}} L_{ik}}, & \text{if } \sum_{k\in\mathcal{N}} L_{ik} > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Thus a financial network is defined by the triple $(\mathcal{N}, z, L)$, and let $\mathcal{F}$ denote the set of all financial networks. The amounts that agents actually pay to one another are determined by the *payment matrix* $P \in \mathcal{M}$, where $\mathcal{M}$ denotes the set of matrices that have zeros on their main diagonal and otherwise contain nonnegative real numbers. P_{ij} gives the amount paid by agent $i \in \mathcal{N}$ to agent $j \in \mathcal{N}$. Given the financial network $(\mathcal{N}, z, L) \in \mathcal{F}$ and the payment matrix $P \in \mathcal{M}$, the *asset value* of agent $i \in \mathcal{N}$ is $a_i(\mathcal{N}, z, P) = z_i + \sum_{j\in\mathcal{N}} P_{ji}$, which is the sum of the initial endowment and the payments received from others.

Definition 2.7. Subtracting the agent's payments from the value of assets yields the agent's *equity*, which for agent $i \in \mathcal{N}$ is $e_i(\mathcal{N}, z, P) = a_i(\mathcal{N}, z, P) - \sum_{j\in\mathcal{N}} P_{ij} = z_i + \sum_{j\in\mathcal{N}}(P_{ji} - P_{ij})$.

Bankruptcy rules assign to a financial network $(\mathcal{N}, z, L) \in \mathcal{F}$ a payment matrix $P \in \mathcal{M}$. They practically specify how much each agent should pay to the others under a given situation.

Definition 2.8. A *bankruptcy rule* is a function $b : \mathcal{F} \to \mathcal{M}$ such that for every $(\mathcal{N}, z, L) \in \mathcal{F}$ we have $b(\mathcal{N}, z, L) \in \mathcal{M}$.

Under the proportional bankruptcy rule $p : \mathcal{F} \to \mathcal{M}$, each agent pays out of its assets in proportion to its liabilities.

Definition 2.9. The *proportional bankruptcy rule* is the function $p : \mathcal{F} \to \mathcal{M}$ such that for every $(\mathcal{N}, z, L) \in \mathcal{F}$ we have $p(\mathcal{N}, z, L) = P$, where the matrix P is the solution to the following system:

$$P_{ij} = \begin{cases} 0, & \text{if } L_{ij} = 0, \\ \min\{\frac{L_{ij}}{\sum_{k\in\mathcal{N}} L_{ik}} a_i(\mathcal{N}, z, P), L_{ij}\}, & \text{otherwise.} \end{cases} \tag{2.3}$$

It follows from Eisenberg and Noe (2001) (Theorem 2) that (2.3) has a unique solution. Each agent repays at most as much as can be covered by its asset side; moreover, if an agent cannot fully meet its obligations, then it must make total payments equal to the value of its assets. Under the proportional bankruptcy rule, it does so by repaying the same fraction of its liabilities to all creditors (both bank and nonbank alike). Throughout this paper we use the proportional bankruptcy rule, but other bankruptcy rules are conceivable; for examples, see Csóka and Kondor (2020).

In what follows we will need the definition of the minimal capital injection. This is the minimal amount of cash that must be distributed among a given set of agents $C \subset \mathcal{N} \setminus \{1\}$ so that all its members remain solvent.

Let $g \in \mathbb{R}^N$ denote the vector of nonnegative cash amounts provided to the individual agents, and let g_i denote the ith element of this vector. The capital injection increases the agents' initial endowments. Hereafter, when we write the expression $\sum_{i\in C}$ (where $C \subseteq \mathcal{N}$ is any nonempty set of agents), we mean summation over the indices $1 \leq i \leq N$ corresponding unambiguously to the agents in C.

Definition 2.10. Given the financial network risk environment $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$ and a coalition of banks $C \subseteq \mathcal{N} \setminus \{1\}$, in state s we define the amount of the *minimal capital injection (henceforth simply capital injection)* required for coalition C, denoted k^s_C, as

$$k^s_C := \min_{g\in\mathbb{R}^N} \sum_{i\in\mathcal{N}} g_i$$

such that

$$p(\mathcal{N}, z^s + g, L^s)_{ij} = L^s_{ij} \quad \forall i \in C,\ \forall j \in \mathcal{N},$$

and

$$g_i \geq 0 \quad \forall i \in C \text{ and } g_j = 0 \quad \forall j \in \mathcal{N} \setminus C.$$

Observe that in k^s_C we focus exclusively on the repayment of members of coalition C; the solvency of the other agents is not a consideration. The definition does not provide a recipe for how to determine the minimal capital injection required in a given state; we return to this question in Section 2.3.1. In Definition 2.10 we defined the minimal amount of money needed to bail out the subset C of agents such that only members of C may receive a nonzero share of the capital injection. Let us temporarily consider a slightly modified definition in which we relax this latter restriction.

Definition 2.11. Given the financial network risk environment $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$ and a coalition of banks $C \subseteq \mathcal{N} \setminus \{1\}$, we define the amount of the *unrestricted minimal capital injection* required for coalition C, denoted κ^s_C, as

$$\kappa^s_C := \min_{g\in\mathbb{R}^N} \sum_{i\in\mathcal{N}} g_i$$

such that

$$p(\mathcal{N}, z^s + g, L^s)_{ij} = L^s_{ij} \quad \forall i \in C,\ \forall j \in \mathcal{N},$$

and

$$g_i \geq 0 \quad \forall i \in \mathcal{N}.$$

This raises the question of whether there can exist a financial network risk environment and a coalition of agents within it for which the unrestricted minimal capital injection is strictly smaller than the minimal capital injection, that is, whether it is cheaper to rescue a group of agents if anyone in the system may receive the injection. Equivalently, start from the minimal capital injection required for a given coalition C (Definition 2.10), and ask whether the additional funds comprising the injection can be reallocated to agents outside the coalition $\mathcal{N} \setminus C$ in such a way that we still guarantee that no member of C is in default, yet the amount of money needed to achieve this decreases. The question is significant in two respects. First, if this is possible, it is revealing about the strength of ripple effects within the system. Second, practical applicability is greatly aided if one need not compute the unrestricted minimal capital injection (Definition 2.11).

Proposition 2.12. *For any $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$ financial network risk environment and any nonempty coalition $C \subset \mathcal{N} \setminus \{1\}$, we have $\kappa^s_C = k^s_C$.*

Proof. Claim 2.13 states that any amount of cash is utilized most effectively when given directly to the members of coalition C, in the sense that this is when the coalition members' aggregate repayment $\left(\sum_{i\in C}\sum_{j\in\mathcal{N}} p(\mathcal{N}, z^s, L^s)_{ij}\right)$ increases the most. This result contradicts the existence of the previously outlined reallocation scheme, thereby proving the claim. $\square$

Proposition 2.13. *Given an amount $k > 0$, when it is optimally distributed among the members of coalition C, the coalition's aggregate repayment $\left(p_C := \sum_{i\in C}\sum_{j\in\mathcal{N}} p(\mathcal{N}, z^s, L^s)_{ij}\right)$ increases by at least as much as if the same amount k were optimally distributed among arbitrary agents in $\mathcal{N} \setminus C$. This also means that any amount of cash that would be given to members of $\mathcal{N} \setminus C$ can be reallocated to members of C in such a way that p_C does not decrease.*

Proof. Let $D \subseteq \mathcal{N} \setminus \{1\}$ denote the set of defaulted agents. Consider the effect of allocating an amount $k > 0$ to an agent $j \in \mathcal{N} \setminus C$. This is only of interest if $j \in D$ as well, that is, if j is in default. Indeed, if j was not in default even before this capital injection, it paid all its creditors in full, and thus its repayments cannot be increased. Suppose that, when this amount k is optimally distributed among the members of C, p_C increases by $a_k > 0$ (which already incorporates all ripple effects). This means that if an amount k reaches any $i \in C$ (either directly or via increased repayments of other agents), then p_C increases by at most a_k. We show that, if the amount k is granted to any agent $j \in \mathcal{N} \setminus C$, the resulting increase in p_C is $b_k \leq a_k$. If we give the amount

k to any such j, the increase in p_C can be bounded from above by the following infinite sum:

$$b_k \leq a_k \sum_{i \in C} \Pi_{ji} + a_k \sum_{k \in D \setminus C} \Pi_{jk} \sum_{i \in C} \Pi_{ki} + a_k \sum_{k \in D \setminus C, f \in D \setminus C} \Pi_{jk} \Pi_{kf} \sum_{i \in C} \Pi_{fi} + \dots$$

$$+ a_k \sum_{j_k \in D \setminus \{i\}, k=1 \dots p-1} \Pi_{jj_1} \Pi_{j_1 j_2} \cdots \sum_{i \in C} \Pi_{j_{p-1} i} + \dots . \tag{2.4}$$

where Π is the relative liabilities matrix (Definition 2.6).

To see that this is indeed an upper bound, let us mentally trace how the amount k flows through the system and what ripple effects arise (as increases in agents' repayments lead to increased repayments of other agents). Certain terms of these sums with nonnegative terms may be smaller or may vanish altogether if some banks leave the set of defaulted banks at some step. The summation ranges only over the agents in $D \setminus C$ because as soon as some part x of k reaches any member of C, the resulting increase in p_c can be upper-bounded by $x a_k$. The latter is also an upper bound for two reasons. First, the funds need not be distributed optimally; second, it may happen that D has shrunk in the meantime, leaving fewer channels to increase repayments, so the effect is at most a_k. Observe that $\sum_{i \in C} \Pi_{ji} \leq 1 - \sum_{k \in D \setminus C} \Pi_{jk}$, and similarly, in the p-th term of the sum, for the last factor of the product, $\sum_{i \in C} \Pi_{j_{p-1} i} \leq 1 - \sum_{j_p \in D \setminus C} \Pi_{j_{p-1} j_p}$. Factor out a_k from all terms:

$$(2.4) \leq a_k \Big((1 - \sum_{k \in D \setminus C} \Pi_{jk}) + \sum_{k \in D \setminus C} \Pi_{jk} (1 - \sum_{f \in D \setminus C} \Pi_{kf}) +$$

$$+ \sum_{k \in D \setminus C, f \in D \setminus C} \Pi_{jk} \Pi_{kf} (1 - \sum_{g \in D \setminus C} \Pi_{fg}) + \dots$$

$$+ \sum_{j_k \in D \setminus C, k=1 \dots p-1} \Pi_{jj_1} \Pi_{j_1 j_2} \dots \Pi_{j_{p-2} j_{p-1}} (1 - \sum_{j_p \in D \setminus C} \Pi_{j_{p-1} j_p}) + \dots \Big). \tag{2.5}$$

For clarity, introduce the following notation.

$$Q_p = \sum_{j_k \in D \setminus C, k=1 \dots p} \Pi_{jj_1} \Pi_{j_1 j_2} \dots \Pi_{j_{p-1} j_p}.$$

Using the same notation

$$(2.5) = a_k \Big((1 - Q_1) + (Q_1 - Q_2) + \dots + (Q_{p-1} - Q_p) + \dots \Big).$$

Since $0 \leq Q_{p+1} \leq Q_p \leq 1 \quad \forall p \leq N$, the telescoping series sums to at most 1, hence $b_k \leq a_k$. The proof is complete, as the above result can be trivially extended to the set of agents outside C by examining, in turn and sequentially, the effect of the amount granted to each agent.

□

In what follows, since with respect to the amount of the capital injection it is irrelevant who may (or may not) receive it, whenever we mention a capital injection we always mean its original version (Definition 2.10). It should be noted that the concept of capital injection we use differs in spirit from that applied by, for example, Demange (2018), where the aim is the optimal allocation of a given available budget. By contrast, Jackson and Pernoud (2024)'s definition of capital injection is very similar, with the difference that for us it is not only the cost of rescuing the entire system that is of interest; we have to determine this amount for every coalition.

2.2.3 Capital allocation and systemic risk indicator

Following Drehmann and Tarashev (2013), we can construct the financial network risk environment as follows, which collects the building blocks necessary to define the systemic risk capital-allocation game.

Definition 2.14. In the $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$ *financial network risk environment*, $\mathcal{N}$ is the set of agents, where agent 1 is the nonbank sector (the real economy) and the other agents are banks; S is the number of states of the world (each occurring with equal probability); z^s is the agents' initial endowment in state s; L^s is the liabilities matrix in state s (which we shall assume henceforth to be identical in every state of the world); and ρ is a coherent risk measure.

In modeling systemic risk, in each state of the world there is a given financial network, which differ in the value of the agents' nonbank assets z (i.e., in z^s). Under the proportional bankruptcy rule, it is unambiguously determined who should pay how much to whom. By choosing the bankruptcy rule, we have therefore also fixed how shocks affecting these assets propagate. For the capital-allocation game we need realization vectors, so the next decision is which parameters of the system we examine to decide how severely these shocks affected the system. For instance, we may look at the losses of nonbank creditors (Definition 2.15) or the minimal capital injection needed to rescue the agents (Definition 2.16).

Definition 2.15. Given the $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$ financial network risk environment and a coalition of banks $C \subset \mathcal{N} \setminus \{1\}$, in state s let the coalition C's *realization defined by nonbank losses*, that is, the element of the realization vector corresponding to scenario s, be

$$X(C)_s = \sum_{i\in C} \left(p(\mathcal{N}, z^s, L^s)_{i1} - L^s_{i1}\right), \tag{2.6}$$

where $L^s_{i1} - p(\mathcal{N}, z^s, L^s)_{i1}$ is the loss to the nonbank sector (the distinguished player 1) caused by bank i. It is natural to consider the loss of the nonbank sector, since measuring total losses would, in line with the reasoning of Drehmann and Tarashev (2013), amount to a kind of multiple counting. Its negative appears in the realization vector, because there negative values represent losses.

Definition 2.16. Given the $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$ financial network risk environment and a coalition of banks $C \subset \mathcal{N} \setminus \{1\}$, in state s let the coalition C's *realization defined by the minimal capital injection required for rescue*, that is, the element of the realization vector corresponding to scenario s, be

$$X(C)_s = -k^s_C, \tag{2.7}$$

where the negative of the capital injection appears, since negative values represent losses.

Definition 2.17. Given the $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$ financial network risk environment, the $(\mathcal{N} \setminus \{1\}, v)$ *systemic risk capital allocation game* has characteristic function $v : 2^{\mathcal{N}\setminus\{1\}} \to \mathbb{R}$ (defining the realization vector either by nonbank losses or by the amount of capital injection) as

$$v(C) = -\rho(X(C)), \quad C \subset \mathcal{N} \setminus \{1\}. \tag{2.8}$$

In what follows, when we refer to a systemic risk capital allocation game, this may mean either the case in which the realization vector is defined by nonbank losses or the case in which it is defined by the amount of capital injection. Wherever it matters, we explicitly state which case we have in mind; otherwise, the statements apply to both versions.

Example 2.18. Consider the $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$ financial network risk environment, where $N = \{1, 2, 3\}$ is the set of agents, agent 1 is the nonbank sector, $S = \{1, 2\}$ is the set of states with $p_1 = p_2 = 0.5$, and ρ is the maximum-loss risk measure (i.e., ES_1, the Expected Shortfall with $k = 1$: the worst single outcome, as defined in Definition 2.3). λ^* denotes the (equilibrium) repayment rate under the proportional bankruptcy rule.

i	z^1_i	L^1_{ij}			λ^*	a_i	e_i	$p(\mathcal{N}, z^1, L^1)_{i1} - L^1_{i1}$
agent 1	0	0	0	0		43	4.3	
agent 2	1.9	1	0	3	0.7	2.8	0	-0.3
agent 3	2.4	4	1	0	0.9	4.5	0	-0.4

Table 2.1: The financial network risk environment in state 1.

Note that in state 1, agent 2 (bank) is in fundamental default[28], whereas agent 3 (bank) defaults due to contagion; it cannot pay only because it does not receive all of its claims.

i	z_i^2	L_{ij}^2			λ^*	a_i	e_i	$p(\mathcal{N}, z^2, L^2)_{i1} - L_{i1}^2$
agent 1	0	0	0	0		4.5	4.5	
agent 2	1.4	1	0	3	0.6	2.4	0	-0.4
agent 3	5	4	1	0	1	6.8	1.8	0

Table 2.2: The financial network risk environment in state 2.

Note that in state 2, agent 2 is in fundamental default, whereas agent 3 is not in default.

Using equation (2.8), we obtain the capital allocation game shown in Table 2.3 (Table 2.4), provided that the realization vector is determined by Definition 2.15 (Definition 2.16). Observe that the game is superadditive—meaning the value of a combined coalition is at least as large as the sum of its parts—in both cases (in terms of systemic risk: the sum of the risks of the parts is greater than the risk of the system as a whole), and the diversification effect that appears can be apportioned using the tools of cooperative game theory.

C	$\{1\}$	$\{2\}$	$\{1,2\}$
$X^1(C)$	-0.3	-0.4	-0.7
$X^2(C)$	-0.4	0	-0.4
$\mathrm{ES}_1(C)$	0.4	0.4	0.7
$v(C)$	-0.4	-0.4	-0.7

Table 2.3: The systemic risk capital allocation game when the realization vector is determined by bank losses.

C	$\{1\}$	$\{2\}$	$\{1,2\}$
$X^1(C)$	-1.1	-0.425	-1.1
$X^2(C)$	-1.6	0.0	-1.6
$\mathrm{ES}_1(C)$	1.6	0.425	1.6
$v(C)$	-1.6	-0.425	-1.6

Table 2.4: The systemic risk capital allocation game when the realization vector is defined via capital injection.

In what follows we show that, when risk is measured via capital injection, a diversification effect arises, that is, the sum of the risks of two disjoint coalitions

[28]The concepts of fundamental default and contagion default are formally defined in Chapter 3. Informally, an agent is in fundamental default if its external assets alone are insufficient to cover its obligations even if all counterparties pay in full; contagion default occurs when an agent would be solvent under full payment but defaults because counterparties fail to pay.

cannot be smaller than their joint risk. Since the characteristic function is the negative of risk, the induced cooperative game is superadditive.

Proposition 2.19. *In the financial network risk environment* $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$*, with a realization vector defined by capital injection, and for disjoint sets of banking agents* $C, T \subset \mathcal{N} \setminus \{1\}$*:*

$$v(C \cup T) \geq v(C) + v(T). \tag{2.9}$$

Proof. We know that

$$k^s_C + k^s_T \geq k^s_{C\cup T} \quad \forall T, C \subset \mathcal{N} \setminus \{1\}, \quad \forall s \in S.$$

That is,

$$X(C \cup T)_s \geq X(C)_s + X(T)_s \quad \forall T, C \subset \mathcal{N} \setminus \{1\}, \quad \forall s \in S.$$

Since this holds in every state (i.e. componentwise), using the monotonicity of a coherent risk measure:

$$\rho(X(C \cup T)) \leq \rho(X(C) + X(T)) \quad \forall T, C \subset \mathcal{N} \setminus \{1\}.$$

Because of the subadditivity of a coherent risk measure:

$$\rho(X(C \cup T)) \leq \rho(X(C) + X(T)) \leq \rho(X(C)) + \rho(X(T)) \quad \forall T, C \subset \mathcal{N} \setminus \{1\}.$$

□

The claim, of course, also holds when the realization vector is defined by nonbank losses; the proof follows trivially.

Having defined the systemic risk game, we still need a method by which we measure each participant's contribution to systemic risk. This is possible by distributing the value of the grand coalition, using a solution concept known from cooperative game theory.

In what follows we will apply the Shapley value (Shapley (1953)). This is the unique allocation principle that satisfies equal treatment, strong monotonicity, and well-definedness, but it is not necessarily in the core, hence coalitional acceptability may be violated when using it. Moreover, it is an allocation, hence individually acceptable: it does not assign to any agent (bank) more risk than what the agent bears alone.

Definition 2.20. Given the financial network risk environment $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$, let $\mathcal{N}_{-1} := \mathcal{N} \setminus \{1\}$ and $k := |\mathcal{N}_{-1}|$. The Shapley value of agent $i \in \mathcal{N}_{-1}$ in the game $(\mathcal{N}_{-1}, v)$ is

$$\phi_i(v) = \sum_{C \subseteq \mathcal{N}_{-1}\setminus\{i\}} \frac{|C|!\,(k - |C| - 1)!}{k!} \big(v(C \cup \{i\}) - v(C)\big).$$

We define the *systemic risk indicator* as

$$\varphi_i := -\phi_i(v),$$

so that $\sum_{i\in\mathcal{N}_{-1}} \varphi_i = \rho(X(\mathcal{N}_{-1}))$.

2.2.4 Relation to official systemic-importance frameworks

The systemic-importance indicator constructed above should be compared with official frameworks only with care. G-SIB and O-SII methodologies are broad structural designation systems. They are designed to be transparent, reproducible and implementable with supervisory data, and they capture institutional features that are absent from the present model. The chapter instead isolates the contribution of network-clearing effects and coalition-level recapitalization needs in a stylized liability network.

Dimension	**Official G-SIB / O-SII logic**	**Chapter 2 logic**
Objective	Resilience of structurally important institutions	Attribution of ex-ante risk of recapitalization needs
Information set	Size, interconnectedness, substitutability or importance, complexity, cross-border activity and supervisory judgment	Liability network, scenario distribution, clearing rule and coherent risk measure
Timing	Ex-ante structural designation	Scenario-wise ex-post injections aggregated ex ante
Mechanism	Broad systemic relevance and implementability	Direct solvency spillovers, clearing effects and coalition effects
Policy role	Official buffer and designation methodology	Conceptual complement; not a replacement

Table 2.5: Comparison between official systemic-importance frameworks and the Chapter 2 indicator.

A bank can be modest in official structural indicators yet appear important in this chapter if its position in the modeled liability network makes it contribute strongly to the coherent risk measure of coalition recapitalization needs. Conversely, a bank may rank highly in official frameworks because of size, substitutability, complexity or cross-border relevance even if this particular direct-solvency channel is less dominant. Different rankings are therefore not a contradiction; they reflect different questions and different information sets.

2.3 Implementation, comparison of systemic risk indicators

In what follows, we describe a detail that is critical for practical applicability. Definition 2.10 did not provide a recipe for computing the minimal capital injection required in each scenario. We will see that if the capital injection is specified according to Definition 2.10, the task is extremely simple. Therefore, Proposition 2.12 is pivotal: regardless of whether we define the capital injection via 2.10 or via 2.11, the total amount required is the same.

2.3.1 Determining the minimally required capital injection

Consider the $(\mathcal{N}, S, \{z^s\}_{s\in S}, \{L^s\}_{s\in S}, \rho)$ financial network risk environment and fix a state $s \in S$. For notational simplicity, write $(\mathcal{N}, z, L) := (\mathcal{N}, z^s, L^s)$ for the realized financial system in this state. Let $C \subset \mathcal{N} \setminus \{1\}$ be a coalition of banking agents; the aim is to determine the minimal amount of cash and its allocation that guarantees the solvency of every agent in C. Under Definition 2.10, no agent outside C may receive any capital injection, hence the injection must be chosen such that every agent in C can meet all its obligations in full. In particular, once the injection is made, each member of C pays its liabilities in full, so that for the purpose of determining repayments by agents outside C it suffices to know that C is solvent, without yet knowing the exact allocation of the injection within C.

Fix the scenario, and modify the financial network $(\mathcal{N}, z, L)$ realized in this scenario as follows:

- In the new system the set of agents is $I^C := \mathcal{N} \setminus C$.
- Define a liability matrix $L^C \in \mathbb{R}^{|I^C|\times|I^C|}$ by first removing the rows and columns of L corresponding to the agents in C, and then reallocating liabilities towards agents in C as if they were owed to the nonbank sector (agent 1):

$$L^C_{ij} := \begin{cases} L_{ij}, & i \in I^C,\ j \in I^C,\ j \neq 1, \\ L_{i1} + \sum_{m\in C} L_{im}, & i \in I^C,\ j = 1. \end{cases}$$

  This reallocation is permissible since repayments made by the real economy are not modeled.

- The new initial endowments are $z^C \in \mathbb{R}^{|I^C|}$, where
$$z_i^C := z_i + \sum_{m \in C} L_{mi} \quad \forall i \in I^C,$$
since all members of C pay in full as a consequence of the capital injection.

The financial network (I^C, z^C, L^C) thus created has the same properties, relevant for the existence and uniqueness of the clearing matrix, as the original financial network; hence a clearing matrix can be computed for it in the same way. Let
$$p^C := p(I^C, z^C, L^C)$$
denote the (unique) clearing matrix corresponding to the proportional bankruptcy rule in the modified financial system.

For each $k \in I^C$, define its total liabilities and total payments in the reduced system by
$$\bar{\ell}_k^C := \sum_{m \in I^C} L_{km}^C, \qquad \bar{p}_k^C := \sum_{m \in I^C} p_{km}^C,$$
and set the repayment fraction
$$\lambda_k := \begin{cases} 1, & \bar{\ell}_k^C = 0, \\ \dfrac{\bar{p}_k^C}{\bar{\ell}_k^C}, & \bar{\ell}_k^C > 0. \end{cases} \tag{2.10}$$
Under the proportional bankruptcy rule, $\lambda_k \in [0, 1]$ and k repays the same fraction λ_k of each of its liabilities. In particular, although the reduced clearing matrix p^C is only defined on indices in I^C, the payment that an agent $k \in I^C$ makes to a creditor $j \in C$ in the original network is
$$\lambda_k L_{kj}.$$

Knowing the repayment fractions $\{\lambda_k\}_{k \in I^C}$, determining the cash amount required by each member of C becomes a direct balance-sheet computation. Indeed, agent $j \in C$ is solvent (and hence pays in full) if and only if its available resources,
$$z_j + x_j + \sum_{l \in C} L_{lj} + \sum_{k \in I^C} \lambda_k L_{kj},$$
cover its total liabilities $\sum_{i \in \mathcal{N}} L_{ji}$. Therefore, the minimally required injection into each $j \in C$ is
$$x_j = \max\Big(\sum_{i \in \mathcal{N}} L_{ji} - z_j - \sum_{k \in I^C} \lambda_k L_{kj} - \sum_{l \in C} L_{lj},\ 0\Big) \qquad \forall j \in C. \tag{2.11}$$
That is, claims against agents in C are received in full (since the injection renders them solvent), whereas claims against agents outside C are received

proportionally according to the repayment fractions $\{\lambda_k\}_{k\in I^C}$ inferred from the reduced clearing problem.

Therefore, the total capital injection required to rescue coalition C is

$$k_C^A = \sum_{j\in C} x_j. \tag{2.12}$$

With this method, determining the optimal (restricted) capital injection is a computational problem comparable to computing a clearing matrix. Naturally, the procedure must be carried out for every state of the world and for every coalition of agents. Under the proportional bankruptcy rule, computing the clearing matrix can in general be solved by a fixed-point iteration, whose convergence can be accelerated, for example, by the method of Irons and Tuck (1969). Another option is to apply the algorithm of Rogers and Veraart (2013).

2.3.2 Simulation example

The purpose of the following example is to anticipate the practical applicability of what has been presented so far, and to illustrate how consequential the choice of method for defining the realization vector can be.

We assume that the only stochastic component in the system is the agents' initial endowment (say, between an arbitrary initial point in time 0 and horizon H, where claims are due. We express this by the decomposition $z_i = e^{Y_i}\bar{z}_i$, where $\bar{z} \in \mathbb{R}_+^N$ is deterministic, and $Y \in \mathbb{R}^N$ denotes the random vector representing shocks to the initial endowment. We generate 200000 scenarios by Monte Carlo simulation (from the joint distribution of the shocks, Y), and then use 1000 bootstrap samples to construct confidence intervals for the relevant quantities.

The system consists of the nonbank sector and 7 banking agents: one central agent, 3 indistinguishable lenders, and 3 indistinguishable borrowers. These agents are connected to each other only through the central agent. Thus, the borrower banks all have liabilities of identical size towards the central agent, who in turn is indebted to the lenders. Expressed via the relative liabilities matrix, the magnitudes of these connections are $\Pi_{borrower,central} = r$ and $\Pi_{central,lender} = r$. Table 2.6 contains the relevant parameters of the system. The effect of increasing connectedness is studied in a similar fashion by Nier et al. (2007).

Regardless of the value of r, the following always hold. Taking the pre-shock values of the initial endowments into account, the agents have a 6% capital adequacy ratio: $\frac{\bar{e}_i}{\bar{e}_i + l_i} = 6\% \quad \forall i \in \mathcal{N}_{-1}$, where $\bar{e}_i = \bar{z}_i + \sum_{j\in\mathcal{N}} L_{ji} - l_i$. The liability towards the nonbank sector is constant, independently of the role an agent plays in the system: $L_{i1} = 87 \quad \forall i \in \mathcal{N}_{-1}$.

	$\bar{z}_i$	l_i	$\sum_j \Pi_{j,i}$	$\sum_j \Pi_{i,j}$	L_{i1}	P_{fund}
central	$(l_c - 0,94 * 3 * l_b)/0,94$	$87/(1 - 3 * r)$	$3r$	$3r$	87	0,05
lender	$l_b/0,94$	$87/(1 - r)$	0	r	87	0,05
borrower	$(l_l - r * l_c)/0,94$	87	r	0	87	0,05

Table 2.6: The system used in the example. Source: authors' own construction

We assume a proportional bankruptcy rule, under which obligations towards all parties (banking agents as well as members of the real economy) have the same seniority. In every scenario $s \in S$ we require the realized values of the initial endowments. For the random variable Y_i representing the log-return of the shock to agent i's asset side, we assume a normal distribution with mean 0 and standard deviation σ_i: $Y_i = \sigma_i \left(\sqrt{\beta}\zeta + \sqrt{1-\beta}\xi_i\right)$, where ξ_i, ζ, the common market and idiosyncratic factors, are standard normal and pairwise independent. We set the volatilities, σ_i, so as to obtain predetermined, fixed fundamental default probabilities (last column of Table 2.6, P_{fund}). In doing so, we assume an investment policy that considers only the riskiness of nonbank assets and ignores counterparty risk on interbank claims, thereby not accounting for the systemic risk arising from agents' interconnectedness.

In this example, the correlation β of the agents' investment returns is 60% in all cases. We define the realization vectors both by the sum of nonbank losses and by the sum of capital injections, and then use the 2% expected shortfall as the risk measure. In the first case, the value of each coalition is the average loss imposed on the nonbank sector by the members of the coalition in the worst 2% of cases. In the second case, it is the average amount of capital injection needed to rescue the coalition in the worst 2% of cases.

Figure 2.1 shows the effect of increasing connectedness (via increasing r) on the systemic risk indicator of the different types of agents (central, lender, borrower). The height of a given bar thus displays the average systemic risk indicator for an agent belonging to the given category (since, due to the simulation, there may be minimal differences among, for example, the three lenders). The figure also includes the 95% bootstrap confidence intervals for the indicators.

The overall risk of the system will, naturally, differ depending on which method we apply to define the realization vector. It is much more interesting to observe how differently the two methods assign systemic risk significance to the central agent.

As connectivity increases, the systemic risk indicator of the central agent also rises; this is common to both methods. However, the magnitude of the increase differs substantially.

As connectivity grows, the central agent is more likely to incur losses on its (creditor-facing) obligations that are large enough to prevent it from meeting

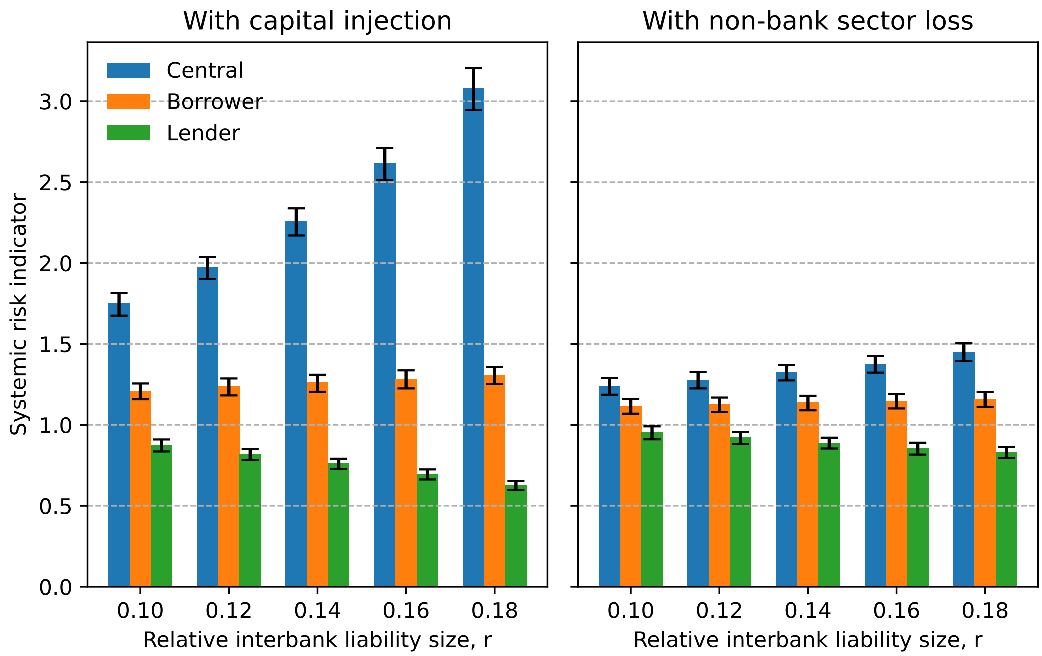

Figure 2.1: systemic risk indicators as a function of r.

all of its debts; the probability of contagion default increases. Since it is the only agent that holds multiple within-system claims, contagion default will affect it the most severely (in the example, the correlation of nonbank asset returns is non-negligible).

Let us first consider the case where the realization vector is defined by nonbank losses. The values of nonbank assets are identical for all agents throughout. Therefore, as we increase connectivity, the weight of nonbank liabilities declines for the central agent. Thus, even though it plays a central role in the propagation of contagion, the rising probability of contagion default can increase its systemic risk indicator only to a small extent.

By contrast, when the realization vector is given by capital injection, increasing connectivity raises the central agent's systemic risk indicator much more dynamically. The larger losses resulting from greater connectivity can be neutralized only by ever larger capital injections.

If we focus exclusively on the losses of nonbank creditors, the central agent's importance is not outstanding, even under high connectivity. If the stability of the system as a whole is at issue (risk measured by capital injection), the central agent's critical role becomes striking, especially at high connectivity.

Although the model is capable of handling circular liabilities, this does not appear in the present example. Instead, the choice of structure is motivated by the markedly different roles of the three types of banking agents and, consequently, by the demonstration of their differing systemic risk significance.

From a supervisory-reading perspective, the divergence between the two indicator families should be interpreted as a difference in the object being measured. Loss-based indicators capture damage imposed on creditors after a failure, while the capital-injection object captures the scenario-specific amount needed to restore solvency. The indicator is nevertheless ex ante, because it forms state-by-state realization vectors of those injections, applies a coherent risk measure, and allocates the resulting systemic-risk game. The simulation example shows that a direct-solvency channel can produce rankings that differ from creditor-loss measures and from official G-SIB or O-SII tools. It does not imply that regulators should mechanically replace existing surcharges with the chapter's indicator, nor that the model calibrates the right level of prudential capital in practice. A real policy implementation would require broader institution-level information than the model contains, including critical functions, substitutability, legal structure, cross-border activity, supervisory judgment and the macroeconomic costs of capital regulation.

2.4 Conclusion and outlook

We have presented a model suitable for measuring systemic risk and the contributions of individual participants, after having set out the building blocks required for this. As an alternative to the usual methods focusing on creditors' losses, capital injections can also be used to measure riskiness. After presenting a result related to this, we demonstrated the difference between the two approaches through a simulation example. The latter result foreshadows the critical decisions to be made in a possible practical application.

The current model contains substantial simplifications at several points. The dynamics of shock propagation and the scope for regulatory intervention are fundamentally different if we allow for the presence of default costs. Another simplifying assumption that may be worth relaxing is the identical seniority level of all liabilities.

It is likewise a simplification that the network of liabilities is exogenous and fixed. Once a method for measuring systemic risk has been fixed, an obvious research direction is to consider interventions that affect the structure of the system. One such intervention may aim, for example, at minimizing systemic risk, with portfolio compression being one possible means. For the latter see D'Errico and Roukny (2021) or Schuldenzucker and Seuken (2020). Since portfolio compression can be reduced to the same graph-theoretic problem as, for example, kidney exchange programs, results related to the latter (see, for instance, Biró et al. (2021)) may significantly aid this line of research. In addition, although in the example presented we held the capital adequacy ratio fixed, our model can also be used to examine how the systemic risk indicator depends on the capital adequacy ratio.

Chapter 3

Correlated external shocks in a stylized financial network: Bayesian network and mean-field approaches

3.1 Introduction

In this chapter, like in other chapters in this dissertation, the model considers obligations due at a single maturity. The distinguishing feature here is the probabilistic, ex-ante perspective in which the stochastic nature of external assets is modeled explicitly. While the previous chapter took asset realizations as given and focused on optimal capital injections once those were given, here we analyze how the correlation structure of shocks interacts with network topology to determine default probabilities.

In a stylized financial system, a network of claims and obligations connects the balance sheets of financial institutions (hereafter referred to as agents). In addition to these claims and obligations, agents are endowed with external assets representing their investment portfolios, which are exposed to random shocks. Whether agents can fulfill their obligations or default therefore depends on two factors: incoming payments from their counterparties and the realization of their external assets. Accordingly, we consider two sources of connectedness that affect agents' default probabilities. Since payment obligations between agents serve as channels of contagion, with default cascades arising from failures to meet payment obligations, the network of payment obligations constitutes the first and most direct form of connection among agents. The second type of connectedness lies in the correlation structure of shocks affecting different agents' external assets; this affects joint default probabilities even in the absence of direct payment obligations. The aim of this chapter is to study the interplay of these two sources of systemic risk in a stylized setting. Assuming a one-factor Gaussian copula model for the shock structure, we focus on how default probabilities vary with the level of pairwise correlation between agents' external assets. We demonstrate how these two sources of connectedness are closely coupled in determining agents' default probabilities and thus the robustness of the system.

The chapter deliberately isolates only two channels: direct default contagion through fixed nominal obligations and common or correlated shocks to external assets. It does not model fire sales, funding runs, margin calls, collateral dynamics, payment-system settlement, endogenous network formation or real-economy feedback. The results should therefore be read as a controlled solvency-network dependence analysis rather than as a full supervisory stress test.

First, we analyze systems consisting of uniform agents under fully connected and ring-like network topologies. By exploiting the special structure of these networks, we introduce analytical calculations for key quantities, enabling us to compute conditional individual and joint default probabilities without relying on simulations. Taking a different perspective, we then examine the mean-field limit, which describes the system's behavior when it consists of infinitely many microscopic agents. These two approaches yield consistent main conclusions. In the homogeneous one-period setting studied in this chapter, agents' default probabilities are non-monotonic in the correlation. Under reasonable parameterizations, a characteristic hump shape appears: default probabilities increase as correlation rises, reaching a peak at relatively high correlation values, and then rapidly decrease as correlation approaches one. The same effect holds for systemic default probabilities (defined in our study as at least two-thirds of agents suffering default) under the analytical approach. Larger interbank connectivity increases the maximal default probability while simultaneously pushing the peak of the hump toward smaller correlation values. Lower expected external asset levels lead to larger fundamental default probabilities, which refer to the default probability of agents in the absence of contagion effects from the network. In addition, lower expected external asset levels produce a more pronounced hump shape with larger maximal default probabilities. The mechanism is the cross-sectional heterogeneity of external asset values at the terminal clearing date: as correlation rises from low levels, diversification across counterparties weakens and default probabilities increase, but when correlation becomes very high, the dispersion that makes pure contagion defaults possible disappears. Under a fully homogeneous setup, those contagion states vanish altogether at $\rho = 1$. This increased sensitivity makes the financial system more susceptible to changes in the correlation structure between the shocks affecting different agents' external assets. As a result, understanding this correlation structure is critical for assessing the stability and resilience of the financial system. To quantify how far this homogeneous benchmark survives once ex-ante symmetry is relaxed, we later complement the analytical results with a Monte Carlo robustness experiment in which bank-specific distances to default are heterogeneous across agents.

Highly correlated asset sides are a common assumption in the literature, but perfect correlation is quite unrealistic in practice. In our model, for most parameterizations, the hump shape is such that default probabilities peak for correlation values around 0.8–0.9. Note that this is precisely the region of typical correlations between banks' assets, as Banerjee and Feinstein (2022) reports. On the other hand, in our homogeneous one-period setup there is a significant and sudden reduction in default probabilities for correlations close to 1, precisely because pure contagion defaults vanish in the comonotonic limit. This result suggests that by assuming perfectly correlated endowments, one might miss an important effect and seriously underestimate default probabilities. Overall, our results highlight the importance of analyzing the effect of correlated endowments and direct connections together.

The chapter is organized as follows. Section 3.2 discusses the relevant literature. In Section 3.3 we present the balance sheet dynamics and other model assumptions. In Section 3.4, we derive formulas for calculating default probabilities using the Bayesian network approach. In Section 3.5 we introduce the mean-field approximation for the large N limit. In Section 3.6 we discuss the results, and Section 3.7 concludes.

3.2 Related literature

Both channels have been studied extensively in the literature, though mostly separately. First, numerous studies examine the effect of network structure on systemic risk. Allen and Gale (2000) find that complete networks are more resilient to the spread of liquidity shocks than less connected ones. Battiston et al. (2012a) find that increasing diversification through greater connectivity does not always enhance stability—for instance, when fundamental default probability is heterogeneous across agents but the average fundamental default probability is low and there is an initial large enough shock.

A large body of literature addresses systemic risk arising from asset-side correlation across agents, without a direct channel of contagion. As our study investigates (joint) default events and their probabilities in a network setup, it connects to the broader field of dependent defaults. For a discussion from a portfolio credit risk perspective, see Frey and McNeil (2003). In the model of Caccioli et al. (2014), as the level of diversification of banks' portfolios increases, the system undergoes two phase transitions, with a region in between where global cascades occur. In the model of Lehar (2005), if banks have highly correlated asset portfolios, the probability of multiple defaults is high, making positive correlation undesirable. In Acharya (2009)'s model, if banks choose the correlation of their investment portfolios endogenously, they have an incentive for risk shifting to the regulator by investing in correlated portfolios, thus increasing aggregate risk across the economy. Additional relevant studies on systemic risk arising from investment similarities include Poledna et al. (2021), Gualdi et al. (2016), Delpini et al. (2019), and Raffestin (2014).

There are also several empirical analyses aimed at supporting the theoretical results in the literature. Examining syndicated loan portfolios, Cai et al. (2018) finds that bank-level diversification is the main driver of interconnectedness in the form of overlapping asset portfolios and that interconnectedness is positively associated with various systemic risk measures. In a similar empirical analysis, Lee et al. (2020) show that high correlation between banks' portfolios increases the systemic risk of the banking sector.

We believe only a few studies focus on the interplay of correlated endowments and direct channels of contagion. Li and Wang (2021) investigates how portfolio diversification affects banking systemic risk while accounting for the interbank network. Depending on the shock structure, when portfolio diversification is low, increasing the level of portfolio diversification might initially increase or decrease banking system resilience. On the other hand, when portfolio diversification exceeds a threshold, further increases enhance the resilience of the banking system. Elliott et al. (2014) study how cascades of failures in a network of interdependent financial organizations depend on network structure. They find that integration (larger interbank exposures) and diversification (more counterparties per organization) have different, non-monotonic effects on the extent of cascades. In a variant of their model, they also consider the case where agents have investments with correlated payoffs. Their result—that increasing correlation increases default probability, with the increase occurring abruptly at a certain level of correlation—contrasts with some of our findings. This is likely attributable to the somewhat different model dynamics they use, which we discuss later.

In the majority of models focusing on the effects of asset correlation, banks, in an effort to create diversified portfolios, end up with highly overlapping, correlated asset holdings. Although individual diversification reduces risk, overlapping portfolios between banks increase systemic risk. Tasca et al. (2014) separates two market regimes: a safe one in which the systemic impact of higher leverage (higher returns but higher default risk) can be mitigated by better diversification, and a risky regime in which this is not possible, resulting in additional systemic risk through correlated asset sides. Diversification through spreading exposure across more and more assets will likely affect asset-side variance as well. To isolate that effect and focus solely on the effect of correlation between different agents' assets, unlike the studies above, we keep endowment variance constant. This separation is made possible by treating correlation between different agents' endowments as an exogenous property rather than an artifact of some investment policy.

Most similar studies employ either Monte Carlo simulations or utilize asymptotic analytical calculations in the large-system limit. We present methods to derive semi-analytical results that can be efficiently calculated for relatively small systems. We wish to employ a proper clearing rule in the spirit of Eisenberg and Noe (2001) within a stochastic endowment setting. The problem of calculating systemic risk in such a stochastic setup first appears in Gourieroux et al. (2012). As Banerjee and Feinstein (2022) points out, tackling this problem without Monte Carlo simulation generally requires computing the measure of 2^N regions for N agents, corresponding to partitioning the endowment space into 2^N possible default scenarios. In a similar setting, Barucca et al. (2020) introduces some approximations to handle the problem and only examines the independent endowment case in actual calculations. Banerjee and Feinstein (2022) also works with a stochastic generalization of the Eisenberg and Noe (2001) model and derives debt and equity valuation formulas under the assumption of comonotonic endowments. They argue that such an assumption is reasonable from a theoretical point of view (if all firms are portfolio optimizers and do not take into account any other firm's investments, the chosen endowments will be comonotonic) and empirically as well (banks' assets exhibit a high degree of rank correlation). For our semi-analytical calculations, we employ a conditional Bayesian network approach (Chong and Kluppelberg, 2018) that exploits the symmetries of certain system structures. We also present results for the large-system limit. Among many general results, Hurd (2016) derives mean-field solutions for the Gai et al. (2011) model in random networks. Hurd et al. (2014) presents more complex balance sheet dynamics and employs mean-field asymptotic calculations. Lorenz et al. (2009) provide a unifying framework for failure cascade models and derive mean-field equations to calculate the total fraction of failed nodes in the system. There is also a growing literature on continuous-time versions of the mean-field approach, ranging from simple interactions as in Carmona et al. (2015) to complex strategic interactions that allow agents to choose their connections as in Amini et al. (2022).

3.3 A model with correlated endowments

We consider a two-period model consisting of time 0 and time $H = 1$ (year). In our stylized setup, the financial system consists of N financial institutions, referred to as agents in what follows, and $I = \{1, 2, \ldots, N\}$ denotes the set of agents. The balance sheet of agent i at time horizon H, following Demange (2016), is as follows.

Balance sheet of agent i at time 1	
assets	liabilities
A_i $\sum_{j=1}^{N} L_{ji}$	ℓ_i $\sum_{j=1}^{N} L_{ij}$ E_i

Definition 3.1. A matrix $L \in \mathbb{R}^{N \times N}$ represents the value of *interbank payment obligations* due at time 1, which are deterministic and known at time 0 as well. These obligations are best interpreted as unsecured interbank loans or deposits, or more cautiously as reduced-form net unsecured bilateral exposures fixed at the modeling horizon. Cross-exposures can give rise to counterparty risk (Tirole (2011)), but the present model does not represent collateralization, close-out netting, margining, settlement conventions, contingent derivative payoffs or CCP novation. Interpretations based on repo, IRS or CDS contracts are therefore outside the current scope except as highly abstract net-exposure summaries (Basel Committee on Banking Supervision and International Organization of Securities Commissions, 2015; Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions, 2012). We have $L_{ij} \geq 0$ for all $i, j \in I$. The entry L_{ij} denotes agent i's liability toward agent j, and $L_{ii} = 0$ for all $i \in I$.

Definition 3.2. $A_i \geq 0$ denotes the realization of agent i's *external assets* (excluding interbank claims). For brevity, external assets are occasionally also referred to as *endowments*. This is the only item through which uncertainty enters the system. One might argue that throughout the period, the agent's endowment is exposed to credit and market risk, modeled by random shocks. Thus, the time 1 value of i's endowment is a random variable[29], denoted by $\mathcal{A}_i$, while a realization of it is A_i.

Definition 3.3. $\ell_i \geq 0$ is agent i's payment obligation at time 1 toward entities other than the agents introduced above. This deterministic item, referred to as *nonbank liabilities* in the following, might consist of, for example, deposits (retail or corporate). We assume that these claims have the same seniority as the agents' interbank liabilities.

The network of interbank claims can be represented by a directed graph. An edge from vertex i to vertex j means agent i has a payment obligation

[29] See Assumption 3.10

due at time 1 toward agent j. These claims are not netted, and in general, directed cycles may exist. We assume that the interbank liability structure is given exogenously. Agents cannot choose their capital structure or the partners to which they are connected. Acemoglu et al. (2015a) and Bluhm et al. (2014) relax this assumption and consider an endogenous interbank lending structure.

Agent i's total payment obligation at time 1 is

$$b_i = \ell_i + \sum_{j=1}^{N} L_{ij}.$$

Agent i's total cash available to satisfy its obligations, assuming that everyone else pays in full, is thus

$$a_i = A_i + \sum_{j=1}^{N} L_{ji}.$$

Definition 3.4. Agent i's *net worth* is given by

$$E_i = a_i - b_i.$$

We assume that agents pay in full or their creditors recover 0 on their claims. Hence, we adopt the zero recovery assumption of Gai and Kapadia (2010), also used by Jackson and Pernoud (2024). This is justified by the potentially short time period considered; on the short run, no liquidity can be hoarded (no overnight loan or repo market). We also assume total loss upon default (zero recovery, i.e., creditors of a defaulting agent lose their entire claim). Therefore, the system (or more precisely, its time 1 structure) is equivalent to that of Rogers and Veraart (2013), with $\alpha = 0$ and $\beta = 0$.

An agent's ability to fulfill its obligations is affected not only by the end-of-period value of A_i but also by other agents' behavior (non-reimbursed liabilities toward i). We distinguish between fundamental and contagion default.

Definition 3.5. The event of *fundamental default* of agent i, denoted F_i, is

$$F_i \;:=\; \big\{E_i = a_i - b_i < 0\big\}, \tag{3.1}$$

that is, the event that agent i would have negative net worth even if every other agent paid in full.

Definition 3.6. Let $p^T = (p_1, p_2, \dots, p_N)$ denote a *survival vector*, where $p_i = 1$ if agent i pays all its obligations and 0 otherwise. Assuming that p describes the survival state of the system, agent i receives $\sum_{j=1}^{N} p_j L_{ji}$ from the other agents in total.

Since agents may have mutual claims on each other and cycles may exist in the system of obligations, we need a clearing algorithm to determine the equilibrium payments, i.e., the equilibrium survival vector p^*. In the following, we rely on Rogers and Veraart (2013) to formalize the definition of this equilibrium payment scheme. We assume this clearing is performed in a centralized way, a common assumption that is relaxed, for example, by Csóka and Herings (2018).

Definition 3.7. We define the *equilibrium survival vector* as a $p^* \in \{0,1\}^N$ vector such that

$$\Phi(p^*) = p^*$$

where $\Phi : \{0,1\}^N \to \{0,1\}^N$ with

$$\Phi(p)_i = \begin{cases} 1 & \text{if } b_i \leq A_i + \sum_{j=1}^{N} p_j L_{ji} \\ 0 & \text{otherwise} \end{cases}$$

Definition 3.8. The event of *contagion default* of agent i, denoted C_i, is

$$C_i \;:=\; \left\{p_i^* = 0\right\} \setminus F_i, \tag{3.2}$$

that is, the event that agent i defaults in the clearing equilibrium ($p_i^* = 0$) despite not being in fundamental default ($E_i \geq 0$). In other words, agent i is unable to repay its obligations solely due to spillover effects from other agents' failures.

Once the endowments $(A_i)_{i\in I}$ are realized, the time-1 clearing problem is a fixed-point problem in the sense of Eisenberg and Noe (2001). In the classical case without default costs (the Eisenberg–Noe model), the payment operator (with non-zero recovery, the quantity of interest is the payment agents finally make, which in our case follows trivially from the survival indicators) is monotone on the complete lattice[30] $[0, b] \subset \mathbb{R}^N$ under the componentwise order. By Tarski's fixed-point theorem[31], the set of clearing payment vectors is nonempty and, in fact, forms a complete lattice. In particular, there exist a smallest and a largest clearing vector. In the Rogers and Veraart (2013) extension, which allows proportional default costs on external assets and on interbank assets via parameters $\alpha, \beta \in [0, 1]$, existence is preserved and the order-theoretic picture remains: there exist least and greatest clearing vectors. In general, however, clearing vectors need not be unique in the presence of default costs; multiple clearing vectors may arise, so working with the greatest element is a natural selection rule.

[30] A complete lattice is a partially ordered set in which every subset has both a least upper bound (supremum) and a greatest lower bound (infimum).

[31] Tarski's theorem: a monotone map on a complete lattice has a complete lattice of fixed points, hence least and greatest fixed points; see, e.g., Tarski (1955).

Note that our binary survival formulation ($p \in \{0,1\}^N$ with the monotone map Φ defined above) is an order-preserving self-map of the finite lattice $\{0,1\}^N$. The same Tarski argument implies that the set of survival equilibria $\{p : \Phi(p) = p\}$ is a nonempty complete lattice with a smallest and a largest element. Once the endowments are realized at time 1, the problem of finding the equilibrium survival vector is the same fixed-point problem already treated by Rogers and Veraart (2013); the fictitious-default procedure we adopt selects the largest survival equilibrium, which corresponds to the greatest clearing vector in their framework.

Throughout, we order payment (or survival) vectors componentwise. For two payment vectors $x, y \in \mathbb{R}^N_{\geq 0}$, we write $x \leq y$ if and only if $x_i \leq y_i$ for every $i \in I$; for survival vectors $p, q \in \{0,1\}^N$, we use the same coordinatewise order $p \leq q$ if and only if $p_i \leq q_i$ for all i. Intuitively, y is larger than x if every creditor of every bank receives at least as much under y as under x (no one is worse off), and q is larger than p if it has (weakly) more survivors. With this order, the set of clearing fixed points is a (complete) lattice: there exists a least clearing vector $x_{\min}$ and a greatest clearing vector $x_{\max}$, where "greatest" means $x \leq x_{\max}$ for every other clearing vector x (componentwise dominance). In particular, $x_{\max}$ yields the largest feasible systemwide recoveries and the largest survival set, while $x_{\min}$ yields the smallest. The order is partial (not total): two payment vectors may be incomparable if one pays more to some creditors but less to others.

There is no strategic interaction among the agents; clearing in default is assumed to be carried out by an outside agent (regulator) with complete information about the system. Moreover, we assume that the greatest equilibrium survival vector is selected. Rogers and Veraart (2013) introduce a fictitious default algorithm to arrive at the greatest clearing (survival in our case) vector. Since we assumed zero recovery, the fictitious default algorithm is very simple. In the first iteration, we check which agents default assuming that everyone can pay their interbank obligations in full. The default set then contains all agents in fundamental default. In the second clearing iteration, we check which of the agents surviving the first round default, knowing that their debtors in fundamental default cannot pay their obligations. The set of defaulted agents is updated with these new, first-round contagion defaults. Then, the same procedure (determining second-, third-, etc. round contagion) is iterated until there is no change in the default set after a clearing round.

In the following, unless otherwise indicated, we assume homogeneity. This means that we make the following additional assumptions about the system.

Assumption 3.9. Each agent credits k others and owes k others. Each interbank loan has a nominal/face value of λ. Thus, the total interbank claims/obligations for each agent is $k\lambda$. We also assume that $\ell_i := l$ for all $i \in I$.

Assumption 3.10. At the period end, each agent has external assets that are normally distributed with expected value μ and standard deviation σ. The pairwise correlation between the external assets of any two agents is ρ. Therefore, we express the time 1 external asset of agent i as the random variable $\mathcal{A}_i = \sigma(\sqrt{\rho}Z + \sqrt{1-\rho}\epsilon_i) + \mu$, where the ϵ_is, the *idiosyncratic factors*, specific to each agent, are standard normal variables, and Z, another standard normal variable, is referred to as the *common factor*. This dependence structure is also commonly referred to as the Gaussian one-factor copula.

This one-factor Gaussian structure is a controlled stylization used to vary dependence cleanly; it is not intended as an empirical model of the full joint dynamics of bank balance sheets, macro-financial variables or supervisory stress-test scenarios. We can choose $H = 1$ (years) and $\sigma = 1$ without loss of generality. More importantly, we investigate the effect of ρ independently of the value of σ. Thus, with increasing correlation, agents do not benefit from diversification, as is typical in the literature, e.g., in Raffestin (2014) or Tasca et al. (2014). Since each agent has the same amount of total interbank liabilities and interbank claims, assuming everyone pays in full, agents still need to pay l to their outside creditors. One might argue that the system is in poor financial health if the agents' expected external asset value is lower than the outside obligations. Such a setup is rather unrealistic. Therefore, in what follows, we limit our attention to $\mu \geq l$. Also, note that in the above approximations there is no dependence on μ and l individually, only on $\mu - l$. Note that since $\mathcal{A}_i$, the external asset values, are random variables, it is meaningful to now take an ex-ante view and consider time-1 balance sheet items as random variables.

3.3.1 Default probabilities

We now introduce the key probabilistic quantities that arise from the clearing mechanism and the stochastic endowments.

Definition 3.11. Let us denote the random outcome of the clearing mechanism by the vector $\mathcal{P} \in \{0,1\}^N$.

Definition 3.12. The *fundamental default probability* of agent i is

$$P_f(i) = \mathbb{P}(\mathcal{A}_i + \sum_{j=1}^{N} L_{ji} - \ell_i - \sum_{j=1}^{N} L_{ij} < 0)$$

Definition 3.13. The (actual) *default probability* of agent i is

$$P(i) = \mathbb{P}(\mathcal{P}_i = 0)$$

Definition 3.14. The *contagion default probability* of agent i is

$$P_c(i) = P(i) - P_f(i)$$

Definition 3.15. Under the above homogeneity assumptions, we will often refer to $\mu - l$ as DD, the *distance to default*, as in the literature on structural credit risk models.

Default in this sense is understood to be fundamental default; under the network symmetry assumption of 3.3, assuming everyone pays in full, the asset value indeed needs to drop by DD to trigger a fundamental default at maturity.

The next result describes under what circumstances comonotonicity leads to the best outcome in terms of individual default probabilities.

Definition 3.16. The random vector $(\mathcal{A}_1, \ldots, \mathcal{A}_N)$ is *comonotonic* (see, e.g., Dhaene et al. (2002)) if there exist nondecreasing functions $f_1, \ldots, f_N : \mathbb{R} \to \mathbb{R}$ and a single random variable Z such that

$$\mathcal{A}_i = f_i(Z) \quad \text{almost surely, for all } i \in I.$$

In the Gaussian one-factor specification of Assumption 3.10, this corresponds to $\rho = 1$, yielding $\mathcal{A}_i = \sigma Z + \mu$ for every i. In words, comonotonicity is the strongest form of positive dependence: all agents' endowments move monotonically with a single common random variable.

Proposition 3.17. *Assume the homogeneous network structure of Assumption 3.3 and identical endowment values with the Gaussian one-factor specification above, so that all agents have the same fundamental default probability $P_f(i)$, $i \in I$.*

Then, among all joint dependence structures of the endowment vector $(\mathcal{A}_i)_{i\in I}$ with these fixed marginals, each agent's individual default probability $P(i)$ is minimized when the endowments are comonotonic. In that case

$$P(i) = P_f(i) \qquad \textit{and hence} \qquad P_c(i) = 0, \quad i \in I.$$

Proof. For each agent $i \in I$, let

$$F_i := \left\{\mathcal{A}_i + \textstyle\sum_{j=1}^N L_{ji} - \ell_i - \sum_{j=1}^N L_{ij} < 0\right\}$$

denote the fundamental default event of i, so that $P_f(i) = \mathbb{P}(F_i)$ by definition. Let $D_i := \{\mathcal{P}_i = 0\}$ be the event that i defaults in the clearing outcome, so that $P(i) = \mathbb{P}(D_i)$ and $P_c(i) = P(i) - P_f(i)$.

First observe that $F_i \subseteq D_i$ for every $i \in I$. Indeed, if i already has negative net worth when all counterparties pay in full (event F_i), then in any equilibrium of the zero-recovery clearing mechanism, i cannot become solvent by receiving less than full interbank payments. Hence fundamental default necessarily implies actual default. Consequently,

$$P(i) = \mathbb{P}(D_i) \ \geq \ \mathbb{P}(F_i) = P_f(i),$$

so the fundamental default probability $P_f(i)$ is a lower bound for $P(i)$ that depends only on the marginal law of $\mathcal{A}_i$ and not on the dependence between endowments.

It therefore suffices to find a dependence structure under which $P(i) = P_f(i)$, i.e., under which there is no pure contagion default. Under the Gaussian one-factor specification with correlation parameter $\rho = 1$, we have

$$\mathcal{A}_i = \sigma Z + \mu, \qquad i \in I,$$

for some standard normal common factor Z. Thus $\mathcal{A}_1 = \cdots = \mathcal{A}_N$ almost surely; the vector $(\mathcal{A}_i)_{i\in I}$ is comonotonic in the sense that all components are (strictly) increasing functions of the same scalar random variable Z.

By the homogeneity assumptions, each agent faces the same total obligations $b_i = b$ and the same total interbank assets (assuming full repayment), so the fundamental default condition (3.1) is identical across agents. Hence there exists a constant threshold $z^\star$ such that

$$F_i = \{Z < z^\star\} \quad \text{for all } i \in I.$$

In particular, either all agents are fundamentally solvent ($Z \geq z^\star$, so F_i does not occur for any i) or all are in fundamental default ($Z < z^\star$, so F_i occurs for every i); there is no state in which some agents are fundamentally solvent while others are fundamentally insolvent.

Now consider the fictitious default algorithm under the zero recovery specification.

If $Z \geq z^\star$, then by definition no agent is in fundamental default, so in the first iteration (where everyone is assumed to pay in full) no bank has negative equity. The initial default set is empty, and the algorithm stops immediately. Hence D_i does not occur for any i.

If $Z < z^\star$, then every agent has negative equity even under full repayment, so every agent is in fundamental default already in the first iteration and thus defaults in every subsequent step. Hence D_i occurs for all i.

In both cases, the actual default event D_i coincides with the fundamental default event F_i for every agent $i \in I$. Therefore

$$P(i) = \mathbb{P}(D_i) = \mathbb{P}(F_i) = P_f(i)$$

and consequently $P_c(i) = P(i) - P_f(i) = 0$ for all i. Since we have already argued that $P(i) \geq P_f(i)$ for any admissible dependence structure with the same marginals, this shows that the comonotonic specification minimizes the individual default probabilities. □

The logic behind Proposition 3.17 is worth emphasizing. The result is as much a symmetry statement as a dependence statement. Comonotonicity rules out contagion defaults only because, under the ex-ante homogeneity assumptions, all agents cross the fundamental-default boundary at the same realization of the common factor. If one introduced ex-ante heterogeneity—for example heterogeneous drifts, means, or distances to default—then even under $\rho = 1$ some agents could remain fundamentally solvent while others were fundamentally insolvent, so losses on interbank claims could again generate contagion. In such an extension, $\rho = 1$ need not minimize default probability. Section 3.6.2 below studies exactly this robustness question by Monte Carlo.

Definition 3.18. Let us denote the probability of exactly m defaults occurring by $P(m)$, and the probability of m defaults occurring conditional on the common factor's realization being z by $P_z(m)$. Similarly, p denotes the probability of default for an agent (same for all agents under our symmetry assumptions), and the default probability of an agent conditional on the common factor's realization being z is denoted by p_z.

For the Bayesian setup, we will show how $P_z(m)$ can be determined under two specific system structures. Once $P_z(m)$ is given, for both system structures, we integrate over the common factor's distribution to obtain the agents' joint default probabilities.

By the law of iterated expectations, the probability of m defaults occurring is given by:

$$P(m) = \mathbb{E}(P_z(m)) = \int_{-\infty}^{\infty} \phi(z) P_z(m)\, dz, \tag{3.3}$$

where $\phi(x)$ denotes the standard normal PDF.

We will also need individual agents' default probability. The agents' default probability, conditional on the common factor's realization z, is given by:

$$p_z = \sum_{l=0}^{N} P_z(l) \frac{l}{N}$$

Analogously to equation 3.3, the probability that a randomly chosen agent will default is:

$$p = \mathbb{E}(p_z) = \int_{-\infty}^{\infty} \phi(z) p_z\, dz, \tag{3.4}$$

where $\phi(x)$ denotes the standard normal PDF. The integrals above can be evaluated numerically.

Because of the zero recovery assumption, agents' default probability is the paramount quantity of interest. In the following, we present two methodologies to derive default probabilities. In the Bayesian network approach of Section 3.4, we focus on determining $P_z(m)$, while the mean-field limit of Section 3.5 only allows us to determine p_z directly.

3.4 Bayesian network approach

For certain network structures, we derive analytical formulas for calculating various probabilities of interest, conditional on the realization of the common factor. The term Bayesian network is used here in the sense of probabilistic graphical models—a directed acyclic graph encoding conditional dependencies—not in the sense of Bayesian statistical inference or updating.

These formulas can be viewed as special cases of the Bayesian network approach of Chong and Kluppelberg (2018), who represent the default of each institution as a node in a directed graph following the pattern of interbank liabilities and show that, when this graph is acyclic (or suitably augmented), the joint default distribution factorizes into local conditional default probabilities depending only on each node's parents, so that contagion is interpreted as the propagation of defaults along the edges of the network. Intuitively, this factorization works because, given the states of a bank's direct borrowers, more distant parts of the network become conditionally independent and can be ignored when computing its own default probability. In our setting, we exploit the same idea under special network structures where this recursion collapses to tractable closed-form expressions, and—through the single common factor assumption above—we allow for correlated endowments, thereby relaxing the independence assumption on endowments in Chong and Kluppelberg (2018) while still retaining a Bayesian-network-style decomposition conditional on the factor.

3.4.1 Fully connected system

We first analyze the simplest topology in which every agent is linked to every other.

Definition 3.19. In a *fully connected* system, each agent is symmetrically connected to all other $N-1$ agents, meaning that the total interbank credits and total interbank obligations of each agent is $k\lambda = (N-1)\lambda$.

Definition 3.20. For an arbitrary, homogeneous, fully connected system, for a given value of common factor z and correlation ρ, assuming j other agents are already in default, let $B_z(j)$ be the largest solution of

$$\sigma(\sqrt{\rho}z + \sqrt{1-\rho}B_z(j)) + \mu + (N-1-j)\lambda = (N-1)\lambda + l.$$

Thus, $B_z(j)$ is the largest realization of the idiosyncratic factor ϵ_i of agent i, still leading to a net worth of 0. If the realization of its idiosyncratic factor is greater than $B_z(j)$, the agent survives; otherwise it defaults.

Let us now fix at $m > 0$ the number of total defaults after the clearing mechanism (under the mild default rule) has converged. Trivially, if $m = 0$:

$$P_z(0) = \big(1 - \Phi(B_z(0))\big)^N$$

In the following, we discuss how to calculate default probabilities when there are nonzero defaults.

Definition 3.21. Assuming $m > 0$, a *default sequence* κ contains the number of new defaults per clearing iteration. If it has one element only, it means there are only fundamental defaults, and no agent suffers contagion default. On the other hand, if it has m elements, it necessarily is an m-length vector of 1s, indicating that there has been a single default in each clearing iteration of the fictitious default algorithm (and that in the $(m+1)$th round, there have been no new defaults, and thus the algorithm converged).

Definition 3.22. Let us denote the set of all possible default sequences leading to $m > 0$ defaults in total under the mild default rule by K_m. Let the jth element of κ be denoted by $\kappa[j]$ and the number of elements in κ by $|\kappa|$. K_m is the set of vectors κ that satisfy the following:

$$|\kappa| \geq 1$$
$$\kappa[i] \in \mathbb{Z}^+ \quad \forall \kappa[i] \in \kappa$$
$$\sum_{i=1}^{|\kappa|} \kappa[i] = m.$$

There can be no 0 in the default sequences: we do not include trailing 0s to indicate no further defaults in later iteration rounds; moreover, no defaults can occur in the fictitious default algorithm once there has been a "0 new default" round.

Definition 3.23. We denote the probability of the default sequence $\kappa \in K_m$ conditional on common factor realization z by $P_{m,z}(\kappa)$.

We derive the probability of m defaults by iterating through all possible default sequences in K_m and summing the respective probabilities. That is, the probability of $0 < m \leq N$ defaults occurring is

$$P_z(m) = \sum_{\kappa \in K_m} P_{m,z}(\kappa)$$

The crux of the problem is to calculate the probability of each possible default sequence. We consider all vectors that satisfy the above criteria and calculate the probability of such a default sequence occurring. We note that while the methodology introduced here is valid for any N, in practice—since we account for and iterate on all possible default sequences—implementation is only feasible for a small number of agents.

As mentioned earlier, we only need to calculate $P_{m,z}(\kappa)$ in the nontrivial, $m > 0$ case. We introduce generic notation used throughout these calculations. Say, out of N agents, a total of $\gamma \geq 0$ have defaulted up to the previous iteration, and we have $\kappa > 0$ new defaults in the current iteration.

Definition 3.24. Let $C(N, \gamma, j, p)$ denote the probability of j new defaulting agents out of the surviving $N - \gamma$ if each default occurs with probability p independently. In this case, the number of defaults is binomially distributed, so

$$C(N, \gamma, j, p) = \binom{N-\gamma}{j} p^j (1-p)^{N-\gamma-j}.$$

The first element of a default sequence contains the number of agents with fundamental default. Conditional on the common factor realization, fundamental defaults are independent and default probabilities depend only on distance to default, thus it is given by 3.5.

Definition 3.25. If $m > 0$, for a given default sequence κ, let the *fundamental default probability part* of the default sequence be

$$P^f_{m,z}(\kappa) = C\Big(N, 0, \kappa[1], \Phi(B_z(0))\Big) \tag{3.5}$$

If the default sequence in question is longer than 1, i.e., there are contagion defaults, we must deal with the contagion default probability part as given by 3.6. Since all agents are connected, conditional on the common factor realization, the probability of a given default sequence corresponds to a particular domain of the idiosyncratic factor realizations. Say we know that $\kappa[1]$ agents suffered fundamental default. We know that the $\kappa[2]$ agents that suffer first-round contagion default have not suffered fundamental default; thus their endowment values must be high enough (higher than $B_z(0)$, they did not suffer fundamental default) but lower than a certain value ($B_z(\kappa[1])$, since they could not withstand the losses incurred by the default of $\kappa[1]$ of their counterparties). The same argument applies to other elements of the default sequence. For the ease of notation, let us introduce a new variable tracking the number of agents already defaulted in the previous rounds.

Definition 3.26. Let $w(i) \in \mathbb{Z}_{\geq 0}$ be defined as

$$w(i) = \begin{cases} 0 & \text{if} \quad i = 1 \\ \kappa[1] & \text{if} \quad i = 2 \\ \sum_{l=1}^{i-1} \kappa[l] & \text{if} \quad i > 2 \end{cases}$$

For the ith, (where $i > 1$) element of the default sequence κ, the agents who default in that round must have larger realizations of idiosyncratic factors than $B_z(w(i-1))$ since they survived the previous clearing round, but smaller than $B_z(w(i))$ since they default now. The binomial coefficient tells us the number of ways in which the newly defaulting $\kappa[i]$ agents can be selected from the $N - w(i)$ that have survived previous rounds.

Definition 3.27. If $m > 0$, for a given default sequence κ with $|\kappa| \geq 2$, let the *contagion probability part* of the default sequence be

$$P^c_{m,z}(\kappa) = \prod_{i=2}^{|\kappa|} C\left(N, w(i), \kappa[i], \frac{\Phi\big(B_z(w(i))\big) - \Phi\big(B_z(w(i-1))\big)}{1 - \Phi\big(B_z(w(i-1))\big)}\right). \tag{3.6}$$

If $m < N$ so there are surviving agents, we need to take the survival probabilities into consideration as well (surviving agents' endowments are sufficiently high); this is given by equation 3.28.

Definition 3.28. If $m > 0$, for a given default sequence κ, and $m < N$, let the *survival probability part* of the default sequence be

$$P^s_{m,z}(\kappa) = \left(\frac{1 - \Phi(B_z(m))}{1 - \Phi(B_z(w(|\kappa|)))}\right)^{N-m}. \tag{3.7}$$

.

Finally, these probabilities need to be multiplied, as each element corresponds to the probability of an agent's endowment being between certain values, and those probabilities are independent conditional on the common factor realization.

$$P_{m,z}(\kappa) = P^f_{m,z}(\kappa) P^c_{m,z}(\kappa) P^s_{m,z}(\kappa).$$

3.4.2 Ring structure system

We now turn to a sparser topology in which each agent is exposed to exactly one counterparty.

Definition 3.29. In a *ring structure* system, each agent credits a single other agent, and these connections form a single closed loop involving all the agents. The total credit/obligation of each agent is $k\lambda = \lambda$.

Definition 3.30. Under the ring structure, an *open path*, OP, stands for a set of agents forming a continuous path of obligations, where the first agent's debtor suffers fundamental default and the last agent's creditor also suffers fundamental default, but none of the agents in the set does. Let $l(OP)$ denote the length of the open path OP: the number of agents that form it.

Contagion defaults can thus occur in the ring structure only along open paths.

Definition 3.31. Under the ring structure, a *fundamental default configuration* FD consists of open paths $\{OP_i\}$ under a given configuration of fundamental defaults. $\mathcal{FD}$ denotes the set of possible fundamental default configurations. Let $d(FD)$ denote the total number of agents with fundamental default under FD. $n(FD)$ denotes the total number of open paths in FD.

A fundamental default configuration can be fully specified with a map (key-value pairs), where the key is a positive natural number c, and the value is the number of open paths of length c:

$$M(FD) := \Big\{c : \Big|\{OP_i \in FD : l(OP_i) = c\}\Big|\Big\}$$

The keys are denoted by c and the values by $M(FD)[c]$. Two fundamental default configurations differ only if they are described by different maps. Note that

$$d(FD) = N - \sum_{OP_i \in FD} l(OP_i)$$

Definition 3.32. Under the ring structure, for a given z realization of the common factor, $PF_z(FD)$ denotes the probability of fundamental default configuration FD.

For any given, finite N, $\mathcal{FD}$ can be determined and the probability of every initial default configuration $PF_z(FD)$ for all $FD \in \mathcal{FD}$ can be computed algorithmically by utilizing the symmetries of the system and the independence of initial defaults (conditional on the realization of common factor z).

Definition 3.33. For a given FD, a *default configuration*[32] DC is a collection of nonnegative integers $\{D_j\}$ of length $n(FD)$, where for each $1 \leq j \leq n(FD)$, D_j is the number of agents suffering contagion default along the open path OP_j. Trivially, $0 \leq D_j \leq l(OP_j)$ for all j. As earlier, let $l(OP_j)$ denote the length of open path OP_j.

Definition 3.34. Let $\mathcal{D}(I, n)$ be the (possibly empty) set of default configurations for a given initial default configuration FD, where $\sum_{j=1}^{n(FD)} D_j = n$. That is, the set of default configurations for a given fundamental default configuration where the total number of agents suffering contagion default is n.

As mentioned earlier, contagion defaults can only occur along open paths.

[32]Not to be confused with the *fundamental default configuration* defined above.

Definition 3.35. Cascading default events on a single, given open path will be referred to as a *default path*, DP, and will be characterized by a sequence of zero or more 1s ending with a 0, with a total length of $|DP| \leq l(OP)$, or a sequence of 1s with a length of $l(OP)$. Each element represents whether the agent at a given index on the open path defaults in the default cascade.

Definition 3.36. Conditional on the common factor realization z, let $PC_z(i,j)$ denote the (conditional) probability that the first i agents default in an open path OP with $l(OP) = j$. It is conditional of all agents being part of an open path, i.e. none of them suffering fundamental default.

Definition 3.37. Since in a ring structure system each agent has exposure to exactly one other agent (counterparty), given the common factor realization z, let $C_z(k)$ denote the solution of

$$\sigma(\sqrt{\rho}z + \sqrt{1-\rho}C_z(k)) + \mu + k\lambda = \lambda + l.$$

k can take the value of 1 (that counterparty survives) or 0 (that counterparty defaults).

Thus, $C_z(1)$ ($C_z(0)$) is the largest realization of the idiosyncratic factor ϵ_i of agent i, still leading to a net worth of 0 if the agent i is exposed to survives (defaults). If the realization of its idiosyncratic factor is greater than $C_z(\cdot)$, the agent survives; otherwise it defaults. By our definition of open path, agents on the open path are not fundamentally defaulting. In our notation that means they satisfy $\epsilon_i \geq C_z(1)$. Given that conditioning, the probability that such an agent fails when its debtor defaults is:

$$q_z = \mathbb{P}(\epsilon < C_z(0) \,|\, \epsilon \geq C_z(1)) = \frac{\Phi\big(C_z(0)\big) - \Phi\big(C_z(1)\big)}{1 - \Phi\big(C_z(1)\big)}.$$

Let us fix the realization of the common factor, z. Let d denote the number of defaults along the default path. We make a distinction between the cases $d = l(OP)$ and $d < l(OP)$.

In the former case,

$$PC_z(j,j) = q_z^j = \left(\frac{\Phi\big(C_z(0)\big) - \Phi\big(C_z(1)\big)}{1 - \Phi\big(C_z(1)\big)}\right)^j. \tag{3.8}$$

In the latter case,

$$PC_z(i,j) = q_z^i(1-q_z) = \left(\frac{\Phi\big(C_z(0)\big) - \Phi\big(C_z(1)\big)}{1 - \Phi\big(C_z(1)\big)}\right)^i \left(1 - \frac{\Phi\big(C_z(0)\big) - \Phi\big(C_z(1)\big)}{1 - \Phi\big(C_z(1)\big)}\right) \tag{3.9}$$

Then,

$$P_z(m) = \sum_{FD \in \mathcal{FD}} PF_z(FD) \sum_{DC \in \mathcal{D}(FD, m-d(FD))} \prod_{i=1}^{n(FD)} PC_z(D_i, l(OP_i)). \tag{3.10}$$

$P_z(m)$ can be calculated computationally efficiently, with far fewer computations than the above formula suggests. To understand why, notice that in Equations 3.8 and 3.9, probabilities of the default paths do not depend on the specific default configuration they are part of, but possibly only on the length of the open path they act on, so they can be precalculated.

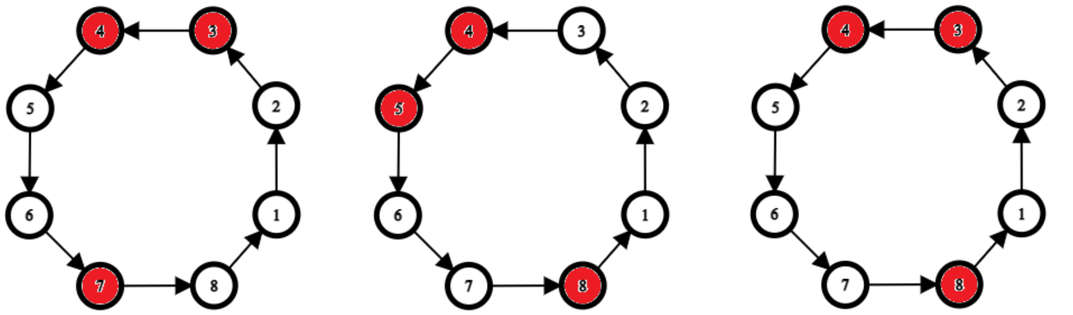

Figure 3.1: Three possible manifestations of the same fundamental default configuration.

Example 3.38. Take the 8-agent system as in Figure 3.1 as an example. The three plots show three possible ways in which three agents can suffer fundamental default. Note that all three outcomes have the same probability. More importantly, these three (and actually many more) outcomes correspond to the exact same fundamental default configuration, as they define the same set of open paths.

In this setup, there are two open paths: one is two agents long, the other one is three agents long. In the following, we arbitrarily choose to use the middle plot as an example. Focusing on the three-long open path, the following are the possible default paths along the open path consisting of agents $\{1, 2, 3\}$:

1. agent 1 does not suffer contagion default, even though agent 8 is in default

2. agent 1 suffers contagion default, but agent 2 does not

3. agents 1 and 2 suffer contagion default, but agent 3 does not

4. in a default cascade, all three agents go into default.

Conditional on the z realization of the common factor and on FD, the probability of each of these events is straightforward to calculate and is given by equation 3.8 in the fourth case and by equation 3.9 in the other three cases. In a similar manner, all possible default paths' probabilities can be calculated on the two-long open path as well. Again, conditional on the fundamental default configuration and on z, the events of given default paths that occur on these two separate open paths are independent.

3.5 Mean-field limit

To obtain a closed-form solution for the individual default probability, for sufficiently large N we take a mean-field approximation[33]. The mean-field (MF) idea is to replace the many-body interaction faced by a representative agent with an average, endogenous environment. To illustrate the self-consistency requirement, consider a simplified example. Suppose each agent survives with some unknown probability s, and an agent's survival depends on how many counterparties pay. In the mean field, instead of tracking who exactly pays, one assumes each counterparty pays independently with probability s. The agent then survives with some probability $g(s)$ that depends on the assumed survival rate. Self-consistency demands $s = g(s)$: the assumed survival probability must equal the survival probability that emerges when every agent uses that same assumption. In general, this is a fixed-point equation for s.

In large, dense, and homogeneous systems, this approximation is justified by propagation of chaos: as $N \to \infty$, suitably scaled and exchangeable interactions make individual agents conditionally independent given a few aggregate variables (Sznitman, 1991). In other words, local randomness averages out and it suffices to follow macroscopic quantities that solve closed equations (Lorenz et al., 2009).

To use the MF toolkit, we require: (i) a large population of homogeneous agents, (ii) dense interconnections so that each agent interacts with many others and connections are scaled as $O(1/N)$ to keep total interbank exposures $O(1)$; and (iii) a regularity condition ensuring that the self-consistency map has (and we can find) a stable solution. These are standard in fully-connected models where MF becomes asymptotically exact (Friedli and Velenik, 2017). In addition, when shocks feature a common component, conditioning on the factor preserves the MF structure: the representative survival probability $s(z)$ is computed for fixed common shock z and then averaged out, as we do below. The MF solution is not appropriate for a ring structure because the network is sparse and low-dimensional: the inflow from all counterparties (only one) cannot be properly represented by an average/expected payment.

Birch and Aste (2014) studies similar networks in the mean-field limit, assuming endowments are independent. We extend their analysis by investigating

[33]Recall from Chapter 1 that a mean-field approximation replaces the detailed interactions among all agents with interaction with an aggregate quantity (typically the average default rate), yielding tractable equations for large systems.

the effect of correlated endowments and how a mean-field approximation can be augmented to handle such cases.

Under homogeneity, dense connectivity, and $O(1/N)$ edge scaling, the interbank inflow of a representative agent equals survival probability times total interbank claim. The random variable corresponding to the amount that an agent recovers on its interbank claims is approximated by its expected value: $s(z)k\lambda$, where $s(z)$ is the common factor-conditional survival probability. The agent survives if its external asset plus expected interbank inflow exceeds its total obligations. This will be the self-consistency equation whose largest solution selects the greatest clearing outcome used later.

3.5.1 Conditional probability of default given the common factor realization

Conditional on the realization z of the common factor, the following mean-field approximation gives the conditional survival probability $s(z) = 1 - p_z$ of an agent.

$$s(z) = \mathbb{P}\big(\sigma(\sqrt{\rho}z + \sqrt{1-\rho}\epsilon) + \mu + s(z)k\lambda \geq k\lambda + l\big)$$

That is,

$$s(z) = 1 - \Phi\big(\frac{k\lambda + l - \mu}{\sigma\sqrt{1-\rho}} - \frac{\sqrt{\rho}z}{\sqrt{1-\rho}} - s(z)\frac{k\lambda}{\sigma\sqrt{1-\rho}}\big) := 1 - \Phi\big(a - bs(z)\big) \quad (3.11)$$

to be solved for $s(z)$.

We characterize the number of solutions of the fixed-point equation

$$s = 1 - \Phi\big(a - bs\big), \qquad s \in [0, 1],$$

as a function of the parameters a and b.

Proposition 3.39. *Let $F(s) := 1 - \Phi(a - bs) - s$ for $s \in [0, 1]$, where Φ and ϕ denote the standard normal CDF and PDF, respectively.*

1. *The equation $F(s) = 0$ always has at least one solution in $[0, 1]$.*
2. *If $b < \sqrt{2\pi}$, then this solution is unique.*
3. *If $b \geq \sqrt{2\pi}$, define $x_b := \sqrt{-2\ln(\sqrt{2\pi}/b)} > 0$. Multiple solutions on $[0, 1]$ can occur only if both turning points*

$$s_- = \frac{a - x_b}{b} \qquad \text{and} \qquad s_+ = \frac{a + x_b}{b}$$

lie in $[0, 1]$ (equivalently, $x_b \leq a \leq b - x_b$) and $F(s_-) < 0 < F(s_+)$. In all other cases, the solution is unique.

Proof. Since F is continuous and

$$F(0) = 1 - \Phi(a) > 0, \qquad F(1) = 1 - \Phi(a-b) - 1 = -\Phi(a-b) < 0,$$

the intermediate value theorem guarantees at least one root, establishing (1).

For the multiplicity analysis, compute

$$F'(s) = b\,\phi(a - bs) - 1.$$

Since $\sup_x \phi(x) = 1/\sqrt{2\pi}$, we have $F'(s) < 0$ for all s when $b < \sqrt{2\pi}$, so F is strictly decreasing and the solution is unique, proving (2).

When $b > \sqrt{2\pi}$, $F'(s) = 0$ if and only if $a - bs = \pm x_b$, giving the two critical points s_- and s_+. At s_- we have $a - bs_- = x_b > 0$, hence $F''(s_-) = b^2(a-bs_-)\phi(a-bs_-) > 0$ and s_- is a local minimum; at s_+ we have $a - bs_+ = -x_b < 0$, hence $F''(s_+) < 0$ and s_+ is a local maximum. For F to cross zero more than once, both critical points must lie in $[0,1]$ and the sign pattern $F(s_-) < 0 < F(s_+)$ must hold, establishing (3). □

Whenever the solution is not unique, we select the largest solution. Under a simulation-based approach, the equilibrium survival vector may not be unique for many realizations; selecting the largest fixed point is intended as a proxy for selecting the largest equilibrium survival vector in each simulated scenario.

We integrate out the conditional probabilities to get the unconditional default probability p as described in Section 3.3.1.

The perfectly correlated/comonotonic case, where a single factor drives all agents' cash flows, requires special attention. Since the system is homogeneous, when $\rho = 1$, there is a particular value of z above which none of the agents defaults, but all agents default below. That particular value is $z_b = \frac{l-\mu}{\sigma}$. Thus, the survival probability s_c in the perfectly correlated case satisfies

$$s_c = 1 - \Phi\left(\frac{l-\mu}{\sigma}\right) \tag{3.12}$$

Proposition 3.40. *In the perfectly correlated case, under the above mean-field approximation, the survival probability is at least as high as in the independent ($\rho = 0$) case.*

Proof. The survival probability s_u in the independent case satisfies

$$s_u = 1 - \Phi\big(\frac{k\lambda + l - \mu}{\sigma} - s_u\frac{k\lambda}{\sigma}\big)$$

Since s_u is positive, using 3.12,

$$s_u = 1 - \Phi\big(\frac{l-\mu}{\sigma} + \frac{k\lambda(1-s_u)}{\sigma}\big) \leq 1 - \Phi\big(\frac{l-\mu}{\sigma}\big) = s_c$$

□

Note that a more general result is already shown in 3.17; here we simply show that the mean-field approximation is consistent with that result.

3.6 Analysis of the model

3.6.1 Bayesian networks

Next, we examine how the probabilities of various events change as the main parameters of the system vary. These events are as follows.

Definition 3.41. *Probability of systemic default*: probability that more than three-quarters of the agents default. In the examples below, when $N = 3$, this amounts to all agents defaulting; when $N = 7$, this amounts to at least 6 agents defaulting; and when $N = 11$, this amounts to at least 9 agents defaulting.

Definition 3.42. *Probability of few defaults*: probability that fewer than one-quarter of the agents default. In the examples below, when $N = 3$, this amounts to all agents surviving; when $N = 7$, this amounts to at most 1 agent defaulting; and when $N = 11$, this amounts to at most 2 agents defaulting.

Definition 3.43. *Individual default probability*: as calculated under Definition 3.13.

Along with the probabilities of these events given various parameter values of the system, we plot the corresponding probabilities in a unconnected system, i.e., one with no interbank loans. That case serves as a reference point to illustrate the effects of contagion, as it also provides the fundamental default probabilities.

As shown in Figures 3.2 and 3.3, larger interbank exposure $k\lambda$ means larger default probability, in general. Both systemic default probability and individual default probability show a characteristic humped shape as a function of asset correlation, ρ. The maximal default probability occurs closer to $\rho = 1$ the smaller the interbank exposures. This applies both to individual default probabilities and systemic default probabilities. At the same time, the maximum default probability increases with increasing $k\lambda$. Similarly, the probability of only a few defaults occurring decreases as interbank exposure increases. The above is true for both network structures. However, there is a remarkable difference between the fully connected and the ring structure when it comes to behavior at small correlation values. While in the fully connected system there is little contagion default probability (and small systemic default probability), in the ring structure there is substantial systemic default probability, which also increases the individual default probability significantly.

The root cause lies in the connection structure: diversification in the form of distributing interbank exposure across many other agents benefits the fully connected system but not the ring structure. Note that even in the ring structure there is a default probability amplification effect as ρ increases, but the hump shape is much less pronounced for sufficiently large $k\lambda$ values. It is also worth highlighting how surprisingly large the probability of contagion default is even in the uncorrelated case, e.g., for $k\lambda = 3.0$ (as shown by the individual probability of default in Figure 3.3), especially compared to the $\rho = 1$ case (when there is trivially no contagion default). To put it differently, a diversified investment strategy (in endowments) can help reduce systemic losses mostly if diversification appears at the interbank connection level as well.

To illustrate the magnitudes: in the fully connected system with $N = 11$ and $DD = 2$, increasing $k\lambda$ from 1 to 3 roughly doubles the peak individual default probability while shifting the peak from around $\rho \approx 0.95$ to $\rho \approx 0.5$. In the ring structure with the same parameters, the peak is less sharply defined but the contagion default probability at $\rho = 0$ is already substantial for $k\lambda = 3$—similar magnitude to the peak value—because each agent's solvency depends entirely on a single counterparty. This contrast underscores that the network topology is at least as important as the correlation structure in determining how severely interbank exposures amplify risk.

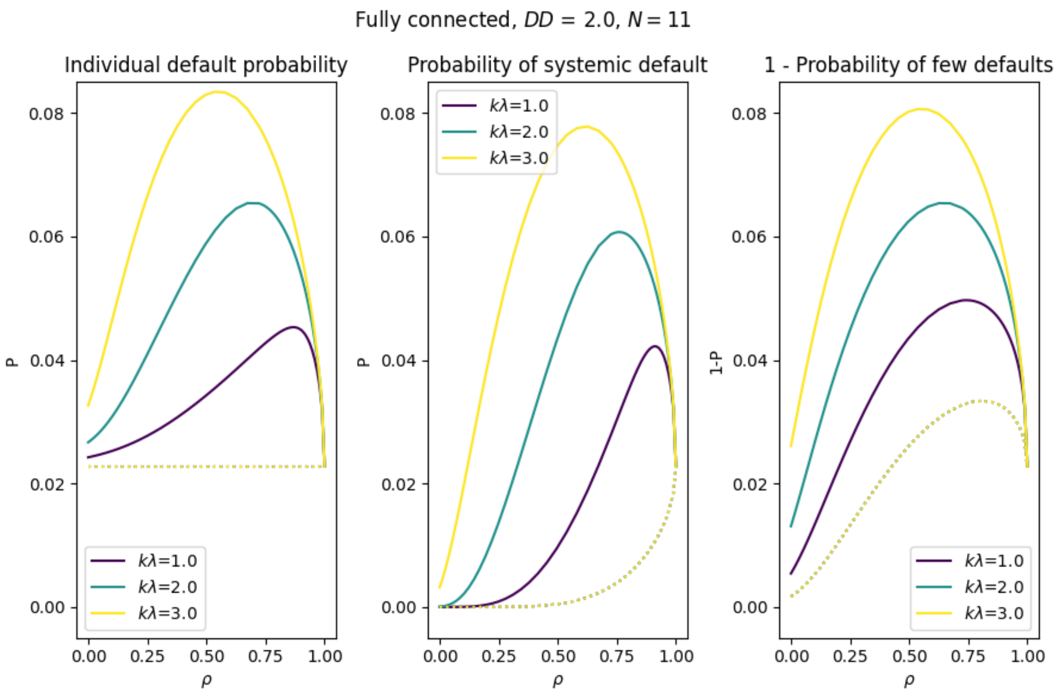

Figure 3.2: Various probabilities as a function of ρ for different $k\lambda$ values: unconnected system results shown with dotted lines.

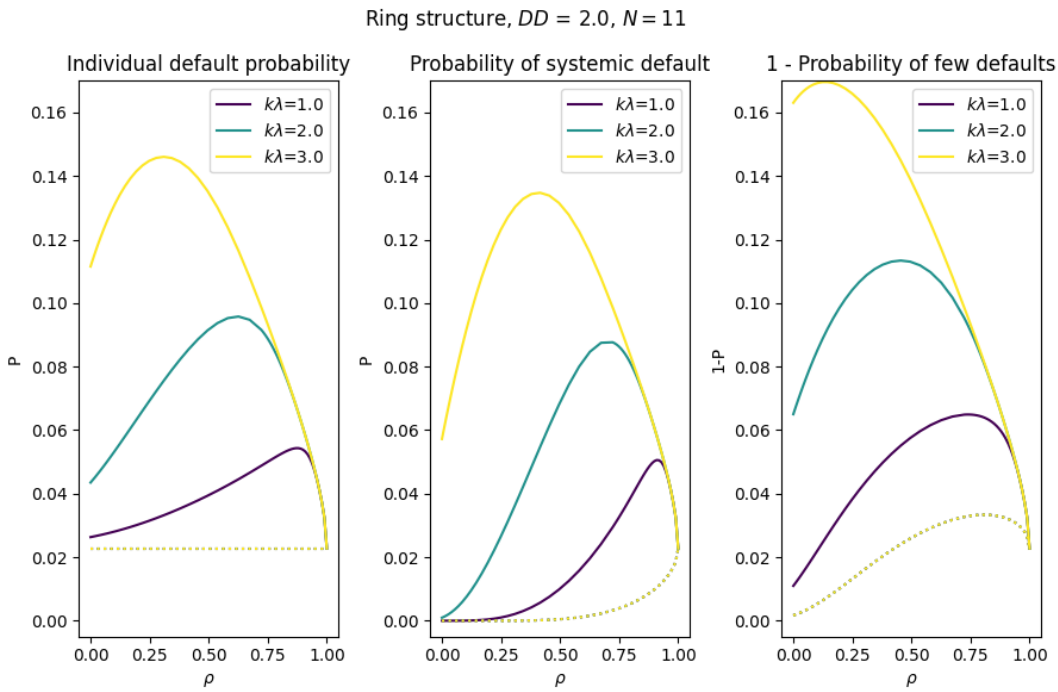

Figure 3.3: Various probabilities as a function of ρ for different $k\lambda$ values: unconnected system results shown with dotted lines.

Next, we illustrate how changing the number of agents affects the probabilities of interest. In Figures 3.4 and 3.5, as before, the dotted lines correspond to the unconnected case. Note that the number of agents affects (joint) fundamental default probabilities as well, but not the individual fundamental default probability. Increasing N, ceteris paribus, makes both the maximal systemic default probability and the maximal individual default probability higher. It typically does so by making the hump shape more pronounced. We note here that the characteristic humped shape can disappear for certain parameter combinations, for example when DD is small and $k\lambda$ is large. In that case, the systemic event probability and individual default probability may become monotonically decreasing, and the few-default probability increasing, as a function of ρ. For both network structures, even in that case, it holds that the maximal systemic/individual default probability is larger the larger N is. This is a somewhat surprising effect that might need further investigation and rigorous proof if possible.

The phenomenon is intriguing, since for small correlation values, whether or not larger N means smaller default probabilities depends on the other parameters as well. If agents are sufficiently far from fundamental default (large DD) and interbank loans are relatively small (small $k\lambda$), small correlation allows for a greater degree of diversification with an increasing number of agents in the system. In the fully connected network structure, this means distributing interbank claims among more counterparties, and thus agents are less exposed to a single partner's default. Perhaps somewhat more surprisingly, this has a similar effect under the ring network structure as well, where larger N does not mean less exposure to each single counterparty, as each agent is exposed to exactly one other. On the other hand, larger N likely increases the length of open paths, making contagion default less likely.

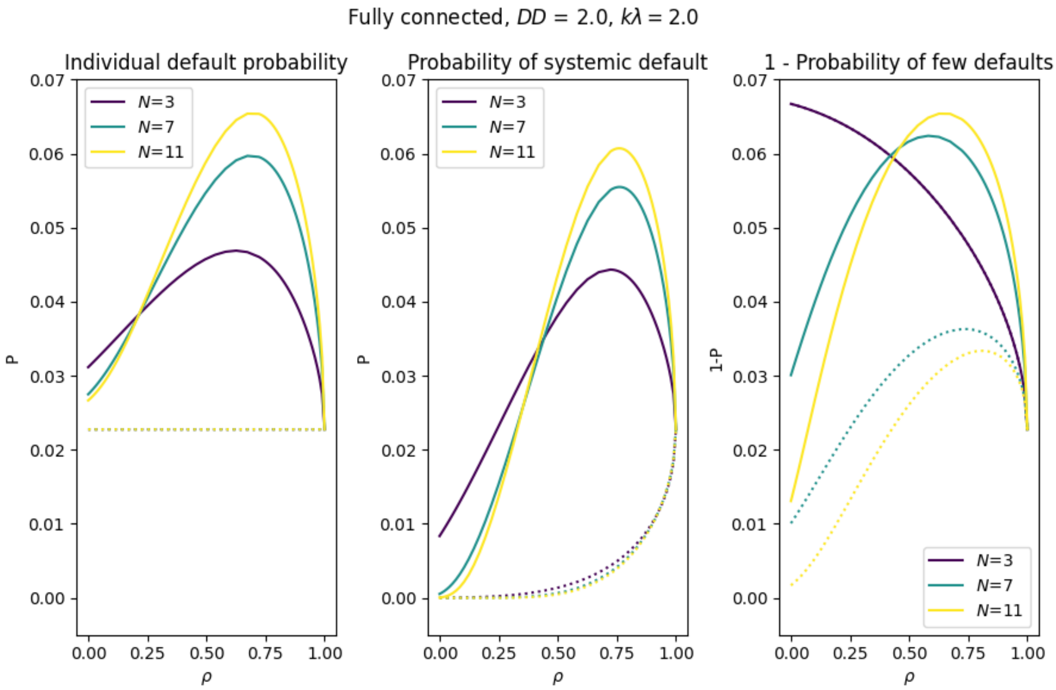

Figure 3.4: Various probabilities as a function of ρ for different N values: unconnected system results shown with dotted lines.

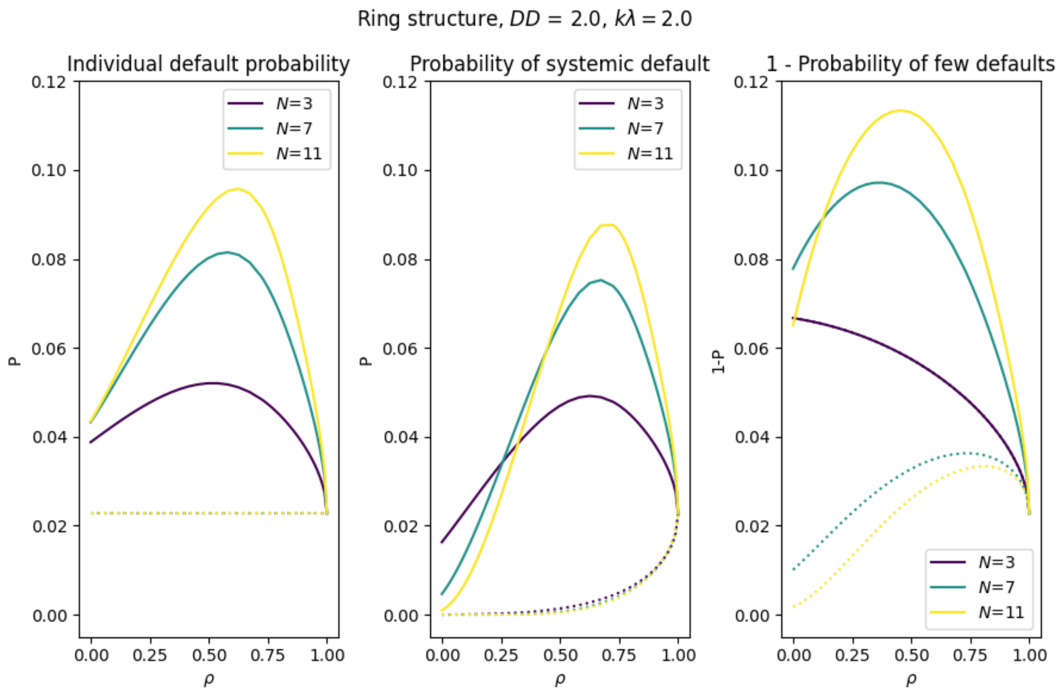

Figure 3.5: Various probabilities as a function of ρ for different N values: unconnected system results shown with dotted lines.

Increasing DD trivially lowers the individual fundamental default probability. At the same time, smaller DD corresponds to a more peaked hump shape in the systemic default probability, with the maximal default probability occurring closer to $\rho = 1$. For instance, at $DD = 3$ the hump is barely visible and default probabilities remain low across all correlation values, whereas at $DD = 1$ the peak default probability can be roughly double the $\rho = 1$ value, illustrating how a modest decline in capitalization dramatically amplifies the sensitivity to the correlation structure. Intuitively, when agents are closer to fundamental default, even a moderate common shock can push several agents below the solvency threshold simultaneously, triggering contagion cascades that would not arise in better-capitalized systems.

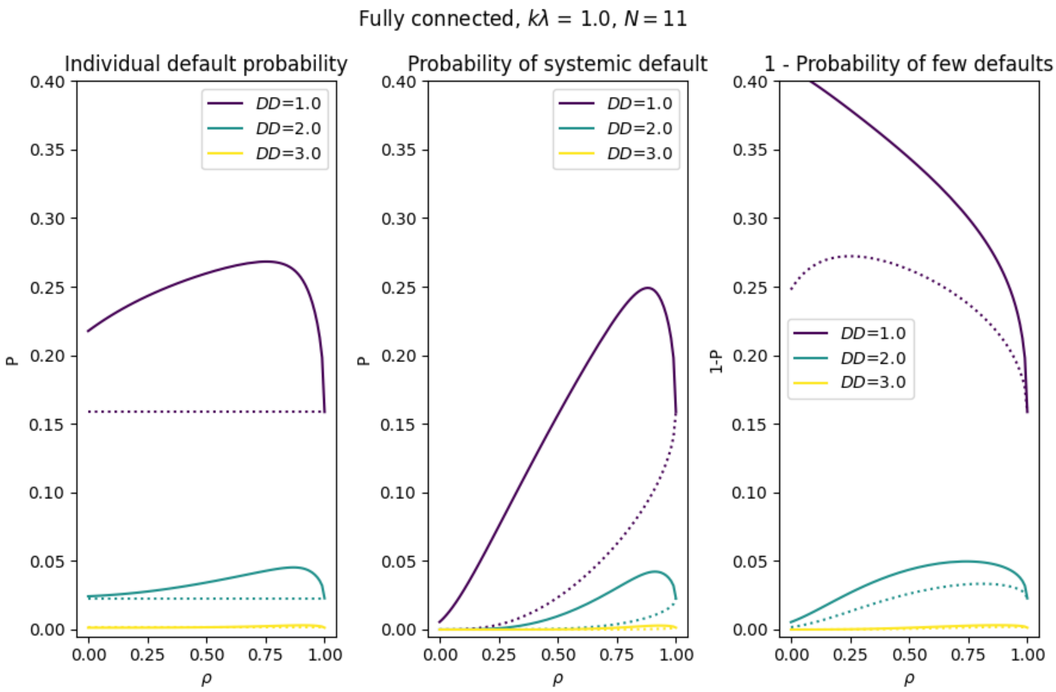

Figure 3.6: Various probabilities as a function of ρ for different DD values: unconnected system results shown with dotted lines.

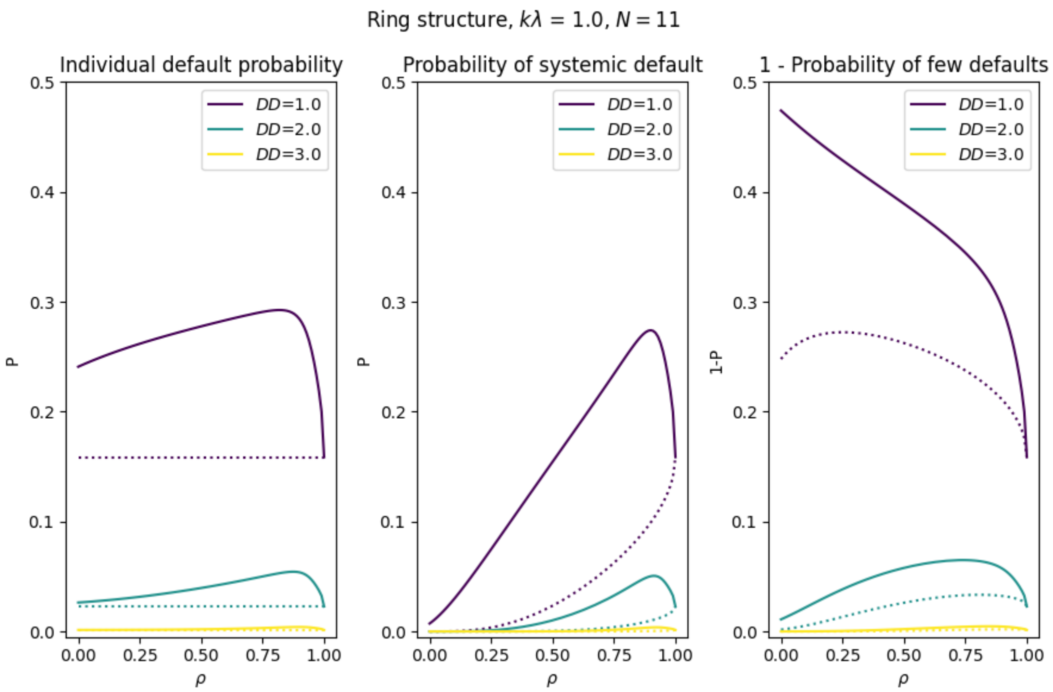

Figure 3.7: Various probabilities as a function of ρ for different DD values: unconnected system results shown with dotted lines.

A unified explanation of the humped shape in this chapter can be formulated in terms of heterogeneity in asset values at the terminal clearing date, that is, at the only date when actual payment obligations are due. This mechanism is the net result of two sub-effects that pull in opposite directions. Starting from low ρ, cross-sectional dispersion in counterparties' endowments provides diversification on the asset side, so default probabilities remain relatively low. As ρ rises, this diversification benefit weakens, which pushes default probabilities upward. However, the same cross-sectional dispersion is also what makes contagion defaults possible: an agent can be pushed into default only in states in which it would survive under full payment, while weaker counterparties fail to pay enough. When ρ becomes very high, such states become rarer. In the limit $\rho = 1$, and under the homogeneity assumptions of Proposition 3.17, all agents are either fundamentally solvent together or fundamentally insolvent together, so contagion defaults disappear. The hump therefore reflects the balance between the loss of diversification and the disappearance of contagion states. This also clarifies why the $\rho = 1$ result is not a generic consequence of correlation alone: it relies on the symmetric setup in which agents lie similarly close to the default boundary. With ex-ante heterogeneity, comonotonicity need not remain optimal. Chapter 4 will build on this logic by adding a second, forward-looking valuation mechanism that is absent from the present one-period setup.

Looking at the Bayesian network results, it also appears that both individual and systemic default probabilities are at least as large in the ring structure system as in the fully connected system. While the result is intuitive, a rigorous proof remains to be done.

3.6.2 Monte Carlo robustness to heterogeneous distances to default

The caveat after Proposition 3.17 suggests a direct robustness question: how far does the homogeneous result that $\rho = 1$ minimizes default probability survive once ex-ante symmetry is relaxed? Because in the one-period setup only $\mu - l$ matters, we keep the normalization and work directly with distances to default. Equivalently, we set $l = 0$ and write

$$\mathcal{A}_i = DD_i + \sqrt{\rho}\, Z + \sqrt{1-\rho}\, \epsilon_i,$$

with $\sigma = 1$ as before.

Heterogeneity is introduced through agent-specific distances to default. In each outer experiment $b = 1, \ldots, B$ we draw a fixed heterogeneous system

$$DD_i^{(b)} = DD + h\, \xi_i^{(b)}, \qquad i = 1, \ldots, N,$$

where the cross-sectional draw is standardized so that

$$\frac{1}{N}\sum_{i=1}^{N} \xi_i^{(b)} = 0, \qquad \frac{1}{N}\sum_{i=1}^{N} \bigl(\xi_i^{(b)}\bigr)^2 = 1.$$

Hence every outer experiment has the same cross-sectional mean distance to default DD, while h is exactly the cross-sectional standard deviation of the DD_i. In the implementation below, the raw draws are sampled from a uniform distribution and then recentered and rescaled. The vector $(DD_i^{(b)})_{i=1}^N$ is held fixed across all inner scenarios of a given outer experiment; this is important because the object of interest is persistent ex-ante heterogeneity across banks. Redrawing DD_i in every scenario would instead add extra idiosyncratic noise and mechanically blur the correlation structure.

Conditional on one fixed outer draw $(DD_i^{(b)})_{i=1}^N$ and one value of ρ, we simulate M scenarios of $(Z, \epsilon_1, \dots, \epsilon_N)$, run the same zero-recovery clearing mechanism as in the homogeneous case, and compute

$$\widehat{\overline{P}}^{(b)}(\rho) = \frac{1}{MN} \sum_{m=1}^{M} \sum_{i=1}^{N} \mathbb{1}\{\mathcal{P}_{i,m}^{(b)}(\rho) = 0\},$$

$$\widehat{\overline{P}}_f^{(b)}(\rho) = \frac{1}{MN} \sum_{m=1}^{M} \sum_{i=1}^{N} \mathbb{1}\{F_{i,m}^{(b)}(\rho)\}, \qquad \widehat{\overline{P}}_c^{(b)}(\rho) = \widehat{\overline{P}}^{(b)}(\rho) - \widehat{\overline{P}}_f^{(b)}(\rho),$$

where $F_{i,m}^{(b)}(\rho)$ denotes the fundamental-default event of bank i in scenario m of outer draw b. We then average over outer draws,

$$\widehat{\overline{P}}_h(\rho) = \frac{1}{B} \sum_{b=1}^{B} \widehat{\overline{P}}^{(b)}(\rho), \qquad \widehat{\overline{P}}_{f,h}(1) = \frac{1}{B} \sum_{b=1}^{B} \widehat{\overline{P}}_f^{(b)}(1), \qquad \widehat{\overline{P}}_{c,h}(1) = \frac{1}{B} \sum_{b=1}^{B} \widehat{\overline{P}}_c^{(b)}(1).$$

To display how much the default-probability curve varies across heterogeneous systems, we also report the empirical 5% and 95% quantiles of $\widehat{\overline{P}}^{(b)}(\rho)$ across outer draws,

$$\widehat{Q}_{0.05,h}(\rho) = \mathrm{Quant}_{0.05}\left(\widehat{\overline{P}}^{(1)}(\rho), \dots, \widehat{\overline{P}}^{(B)}(\rho)\right)$$

$$\widehat{Q}_{0.95,h}(\rho) = \mathrm{Quant}_{0.95}\left(\widehat{\overline{P}}^{(1)}(\rho), \dots, \widehat{\overline{P}}^{(B)}(\rho)\right)$$

Thus $h = 0$ recovers the homogeneous benchmark exactly, while $h > 0$ measures how far the system has moved away from the symmetry argument underlying Proposition 3.17.

In the figure below, the left panel reports the mean curve $\widehat{\overline{P}}_h(\rho)$ together with the pointwise 5% and 95% percentiles of $\widehat{\overline{P}}^{(b)}(\rho)$ across outer draws. This makes the dispersion caused by persistent bank heterogeneity directly visible.

This design targets three questions. First, does the humped shape of $\rho \mapsto \widehat{\overline{P}}_h(\rho)$ survive for small and moderate values of h? Second, does $\rho = 1$ continue to minimize average default probability, or does the minimum move away from one once some banks are persistently closer to the default boundary than others? Third, and most directly, does contagion reappear at $\rho = 1$? The last

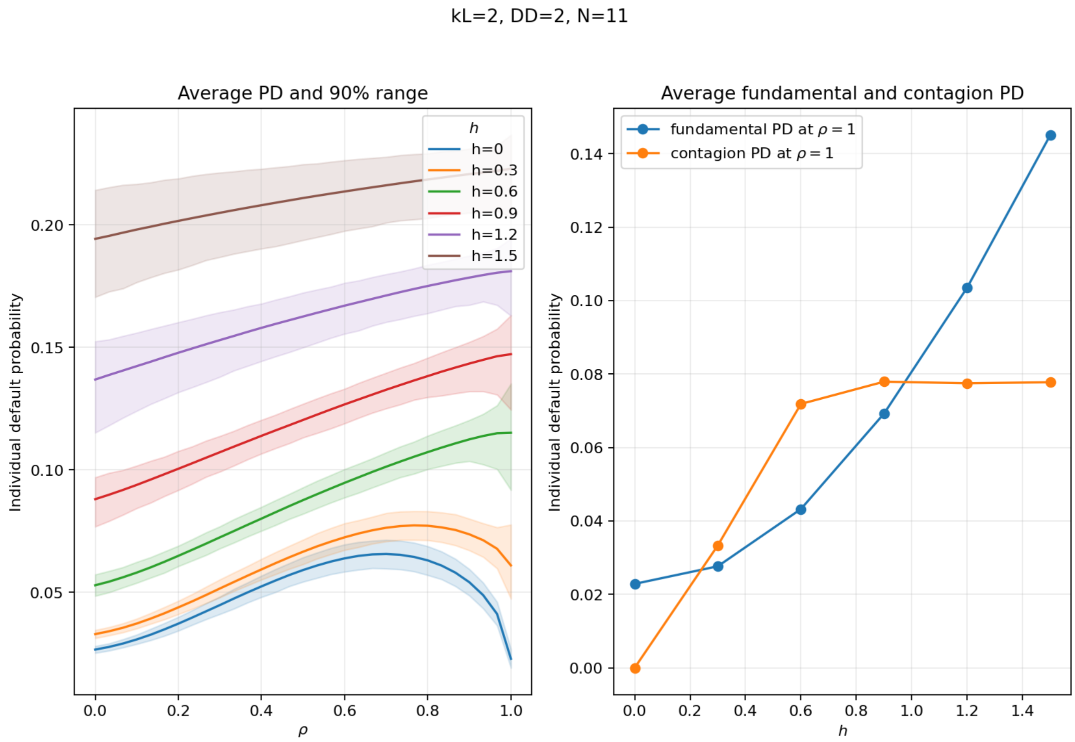

Figure 3.8: Monte Carlo robustness experiment with heterogeneous distances to default. The left panel plots $\hat{\bar{P}}_h(\rho)$ together with the pointwise 5% and 95% percentiles of $\hat{\bar{P}}^{(b)}(\rho)$ across outer experiments for selected values of h. These bands visualize cross-system dispersion induced by heterogeneous DD_i, not Monte Carlo confidence intervals for the mean curve. The right panel overlays $\hat{\bar{P}}_{f,h}(1)$ and $\hat{\bar{P}}_{c,h}(1)$, that is, the outer-draw averages of the fundamental and contagion components at $\rho = 1$, as functions of h. The benchmark $h = 0$ corresponds to the homogeneous system.

question is the sharpest diagnostic because under the homogeneous benchmark the contagion component vanishes exactly at $\rho = 1$.

Figure 3.8 should be read as a robustness check of the I) the hump shape, II) the symmetry-based $\rho = 1$ result. The shaded region in the left panel is a percentile envelope across heterogeneous systems[34]. For small h, one expects the homogeneous logic to remain visible: the hump should persist and the contagion component at $\rho = 1$ should remain limited. As h increases, banks cease to cross the default boundary together, so contagion at $\rho = 1$ reappears and the minimizing correlation need no longer be exactly one. Moreover, even the hump shape can disappear.

The percentile band widens as heterogeneity increases. When h is small, each outer draw is only a mild perturbation of the homogeneous setup, so the resulting curves remain relatively similar. As h increases, the sampled systems become much more different from one another: some draws contain only modest

[34]so not a confidence interval for the estimated mean curve

asymmetries across agents, while others generate a much more pronounced lower tail of weak institutions. Since the shape of the default curve is very sensitive to how many agents are close to the default boundary, this cross-draw variation becomes larger, and the percentile envelope widens. The widening band is not a sign that the simulation has become less precise. It is evidence that, once heterogeneity is sufficiently strong, the correlation effect depends much more on which particular heterogeneous banking system is realized.

Next, following the right hand side figure, we focus on what happens at $\rho = 1$. In the homogeneous setup, $\rho = 1$ aligns all banks along the same common path. Under that symmetry, an agent cannot be dragged into default by others unless it was already in essentially the same position itself, so contagion at $\rho = 1$ is shut down. Heterogeneity breaks exactly this symmetry. Once agents differ in their distances to default, some institutions become systematically weaker than others. Then there are states in which the weak ones fail on fundamentals, while other banks would still survive if all claims were paid in full. Those are precisely the states in which missed interbank payments can generate contagion. As h rises, such asymmetric states become more common, so the contagion default probability at $\rho = 1$ increases.

Then, remarkably, contagion eventually flattens out. The increase in contagion does not continue indefinitely because very large heterogeneity changes the composition of bad outcomes. Once the weakest agents are sufficiently far below the rest, many adverse states cease to be "near-miss" contagion events and instead become outright fundamental default states for those banks. Put differently, heterogeneity first creates the asymmetry needed for contagion, but beyond some point additional heterogeneity increasingly pushes weak banks directly into fundamental failure rather than merely making them likely to transmit losses. At the same time, the stronger banks in the system are often far enough from the boundary that further increases in heterogeneity do not proportionally enlarge the set of states in which they can be tipped over by unpaid claims. The result is that average contagion default at $\rho = 1$ rises strongly at first and then tends to level off, even though overall fragility continues to increase.

At the same time, fundamental default probability starts rising more sharply as heterogeneity increases. A rise in h spreads agents further apart in terms of their distances to default while keeping the average distance to default fixed. The key intuition is that default risk is not linear in this distance. Making an already weak agent even weaker raises its fundamental default probability by more than making a strong agent equally stronger lowers its own. As heterogeneity increases, the lower tail of the cross-sectional distribution moves further into the region where agents are highly vulnerable, while the strongest banks are already so far from the boundary that additional improvements have little effect on their already small default probabilities. The result is that the weak tail contributes more and more to the average, so the average fundamental default probability not only rises with h but, in the relevant parameter range, can do so at an increasing rate.

3.6.3 Mean-field limit

First, we show how the mean-field limit relates to the results determined via the Bayesian network approach in Section 3.6.1. Figure 3.9 shows this relationship for a canonical parameter combination.

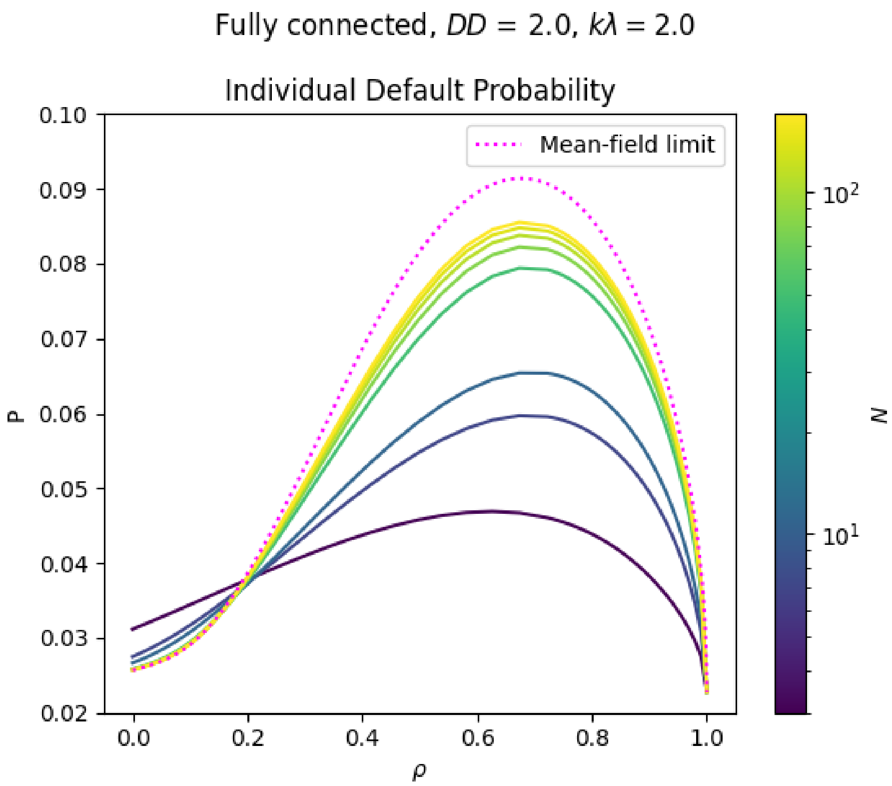

Figure 3.9: Default probabilities as a function of ρ for various N values and the mean-field limit.

Apparently, the Bayesian system results haven't converged to the MF limit even at N as large as 140. The individual default probability p is obtained by integrating the conditional quantity p_z over the common factor. The conditional default probability p_z is itself computed from the joint default-count distribution $P_z(m)$ of the N-agent Bayesian network system. For a fixed common factor realization z, the convergence of p_z as $N \to \infty$ is not uniform across z. In particular, for z values where the fundamental default probability is neither very small nor very close to one—that is, in the intermediate regime where a non-trivial fraction of agents default fundamentally and the contagion cascade is most sensitive to the system size—the conditional distribution $P_z(m)$

changes shape substantially as N grows, and p_z converges slowly. These slowly converging z values dominate the convergence of the integrated quantity p as well. Figure 3.10 illustrates this effect: while p_z stabilizes quickly for extreme z (where either almost no agent or almost every agent defaults fundamentally), the curves for intermediate z are still visibly moving at $N = 500$.

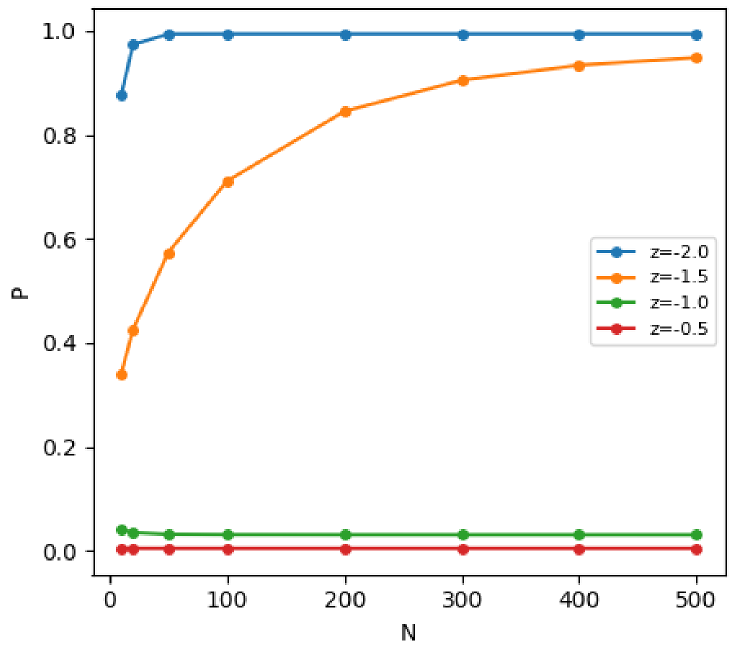

Figure 3.10: Default probabilities for fixed z values, p_z. In this example, $\rho = 0.66$, $k\lambda = 2$, $DD = 2$.

Since the novelty of the approach lies in introducing correlated asset sides in the mean-field-based approximation, we focus on the effect of correlation on individual default probability and on how various other model parameters affect this relationship.

Figures 3.12 and 3.11 illustrate the effect of ρ on the default probability $p = (1 - s)$ for various values of $k\lambda$ and DD. Similarly to the results of the previous section, increasing $k\lambda$ typically creates a “hump” in default probabilities peaking for relatively large correlations. The position of this peak shifts to the left as $k\lambda$ increases, as does the height of the peak.

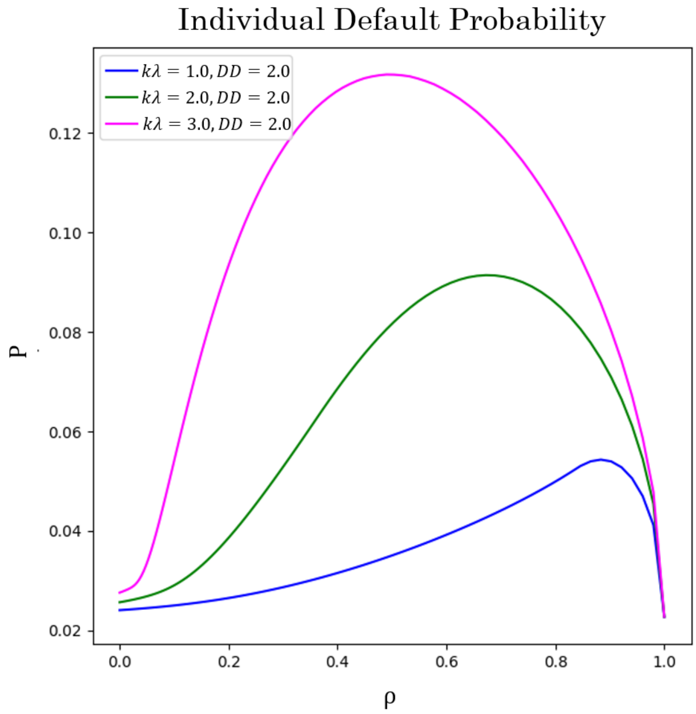

Figure 3.11: Default probabilities from the mean-field approximation as a function of ρ for various $k\lambda$ values.

Decreasing DD (or equivalently, also denoted by DD) has a similar effect in that it also produces a peak in default probabilities close to $\rho = 1$. On the other hand, the position of the peak is less easily shifted to the left, and decreasing DD much more heavily affects $\rho = 0$ default probabilities as well.

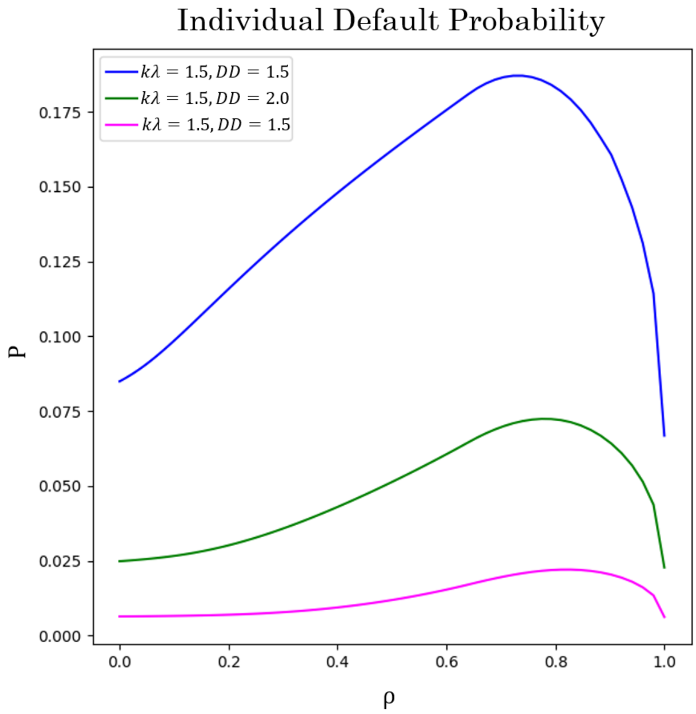

Figure 3.12: Default probabilities from the mean-field approximation as a function of ρ for various DD values.

Thus, within this model, default probabilities are lowest either under nearly independent cash flows or under perfectly correlated endowments, whereas high but not sufficiently high correlation can be the worst case. Note that there are extreme parameter combinations[35] under which the uncorrelated case is the least favorable, but for reasonable parameterizations, this characterization seems to hold.

These results are illustrated in the following figures. Figure 3.13's right panel shows the maximal default probability values, and Figure 3.13's left panel depicts at which ρ value that maximal default probability occurs. If interbank connections are small enough, the maximal default probability occurs at correlation values close to 1, and its value increases with decreasing DD. Initially, the maximum's location barely depends on DD, with the relationship becoming more pronounced as $k\lambda$ increases. As interbank connections become even larger, the DD-dependence of $\arg\max_\rho(P)$ becomes more pronounced, with a clearly discernible domain of parameters appearing. This domain, which appears at small DD (large fundamental default probabilities) and large $k\lambda$ (significant interbank connections), is characterized by remarkably large default probabilities that occur at low correlation values. Also note that the transition between this and the "normal" domain is nonsmooth, with sudden

[35]Unreasonable, e.g., for leading to excessively high default probabilities independently of the asset correlation, thus unexpected to arise in any practical application: very small DD with interbank claims dominating the balance sheet.

jumps in both $\max(P)$ and $\arg\max_\rho(P)$. The abruptness of the transition suggests that a financial system close to this boundary could experience a dramatic increase in systemic risk from a seemingly small deterioration in capitalisation or a modest increase in interbank exposures.

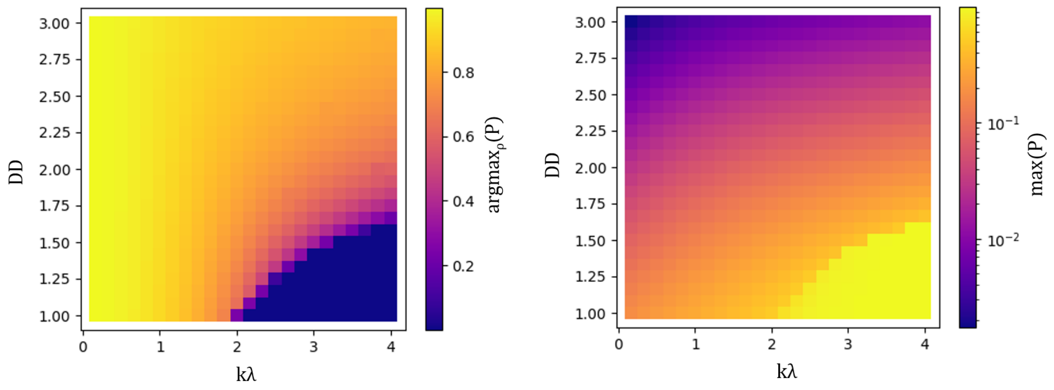

Figure 3.13: Heatmaps over $(k\lambda, DD)$: left shows the maximizing ρ, right shows the corresponding maximal default probability.

As discussed above, the ρ value leading to the minimal default probability is always 1 in homogeneous systems, with the minimal (fundamental) default probability having a trivial dependence on DD (the smaller DD is, the larger the minimal default probability) and no dependence on $k\lambda$.

Next, we focus on the "normal" domain of parameters, specifically when DD is sufficiently large, and focus on the effect that smaller $\arg\max_\rho(P)$ is typically associated with larger $\max(P)$. On scatter plots 3.14, results are generated from a grid between $2 \leq DD \leq 3$ and $0 \leq k\lambda \leq 4$, with one dot corresponding to results for one pair of parameters. The two plots differ only in coloring. For smaller DD values, as $k\lambda$ increases, $\max(P)$ increases more dramatically with $\arg\max_\rho(P)$, with the range of achievable $\arg\max_\rho(P)$ also increasing.

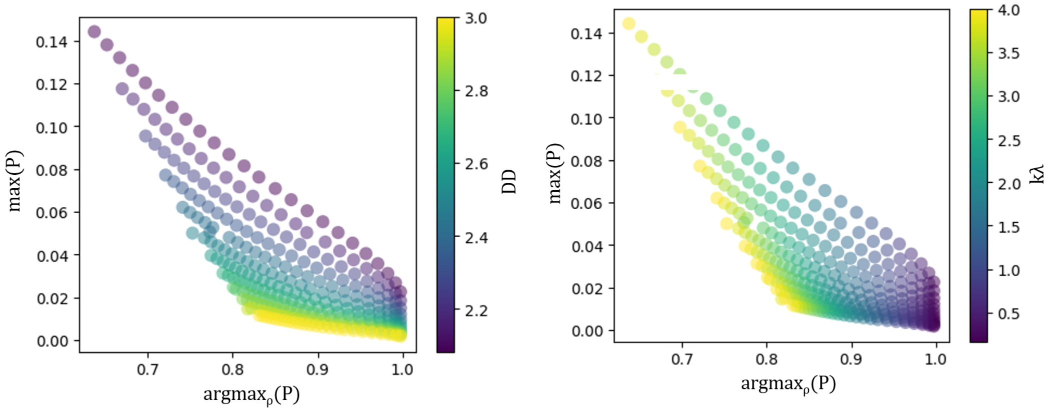

Figure 3.14: Relationship between the ρ value achieving the maximal default probability and the maximal default probability. Left: colors correspond to different DD values. Right: colors correspond to different $k\lambda$ values.

3.7 Conclusion and outlook

First, we analyzed (for practical reasons, small) systems via a Bayesian network approach, focusing on systems with special structures: fully connected and ring-like ones. By exploiting the special structure of these systems, we introduced analytical calculations for key quantities, enabling us to compute conditional default probabilities analytically without relying on simulations. We then examined the mean-field limit, the system's behavior when it consists of infinitely many microscopic agents. These two approaches turned out to be consistent in their main conclusions. Agents' default probabilities are non-monotonic in the correlation. Under reasonable parameterizations, a characteristic hump shape appears: default probabilities increase as correlation increases, reaching a peak at relatively high correlation values, and then rapidly decrease as correlation approaches one. The same effect holds for systemic default probabilities (defined in our study as at least two-thirds of the agents suffering default) in the Bayesian network approach. Larger interbank connectivity increases the maximal default probability and, at the same time, pushes the peak of the hump toward smaller correlation values. Lower expected external asset levels lead not only to larger fundamental default probabilities but also to a more pronounced hump shape, meaning larger maximal default probabilities, making the system more sensitive to changes in the correlation.

Highly correlated asset sides are a common assumption in the literature, but perfect correlation is quite unrealistic in practice. In our model, for most parameterizations, the hump shape is such that default probabilities peak for correlation values around 0.8–0.9. Note that this is precisely the region of typical rank correlations between banks' assets, as Banerjee and Feinstein (2022) reports. On the other hand, in this homogeneous one-period setup there is a significant and sudden reduction in default probabilities for correlations close to one. This result suggests that by assuming perfectly correlated endowments, one might miss an important effect and seriously underestimate default probabilities.

The stress-testing interpretation should be stated more narrowly. Within this stylized solvency-only network, perfectly correlated endowments need not generate the highest default probability. Dependence assumptions therefore deserve explicit sensitivity analysis rather than automatic collapse to a single "worst-case" convention. This does not by itself establish a general supervisory stress-testing rule: liquidity channels, fire sales, collateral dynamics, contract-specific market mechanics and real-economy feedback are outside the present scope.

Within the homogeneous one-period environment studied here, a broad conclusion is that default probabilities are minimized in the perfectly correlated case. The reason is not that high correlation is generically stabilizing. Rather, perfect correlation removes the cross-sectional asset heterogeneity needed for pure contagion defaults. The result is therefore as much a symmetry statement as a dependence statement. This contrasts with many results in the literature, where highly correlated asset sides are the major source of systemic risk, such as in Elliott et al. (2014) or Lehar (2005). Our conclusions are perhaps closest to Li and Wang (2021), who also focus on the interaction of asset side overlap and direct connections. It is important to highlight that our setup has certain assumptions that likely affect the result on the perfectly correlated case being optimal. In particular, we assumed uniform distances to default and hence uniform fundamental default probabilities. If one allows ex-ante heterogeneity—for example heterogeneous drifts or initial distances to default—then even under $\rho = 1$ some agents may remain above the default boundary while others fall below it, so contagion can reappear and the descending part of the hump may weaken or disappear. Section 3.6.2 investigates this through a Monte Carlo robustness exercise by letting $DD_i = DD + h\xi_i$ while keeping the cross-sectional mean of DD_i fixed. The key diagnostic is whether contagion default probability at $\rho = 1$ becomes non-zero once agents are no longer equally close to the default boundary. This observation helps reconcile the present chapter with Chapter 4: the current chapter isolates only the clearing-date heterogeneity mechanism, whereas Chapter 4 adds an additional forward-looking valuation channel. In Battiston et al. (2012a), when the initial fundamental default probability is heterogeneous across agents, increasing connectivity does not always enhance stability; the Monte Carlo robustness exercise shows how this issue already arises within the present one-period framework once exact homogeneity is relaxed.

There are numerous additional possibilities for model enhancements, shifting focus to different aspects of the problem. In our setup, interbank links only serve as a channel of contagion. Thus, larger interbank exposures always increase default probabilities. Introducing more complex interactions between agents would likely affect this relationship. Moreover, in the present study, correlation has been chosen independently of endowment variance. Another possible model enhancement could be to explicitly model the investment universe and thus take the interplay between correlation and portfolio variance into consideration. Under the large-system asymptotic, we did not consider random networks or core-periphery structures as in Rombach et al. (2014).

Chapter 4

Forward-looking network valuation using machine learning

4.1 Introduction

This chapter develops a forward-looking network valuation framework in which interbank claims are marked to market using default probabilities that are endogenously determined by the same dynamics that they help to shape. Building on the balance sheet setup and dynamic clearing mechanism described below, we model agents' solvency in discrete time and define haircuts as conditional to-maturity [36] default probabilities of counterparties. The key requirement is *strong self-consistency*: the probabilities used to discount a claim on agent j at time t must coincide with the actual probability that j defaults by T under the network dynamics induced by those very discounts.

To convey the intuition, consider a simple three-agent revaluation spiral (see Figure 4.1). Agents A, B, and C each have 60 in external assets and owe 50 to outside creditors. Interbank liabilities form a directed cycle:

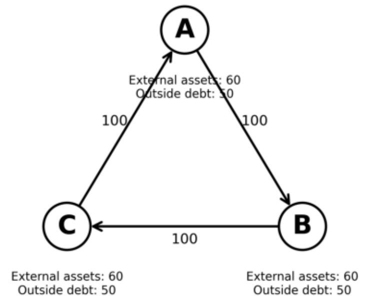

Figure 4.1: Asset and liability sides of agents in the example. Edges point from debtors to creditors.

[36] All claims are due at a single maturity, denoted by T hereafter.

If all interbank receivables are carried at face value, each agent has assets $60 + 100 = 160$ and liabilities $50 + 100 = 150$, so equity is 10 (indeed, all three agents' equities are $60 + 100 - 50 - 100 = 10$) and all three agents are solvent.

Now suppose that, for any reason (a downgrade, a rumour, a tightening of funding conditions), the market applies a 20% write-down to claims on A at the interim revaluation date. Agent B therefore marks its 100 receivable from A down to 80, so its equity becomes $60 + 80 - 50 - 100 = -10$. This purely anticipatory markdown can thus force B into default immediately under mark-to-market accounting or capital constraints. Once B is regarded as defaulted, its creditor C must also write down its claim on B (for instance to 80 as well[37]), which yields $60 + 80 - 50 - 100 = -10$ for C's equity and can in turn push C into default. Finally, A marks down its claim on C (again, say to 80), so its equity becomes $60 + 80 - 50 - 100 = -10$ and A itself can become insolvent even though its external assets never changed. In this way, an initially "small" forward-looking write-down can generate a self-fulfilling spiral: markdowns today trigger earlier defaults, which reduce future payments and thereby justify the markdowns.

A strongly self-consistent[38] valuation is one in which this loop closes exactly: the discounts used today are consistent with the future payments implied by the system dynamics that those very discounts generate. Achieving this fixed point is the central computational challenge of the chapter.

Like the other chapters in this dissertation, the present model considers obligations due at a single maturity T. However, it introduces multiple interim revaluation dates before T, creating a genuinely multi-period setting. This chapter extends the analysis of the previous chapters by allowing anticipated future defaults to trigger current defaults through mark-to-market revaluation of interbank claims.

The practical interpretation is deliberately narrower than the full universe of bank exposures. The chapter models a forward-looking economic valuation mechanism for bilateral claims in an interbank network. The cleanest interpretation is uncollateralized or unsecured interbank claims, or reduced-form bilateral credit-sensitive exposures. It is not a literal model of all accounting treatments of interbank loans, nor of derivative collateral, margining, repo or CCP systems.

The main technical contribution is a learning-based scheme to approximate this self-consistent probability surface in higher dimensions. The computational difficulty is that the object we seek is simultaneously forward-driving

[37] We only wish to illustrate how applying haircuts on claims can push agents into default even before maturity. The choice of haircuts in this examples is arbitrary; in fact, the haircut C applies on its claim on B will be overly optimistic and misaligned with some assumptions we make later this chapter.

[38] A concept we will formally introduce later, see Definition 4.16.

and backward-looking. On the one hand, any candidate surface of to-maturity default probabilities (or, more generally, discount factors for interbank claims) determines today's mark-to-market equities and therefore who survives to the next revaluation date—this is the forward evolution of the system. On the other hand, those same probabilities are conditional expectations about events between t and T, so they are naturally computed by a backward recursion from the terminal condition at maturity. Strong self-consistency therefore amounts to a forward–backward fixed point: probabilities are required to value balance sheets forward in time, but the probabilities themselves must be obtained by working backward given the very default dynamics induced by those valuations.

Feinstein and Sojmark (2022) make this structure explicit on a discretized multinomial tree for the exogenous risk factors. Conceptually, one alternates between (i) a forward sweep that propagates mark-to-market balance sheets and solvency indicators through the tree under the current pricing rule, and (ii) a backward sweep that recomputes conditional to-maturity solvency/default probabilities at each node via dynamic programming. Iterating these two passes closes the fixed point, but the number of tree nodes grows exponentially in the dimension of the state, quickly becoming infeasible for larger networks.

Instead of solving this forward–backward fixed point on a discretized state space, we construct a family of time-indexed neural networks that map current balance sheet states to to-maturity default probabilities. These predictors are trained backwards in time via a temporal-difference (TD) style scheme (see Footnote 18). At each date, the raw default probability predictions are passed through a clearing step that enforces weak self-consistency (claims on defaulted agents should be valued at 0); the resulting post-clearing survival indicators define a one-step bootstrap target for the previous time point. The TD construction itself does not require the system to be symmetric. On the other hand, in the numerical experiments we do assume symmetric systems. Under that homogeneous fully connected specification, the labelled bank ordering is arbitrary, so in some experiments we also utilize a permutation-equivariant DeepSets[39] parametrization of the neural predictor; see Appendix D.

To put the fully state-dependent valuation into perspective, we also introduce three benchmarks. The first is a state-independent haircut scheme (Section 4.4) in which claims are discounted using default probabilities conditional only on survival up to time t, but not on the realised path of assets. This yields a weakly self-consistent yet optimistic valuation that serves as a lower bound on contagion effects. The second is a comonotonic benchmark

[39]DeepSets (Zaheer et al., 2017) is a neural network architecture designed for set-valued (unordered) inputs. It processes each element through a shared encoder, aggregates the encoded representations (e.g., by averaging), and applies a decoder to produce outputs that are equivariant to permutations of the input elements.

(Section 4.6.5), where all external assets move in lockstep so that the multi-agent problem collapses to a scalar recursion in the common asset level; this provides a tractable consistency check for the learned TD model in the fully correlated case. The third benchmark (Section 4.6.6) keeps the same network revaluation / clearing dynamics as in training, but replaces the neural network predictor of the probability of default with a static barrier-crossing probability for the external assets, assuming interbank claim valuations remain constant at the given level.

Numerical experiments under symmetric, fully connected systems (Assumption 4.17) reveal several qualitative features. First, default probabilities exhibit a characteristic hump-shaped dependence on the asset correlation parameter ρ: intermediate correlations often maximise the likelihood of default, while very low or very high correlations can be safer. Second, the contrast with Chapter 3 can be organized around two mechanisms. The first is the same clearing-date heterogeneity mechanism studied there: as correlation rises, diversification of interbank asset values weakens, but very high correlation also eliminates the cross-sectional dispersion needed for pure contagion defaults. The second is new to the present chapter: under multi-period, strongly self-consistent, state-dependent haircuts, bad common paths generate synchronous markdowns of interbank assets before default is realized. Because this second mechanism is strongest at high correlation, the minimal default probability no longer necessarily occurs at $\rho = 1$; instead, the best-case correlation can shift to an interior value or zero. Third, increasing either the number of banks or the size of interbank exposures magnifies this hump-shape characteristic, very similarly to the previous chapter's results.

Taken together, the results show that a TD-style learning approach can deliver strongly self-consistent forward-looking valuations on networks that can sidestep the curse of dimensionality[40] of the tree-based method of Feinstein and Sojmark (2022) while capturing the endogenous feedback that drives forward-looking contagion.

The remainder of this chapter is organized as follows. Section 4.2 introduces the related literature. Section 4.3 presents the balance sheet dynamics and clearing mechanism. Section 4.4 develops the state-independent benchmark. Section 4.6 introduces the TD-learning approach for state-dependent valuations. Section 4.7 reports numerical experiments, and Section 4.8 concludes. Appendix C describes the fitting scheme in pseudocode terms while Appendix D records the DeepSets implementation used for the symmetric experiments.

[40]The curse of dimensionality refers to the exponential growth in computational cost or data requirements as the number of state variables increases, making grid-based or tree-based methods infeasible for high-dimensional problems.

The Python code for this chapter is available in the following repository: `https://github.com/sasszol/NetworkValuationML`

4.2 Related literature

While Chapter 1 provided a unified literature review on financial networks and systemic risk, here we focus on the specific literature most relevant to forward-looking valuation in network settings. The literature on network valuation has progressed through three stages. In the foundational Eisenberg and Noe (2001) model and its extensions, interbank liabilities are carried at face value until a counterparty actually defaults. This historical-cost approach captures default cascades at the moment of settlement but precludes anticipatory adjustments: creditors do not mark down claims even as a borrower's solvency deteriorates, potentially underestimating the buildup of systemic stress before an actual default event.

A second generation of models introduces ex-ante valuation adjustments. Bardoscia et al. (2017, 2019) and Barucca et al. (2020) propose frameworks in which a claim on a counterparty is valued not at par but at an amount reflecting the probability that the counterparty will remain solvent by maturity. Expected future losses thereby depress each bank's equity today, even without any current default—capturing the important intuition that mere concerns about a counterparty's creditworthiness can create informational contagion. However, the default probabilities driving these haircuts are exogenously specified (e.g. from standalone asset dynamics), so there is no feedback from the act of marking down assets to the network's evolution. The adjustments are therefore only partially self-consistent.

The third stage endogenizes this process. Feinstein and Sojmark (2022) develop a dynamic model in which mark-to-market values and solvency probabilities are jointly determined in equilibrium via a forward-backward iteration. This captures powerful feedback loops—a perceived increase in one bank's default risk triggers markdowns at its creditors, weakening their balance sheets and potentially raising their own default risk, which feeds back to the first bank. Among the qualitative novelties of this fully endogenous approach are contagion without actual default (anticipation alone can trigger insolvency), endogenous volatility clustering in banks' equity trajectories, and the emergence of a term structure of network risk when multiple maturities are present.

Computing such self-consistent, forward-looking valuation of interbank liabilities is demanding because it requires solving the forward–backward fixed point described above over the joint state of all banks' assets and liabilities. Two strategies have emerged to cope with this complexity. The first is state-space discretisation: Feinstein and Sojmark (2022) construct a multinomial tree (lattice) for the joint external-asset process and iterate between a forward propagation of solvency/default states and a backward recursion for conditional to-maturity solvency probabilities. While general, the number of nodes grows exponentially with the dimension, manifesting the classic "curse of dimensionality." High-dimensional tree methods are well known to become prohibitively memory-intensive even for ten underlying factors in option-pricing contexts, a benchmark that translates into similar infeasibility for networks with more than a handful of banks (Chen and Goldberg, 2018).

The second assumes perfectly comonotonic external assets, collapsing the state space to a single factor and enabling closed-form debt and equity prices (Banerjee and Feinstein, 2022). Although comonotonic endowments deliver sharp analytical bounds, they capture only an extreme dependence structure and may overstate contagion whenever asset returns are less than perfectly correlated, limiting their applicability to broader stress-testing scenarios. Full discretisation remains universally applicable but its computational burden motivates research into dimension-reduction techniques and machine-learning surrogates for large-scale networks (Ghosh et al. (2023); Fernández-Villaverde et al. (2024)). This chapter contributes to this effort by developing a temporal-difference learning approach (Sutton and Barto, 2018) that can handle larger networks while preserving the endogenous feedback structure.

A related body of literature extends the Eisenberg-Noe mechanism to dynamic settings while retaining historical-cost accounting. Capponi and Chen (2015) formulate a multi-period clearing system with lender-of-last-resort liquidity; Banerjee et al. (2025) push clearing to continuous time via cash-flow SDEs; Calafiore et al. (2023) cast dynamic clearing as a sequence of linear programmes admitting temporary pseudo-defaults; and Sonin and Sonin (2025) models money as fluid flowing through pipes. All four approaches deliver richer default timing than static models, but none incorporates mark-to-market haircuts—creditors' valuations remain insensitive to rising default probabilities until the moment of failure, precisely the limitation that the present chapter addresses. This literature is used here as an economic valuation and distress-contagion lens. Although the mechanism is related to debates about fair-value measurement and procyclicality, the chapter does not propose a ready-made supervisory accounting rule or prudential valuation adjustment.

4.3 Balance sheet dynamics and forward-looking solvency contagion

We adopt a discretized time grid for two principal reasons. First, continuous-time revaluations are infeasible from a practical standpoint. Second, our framework relies on machine learning and numerical methods, which are naturally formulated in discrete time.

In our stylized setting, the financial system comprises N financial institutions, which we shall refer to as agents. The index set of agents is denoted by $I = \{1, 2, \ldots, N\}$.

Let $\mathcal{T} = \{0, 1, \ldots, T\}$ denote the set of discrete revaluation indices over a fixed calendar horizon $H > 0$, and define the time step $\Delta t = H/T$. The index $t \in \mathcal{T}$ corresponds to calendar time $t\,\Delta t$; throughout we use the shorthand $Z(t)$ for the value of a process at grid index t. We work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ endowed with a filtration[41] $\{\mathcal{F}(t)\}_{t\in\mathcal{T}}$. Whenever we appeal to continuous-time Brownian scaling (e.g. in barrier-crossing approximations), the relevant calendar time remaining from t to T is $\tau_{\text{rem}}(t) = (T - t)\Delta t$.

Definition 4.1. The external asset processes are discrete-time stochastic processes adapted to this filtration. At time t, the external asset of agent i—also referred to as nonbank assets or *endowments*—is denoted by $A_i(t)$, an $\mathcal{F}(t)$-measurable random variable.

These endowments represent the sole source of uncertainty in the system. One may interpret them as being subject, up to terminal time T, to both credit and market risk, modeled through stochastic shocks.

The balance sheet of agent i at time t is given below.

Balance sheet of agent i at time t	
assets	liabilities
$A_i(t)$ $\sum_{j=1}^{N} V_j(t) L_{ji}$	ℓ_i $\sum_{j=1}^{N} L_{ij}$ $E_i(t)$

Definition 4.2. An $L \in \mathbb{R}^{N\times N}$ matrix represents the system of interbank payment obligations due at time T, which are deterministic and already known at the initial time 0.

These obligations are best interpreted as fixed nominal interbank obligations, such as unsecured loans or deposits, or as reduced-form bilateral credit-sensitive exposures measured at the modeling horizon. Such exposures generate counterparty risk (see, e.g., Tirole (2011)), but derivative, repo or CCP-cleared interpretations would require explicit netting, collateral, variation or initial margin, close-out and default-waterfall mechanisms, which are omitted here (Basel Committee on Banking Supervision and International Organization of

[41] A filtration $\{\mathcal{F}(t)\}$ is an increasing family of information sets (sigma-algebras) representing what is observable at each time t. Intuitively, $\mathcal{F}(t)$ captures all information available to market participants up to and including time t.

Securities Commissions, 2015; Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions, 2012). By construction, $L_{ij} \geq 0$ for all $i, j \in I$, where L_{ij} denotes agent i's liability to agent j, and $L_{ii} = 0$ for all $i \in I$. Importantly, liabilities are carried on the balance sheet at face value.

Definition 4.3. For each $i \in I$, let $\ell_i \geq 0$ denote agent i's total payment obligation at time T to entities external to the system of agents. These deterministic *non-bank liabilities* may consist, for instance, of deposits (retail or corporate). We assume that such claims have the same seniority as the agents' interbank liabilities.

The network of interbank obligations may be represented by a directed graph. An edge from i to j indicates that agent i has a payment obligation to agent j at time T. These obligations are not netted, and directed cycles may occur in general. We assume that the interbank liability structure is exogenous: agents cannot alter their capital structure nor choose their counterparties.

Definition 4.4. Agent i's *total payment obligation* at terminal time T is given by

$$b_i = \ell_i + \sum_{j=1}^{N} L_{ij}.$$

While liabilities are always accounted for at par, agents mark their interbank assets to market via some valuation methodology. We express this in form of valuation functions for the interbank claims.

Definition 4.5. Let $V_j(t)$, an $\mathcal{F}(t)$-adapted stochastic process denote the *valuation function* agent j's creditors use to mark j's obligation to market at time t. $1 - V_j(t)$ will be called the corresponding *haircut*[42].

This is to express that valuations trivially depend on the nominal value of the payment obligation, and additionally might depend on other parameters or variables. Importantly, any t-valuation can only depend on $\mathcal{F}(t)$-measurable random variables. We exclude information asymmetry in how agents view each other's solvency, therefore two agents crediting a common third will both use the same valuation function in marking these contracts to market, resulting in the same haircuts.

Before we specify the natural requirements that haircuts must satisfy and introduce some possible ways of defining these haircuts, we must first define default.

[42] The term haircut is used here in a model-specific sense: the complement of the valuation function, equal to the anticipated probability that the counterparty defaults by maturity. This is distinct from the collateral or repo haircuts used in market practice.

Definition 4.6. The *equity* $E_i(t)$ of agent i at time t is the difference between its asset and liability sides:

$$E_i(t) = A_i(t) - \ell_i + \sum_{j=1}^{N} V_j(t) L_{ji} - \sum_{j=1}^{N} L_{ij} \quad \forall\, t \in \mathcal{T},\ t < T. \tag{4.1}$$

In calculating the equity, it is assumed that interbank obligations at maturity towards already defaulted counterparties are still accounted for. This is in line with the typical assumptions in the literature and with the one-period model of Chapter 3.

Definition 4.7. The *survival indicator process* $S_i(t)$ of agent i is a $\{0,1\}$-valued stochastic process adapted to $\mathcal{F}(t)$. $S_i(t)$ represents whether agent i survived up to and including the revaluation at time step t. We construct the survival indicators by re-evaluating the balance sheets at each grid point. For $t = 0$,

$$S_i(0) = \mathbb{1}\{E_i(0) \geq 0\} \tag{4.2}$$

While for later t, we also need to take past defaults into account.

$$S_i(t) = S_i(t-1)\, \mathbb{1}\{E_i(t) \geq 0\} \quad \forall\, t \in \mathcal{T} \setminus \{0, T\}. \tag{4.3}$$

The terminal date T requires special treatment, because payment obligations must be cleared. For that, similarly to Chapter 3, we assume zero recovery on interbank claims.

Definition 4.8. Let $z^T = (z_1, z_2, \ldots, z_N)$ denote the *terminal clearing vector*, where $z_i = 1$ if agent i meets all obligations and $z_i = 0$ otherwise. Given z^T, agent i receives $\sum_{j=1}^{N} z_j L_{ji} S_j(T-1)$ from the other agents in total.

The factor $S_j(T-1)$ ensures that only agents who survived the previous round can make payments in the final clearing. We follow Rogers and Veraart (2013) in formalizing the equilibrium payment scheme.

Definition 4.9. A *terminal equilibrium survival vector* is a vector $z^* \in \{0,1\}^N$ satisfying

$$\Phi(z^*) = z^*,$$

where $\Phi : \{0,1\}^N \to \{0,1\}^N$ with

$$\Phi(z)_i = \begin{cases} 1 & \text{if } b_i \leq A_i(T) + \sum_{j=1}^{N} z_j L_{ji} S_j(T-1), \\ 0 & \text{otherwise.} \end{cases}$$

Rogers and Veraart (2013) prove the existence of such an equilibrium vector in this setup; uniqueness is not generally guaranteed. Let us observe that the system's behavior during terminal clearing is exactly the same as in Chapter

3, with the caveat that some of the incoming payments L_{ij} are now zeroed out due to earlier defaults. For a more detailed discussion of existence and uniqueness, we refer the reader to Section 3.3 thereof. Similarly to the earlier chapter, we always choose the greatest terminal clearing vector. We can now complete the definition of the survival indicator process.

Definition 4.10. The *survival indicator process* at terminal time T for agent i is given by combining the terminal payment indicator and the $T-1$ survival indicator:

$$S_i(T) = S_i(T-1)\, z_i^*. \tag{4.4}$$

Definition 4.11. The *default time* of agent i is the stopping time

$$\tau_i = \inf\{t \in \mathcal{T} : S_i(t) = 0\}.$$

Note the iterative construction of the survival indicator process in Equations 4.2, 4.3 and 4.4. Therefore, by construction, there is no escape from default:

$$S_i(t) = 0 \quad \text{for } t > \tau_i. \tag{4.5}$$

Consequently, any reasonable valuation function must satisfy $V_i(t) = 0$ for all $t > \tau_i$. Furthermore, a default at τ_i should immediately affect the counterparties. This is formalized by the notion of weak self-consistency.

Definition 4.12. A valuation function is called *weakly self-consistent* if

$$V_i(t) = 0 \quad \forall\, t \geq \tau_i.$$

The importance of weak self-consistency in avoiding pathological behavior is illustrated below; the provenance of the valuation function should be ignored at this stage—we focus only on the consistency requirement.

Example 4.13. Assume $N = 2$, $\ell_1 = \ell_2 = 1$, $\mathcal{T} = \{0, 1, 2\}$,

$$L = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

and deterministic asset processes

$$A_1 = (2, 1, 2), \qquad A_2 = (2, 1, 2).$$

Let $V_2(1) = 0.5$. Then using Definition 4.6, $E_1(1) = A_1(1) + V_2(1)L_{2,1} - L_{1,2} - \ell_1 = 1 + 0.5 - 1 - 1 = -0.5$, so agent 1 defaults at $t = 1$. Choose a valuation function for agent 1 that is not weakly self-consistent, say $V_1(1) = 1$. Agent 2 therefore avoids default at $t = 1$ with $E_2(1) = 1 + 1 - 1 - 1 = 0$. Because there is no escape from default, $S_1(2) = 0$, implying $V_1(2) = 0$. Since agent 2's external assets are high enough at $t = 2$, it survives with $E_2(2) = 0$. Had weak consistency been imposed, $V_1(1)$ would have been forced to zero at $t = 1$, inducing an immediate default of agent 2 as well.

As in earlier chapters, we again distinguish fundamental default and contagion default.

Definition 4.14. Agent i suffers *fundamental default* at time $\tau \leq T$ if $S_i(\tau) = 0$ while $S_i(\tau - 1) = 1$ and $A_i(\tau) - \ell_i + \sum_{j=1}^N L_{ji} - \sum_{j=1}^N L_{ij} < 0$.

Intuitively, fundamental default happens if equity gets negative even if all haircuts are 0 (or equivalently, all valuation functions are 1) before T and all claims are received at T. In contrast, contagion default means default due to haircuts being larger than 0 before maturity, or interbank inflows being smaller than the face values at maturity while an agent would not have suffered fundamental default.

Definition 4.15. Agent i suffers *contagion default* at time $\tau \leq T$ if $S_i(\tau) = 0$ while $S_i(\tau - 1) = 1$ and $A_i(\tau) - \ell_i + \sum_{j=1}^N L_{ji} - \sum_{j=1}^N L_{ij} \geq 0$.

Self-consistency naturally leads to a particular specification of the haircuts in a zero-riskless-rate world.

Definition 4.16. The *strongly self-consistent* valuation function uses counterparties' $\mathcal{F}(t)$-conditional default probabilities[43] as haircuts (or equivalently, $\mathcal{F}(t)$-conditional survival probabilities as valuation functions):

$$V_j^{ssc}(t) = \mathbb{P}\big[S_j(T) = 1 \,\big|\, \mathcal{F}(t)\big] \quad \forall\, j \in I.$$

Assuming zero recovery (a typical choice for similar models, see e.g. Gai et al. (2011)) and zero interest rates, this choice of valuation function means interbank claims are priced at their expected terminal payments at maturity. It is self-consistent by design: the default probabilities embedded in the valuation functions are exactly those realised by the system's dynamics. Computing or approximating the functions $V_j^{ssc}(t)$ is therefore the main task. Note that this function is also weakly self-consistent.

Throughout the chapter, we will apply a number of simplifying assumptions on the system's parameters.

Assumption 4.17. The network is fully connected with homogeneous exposures: each agent owes a total amount of $k\lambda$ in interbank liabilities and holds the same face value of interbank assets. Here $k = N - 1$, and each interbank obligation has face value λ (so $L_{ij} = \lambda$ for $i \neq j$ and $L_{ii} = 0$). Each agent starts with the same external asset value $A_i(0) = a_0$ and external assets evolve

[43] Throughout, we work under the real-world measure.

on the grid as driftless arithmetic Brownian motions (ABMs) via the following common market-factor representation: for $t = 0, \dots, T-1$,

$$A_i(t+1) = A_i(t) + \sigma\sqrt{\Delta t}\left(\sqrt{\rho}\, M(t) + \sqrt{1-\rho}\,\varepsilon_i(t)\right),$$

where $\{M(t)\}_{t=0}^{T-1}$ and $\{\varepsilon_i(t)\}_{i=1,\dots,N;\, t=0,\dots,T-1}$ are i.i.d. standard normal random variables, independent across time and mutually independent. In particular, $\mathrm{Corr}(A_i(t+1) - A_i(t),\, A_j(t+1) - A_j(t)) = \rho$ for all $i \neq j$.

The ABM specification allows us to make a number of simplifying assumptions without the loss of generality. The external liability vector ℓ is incorporated into the assets (so $\ell = 0$). The time horizon is $H = 1$ (year) and the grid has T steps, so $\Delta t = H/T$. The asset volatility parameter is $\sigma = 1$ in our numerical experiments. As is clear from Definition 4.14, in fully connected and symmetric systems $A_i(t) - \ell_i$ is the key quantity in driving fundamental defaults. Therefore, since $\ell = 0$ we will also denote a_0 as DD (distance-to-default). This assumption is maintained in the numerical experiments and in the computational shortcuts described later; the definition of strong self-consistency and the TD training recursion do not depend on it.

4.4 A benchmark model using state-independent haircuts

Next we define a valuation function that will serve as a benchmark for the later models. It is a simple, and optimistic proxy that is weakly, but not strongly self-consistent.

Definition 4.18. The *state-independent survival probability based valuation function* uses counterparties' survival probabilities—conditional on the counterparty survives up to that point—as haircuts:

$$V_j^{si}(t) = \mathbb{P}\big[S_j(T) = 1 \,\big|\, S_j(t) = 1\big] \quad \forall\, j \in I.$$

The major difference compared to Definition 4.16 is the lack of dependence on the state information. We only condition on agent j being alive at time t. The corresponding haircut is, in fact, a deterministic time-dependent function for all agents.

Algorithm 1 computes a deterministic matrix of state-independent (time-dependent, but state-averaged) default probabilities, i.e. a PD surface $\widehat{P} \in [0,1]^{(T+1)\times N}$. The key idea is to view such a PD surface as an input that defines deterministic haircuts $\widehat{V}_j^{\widehat{P}}(t) = 1 - \widehat{P}_{t,j}$, run the balance-sheet dynamics and embedded clearing with these haircuts fixed, and then measure the conditional-to-maturity default probabilities that the resulting system dynamics generates.

Consistency, in this benchmark, means that the PD surface used to haircut claims is the same as the PD surface implied by the realised (state-averaged) survivals under that choice.

Formally, fix a candidate $\widehat{P} \in [0,1]^{(T+1)\times N}$ and let $S_i^{\widehat{P}}(t) \in \{0,1\}$ denote the survival indicator produced by the discrete-time dynamics (with the same within-period revaluation and monotone clearing logic as in the inner loop of Algorithm 1). Write the unconditional survival probabilities as

$$\pi_{t,i}^{\widehat{P}} = \mathbb{E}\Big[S_i^{\widehat{P}}(t)\Big], \qquad t \in \mathcal{T},\ i \in I.$$

Whenever $\pi_{t,i}^{\widehat{P}} > 0$, the induced *state-independent* conditional default probability from t to T is

$$\mathbb{P}\Big[S_i^{\widehat{P}}(T) = 0 \,\Big|\, S_i^{\widehat{P}}(t) = 1\Big] = 1 - \frac{\pi_{T,i}^{\widehat{P}}}{\pi_{t,i}^{\widehat{P}}}.$$

This motivates a (population) operator Θ on PD surfaces, defined componentwise by

$$(\Theta(\widehat{P}))_{t,i} = 1 - \frac{\pi_{T,i}^{\widehat{P}}}{\pi_{t,i}^{\widehat{P}}}, \qquad t = 0,\dots,T, \quad i = 1,\dots,N, \tag{4.6}$$

(with the natural convention $(\Theta(\widehat{P}))_{t,i} = 1$ when $\pi_{t,i}^{\widehat{P}} = 0$). A state-independent benchmark surface is then a fixed point of this map,

$$\widehat{P}^\star = \Theta(\widehat{P}^\star),$$

which is the discrete-time counterpart of Definition 4.18.

In general, Θ is not available in closed form because the expectation in $\pi_{t,i}^{\widehat{P}}$ is taken over the full coupled network dynamics. Algorithm 1 therefore replaces Θ by its Monte Carlo analogue: for K simulated scenarios $\omega_1,\dots,\omega_K$, compute the empirical survival frequencies $\widehat{\pi}_{t,i}^{\widehat{P}} = \sum_{k=1}^K S_i^{\widehat{P}}(t;\omega_k)$ and update via

$$(\Theta_K(\widehat{P}))_{t,i} = 1 - \frac{\sum_{k=1}^K S_i^{\widehat{P}}(T;\omega_k)}{\sum_{k=1}^K S_i^{\widehat{P}}(t;\omega_k)},$$

again setting $(\Theta_K(\widehat{P}))_{t,i} = 1$ if the denominator vanishes. The outer loop is a Picard-type fixed-point iteration[44] $\widehat{P}^{(m+1)} = \Theta_K(\widehat{P}^{(m)})$, initialized at the most optimistic surface $\widehat{P}^{(0)} = \mathbf{0}$ and terminated when $\|\widehat{P}^{(m+1)} - \widehat{P}^{(m)}\|_2$ falls below a tolerance.

We do not attempt to establish existence/uniqueness of fixed points of Θ, nor a general convergence or stability guarantee for the induced iteration

[44] A Picard iteration (also called successive approximation) is a method for finding fixed points of a map T: starting from an initial guess $x^{(0)}$, one iterates $x^{(m+1)} = T(x^{(m)})$ until convergence.

(either at the population level or for its Monte Carlo approximation). The mapping is highly nonlinear because defaults feed back into valuations, and in principle one can imagine parameter regimes where multiple fixed points exist or where naive fixed-point iteration fails to converge. For the purposes of this chapter, we therefore treat Algorithm 1 as a practical computational procedure: we monitor convergence directly via the stopping criterion above. Empirically, across all numerical experiments reported later, the iteration always stabilised from the optimistic initialisation and produced a time-dependent PD surface that was robust to further iterations, which we take as our state-independent benchmark for comparison with the state-dependent (and more strongly self-consistent) methods.

4.5 A naive machine learning method to estimate strongly self-consistent haircuts

A first, very direct way to estimate self-consistent haircuts is to define a default probability surface $V_\theta(a,s,t) \in [0,1]^N$ as an ordinary regression object. The surface—once fitted—is supposed to be used for reading off PD predictions for a given parameter vector. Given a vector of candidate parameters θ and a collection of initial states (a,s,t) (asset levels and survival indicators for all agents, plus time indices) we can simulate the system from the beginning to maturity T in the current model (using $V_\theta(a,s,t)$ to calculate haircuts when needed in the revaluation mechanisms at $t<T$) and the clearing mechanism at T to generate terminal survival indicators $S_T^i \in \{0,1\}$ for each agent i for each initial state. Repeating this procedure over many Monte Carlo scenarios $\omega_1,\ldots,\omega_K$ yields an empirical estimate of the conditional default probability at horizon T,

$$\widehat{p}^{\,i,\mathrm{MC}}(t,a,s) \;=\; \frac{1}{K}\sum_{k=1}^{K} \mathbb{1}\big\{S_T^i(t,a,s;\,\omega_k)=0\big\},$$

for each agent i and each training state (t,a,s) where K denotes the total number of scenarios. A neural network with inputs $(a\in\mathbb{R}^N, s\in\mathbb{R}^N, t\in\mathbb{R})$ is then trained, to regress onto these Monte Carlo estimates by minimizing a loss of the form $\sum_i \big(V_\theta^i(a,s,t)-\widehat{p}^{\,i,\mathrm{MC}}(t,a,s)\big)^2$, aggregated over a buffer of size B of sampled states.

In pseudocode terms, the algorithm alternates between (i) generating fresh Monte Carlo estimates of realised default probabilities for each sampled state by running the full forward simulation, and (ii) taking a gradient step so that the current network output matches these realised frequencies. In practice, one can train a set of neural networks, one for each $t\in\mathcal{T}\setminus\{T\}$. Starting from the $T-1$ network (there is nothing to fit at T, the outcome is computed by performing the terminal clearing), progressing backwards, we fit one network at a time, keeping later steps' already fitted networks frozen. Each network now has inputs $(a \in \mathbb{R}^N, s \in \mathbb{R}^N)$. The sample batch then only contains initial states starting at a given t, and all $\tau > t$ networks are kept frozen at their fitted parameter values θ in the forward simulation.

While conceptually simple, this naive Monte Carlo regression has several serious drawbacks. First, it is extremely sample– and computationally inefficient. For every gradient update we must, for each training state and independently of the time–to–maturity level, run a full forward simulation with scenarios K of the network over all remaining time steps and across many return scenarios.

Second, the targets $\widehat{p}^{i,\mathrm{MC}}(a, s)$ are noisy Monte Carlo estimates of an expectation, and they are re–computed from scratch after every parameter update. Because the environment itself depends on the current PD surface (via the forward–looking revaluations), each change in θ alters the entire mapping from state to terminal default probabilities. This makes the regression target highly non–stationary: the network is chasing a moving Monte Carlo estimate that it itself helps to define.

A further issue is that the naive approach effectively treats each (a, s) as an independent supervised–learning problem with a long–horizon label. The Markovian structure of the dynamics is not exploited. All information about the intermediate states $((A_{t+1}, S_{t+1}), (A_{t+2}, S_{t+2}), \dots)$ along a simulated path is discarded; each trajectory only contributes a single terminal default indicator to the training target for its initial state. Errors in the PD surface at future times, say $V_\theta(\cdot, t+1)$, only influence learning at time t indirectly, via their effect on the Monte Carlo estimate of the terminal default frequency. There is no local, one–step notion of temporal credit assignment that would tell the model an over– or under–estimate at time $t+1$ should feed back into the prediction at time t. In a strongly coupled network, this lack of structured feedback can lead to slow convergence, oscillations between incompatible PD surfaces, or even apparent convergence to a self–inconsistent fixed point.

4.6 Learning state-dependent default probabilities via a TD analogue

The limitations discussed in Section 4.5 motivate moving away from pure Monte Carlo regression towards a temporal-difference (see (Sutton and Barto, 2018); TD in the following) scheme. Instead of trying to learn the till-maturity default probability in one shot from a noisy Monte Carlo target, a TD learner exploits the dynamic programming structure of the problem: default probabilities at time t are related to those at time $t+1$ through the one–period transition of the assets and the periodic mark-to-market revaluation. By bootstrapping [45] on the network's own prediction at the next time step (as opposed to running all the way to T and averaging), we obtain lower–variance targets, can reuse every step of every simulated path for learning, and impose self–consistency locally in time. The TD scheme in this section therefore provides a much more efficient and stable way to drive the PD surface towards a fixed point that is consistent with the underlying network dynamics.

Previously, we only talked about clearing at T. In the following, to make sure that the valuation function is weakly self-consistent, we will perform a very similar clearing step at each $t < T$ as well. In particular, we first get probability of default predictions from the learner. Then, these raw default probability predictions will be run through a clearing mechanism which is essentially equivalent to the one in Definition 4.9, and with that post-processing step we guarantee avoiding the pathological cases illustrated in Example 4.13.

As earlier, we still consider a discrete grid $\mathcal{T} = \{0, 1, \ldots, T\}$ but model the within-period timing at each $t < T$ explicitly. Two survivorship notions are distinguished.

Definition 4.19. The *pre-clearing indicators* $S^-(t) \in \{0,1\}^N$ encode who is considered alive at the start of step t, before any mark-to-market revaluation and clearing at t.

Definition 4.20. The *post-clearing indicators* $S^+(t) \in \{0,1\}^N$ are obtained after valuing interbank claims using counterparties' to-maturity survival probabilities (so applying the haircut) and running the balance-sheet clearing map at time t to guarantee weak self-consistency.

Defaults revealed by this valuation-and-clearing stage are time-t events and are reflected in $S^+(t)$. By construction $S^+(t) \leq S^-(t)$ componentwise, $S^-(t) = S^+(t-1)$ for $t \geq 1$, and default is absorbing: once $S_i^+(t) = 0$, then $S_i^-(t') = S_i^+(t') = 0$ for all $t' > t$.

The learning model uses only information measurable at the start of the step. Concretely, the predictor at time t receives as covariates the pre-clearing external asset vector $A(t) \in \mathbb{R}^N$. Survivorship indicators are not passed as exogenous inputs; agents already defaulted prior to t are encoded in $A(t)$ by masking their external asset value with a dedicated low value a_{dead}, so that pre-clearing survivorship $S^-(t)$ can be inferred from assets.

[45] In this chapter, bootstrapping is used in the reinforcement-learning sense: using the model's own next-step prediction as a training target, rather than waiting for the final outcome. This differs from the statistical bootstrap, which involves resampling with replacement.

Definition 4.21. Let $P^{\mathrm{raw}}(t) \in [0,1]^N, \quad t < T$ denote the vector of raw till-maturity default probabilities predicted from the pre-clearing inputs at t.

This is only an ingredient for deriving the haircuts; we enforce weak self-consistency of haircuts by defining a clearing map on survivorship indicators similar to the terminal one. Then, feeding in the pre-clearing survivorship indicator as the first iterate we run a clearing algorithm at the same information time t until no further defaults arise.

Definition 4.22. Fix $t < T$. Let $S^{-}(t) \in \{0,1\}^N$ denote the pre-clearing survivorship indicator and $P^{\mathrm{raw}}(t) \in [0,1]^N$ the raw (to-maturity) default probability estimates. Define the interim clearing map $\Phi_t : \{0,1\}^N \to \{0,1\}^N$ by

$$(\Phi_t(S))_i \;=\; S_i^{-}(t)\,\mathbb{1}\left\{ b_i \le A_i(t) + \sum_{j=1}^{N} \big(1 - P_j^{\mathrm{raw}}(t)\big)\, S_j^{-}(t)\, S_j\, L_{ji} \right\}. \tag{4.7}$$

Let $\mathcal{L}_t = \{S \in \{0,1\}^N : \ S \le S^{-}(t)\}$ (componentwise order). The post-clearing survivorship indicator $S^{+}(t)$ is defined as the greatest fixed point of Φ_t on $\mathcal{L}_t$, i.e.

$$S^{+}(t) \;=\; \max\{S \in \mathcal{L}_t : \ \Phi_t(S) = S\}.$$

Initialize $S^{(0)} = S^{-}(t)$ and iterate

$$S^{(k+1)} = \Phi_t\big(S^{(k)}\big)\,, \qquad k = 0,1,2,\ldots$$

Since all coefficients in the sum are nonnegative, Φ_t is monotone on $\mathcal{L}_t$. Moreover, $S^{(1)} = \Phi_t(S^{(0)}) \le S^{(0)}$, hence by monotonicity $S^{(k+1)} \le S^{(k)}$ for all k, so the sequence is componentwise nonincreasing and cannot resurrect defaulted agents. Consequently it stabilizes after at most $\sum_{i=1}^{N} S_i^{-}(t) \le N$ iterations, and its limit equals the greatest fixed point $S^{+}(t)$.

Definition 4.23. The *post-clearing probabilities* are then given as

$$P^{\mathrm{post}}(t) \;=\; \big(1 - S^{+}(t)\big) \;+\; S^{+}(t)\,P^{\mathrm{raw}}(t), \tag{4.8}$$

i.e., agents that fail at t are set to probability one immediately; for survivors the raw estimate is retained.

Importantly, we do not re-evaluate the network inside the clearing loop: we only fetch $P^{\mathrm{post}}(t)$ once, all valuations in Equation 4.7 rely only on the very same estimates. Clearing is solved to its own fixed point conditional on those default probability estimates, and the result is used to build the learning target and to propagate the system forward. Weak self-consistency is guaranteed without having to assume monotonicity in the learning model. Current predictions are taken seriously for valuation, and valuation is solved to consistency at t, but the predictor is not iterated in the inner loop. In the following we describe learning in detail.

4.6.1 The quantity to learn and its one-step recursion

We work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with discrete filtration $\{\mathcal{F}(t)\}_{t=0}^{T}$. At each $t < T$ we first run the time-t valuation and clearing above using $\mathcal{F}(t)$-measurable inputs (assets $A(t)$, survival indicators $S^-(t)$, the non-random parameters of the system, and the haircuts computed from the predictor at t). This produces the post-clearing survivorship vector $S^+(t) \in \{0,1\}^N$. Because the clearing is a deterministic, monotone map of $\mathcal{F}(t)$-measurable inputs, $S^+(t)$ itself is $\mathcal{F}(t)$-measurable.

Definition 4.24. For each agent i and $t < T$, the *to-maturity default event* is

$$Y_i(t) \;=\; \mathbb{1}\{S_i(T) = 0\},$$

viewed as a random variable on $(\Omega, \mathcal{F}, \mathbb{P})$.

Since $T > t$, $Y_i(t)$ is not $\mathcal{F}(t)$-measurable. It is crucial that $Y_i(t)$ is not a one-step default event; it is the indicator of default by T. On the other hand, one-step events will appear as follows, only as a device to decompose $Y_i(t)$.

Definition 4.25. Let $P_i(t)$ denote the Bernoulli parameter (conditional default probability) at time t.

$$P_i(t) \;=\; \mathbb{P}\big[S_i(T) = 0 \,\big|\, \mathcal{F}(t)\big] \;=\; \mathbb{E}\big[Y_i(t) \mid \mathcal{F}(t)\big] \in [0,1]$$

Ultimately, that is the conditional default probability we would like our predictor to learn, and use it as a strongly self-consistent valuation's haircut.

Let $\tau_i = \inf\{s \in \{0, \dots, T\} : S_i^+(s) = 0\}$ be the default time (absorbing), slightly modifying its previous definition to explicitly account for pre- and post clearing defaults. Default at t is the $\mathcal{F}(t)$-measurable event $\{\tau_i = t\} = \{S_i^+(t) = 0\}$. Default in the next step is the event $\{\tau_i = t{+}1\} = \{S_i^+(t) = 1,\ S_i^+(t{+}1) = 0\}$, where the second indicator is determined by the time-$(t{+}1)$ clearing. Thus, one-step default events will involve the next clearing; in contrast, the decomposition we use below separates default at t (known after the time-t clearing) from survive t and then face the to-maturity probability at $t{+}1$ —a continuation quantity that already includes the next step's clearing because it is evaluated at the next information time.

Proposition 4.26. *$S_i^+(t)$ is $\mathcal{F}(t)$-measurable; hence*

$$\mathbb{P}\big[S_i^+(t) = 0 \mid \mathcal{F}(t)\big] = 1 - S_i^+(t), \qquad \mathbb{E}\big[S_i^+(t)\, Z \mid \mathcal{F}(t)\big] = S_i^+(t)\, \mathbb{E}[Z \mid \mathcal{F}(t)]$$

for any integrable Z.

Proof. Clearing at t is a deterministic function of $\mathcal{F}(t)$-measurable inputs; thus $S^+(t)$ is $\mathcal{F}(t)$-measurable. The identities follow from properties of conditional expectation of $\mathcal{F}(t)$-measurable random variables. $\square$

Next we introduce the law-of-total-probability recursion that will help us define targets for learning. Decompose the to-maturity default event by whether default happens already at t or later:

$$\{S_i(T) = 0\} = \{\tau_i = t\} \,\cup\, \big(\{\tau_i > t\} \cap \{S_i(T) = 0\}\big), \quad \text{disjoint.}$$

Taking conditional expectation given $\mathcal{F}(t)$,

$$P_i(t) \;=\; \underbrace{\mathbb{P}\big[S_i^+(t) = 0 \mid \mathcal{F}(t)\big]}_{\text{default at } t, \text{ known after clearing}} \;+\; \underbrace{\mathbb{E}\big[S_i^+(t)\, P_i(t{+}1) \mid \mathcal{F}(t)\big]}_{\text{survive } t \text{ then continue at } t+1}. \tag{4.9}$$

By Proposition 4.26 this simplifies to the degenerate form

$$P_i(t) \;=\; \big(1 - S_i^+(t)\big) \;+\; S_i^+(t)\,\mathbb{E}\big[P_i(t{+}1) \mid \mathcal{F}(t)\big]. \tag{4.10}$$

Equations (4.9)–(4.10) are the discrete-time survival analogue of risk-set conditioning: $S_i^+(t)$ is the indicator that i is in the risk set after the time-t clearing, so only those paths contribute to continuation.

Switching to estimates, our network at t outputs a raw to-maturity probability vector $P^{\text{raw}}(t)$ pre-clearing. We then apply the time-t clearing to obtain the post-clearing prediction $P^{\text{post}}(t)$ and the post-clearing indicators $S^+(t)$ (defaults at t are set to probability 1). For the continuation we evaluate the frozen (already fitted) next-step predictor at $t{+}1$ on pre-clearing inputs, and we again apply the time-($t{+}1$) clearing to obtain $P^{\text{post}}(t{+}1)$.

Definition 4.27. The *one-step TD target* is

$$\widetilde{Y}_i(t) \;=\; \big(1 - S_i^+(t)\big) \;+\; S_i^+(t)\, P_i^{\text{post}}(t{+}1). \tag{4.11}$$

In (4.11), the first term is the time-t default determined by the time-t clearing; the second term is the to-maturity continuation *after* the time-($t{+}1$) clearing. The next proposition states the conditional unbiasedness of the TD target.

Proposition 4.28. *If* $P^{\text{post}}(t{+}1) = P(t{+}1)$ *almost surely, then*

$$\mathbb{E}\big[\widetilde{Y}_i(t) \mid \mathcal{F}(t)\big] = P_i(t).$$

Proof. By Equation 4.11 and Proposition 4.26,

$$\mathbb{E}\big[\widetilde{Y}_i(t) \mid \mathcal{F}(t)\big] = \mathbb{E}\big[1 - S_i^+(t) \mid \mathcal{F}(t)\big] + \mathbb{E}\big[S_i^+(t)\, P_i(t{+}1) \mid \mathcal{F}(t)\big] = P_i(t)$$

by Equation 4.9. This is the TD identity with discount factor of 1 and absorbing default; see (Sutton and Barto, 2018; Bertsekas and Tsitsiklis, 1995). □

4.6.2 Loss function

We now specify the objective used to train the neural network predictors.

Definition 4.29. For fixed (i, t) the Bernoulli observation $Y_i(t) \in \{0, 1\}$ has parameter $P_i(t)$. The binary cross-entropy (negative log-likelihood) for a single Bernoulli trial with prediction $q \in (0, 1)$ is

$$\ell\big(Y_i(t), q\big) \;=\; -\,Y_i(t)\,\log q \;-\; \big(1 - Y_i(t)\big)\,\log\big(1 - q\big). \tag{4.12}$$

The logarithmic score $\ell(Y, q)$ is a strictly proper scoring rule[46] for Bernoulli events: for any $\mathcal{F}(t)$-measurable $P_i(t) \in [0, 1]$, the conditional expectation $\mathbb{E}[\ell(Y_i(t), q) \mid \mathcal{F}(t)]$ is uniquely minimized at $q = P_i(t)$ (Gneiting and Raftery, 2007; Dawid and Musio, 2014).

Properness is satisfied pointwise and therefore remains valid for heterogeneous populations: if different agents and scenarios have different true probabilities $P_{i,k}(t)$, the sum $\sum_i \mathbb{E}[\ell(Y_{i,k}(t), q_{i,k}) \mid \mathcal{F}(t)]$ is minimized by choosing $q_{i,k} = P_{i,k}(t)$ for each (i, k) pair separately Gneiting and Raftery (2007). In particular, replacing the intractable indicator $Y_i(t)$ by any conditionally unbiased predictor leaves the minimizer unchanged. Using the TD target of Equation 4.11 in Definition 4.28, the stochastic objective [47]

$$C_t(\theta) \;=\; -\sum_{i=1}^{N}\Big\{\widetilde{Y}_i(t)\,\log P_i(t;\theta) + \big(1 - \widetilde{Y}_i(t)\big)\,\log\big(1 - P_i(t;\theta)\big)\Big\}$$

is an unbiased Monte-Carlo estimator of the exact negative log-likelihood $\sum_i \ell(Y_i(t), P_i(t;\theta))$ and has its conditional expectation minimized at the truth. Freezing the continuation $P(t{+}1)$ during the inner loop at t is the standard stabilizing device in TD methods with function approximation (Sutton and Barto, 2018; Bertsekas and Tsitsiklis, 1995). This justifies our choice of loss function as described above. For simplicity, omitted in notation, but in loss calculation P is to be understood as P^{post}.

Next we take a more algorithmic view of how the models are trained.

4.6.3 Backward training schedule and algorithmic outline

We train backwards in time $t = T{-}1, T{-}2, \ldots, 0$, freezing θ_{t+1} while fitting θ_t. Freezing θ_{t+1} keeps the bootstrap target stationary while fitting θ_t, the

[46] A scoring rule assigns a numerical score to a probabilistic forecast given the observed outcome. It is proper if the expected score is optimized when the forecaster reports the true probability, and strictly proper if this optimum is unique. Binary cross-entropy (the logarithmic score) is a standard strictly proper scoring rule for binary events, ensuring that minimizing the training loss drives the neural network's output toward the true conditional default probability.(Gneiting and Raftery, 2007; Dawid and Musio, 2014)

[47] Here we denote the θ-dependence of $P_i(t;\theta)$ explicitly

standard TD stability trick. Each training step uses pre-clearing inputs and a clearing iteration to enforce weak self-consistency at the date. In the numerical implementation, gradients through this discontinuous clearing step are handled by a straight-through estimator, while the probabilistic TD identity itself relies only on measurability and the absorbing default convention. A typical iteration proceeds as follows.

1. Draw a buffer of size B of pre-clearing assets $A(t)$. Agents already defaulted prior to t are encoded in $A(t)$ by a_{dead}.
2. Evaluate the time-t network on $A(t)$ to obtain raw probabilities $P^{\text{raw}}(t)$, then apply the clearing at t to compute post-clearing probabilities $P^{\text{post}}(t)$ and the post-clearing survivorship $S^{+}(t)$.
3. Shock surviving assets to build the next pre-clearing assets $A(t{+}1)$; carry forward defaulted agents at a_{dead}.
4. Evaluate the frozen next-step network at $t{+}1$ on the pre-clearing $A(t{+}1)$, apply the clearing at $t{+}1$ to obtain $P^{\text{post}}(t{+}1)$, and form the bootstrap target
$$\widetilde{Y}(t) \;=\; \big(1 - S^{+}(t)\big) \;+\; S^{+}(t)\, P^{\text{post}}(t{+}1).$$
5. Update θ_t by minimizing the cross-entropy between $P^{\text{post}}(t)$ and $\widetilde{Y}(t)$, optionally restricting the loss to the risk set $\{S^{+}(t) = 1\}$. With n agents and a batch of B scenarios, we minimize binary cross-entropy over all agent–scenario pairs. If we mask to the not yet defaulted agents (a choice we typically make), weights are $w_{b,i}(t) = S^{+}_{b,i}(t)$:
$$\begin{aligned} \widehat{C}_t(\theta) &= \frac{1}{\sum_{b=1}^{B}\sum_{i=1}^{N} w_{b,i}(t)} \sum_{b=1}^{B}\sum_{i=1}^{N} w_{b,i}(t)\, \ell\Big(\widetilde{Y}_{b,i}(t),\; P_{b,i}(t;\theta)\Big), \\ \ell(y,p) &= -y\log p - (1-y)\log(1-p). \end{aligned} \tag{4.13}$$
6. Periodically refresh $\widetilde{Y}(t)$ by re-simulating shocks and re-evaluating the frozen next-step predictor[48], potentially averaging over independent draws for variance reduction.

A more detailed pseudo–code representation can be found in Appendix C

[48] A typical value in our numerical experiment is periodic refresh after $20 - 40$ updates of θ_t

4.6.4 Notes on implementation

On the choice of a_{dead}. Since the interbank liability matrix is symmetric, each agent in the best case recovers as much on interbank assets as it needs to pay. That means the difference between interbank assets and external liabilities defines a barrier below which default upon clearing is imminent. In particular, a $DD < 0$ means certain default, defining a convenient $a_{\text{dead}} = 0$.

Predictor architecture. The TD scheme only requires a predictor that maps the pre-clearing asset vector $A(t)$ to one raw logit per agent. The generic option is a flattened multilayer perceptron: it consumes node-wise features in a flattened block and outputs per-agent logits (later mapped to probabilities and post-processed by clearing). Concretely, with N agents, the input has size N; the MLP applies three ReLU layers of widths $64 \to 64 \to 32$ followed by a linear head of size N. Using a shallow MLP is consistent with universal-approximation results for feed-forward networks and keeps training stable in our backward, TD-style schedule (Nishijima, 2021). The final layer bias is initialised at typically $\log(0.1/0.9)$ (a 10% prior) so that untrained outputs start in a conservative range. Also see the next point.

In the symmetric experiments for larger N, however, we use the DeepSets option. It applies a shared encoder to each scalar asset value, forms for each agent a leave-one-out average of the other agents' embeddings, and applies a shared decoder to the pair consisting of the agent's own asset and this counterparty summary. The result is again one logit per agent, so the clearing map, TD target, loss function, zero-warmstart and terminal clearing module are unchanged. Similarly to the MLP architecture, the final bias is initialized at $\log(0.02/0.98)$, corresponding to a raw PD near 2% before clearing. Appendix D gives the exact parametrization and explains why this restriction is appropriate for Assumption 4.17 but not required by the TD methodology.

Finding the best-case solution via zero-warmstart. On each revaluation, the post-processing clearing map is monotone in counterparties' valuations, so starting from larger valuations selects a more optimistic (lower-PD) solution. To bias the global backward fit toward this outcome, simulation path for initial TD targets can be warm-started [49] with a false network that returns near-zero probabilities for all learnable steps. Therefore, in each step, before the training of the model starts, a fake model emits a large negative logit (≈ -50), so valuation functions $V = 1 - \text{PD}$ are initially near 1 and the first target pass reflects the optimistic payments compatible with assets. Subsequent refreshes

[49] Warm-starting means initializing the iterative procedure from a particular solution rather than from scratch, in order to bias convergence toward a desired fixed point.

and re-training then move from that baseline toward a strongly self-consistent solution under the learned predictor class.

Special last-step clearing. At the final step there is no mark-to-market valuation: we replace the learnable predictor by a deterministic module that only clears and returns the binary survival/default indicators. That step takes the current state (survivorship is inferred from assets via the sentinel rule) and iterates payments up to a fixed point; it returns logits for PD $= 1 - S^+$, i.e., the hard Bernoulli label determined entirely by the terminal clearing, which is then used both during training of earlier steps and at inference. This ensures that the terminal label is exact and that all earlier TD targets are anchored to consistent end-of-horizon outcomes.

Validation loss and early stopping. While the parameters of the models are updated by minimizing the binary cross-entropy against the TD bootstrap targets $\widetilde{Y}(t)$, we monitor convergence and decide when to stop training at each time t using a separate validation mean-squared error (MSE). Concretely, we keep aside a held-out subset $\mathcal{V}$ of the simulation buffer that is never used for gradient steps, and whenever the TD targets are refreshed we recompute validation targets on $\mathcal{V}$ with a larger Monte Carlo sample size $k_{\mathrm{MC}}^{\mathrm{val}} \geq k_{\mathrm{MC}}^{\mathrm{train}}$ to reduce sampling noise; we then evaluate

$$\mathrm{MSE}_t \;=\; \frac{1}{|\mathcal{V}|} \sum_{(b,i)\in\mathcal{V}} \Big(P_{b,i}^{\mathrm{post}}(t) - \widetilde{Y}_{b,i}(t)\Big)^2,$$

retain the parameter iterate achieving the lowest MSE_t, and early-stop once MSE_t fails to improve for a fixed patience window. Using MSE for early stopping is particularly convenient in our setting because the TD targets are probabilities in $(0,1)$ and not binary indicators: although cross-entropy is the principled training objective (proper scoring rule), its attainable minimum is not 0 but the (unknown) Bernoulli entropy of the target surface, $-\widetilde{Y}\log\widetilde{Y} - (1-\widetilde{Y})\log(1-\widetilde{Y})$, so interpreting absolute BCE values would require estimating this moving baseline. In contrast, MSE has a universal optimum of 0 when $P^{\mathrm{post}}(t)$ matches $\widetilde{Y}(t)$, and—combined with the larger validation Monte Carlo sample—provides a lower-variance, out-of-sample signal that prevents overfitting to simulation noise and avoids spending epochs chasing stochastic fluctuations in the refreshed targets. We retain the parameter iterate achieving the lowest validation MSE and stop once improvements fail to materialize over a patience window. In the 10-step discretization used in our experiments, the typical converged validation MSE was around 4×10^{-5} for the last learnable step ($t = T-1$) and decreased gradually for earlier times, reaching around 10^{-5} for the step-0 fit; this magnitude and downward trend were observed consistently across sweeps in the number of agents, correlation ρ, and other parameter choices.

4.6.5 Benchmark for the comonotonic case

Let us assume (as in all of our examples) that all agents start from the same asset buffer a_0 and are linked by a symmetric interbank network. They all have total liabilities b and total interbank assets with face value $\sum_{i=1}^{N} L_{ij}$. Let us further assume that the common external asset evolves as a driftless arithmetic Brownian motion sampled on the grid with step Δt and volatility parameter σ,

$$A(t+1) \;=\; A(t) + \sigma\sqrt{\Delta t}\, Z(t), \qquad Z(t) \sim \mathcal{N}(0,1).$$

When the asset correlation is $\rho = 1$, all agents experience the same shock $Z(t)$ each step (comonotonic evolution). With identical initial conditions and a symmetric liability structure, the state of the system is one dimensional: at each date all agents share the same asset level $A(t)$ and either survive together or default together. Hence the multi-agent problem collapses to a scalar recursion in A.

Definition 4.30. Let $p(A,\tau)$ denote the post-clearing probability of default by final maturity when the common asset level is A and τ time remains.

This is exactly the strongly self-consistent valuation methodology's haircut (Definition 4.16) in a slightly different notation. We introduce an auxiliary, naive default probability to help calculate p.

Definition 4.31. We define the raw next-step default probability

$$r(A,\tau) \;=\; \mathbb{E}\left[p\Big(A + \sigma\sqrt{\Delta t}\, Z,\; \tau - \Delta t\Big)\right], \qquad Z \sim \mathcal{N}(0,1), \tag{4.14}$$

which is the conditional expectation of the next slice $p(\cdot, \tau - \Delta t)$ under the common Gaussian shock.

Raw next-step default probabilities are not yet weakly self-consistent, we need to run them through the clearing first. On the other hand, they connect this step's default probabilities to the next step's default probabilities. Interim clearing at time t is performed using the usual valuation function. Under symmetry, and using the new notation, the survival test for each agent reduces to the scalar inequality using the raw next-step

$$\underbrace{A}_{\text{own assets}} \;+\; \underbrace{(1-r)\,k\lambda}_{\text{interbank assets times valuation function}} \;\geq\; \underbrace{b}_{\text{liabilities}}\;.$$

If this inequality fails, the system is in default now; if it holds, the system continues with default probability r at the next step. Therefore the post–clearing default probability satisfies

$$p(A,\tau) \;=\; \{A+(1-r)k\lambda<b\}+\{A+(1-r)k\lambda\geq b\}\,r(A,\tau),$$

with terminal condition $p(A,0)=\{A+k\lambda<b\}$ (final deterministic settlement).

The expectation in (4.14) is one-dimensional and admits an efficient quadrature. Let us write

$$\mathbb{E}[f(Z)] \;=\; \int_{-\infty}^{\infty} f(z)\,\frac{1}{\sqrt{2\pi}}e^{-z^2/2}\,dz.$$

With the substitution $z=\sqrt{2}\,x$ (so $dz=\sqrt{2}\,dx$), this becomes

$$\mathbb{E}[f(Z)] \;=\; \frac{1}{\sqrt{\pi}}\int_{-\infty}^{\infty} e^{-x^2}\,f\big(\sqrt{2}\,x\big)\,dx. \tag{4.15}$$

The Gauss–Hermite (GH) quadrature is the m-point Gaussian rule for integrals of the form $\int_{-\infty}^{\infty} e^{-x^2}g(x)\,dx$, with nodes x_k equal to the zeros of a Hermite polynomial and weights $w_k>0$ chosen so that the rule is exact for all g that are polynomials up to degree $2m-1$ (Davis and Rabinowitz, 2007). Applying GH to (4.15) gives the practical formula

$$\mathbb{E}[f(Z)] \;\approx\; \frac{1}{\sqrt{\pi}}\sum_{k=1}^{m} w_k\,f\big(\sqrt{2}\,x_k\big), \tag{4.16}$$

and, for our step (4.14),

$$r(A,\tau) \;\approx\; \frac{1}{\sqrt{\pi}}\sum_{k=1}^{m} w_k\,p\Big(A+\sigma\sqrt{2\,\Delta t}\,x_k,\ \tau-\Delta t\Big). \tag{4.17}$$

In the examples presented below, we do not plot the results with this comonotonic benchmark model, but we used the results to verify the TD model's results at $\rho=1$.

4.6.6 Training buffer via a State–dependent static–barrier benchmark

The TD learner at any time t needs samples that are jointly representative of the evolution of assets $A(t)$, the next–step increments $\Delta A(t)=A(t{+}1)-A(t)$ that advance the system to $t{+}1$, and the survival indicators $S(t)$.

We generate multi–step asset paths with additive Brownian increments and equicorrelation via a single common factor (same Gaussian copula formulation as in the Chapter 3):

$$\Delta A_i(t) \;=\; A_i(t+1)-A_i(t) \;=\; \sigma\sqrt{\Delta t}\,\big(\sqrt{\rho}\,M(t)+\sqrt{1-\rho}\,\varepsilon_i(t)\big), \tag{4.18}$$

where $\{M(t)\}_{t=0}^{T-1}$ is a common market shock and $\{\varepsilon_i(t)\}$ is idiosyncratic noise for agent i, all i.i.d. $\mathcal{N}(0,1)$. For each scenario we initialize an asset vector $A(0)$

and simulate the system by 4.18. The advantage of this construction is that the resulting asset values are fully consistent with the later system dynamics used for training.

The more challenging task is to account for the interim defaults due to valuation of the interbank claims. Trivially, there is a default barrier determined by asset value. If any agent's asset value is below $b_i - \sum_{j=1}^{N} L_{ji}$, default at the revaluation is imminent. that is, because even in the best case scenario, assuming $V_j(t) = 1 \quad \forall j$, equity cannot be greater than zero. We guarantee no resurrection by freezing such asset values. Whenever, for any agent, $A_i(t) \leq 0$ we set $\Delta A_i(s) = 0$ for all $s \geq t$ (so the asset value is frozen after crossing this hard lower bound). Assuming default only happens at such low asset values, however, underestimates the default probability and thus the frequency of already defaulted agents the samples realistically should have. For that purpose we introduce a state-dependent static–barrier proxy to approximate the effect of periodic revaluations inducing new, interim defaults.

The methodology keeps the same network revaluation / clearing dynamics as in training, but replaces the neural network predictor of probability of default by maturity with the classical static barrier crossing probability for an arithmetic Brownian motion.

Fix a time t and remaining calendar time $\tau_{\text{rem}} = (T - t)\Delta t$.

Definition 4.32. For a vector of forward-looking till-maturity survival probabilities $p = (p_1, \ldots, p_N) \in [0, 1]^N$, the effective default barrier for agent i is

$$H_i(p) \;=\; b_i - \sum_{j=1}^{N} p_j \, L_{ji},$$

where as earlier, b_i is the total liability level for i.

Assuming the barrier stays constant over the remaining calendar-time interval of length τ_{rem}, the ABM survival probability of agent i above H_i admits the reflection–principle formula

$$\mathbb{P}\left[\min_{0 \leq s \leq \tau_{\text{rem}}} (A_i(t) + \sigma W(s)) > H_i\right] \;=\; 2\,\Phi\left(\frac{A_i(t) - H_i}{\sigma\sqrt{\tau_{\text{rem}}}}\right) - 1,$$

(Karatzas and Shreve, 2014).

To obtain a weakly self-consistent prediction, we iterate the Picard map on survival probabilities, starting from "alive pays in full":

$$p_i^{(0)} = S_i^-(t),$$

relying on the notations introduced earlier and for $m = 0, 1, \ldots$ define

$$
\begin{aligned}
H_i^{(m)} &= b_i - \sum_{j=1}^{N} p_j^{(m)} L_{ji}, \\
\phi_i^{(m+1)} &= S_i^-(t) \max\left(0, 2\,\Phi\left(\frac{A_i(t) - H_i^{(m)}}{\sigma\sqrt{\tau_{\text{rem}}}}\right) - 1\right), \\
p_i^{(m+1)} &= \begin{cases} \phi_i^{(m+1)} & \text{if } A_i(t) + \sum_{j=1}^{N} \phi_j^{(m+1)} L_{ji} - b_i \geq 0, \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}
$$

We run this until convergence, i.e. until $\|p^{(m+1)} - p^{(m)}\|_\infty < \varepsilon$.

Definition 4.33. The *state-dependent static barrier crossing-based valuation function* (SDB) uses these converged survival probabilities as valuations:

$$
V_i^{sdb}(t) = p_i(t) \qquad p_i(t) = p_i^{(m_\star)} \quad \forall\, i \in I.
$$

Note that the built-in clearing iteration guarantees weakly self-consistent predictions by construction. On the other hand, nothing guarantees strong self-consistency; the forward-looking predictions used a naive, barrier–crossing based prediction and not the actual system dynamics.

For generating samples, we run the same discrete–time revaluation / clearing dynamics as in training, but at each grid time t we set the "predicted" till–maturity survival vector to $p(t)$ obtained by the converged Picard iteration above. Concretely, for each scenario we evolve the external asset process with the same stepwise ABM shocks as in training, and at each time t:

1. compute $p(t)$ from the current state $(A(t), S^-(t))$
2. apply the corresponding clearing to obtain the post–clearing survival indicator $S^+(t)$ (defaults are absorbing)
3. propagate survivors to $t+1$ by adding the next external shock increment.

We can use this forward simulation to generate default paths and stepwise state buffers for training. Fix M scenarios. For each scenario $m \in \{1, \ldots, M\}$ we obtain a discrete default path

$$
\big(A^{(m)}(t),\, S^{+,(m)}(t)\big)_{t=0,\ldots,T},
$$

where $S^{+,(m)}(t)$ is the post–clearing survival indicator produced by the SDB rule at time t. The time-t state buffer is then formed by stacking these post–clearing states across scenarios,

$$
\mathcal{B}(t)^{\text{sdb}} \;=\; \left\{\big(A^{(m)}(t),\, S^{+,(m)}(t)\big) :\; m = 1, \ldots, M\right\}, \tag{4.19}
$$

and used as the initial state distribution for any downstream procedure that requires samples at time t (e.g. benchmarking alternative valuation rules or fitting a predictor on the induced state distribution). Because $S^{+,(m)}(t)$ is computed after clearing, the buffer explicitly encodes the endogenous default cascade up to time t, and the resulting collection $\{\mathcal{B}(t)^{\mathrm{sdb}}\}_{t=0}^{T}$ provides a consistent set of "initial configurations" across time steps aligned with the same network revaluation dynamics, differing only in the substitution of the neural predictor by the static–barrier proxy.

We can also use the valuation rule in Definition 4.33 to sanity–check the learned network–consistent default probabilities by using it as a benchmark. Since the valuation rule in Definition 4.33 is in general not strongly self–consistent, the quantity $1 - p_i(t)$ should be interpreted as a proxy for the model–implied probability of default by maturity rather than the true conditional default probability under the induced dynamics. For benchmarking, we evaluate the methodology by simulating the system using the proxy valuation. At maturity, the default event for agent i is $\{S_i^+(T) = 0\}$, and we define the SDB model–implied default probability by

$$PD_i^{\mathrm{sdb}} \;=\; \mathbb{P}^{\mathrm{sdb}}\big[S_i^+(T) = 0\big]\,. \tag{4.20}$$

estimated in practice by Monte Carlo frequencies across scenarios. This "simulate–then–count" benchmark makes the comparison to ML outputs unambiguous: both are assessed by the realized default frequency under their respective induced dynamics.

4.7 Analysis of the model

Next we look at some results primarily to characterize the realized default probabilities using the strongly self-consistent state-dependent valuation functions of the TD-learner (Section 4.6). Furthermore, we also investigate, what differences would emerge if one were to use the state-independent valuation functions of Section 4.4 or the state-dependent but not self-consistent, static-barrier-based valuation function of Definition 4.33 instead. Throughout this discussion we maintain Assumption 4.17. Looking at the default probabilities, the most striking difference compared to the results in Chapter 3 is that the minimal default probability no longer necessarily occurs at $\rho = 1$. So before looking at the results, we investigate why that is the case.

All numerical statements in this section are conditional on the stylized environment used for the experiment: homogeneous fully connected networks, ABM external-asset dynamics, zero recovery, deterministic fixed nominal liabilities, and no collateral, margining, liquidity hoarding, fire sales, close-out netting or CCP mechanics.

4.7.1 When does $\rho = 1$ lead to minimal default probability?

The results of Chapters 3 and 4 can be organized using two unified mechanisms. Chapter 3 contains only the first one, because obligations are settled only once and there are no interim forward-looking haircuts. The present chapter contains the same first mechanism at each clearing date and, in addition, a second mechanism generated by interim state-dependent revaluations.

Definition 4.34. *Heterogeneity in asset values at clearing (HAVC)* This mechanism arises from cross-sectional dispersion in realized external asset values at a date when solvency is determined by a clearing step. It is the net effect of two forces that work in opposite directions. Starting from low correlation, dispersion across counterparties provides diversification, so default probabilities are relatively low. As ρ rises, this diversification benefit weakens and default probabilities tend to increase. But the same dispersion is also what makes contagion defaults possible: an agent can be pushed into default only in states in which it would survive under full payment, while weaker counterparties fail to pay enough. When ρ becomes very high, such states become rare. Under the homogeneous symmetric setup, they disappear at $\rho = 1$, because agents become fundamentally solvent or insolvent together. Hence, when this is the only mechanism, default probability is minimized at $\rho = 1$ and a hump shape emerges from the trade-off between lost diversification and disappearing contagion states.

This is therefore not a pure correlation result. It also relies on the symmetry of the setup. If agents are ex-ante heterogeneous—for example because drifts, means or distances to default differ across banks—then even under comonotonic shocks some banks may remain above the default boundary while others fall below it. In that case, $\rho = 1$ need not eliminate contagion and need not minimize default probability.

Definition 4.35. *Synchronous valuation–spiral (SVS)* This mechanism arises only when interbank claims are repriced before maturity using state-dependent haircuts. A receivable is marked down by the counterparty's conditional default probability, which is a continuous increasing function of how weak the counterparty's current balance-sheet position is. Hence claims lose value already at interim dates, before terminal default is realized. Importantly, this channel does not require cross-sectional heterogeneity in external assets to remain active. Even if external assets are perfectly correlated and the system starts from a homogeneous state, a bad common path lowers all counterparties' asset values together, so every receivable is marked down simultaneously. Higher ρ reduces diversification in these markdowns, making the effect strongest as

$\rho \to 1$. This mechanism is absent in Chapter 3 and in the state-independent benchmark of Section 4.4; it appears only in the multi-period, state-dependent valuation model.

Taken together, the two mechanisms explain both the hump shape and the location of the minimum default probability. HAVC by itself pushes the minimum toward $\rho = 1$. SVS pushes in the opposite direction, because synchronous markdowns are most severe at high correlation. Which force dominates near $\rho = 1$ determines whether the safest correlation remains $\rho = 1$ or shifts to a smaller value.

Multi-period model

State-independent haircut. With Definition 4.18 the haircut can be expressed as

$$V_j^{si}(t) = 1 - \mathbb{P}\big[S_j(T) = 0 | S_j(t) = 1\big].$$

However, by symmetry, the haircuts are deterministic functions of t and identical across agents. The equity simplifies to

$$E_i(t) = A_i(t) - (1 - V_i^{si}(t))k\lambda \quad \forall\, t \in \mathcal{T},\ t < T. \tag{4.21}$$

Therefore, the survival condition reduces to a common-threshold problem. SVS is absent, because interim markdowns do not react to the realized cross-section of states. The only remaining channel for cascades is HAVC, namely cross-sectional heterogeneity in $A(t)$ at the clearing dates. Driving $\rho \to 1$ removes that heterogeneity date by date and, in particular, at terminal clearing, thereby reproducing the one-period logic over the whole horizon. Hence the smallest default probability is attained at $\rho = 1$.

Strongly self–consistent, state-dependent haircut. With strong self–consistency, the haircuts react to the realized state:

$$1 - V_j(t) = \mathbb{P}\big[S_j(T) = 0 \,\big|\, \mathcal{F}_t\big].$$

For agent i the equity can be written as

$$E_i(t) \;=\; A_i(t) - k\lambda + \lambda \sum_{j \neq i} V_j(t) \;=\; A_i(t) - \lambda \sum_{j \neq i} \big(1 - V_j(t)\big) \tag{4.22}$$

In this case, both mechanisms are active. HAVC still operates whenever a clearing step is applied to heterogeneous realized asset values; in particular, at the terminal date T there are no further ex-ante forward-looking haircuts, so the final clearing is governed only by HAVC, exactly as in Chapter 3 apart from the fact that some counterparties may already have defaulted earlier.

Interim dates, however, also feature SVS: state-dependent haircuts co-move with counterparties' realized asset positions and can mark down all interbank assets simultaneously even under $\rho = 1$. The first mechanism therefore pushes the minimum toward $\rho = 1$, while the second pushes it away from $\rho = 1$. If the second dominates near $\rho = 1$, the smallest default probability occurs at a lower correlation, often close to $\rho = 0$ in our numerical examples.

One-period model

In a one–period model ($\mathcal{T} = \{0, 1\}$), valuation takes place only at time 0, while at time 1 only the final clearing occurs and no valuation is performed. Under Assumption 4.17, even with the proper state–dependent valuation, the time–0 valuation haircut is necessarily the same for all agents. Hence SVS is inactive: there is no interim repricing after the state has become heterogeneous. At the time–1 clearing, the only remaining channel for cascades is HAVC, that is, cross-sectional heterogeneity in the terminal $A(1)$. That channel is eliminated if $\rho = 1$, which means the default probability is minimized at that value.

4.7.2 The effect of various parameters

First, we discuss how changing the number of agents affects the default probability. It is illustrated with Figure 4.2 for a $T = 10$ period setup. Increasing N, ceteris paribus, makes the maximal individual default probability higher. It typically does so by making the hump shape more pronounced and the maximum occurring closer to $\rho = 1$. For small correlation values, larger N means smaller default probabilities. Small correlation allows for a greater degree of diversification as the number of agents increases. Since we are investigating the fully connected network structure, larger N directly translates into more counterparties for a given total exposure, which allows interbank claims to be distributed more widely. Hence agents are less exposed to the weakness of any single partner. As ρ rises from low values, that diversification benefit fades and default probabilities increase. In the multi-period, state-dependent model this rise is reinforced by the synchronous valuation–spiral of Definition 4.35. By contrast, very close to $\rho = 1$ the heterogeneity-in-asset-values mechanism of Definition 4.34 becomes weaker because cross-sectional dispersion vanishes. In the exact $\rho = 1$ limit, the remaining default probability is driven by common shocks and synchronous interim markdowns, and it does not depend on N; this limit can also be derived using the comonotonic benchmark described in Section 4.6.5.

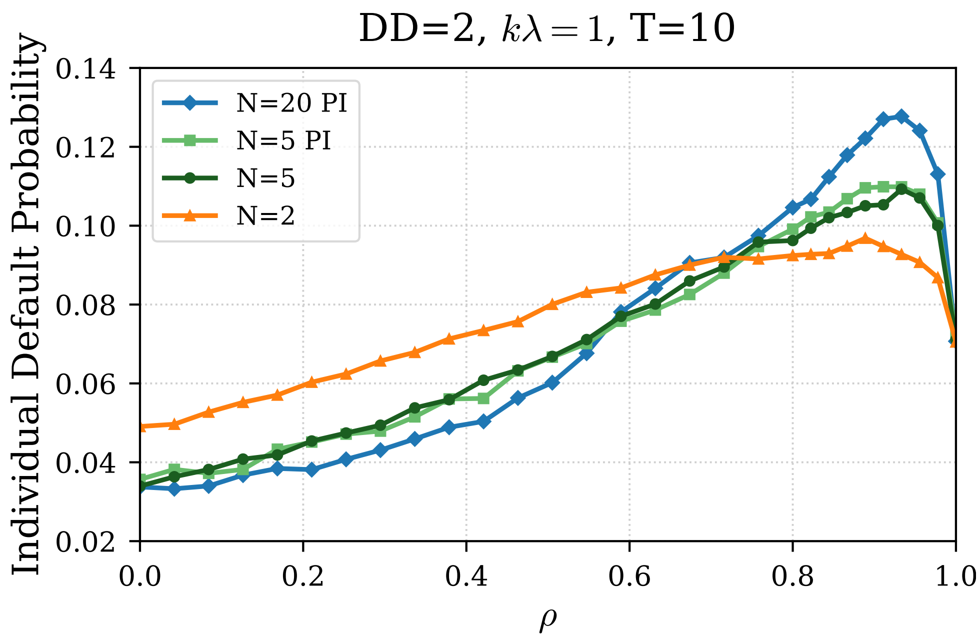

Figure 4.2: Default probabilities for various N values in the 10-period system, while keeping other parameters fixed. PI means the DeepSets architecture were used in the experiment, other results were generated using the MLP architecture. For $N = 5$, equivalence (up to numerical noise) of the two methodologies are also illustrated.

Decreasing DD trivially increases the individual fundamental default probability, also increasing the total default probability. At the same time, smaller DD makes both mechanisms more potent: under HAVC, smaller payment shortfalls are sufficient to push an otherwise surviving bank below the clearing threshold, while under SVS a common markdown of interbank assets more easily depresses equity below zero. Therefore, smaller DD corresponds to a more pronounced hump shape, with a larger difference between maximal and minimal default probabilities across ρ values. This characteristic again resembles how changing DD affects individual default probabilities in the one-period systems of Chapter 3 without mark-to-market interbank claim valuations. Figure 4.3 illustrates this effect through an example.

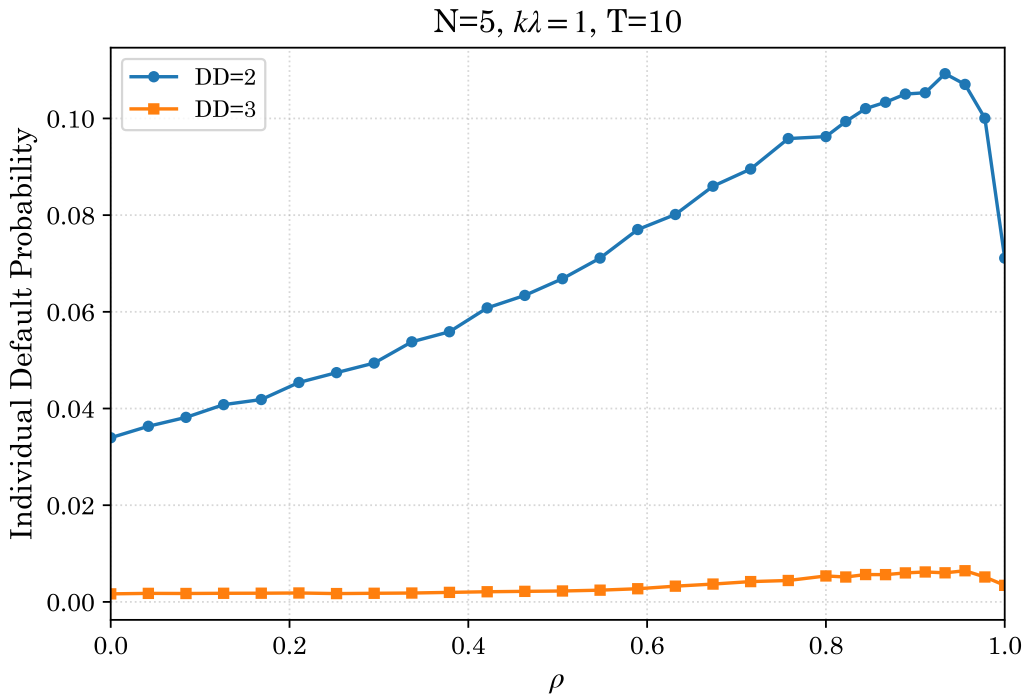

Figure 4.3: Default probabilities for various DD values in the 10-period system, while keeping other parameters fixed.

As shown in Figure 4.4, larger interbank exposure $k\lambda$ means larger default probability, in general. Individual default probabilities show a characteristic humped shape as a function of asset correlation, ρ. The maximal default probability occurs the closer to $\rho = 1$, the smaller the interbank exposures are. Larger $k\lambda$ leads to a more pronounced hump shape, with a larger difference between maximal and minimal default probabilities across ρ values. This is natural in the unified framework: larger interbank exposures make both mechanisms more influential, because they increase the amount of solvency support that depends on counterparties at clearing and, in the state-dependent model, they also increase the balance-sheet impact of synchronous markdowns. These stylized facts—again—resemble many of our conclusions in the one-period systems of Chapter 3 without mark-to-market interbank claim valuations. However, an important quantitative difference is that the mark-to-market framework generally produces higher peak default probabilities than the one-period model for the same parameter values, because the valuation-spiral channel amplifies stress at each interim revaluation date. The gap between the $k\lambda = 0$ (non-connected) and $k\lambda > 0$ curves directly quantifies the contribution of interbank contagion: for larger $k\lambda$, this gap widens substantially, especially near the peak of the hump, indicating that forward-looking contagion is most damaging precisely in the correlation range where default risk is already elevated.

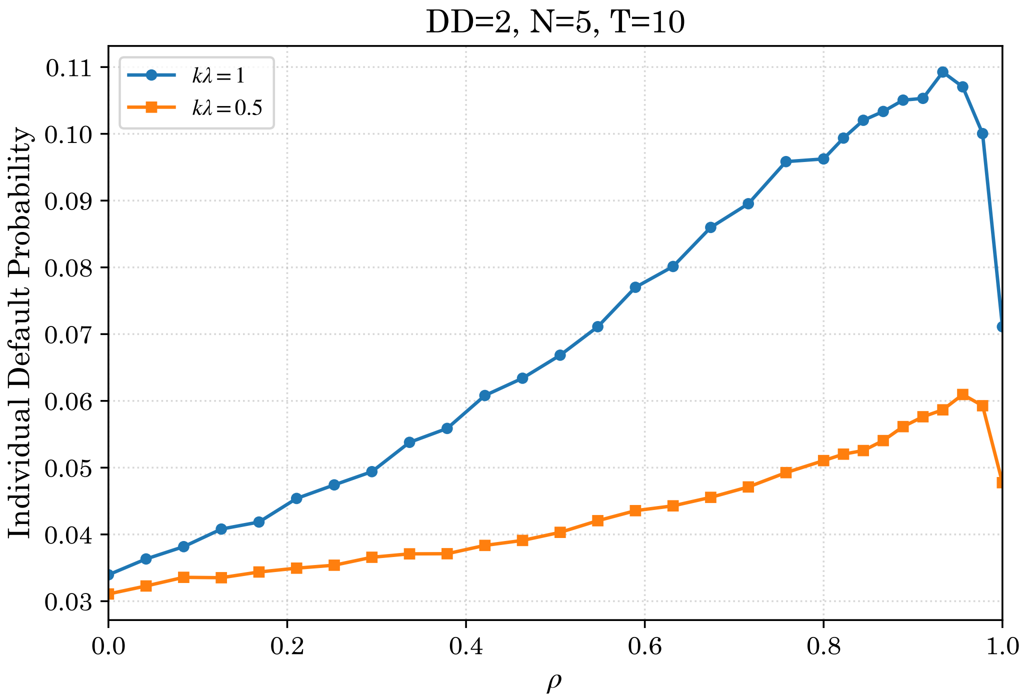

Figure 4.4: Default probabilities for various $k\lambda$ values in the 10-period system, while keeping other parameters fixed.

4.7.3 Comparing state dependent and state independent haircuts

In the following (Figures 4.5 – 4.7) we compare how the model of Section 4.6 with state dependent haircuts relate to the state-independent benchmark model of Section 4.4. For completeness, we also plot the state-independent model's result for the unconnected systems ($k\lambda = 0$). Note that these values should be the same as the state-dependent model's default probabilities with differences solely due to numerical noise.

The characteristic hump shape appears in both the state-independent and the state-dependent results, but the mechanism behind the minima differs. In the state-independent benchmark, only HAVC operates, so the logic of Chapter 3 carries over and the minimal default probability occurs at $\rho = 1$. In the state-dependent model, HAVC is still present—especially at terminal clearing—but interim repricing activates SVS. This additional force is strongest at high correlation, so the minimum no longer needs to be at $\rho = 1$; in several of our parameterizations it moves close to $\rho = 0$. We also do not expect the three lines to meet at $\rho = 1$: even in the comonotonic case, forward-looking haircuts can depress valuations before maturity.

The vertical gap between the state-dependent and state-independent curves at a given ρ therefore quantifies the additional default risk attributable specifically to SVS, that is, endogenous path-dependent revaluation feedback. This gap can be negligible at $\rho = 0$ when DD is sufficiently high, as in Figures 4.6 and 4.7, where cross-sectional asset heterogeneity provides natural insulation against synchronous markdowns. It grows substantially at intermediate-to-high correlations in the multi-period setting. For the parameter combinations shown in the figures, the state-dependent model can produce default probabilities that are several percentage points higher than the state-independent benchmark near the hump's peak. From a supervisory-reading perspective, this is a model-dependent warning about the economic interpretation of valuation contagion: state-averaged or unconditional haircuts can understate forward-looking feedback in this stylized unsecured-exposure setting. It is not a claim that a particular accounting rule or prudential buffer is miscalibrated.

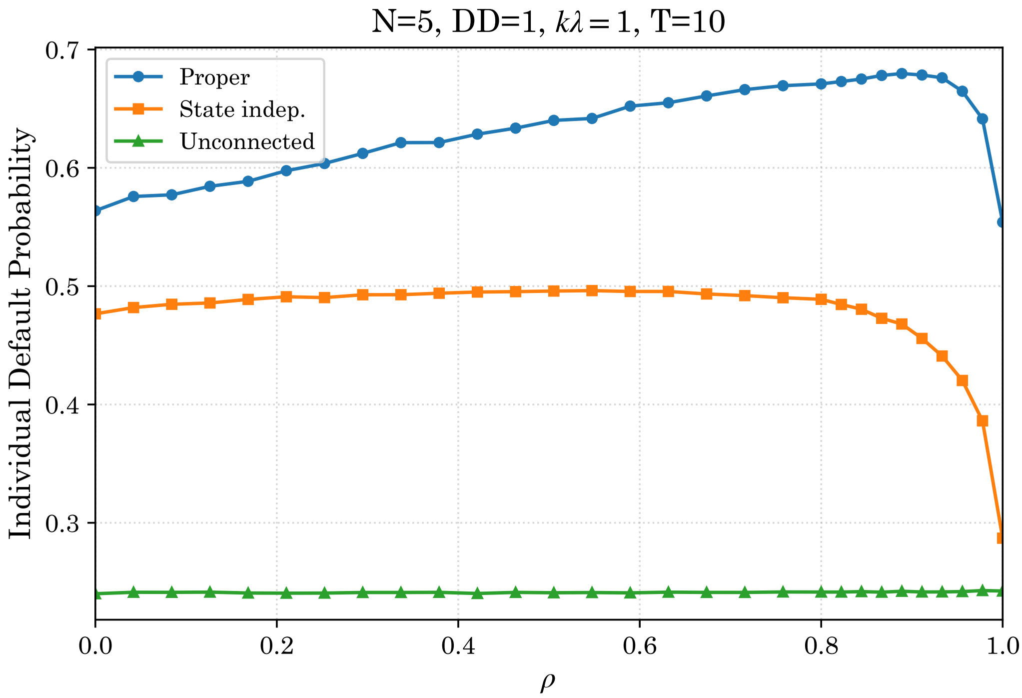

Figure 4.5: Individual default probabilities in a 10-period system. Proper: state-dependent, using the TD-learner.

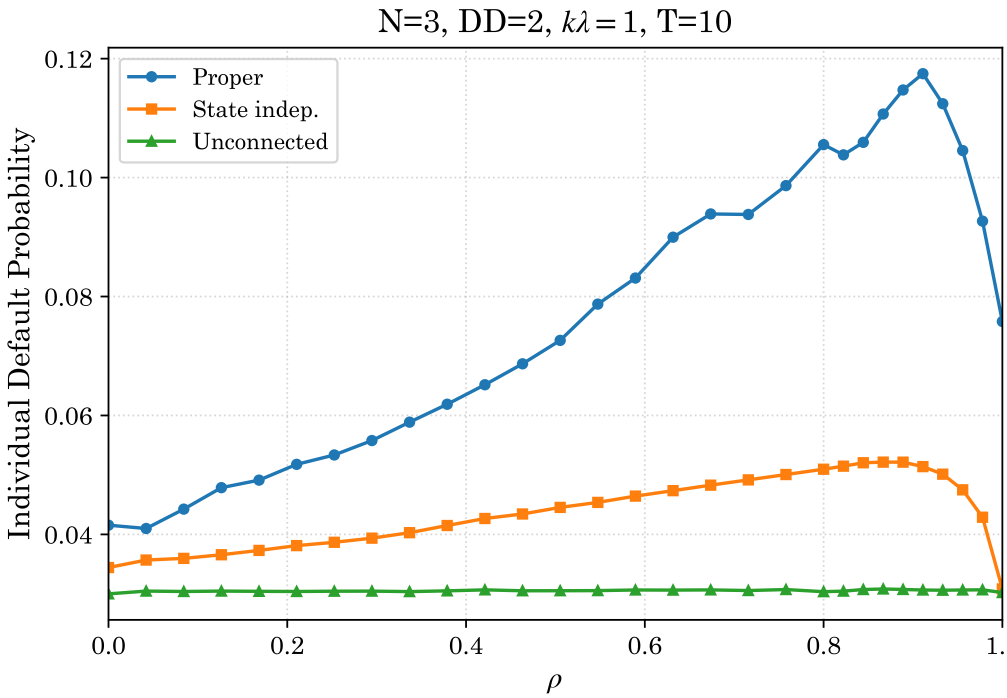

Figure 4.6: Individual default probabilities in a 10-period system. Proper: state-dependent, using the TD-learner.

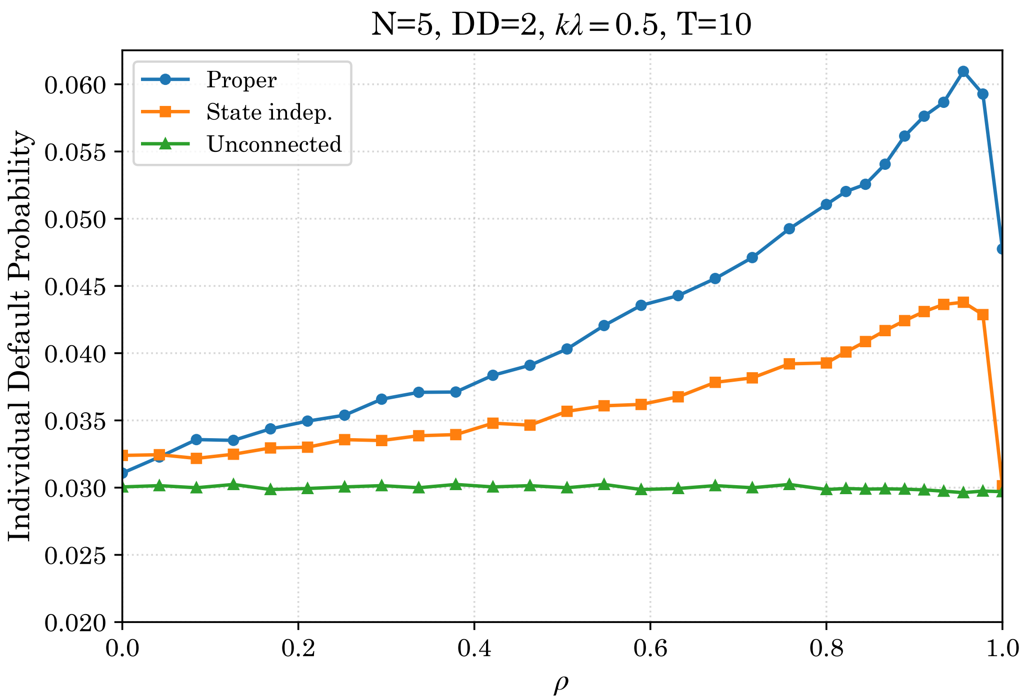

Figure 4.7: Individual default probabilities in a 10-period system. Proper: state-dependent, using the TD-learner.

The one-period system deserves special attention because it shuts down SVS by construction. Revaluation happens only at time 0, where symmetry implies that the valuation haircut is identical across agents; after that there is only the terminal clearing at time 1. Hence state-dependent and state-independent haircuts should lead to the same default probabilities, and Figure 4.8 supports this statement. The differences between state-dependent and state-independent haircuts are only due to numerical noise: as an illustration, we also plot the standard error bounds calculated from the MSE loss of the state-dependent model. Standard error expresses the possible uncertainty in the default probabilities. The only remaining contagion channel is therefore HAVC, namely heterogeneity in asset values at clearing (Definition 4.34). Setting $\rho = 1$ switches that channel off in the homogeneous benchmark, so the default probability is minimized at $\rho = 1$. Note that this default probability is still slightly higher than in the non-connected system, due to the initial interbank valuation at $t = 0$. All agents still default or survive together, but factoring in that probability at the initial valuation brings this joint default closer.

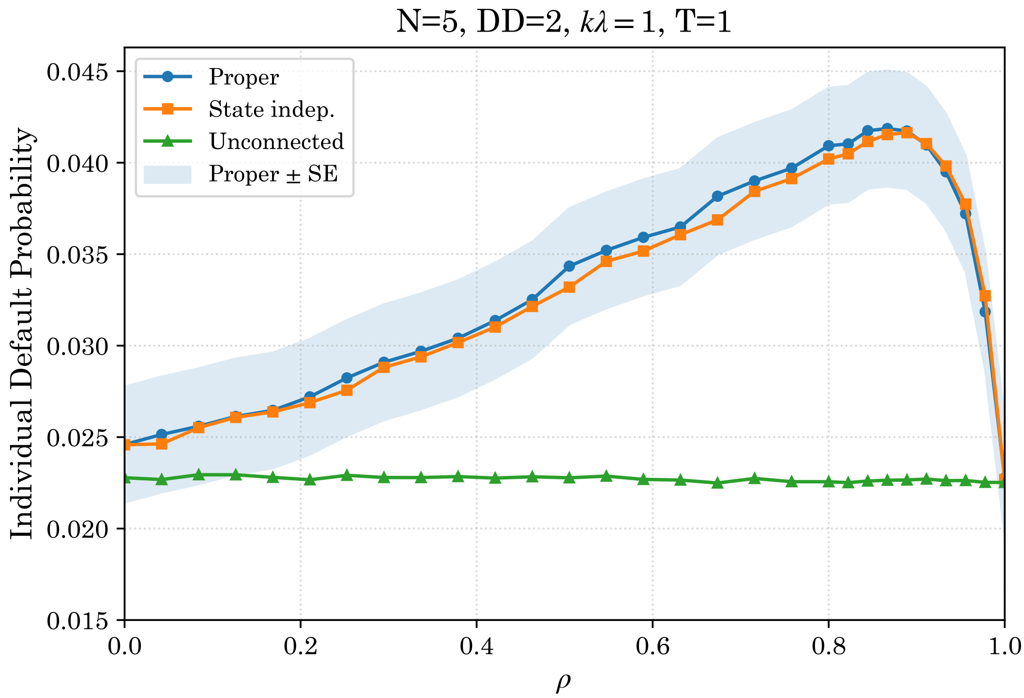

Figure 4.8: Individual default probabilities in a 1-period system.

4.7.4 Comparison to the static–barrier benchmark

Figure 4.9 compares the default probabilities using the state-dependent haircuts to the ones realized if one uses the SDB benchmark of Definition 4.33.

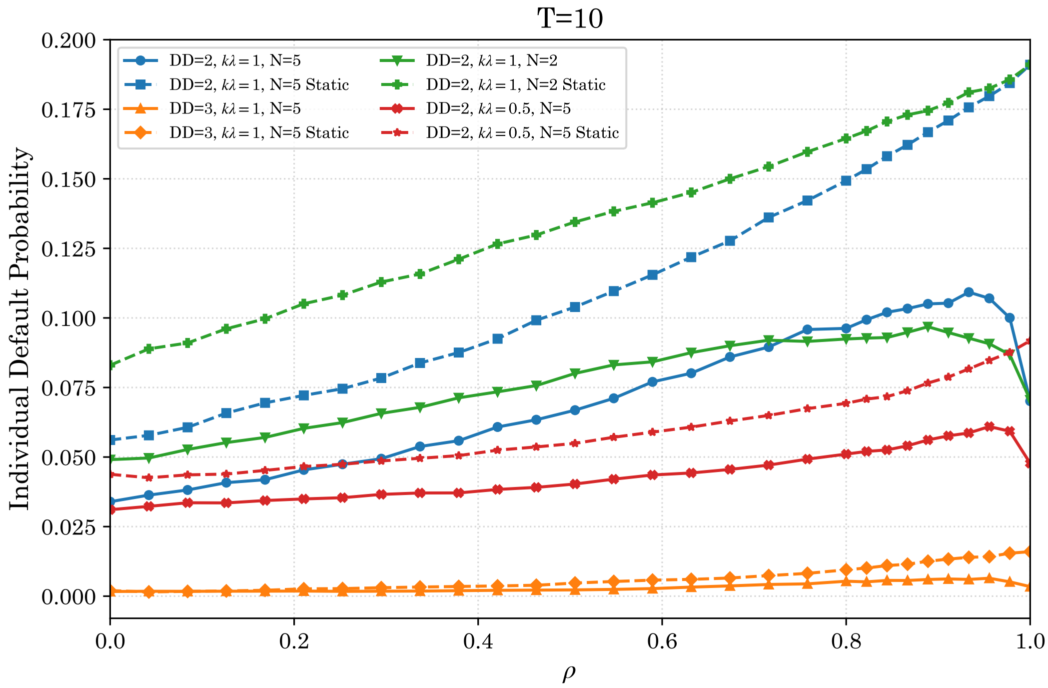

Figure 4.9: Default probabilities for various parameter combinations in the 10-period system, comparing the TD learner and the static-barrier benchmark.

In our experiments the SDB benchmark often produces higher default probabilities than the learned, strongly self–consistent TD model. One reason is that the proxy computes a continuous–time first–passage probability for an ABM against a barrier H_i that is held fixed over $[t, T]$; this counts any intra–period excursion below H_i and therefore tends to overstate default risk relative to the discretely monitored revaluation dynamics. In addition, freezing the barrier suppresses the endogenous evolution of haircuts and incoming payments: conditional on surviving intermediate clearings, the effective barrier can shift downward and valuations can relax in the full model, an effect the proxy cannot condition on. Finally, because the proxy PDs are fed back into the forward simulation that generates $\mathcal{B}(t)^{\text{sdb}}$, overly pessimistic haircuts can become partially self–fulfilling by inducing extra interim defaults, pushing PD^{sdb} upward. Taken together, these effects explain why the SDB benchmark provides a conservative upper bound on default probabilities across most parameter combinations. For a practitioner, this observation has a dual interpretation. On the one hand, the static-barrier proxy is cheap to compute and could serve as a quick, conservative screening tool: if the proxy-implied default probabilities are acceptable, the true self-consistent probabilities are likely even lower. On the other hand, the degree of overstatement varies across parameter regimes—it tends to be larger for small ρ and large DD—so the proxy cannot simply be "deflated" by a constant factor.

4.7.5 Robustness checks

This subsection examines how sensitive the reported self-consistent default probabilities are to modelling and numerical choices that are treated as fixed in the baseline experiments. Two potential concerns are particularly relevant in our setting: (i) the state distribution induced by the sampling / buffer-generation procedure used during training, and (ii) the possibility of multiple self-consistent solutions, in which case the backward TD fit could in principle converge to different equilibria depending on the warm-start used to form the first set of TD targets.

First, one might be interested in checking robustness to the training-buffer generator. All results presented so far use the state-dependent static-barrier buffer generator (SDB) described in Section 4.6.6; see Equation 4.19. Since the neural networks are fitted on samples drawn from these buffers, it is natural to ask whether the learned fixed point is an artifact of this particular data-generating mechanism, especially because interim defaults affect the future distribution of states. To probe this, we repeat the full training procedure using the alternative multi-step rollout described in Appendix B. This alternative generator (ED) simulates correlated external-asset paths and structural barrier defaults, and then superimposes additional interim defaults through an external-default overlay driven by the network stress feature (and correlated across agents). Importantly, ED is not meant to be a valuation proxy; rather, it is a deliberately different but still structured mechanism for generating default paths and survivor sets.

Figure 4.10 shows that the two generators can induce markedly different default frequencies by maturity, and hence materially different empirical training distributions. Nevertheless, once the TD learner is trained to convergence on either buffer, the resulting self-consistent default probabilities are very close. In other words, within this canonical configuration the equilibrium learned by the TD procedure appears stable to substantial perturbations of the sampling scheme, suggesting that the reported results are not driven by a finely tuned choice of buffer generator.

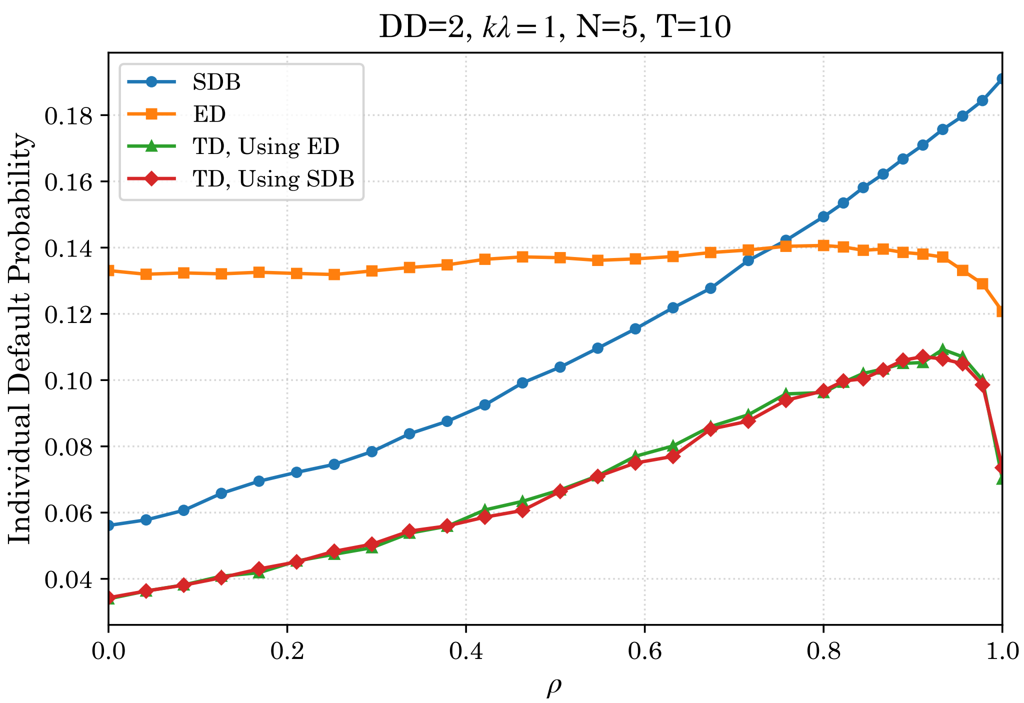

Figure 4.10: Robustness to the buffer generator (canonical parameter set). “SDB” reports the default probabilities obtained from the static-barrier benchmark of Section 4.6.6. “ED” reports the default frequencies produced by the mechanism incorporating the external-default overlay of Appendix B (which is not a valuation rule but a path generator). “TD (trained on ED)” and “TD (trained on SDB)” denote the self-consistent default probabilities produced by the TD learner after training on the corresponding buffers. Although the raw default rates differ substantially across SDB and ED, the learned self-consistent probabilities remain very similar.

A second concern is that the coupled clearing/valuation problem need not admit a unique fixed point. As discussed in the implementation notes (zero-warmstart), our baseline procedure biases the first TD target pass toward the most optimistic solution by warm-starting with a “false” network that outputs near-zero default probabilities (equivalently, near-full valuations). To check whether this choice materially affects the final outcome, we rerun the backward TD training while replacing the zero-warmstart by deliberately pessimistic initial predictors that output substantially larger default probabilities when first queried (roughly around 50% and close to 100%).

The resulting comparison is shown in Figure 4.11. While these warm-starts can change the early learning dynamics, the fitted models converge to essentially the same default-probability profile once training is allowed to run sufficiently long. This does not prove global uniqueness of the equilibrium in all parameter regimes, but it provides evidence that—for the canonical configuration considered here—the reported solution is stable and lies in a large basin of attraction for the learning dynamics. Note that in these experiments we also vary the buffer generating procedure and how frequently we perform the periodic refresh of $\widetilde{Y}(t)$ by re-simulating shocks and re-evaluating the frozen next-step predictor.

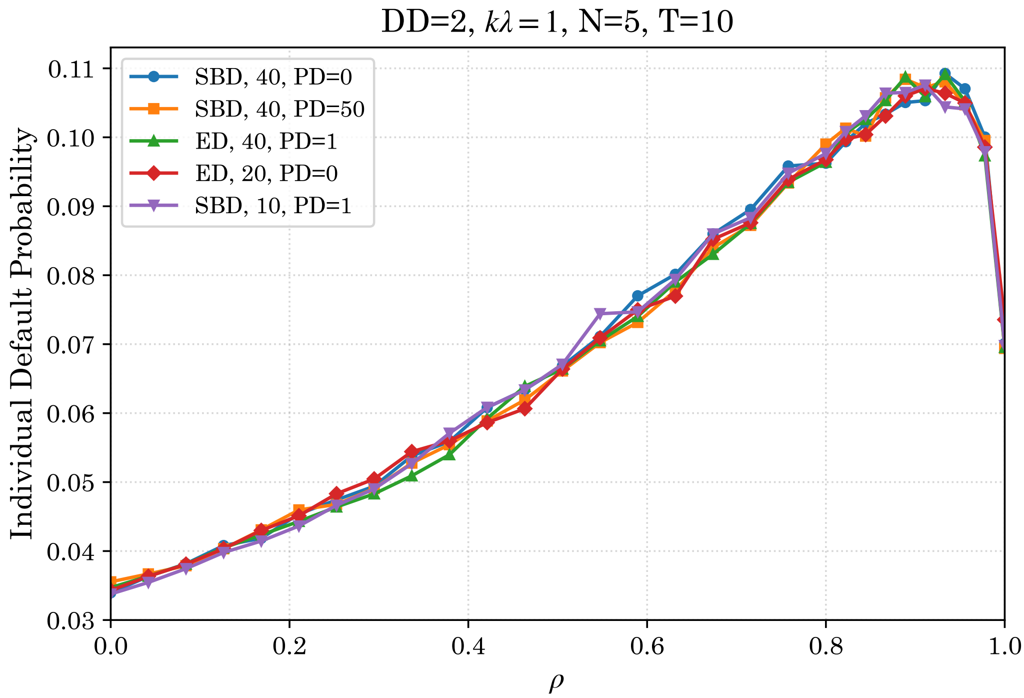

Figure 4.11: Robustness to warm-starting, refresh frequency and training buffer generation method. We compare the baseline zero-warmstart against pessimistic warm-starts in which the initial "teacher" outputs default probabilities around 50% and close to 100% before any optimization. The numbers 10, 20 or 40 refers to the number of optimizer steps before a refresh of $\widetilde{Y}(t)$. After training to convergence, all runs yield almost identical self-consistent default probabilities.

4.8 Conclusion and outlook

This chapter has proposed a discrete-time framework for forward-looking solvency contagion in interbank networks, in which claims are valued under strongly self-consistent, endogenous default probabilities. Starting from a standard Eisenberg-Noe style balance sheet with mark-to-market interbank assets, we introduced equity and a survival indicator process that evolves through repeated interim clearing steps and a final terminal clearing. Within this setting, haircuts are defined as conditional to-maturity default probabilities, ensuring that the valuation of a claim on a counterparty is anchored in the counterparty's actual risk of failure under the same dynamics.

To make this concept computationally tractable in higher dimensions, we developed a temporal-difference (TD) analogue for learning the entire surface of state-dependent default probabilities. A naive Monte Carlo regression approach was shown to be conceptually straightforward but computationally wasteful and numerically unstable, as it treats each state as a separate long-horizon prediction problem and continually chases its own moving Monte Carlo targets. The TD scheme overcomes these limitations by exploiting the law-of-total-probability recursion for to-maturity default, using post-clearing default/survival at time t and post-clearing probabilities at time $t+1$ to build low-variance, locally consistent targets. Crucially, clearing is enforced at each step so that weak self-consistency is guaranteed at the balance-sheet level even before training converges.

We contrasted the fully state-dependent valuation with two benchmarks: a state-independent haircut model that conditions only on survival up to t, and a comonotonic setting in which all external assets move in lockstep. Under symmetric, fully connected networks the results can be understood through the same two-mechanism framework developed in Section 4.7.1. HAVC is heterogeneity in asset values at clearing: by itself it generates the hump shape of Chapter 3 and implies that, in homogeneous systems, perfect correlation minimizes default probability because it removes pure contagion states. SVS is the synchronous valuation–spiral created by interim state-dependent haircuts: bad common paths lower all interbank valuations together even when external assets move in lockstep. The state-independent and one-period settings contain only HAVC, so their minimum remains at comonotonicity. The multi-period strongly self-consistent model contains both mechanisms. Although HAVC is still present at the terminal clearing, SVS can dominate near high correlation and shift the safest value away from $\rho = 1$, in our numerical examples often toward low correlations including $\rho = 0$.

From a methodological perspective, the TD-based learning approach offers a promising route to approximate fully forward-looking network valuations in settings where tree methods or closed-form comonotonic bounds are either infeasible or too restrictive. It accommodates heterogeneous balance sheets, rich liability structures and flexible shock distributions, provided one can simulate asset dynamics and run the clearing map. The DeepSets predictor used in some of our symmetric numerical experiments is therefore best viewed as a parameter-sharing and stabilization device compatible with Assumption 4.17, enabling running experiments for larger N values.

From a prudential interpretation perspective, the chapter shows that forward-looking economic valuation feedback can amplify stress relative to default-only contagion models. This matters for interpreting market-based or economic-valuation contagion under stylized uncollateralized bilateral interbank exposures. The chapter does not, however, derive a supervisory rule for fair-value accounting, nor does it calibrate countercyclical buffers, valuation adjustments or capital surcharges.

Several extensions are left for future work. These include relaxing the symmetry assumptions of Assumption 4.17, introducing multiple seniority classes and recovery rates, incorporating liquidity-driven fire-sale mechanisms alongside solvency contagion, and combining the TD-style surrogate with scenario-based stress testing under alternative macro-financial paths. Regulatory interventions, such as capital surcharges, would require empirical and regulatory calibration before they could be interpreted as policy implications of the model. Dynamic margining would be relevant only in an extended model for collateralized or derivative exposures with explicit margin accounts. More broadly, the approach illustrates how tools from reinforcement learning and function approximation can be integrated with classical clearing models to analyze forward-looking systemic risk in complex financial networks.

Appendix A

Algorithmic outline for the state-independent haircuts model of Section 4.4

Algorithm 1 Fixed-Point Iteration for State-averaged Default Probabilities

Require: total steps T, agents N, scenarios k; initial assets $a_{\text{init}} \in \mathbb{R}^N$; initial survival flags $s_{\text{init}} \in \{0,1\}^n$; interbank liability matrix $L \in \mathbb{R}^{N\times N}$, total liabilities $b \in \mathbb{R}^N$; tolerance `tol` > 0, maximum iterations `max_iter`;
Ensure: $\widehat{P} \in \mathbb{R}^{(T+1)\times n}$ (state-averaged default probabilities)

Initialize $\widehat{P}^{(0)} \leftarrow \mathbf{0}_{(T+1)\times n}$; $A \leftarrow \mathbf{NaN}_{(T+1)\times k\times N}$; $S \leftarrow \mathbf{NaN}_{(T+1)\times k\times N}$; dist $\leftarrow \infty$; iter $\leftarrow 0$. ▷ start from the most optimistic initialization
while dist $>$ `tol` **and** iter $<$ `max_iter` **do**
  iter $\leftarrow$ iter $+ 1$
  (Path simulation with embedded clearing)
  A$[0,:,:] \leftarrow a_{\text{init}}$
  for $t = 0,\ldots,T$ **do** ▷ handle all scenarios in parallel
    $\text{surv} \leftarrow \begin{cases} s_{\text{init}}, & t = 0 \\ S[t-1,:,:], & t > 0 \end{cases}$
    for $r = 1,\ldots,N$ **do** ▷ iterate clearing up to N rounds (worst-case one new default per round)
      $p \leftarrow \widehat{P}^{(\text{iter}-1)}[t,:]$ ▷ PD vector at time t from previous outer iterate
      $s \leftarrow \min\{\text{surv}, \mathbf{1} - p\}$
      receive $\leftarrow s\,L + A[t,:,:]$
      survivors_step $\leftarrow \mathbf{1}\{\text{receive} > b\}$
      surv $\leftarrow \min\{\text{surv, survivors\_step}\}$ ▷ monotone (no resurrection)
      if *all entries of* surv *stabilised since previous round* **then**
        break ▷ early exit if clearing fixed point reached, will happen in at most n iterations
      end if
    end for
    $S[t,:,:] \leftarrow$ surv
    if $t < T$ **then**
      Propagate assets to next step
    end if
  end for
  (Aggregate across scenarios to form unconditional PDs)
  for $t = 0,\ldots,T$ **do**
    for $i = 1,\ldots,N$ **do**
      denom $\leftarrow \sum_k S[t,k,i]$
      if denom > 0 **then**
        $\widehat{P}^{(\text{iter})}[t,i] \leftarrow 1 - \dfrac{\sum_k S[T,k,i]}{\text{denom}}$
      else
        $\widehat{P}^{(\text{iter})}[t,i] \leftarrow 1$
      end if
    end for
  end for
  dist $\leftarrow \left\lVert \widehat{P}^{(\text{iter})} - \widehat{P}^{(\text{iter}-1)} \right\rVert_2$ ▷ global convergence check
end while
return $\widehat{P}^{(\text{iter})}$

Appendix B

An alternative way of choosing the training buffer

Below we describe an alternative way to generate the training buffer. In code, This subsection documents the path-consistent buffer generator used when `USE_ROLLOUT_PATHS=True` and `NAIVE_PATHS_METHOD="sc_data"` (i.e. the routine `multistep_asset_buffers` in the code). It produces full multi–step paths of external assets and default indicators on the discrete grid $t_k := k\Delta t$, $k = 0, \ldots, T$, with total horizon $T_{\text{total}} = T\Delta t$. The per–step training buffers used by the TD learner are then obtained by slicing these same paths: for each t_k we take the cross–section $\{A(t_k)\}_{\text{scenarios}}$ and the corresponding one–step shock increments $\{dA(t_k)\}_{\text{scenarios}}$. This ensures that the buffers $\{\mathcal{B}(t_k)\}_{k=0}^{T-1}$ are mutually consistent in time (they come from the same simulated rollouts), rather than being sampled independently at each t_k.

Throughout, we simulate S scenarios of an N–agent system. For a fixed scenario, we write $A_i(t_k)$ for the (distance–to–default) external asset of agent i at time t_k, and we track survivorship via an indicator $S_i(t_k) \in \{0, 1\}$ (defaults are absorbing).

Asset increments and initialisation

We generate multi–step asset paths with additive Brownian increments and equicorrelation via a single common factor. For each time step k and scenario we draw a market factor $M_k \sim \mathcal{N}(0, 1)$ and independent idiosyncratic shocks $\varepsilon_{i,k} \sim \mathcal{N}(0, 1)$, and set

$$dA_i(t_k) \;=\; \sigma\sqrt{\Delta t}\left(\sqrt{\rho}\, M_k + \sqrt{1-\rho}\,\varepsilon_{i,k}\right) \;+\; \zeta_{i,k}, \tag{B.1}$$

where $\rho \in [0, 1]$ is the one–factor correlation and $\sigma > 0$ the volatility. The optional term $\zeta_{i,k} \sim \mathcal{N}(0, s_{\text{step}}^2)$ is a small extra idiosyncratic jitter (parameter `STEP_JITTER_SD`: s_{step} = `STEP_JITTER_SD`); setting $s_{\text{step}} = 0$ recovers the pure one–factor model. All shocks are taken independent across time steps and scenarios.

To initialise $A(0)$ in a way that is consistent with the correlation level, we first draw a scenario base $\bar{A}$ around the target initial distance–to–default INIT_DD, and then perturb each agent by a small idiosyncratic term whose dispersion shrinks as $\rho \to 1$. Concretely, we draw

$$\bar{A} \sim \mathrm{Unif}\big(\mathrm{INIT\_DD} - w_0,\ \mathrm{INIT\_DD} + w_0\big), \qquad w_0 := m\Big(0.4 + 0.1\,\sigma\sqrt{T\,\Delta t}\Big),$$

where $m > 0$ is a bandwidth multiplier (`BAND_MULT`) and $T\Delta t$ is the total maturity horizon. Optionally, with probability p_{stress} (`SCENARIO_STRESS_PROB`) we shift the entire scenario base by a constant Δ_{stress} (`SCENARIO_STRESS_SHIFT`), producing a simple stressed mixture of initial conditions. Finally, for each agent i we set

$$A_i(0) = \bar{A} + \eta_i, \qquad \eta_i \sim \mathcal{N}\big(0,\ (\alpha\, w_0 \sqrt{1-\rho})^2\big),$$

where $\alpha \geq 0$ is the per-agent jitter coefficient (`INIT_JITTER`). Thus, if $\alpha = 0$ all agents start identically within a scenario ($A_1(0) = \cdots = A_N(0) = \bar{A}$), while for ρ close to one the initial cross–sectional dispersion is automatically dampened.

Structural barrier defaults within a step (Brownian bridge).

For each step $k \in \{0, \ldots, T-1\}$ we propose a "free" update

$$A_i^{\text{free}}(t_{k+1}) = A_i(t_k) + dA_i(t_k).$$

We impose a structural default barrier at 0 and encode defaulted agents by a fixed sentinel value $a_{\text{dead}} \leq 0$. An agent alive at t_k can default over the step $[t_k, t_{k+1}]$ in two ways: (i) endpoint hit if $A_i^{\text{free}}(t_{k+1}) \leq 0$; and (ii) intra-step hit even when both endpoints are positive. Conditional on $A_i(t_k) > 0$ and $A_i^{\text{free}}(t_{k+1}) > 0$, the probability that an arithmetic Brownian motion bridge crosses 0 within the step is

$$p_{\text{BB}}\big(A_i(t_k), A_i^{\text{free}}(t_{k+1})\big) = \exp\left(-\frac{2\,A_i(t_k)\,A_i^{\text{free}}(t_{k+1})}{\sigma^2\,\Delta t}\right),$$

and we sample an intra-step default whenever $U_{i,k} \leq p_{\text{BB}}$ for an independent $U_{i,k} \sim \mathrm{Unif}(0,1)$. Putting endpoint and intra-step hits together, the structural step-default indicator is

$$D_i^{\text{str}}(t_k) := \mathbb{1}\big\{A_i^{\text{free}}(t_{k+1}) \leq 0\big\} + \mathbb{1}\big\{A_i(t_k) > 0,\ A_i^{\text{free}}(t_{k+1}) > 0,\ U_{i,k} \leq p_{\text{BB}}\big\}.$$

We then update survivorship and assets as

$$S_i(t_{k+1}) = S_i(t_k)\big(1 - D_i^{\text{str}}(t_k)\big), \qquad A_i(t_{k+1}) = \begin{cases} a_{\text{dead}}, & S_i(t_{k+1}) = 0, \\ A_i^{\text{free}}(t_{k+1}), & S_i(t_{k+1}) = 1, \end{cases}$$

so that once a agent defaults it is marked by a_{dead} and remains out of the risk set at later dates. This produces path-consistent pairs $\big(A(t_k), dA(t_k)\big)$ for all k, together with a survival mask $S(t_k)$.

The rollout simulator internally produces tensors

$$A_{\text{path}} \in \mathbb{R}^{S\times(T+1)\times N}, \qquad dA_{\text{all}} \in \mathbb{R}^{S\times T\times N},$$

containing the full asset paths and the one-step increments used at each time step. The actual TD training code uses the slices

$$A_{\text{buf}}(t_k) := A_{\text{path}}[:, k, :] \in \mathbb{R}^{S\times N}, \qquad dA_{\text{buf}}(t_k) := dA_{\text{all}}[:, k, :] \in \mathbb{R}^{S\times N},$$

as the time-t_k buffer inputs for TD0BatchTargets. In particular, the same scenario index is used consistently across k (the buffers are not reshuffled independently at each time).

The more challenging task is to account for interim defaults due to valuation of the interagent claims. Trivially, there is a default barrier determined by asset value. If any agent's asset value is below $b_i - \sum_{j=1}^{N} L_{ji}$, default at the revaluation is imminent, because even in the best case scenario (assuming $V_j(t) = 1$ for all j), equity cannot be greater than zero. The structural barrier simulation above only captures defaults driven directly by the external assets; it tends to underestimate early defaults triggered by deteriorating network conditions. For that purpose we introduce the following variable to proxy the proximity to default once expected counterparty payments are taken into account. This feature will be denoted by $\psi_i(t) \in [0, 1]$ for agent i at time t. This variable will be the driver of the extra stochastic defaults in the training buffer.

Adding additional (network-driven) defaults

Fix a time $t = t_k < T$ and let $A(t) = (A_1(t), \ldots, A_N(t))$ be the current external asset vector. Denote by $T_{\text{rem}} := T_{\text{total}} - t$ the remaining time to maturity (in the implementation: $T_{\text{rem}} = \max\{T_{\text{total}} - k\Delta t, \varepsilon\}$). As in the training, let $a_{\text{dead}} \le 0$ be the sentinel asset value used to encode defaulted agents in the simulation, and define the pre-clearing survivorship at time t as

$$S_i(t) \;:=\; \mathbb{1}\{A_i(t) > a_{\text{dead}}\}, \qquad i = 1, \ldots, N.$$

(With the structural simulator above, $A_i(t)$ is either strictly positive or equal to a_{dead}, so this indicator is equivalent to $\mathbb{1}\{A_i(t) > 0\}$.) Only agents with $S_i(t) = 1$ can contribute to future payments.

For a single agent with current asset x and a flat default barrier h, we model external assets between t and T_{total} as an arithmetic Brownian motion

$$X(s) = x + \sigma W_s, \qquad s \in [0, T_{\text{rem}}],$$

with volatility parameter $\sigma > 0$ and zero drift. The probability that this process does not cross the barrier h before maturity is given by the classical reflection-principle formula

$$\mathbb{P}\Big[\min_{0 \le s \le T_{\text{rem}}} X(s) > h \,\Big|\, X(0) = x\Big] \;=\; 2\,\Phi\!\left(\frac{x-h}{\sigma\sqrt{T_{\text{rem}}}}\right) - 1,$$

where Φ is the standard normal distribution function. We clip this value to $[0,1]$ and set it to zero whenever the agent is already defaulted.

In the network, each agent i faces an effective default barrier that depends on how much it expects to receive from its counterparties. Let $p_j \in [0,1]$ denote a payment probability for counterparty j (on each of its interagent liabilities). For a given vector $p = (p_1, \ldots, p_N)$, the expected incoming interagent cash flow to agent i is

$$\sum_{j=1}^{N} p_j L_{ji},$$

so its effective liability "wall" becomes

$$H_i(p) \;:=\; b_i - \sum_{j=1}^{N} p_j L_{ji}.$$

If counterparties' default probability increases, expected payments shrink, $H_i(p)$ increases and the agent is closer to default.

The same way as in training, we guarantee weak self-consistency by approximating network-consistent survival probabilities $\phi_i(t) \in [0,1]$ via iterating a fixed-point map on the vector of expected payment fractions. Starting from the initial guess that each alive agent pays in full,

$$p_i^{(0)} \;=\; S_i(t), \qquad i = 1, \ldots, N,$$

we perform a fixed number k of iterations (parameter `K_SURV_ITERS`) of

$$\begin{aligned}
H_i^{(m)} &:= b_i - \sum_{j=1}^{N} p_j^{(m)} L_{ji},\\
\phi_i^{(m+1)} &:= S_i(t)\,\Big(2\,\Phi\Big(\frac{A_i(t) - H_i^{(m)}}{\sigma\sqrt{T_{\text{rem}}}}\Big) - 1\Big),\\
p_i^{(m+1)} &:= \min\big\{S_i(t),\, \phi_i^{(m+1)}\big\}, \qquad i = 1, \ldots, N.
\end{aligned}$$

The factor $S_i(t)$ ensures that already defaulted agents ($A_i(t) \le a_{\text{dead}}$) have zero survival and never resurrect.

After k iterations we denote

$$\phi_i(t) \;:=\; \phi_i^{(k)}, \qquad i = 1, \ldots, N,$$

and interpret $\phi_i(t)$ as a counterparty-aware, network-implied probability that agent i survives to maturity without hitting its effective barrier, under the current asset configuration $A(t)$. Note that this is still an imperfect estimate of the actual survival probability: at each state it checks the probability of an agent's asset side hitting a constant default barrier that depends on the counterparties' current financial health. On the other hand, it ignores that the barrier itself is random and will evolve with future defaults.

Definition B.1. The *proximity to default* for agent i at time t is defined as the complement of the iterated survival probability,

$$\psi_i(t) \;:=\; 1 - \phi_i(t) \;\in\; [0, 1].$$

Low values of $\psi_i(t)$ correspond to safe, well-capitalised agents whose assets remain comfortably above their network-implied barrier even if some counterparties weaken. High values of $\psi_i(t)$ identify fragile agents whose survival probability $\phi_i(t)$ is small once counterparty risk is accounted for. Next we show how the proximity-to-default feature is overlaid with the generated asset paths to yield the training buffer.

We take one-step asset increments generated as described above and superimpose an additional layer of stochastic defaults that are correlated both with the current asset level and with the liability network. In the implementation this overlay is optional (`EXTRA_DEFAULTS_ON`) and is applied after the ABM increment has produced the provisional next-step asset $A^{\text{free}}(t+\Delta t)$; the stored increment $dA(t)$ itself is left unchanged.

The extra-default mechanism turns $\psi_i(t)$ into a per-step default probability by an affine rescaling with saturation. Let

$$\lambda_{\text{extra}} > 0 \qquad \text{and} \qquad \overline{p}_{\text{extra}} \in (0, 1)$$

denote the extra intensity and extra cap parameters (`EXTRA_INTENSITY` and `EXTRA_CAP`), respectively. Conditionally on being alive at time t, agent i receives an additional default probability

$$p_i^{\text{extra}}(t) \;=\; \min\Big\{ \lambda_{\text{extra}}\, \psi_i(t),\; \overline{p}_{\text{extra}} \Big\}, \qquad \text{and} \;\; p_i^{\text{extra}}(t) = 0 \text{ if } A_i(t) \leq 0. \tag{B.2}$$

λ_{extra} determines how aggressively network stress is converted into immediate default probability in a single step, while $\overline{p}_{\text{extra}}$ is a ceiling, preventing per-step extra default probabilities from becoming unrealistically close to one even when $\psi_i(t)$ is near unity.

Cross-sectional dependence of these extra defaults is introduced through a one-factor Gaussian copula. For each time t we draw a common market shock $M(t) \sim \mathcal{N}(0,1)$ and idiosyncratic shocks $\varepsilon_i(t) \sim \mathcal{N}(0,1)$, and set

$$Z_{i,t} \;=\; \sqrt{\rho_{\text{extra}}}\,M(t) + \sqrt{1-\rho_{\text{extra}}}\,\varepsilon_i(t), \qquad U_i(t) \;=\; \Phi(Z_{i,t}),$$

where $\rho_{\text{extra}} \in [0,1)$ is an overlay correlation parameter (`EXTRA_RHO`); if not specified we take $\rho_{\text{extra}} = \rho$. Conditional on $A_i(t) > 0$, an extra default of i in step t occurs whenever

$$D_i^{\text{extra}}(t) \;:=\; \mathbb{1}\big\{U_i(t) \le p_i^{\text{extra}}(t)\big\} \;=\; 1.$$

These extra defaults are then applied after the ABM increment from t to $t{+}\Delta t$ has been simulated: if $D_i^{\text{extra}}(t) = 1$, we overwrite the next asset value by the sentinel and mark the agent as defaulted,

$$A_i(t{+}\Delta t) \;\leftarrow\; a_{\text{dead}}, \qquad S_i(t{+}\Delta t) \;=\; 0,$$

so that default is absorbing and the agent cannot "resurrect" at later dates.

In summary, the simulation of training paths proceeds in two layers. First, a structural ABM with barrier at 0 generates the raw asset values and defaults driven purely by external assets (including intra-step barrier hits via the Brownian-bridge rule, and optional per-step jitter in the increments). Second, conditional on the current state, the overlay mechanism (B.2) injects additional, network-correlated defaults with probability proportional to the breach feature $\psi_i(t)$, scaled by λ_{extra} and truncated at $\overline{p}_{\text{extra}}$. This allows the training distribution to contain early defaults that are triggered by deteriorating network conditions rather than by a hard crossing of the external-asset barrier alone.

Appendix C

Algorithmic outline of the TD learner of Section 4.6

Algorithm 2 TRAIN - backward TD training of state-dependent PD in networks for given parameters

Require: agents N, time grid $\mathcal{T} = \{0, \dots, T\}$, interbank liability matrix $L \in \mathbb{R}^{N\times N}$, total liabilities $b \in \mathbb{R}^N$; asset–dynamics parameters $(\rho, \sigma, \Delta t, a_{\text{dead}})$; buffer size B, epochs $N_t^{(\text{ep})}$ for each t, learning rate schedule $\eta(\text{epoch})$; number of training and validation Monte Carlo replicas $k_{\text{MC}}^{\text{train}}$, $k_{\text{MC}}^{\text{val}}$; refresh schedule `REFRESH_EVERY`$(\cdot)$

Ensure: fitted predictors $\{Net_t\}_{t=0}^{T-1}$ and calibrated step-0 valuations $\widehat{V}^{\text{post}}(0, \cdot)$

Initialize parametric networks $\{Net_t\}_{t=0}^{T-1}$ (e.g. MLPs) with a conservative prior on default

Define a deterministic terminal module Net_T that maps any pre-clearing state at T to $P^{\text{post}}(T) = \mathbb{1}_{\{S^+(T)=0\}}$ via terminal clearing

for $t \leftarrow T-1, T-2, \dots, 0$ **do** ▷ backward induction in t

  Freeze $\{Net_u\}_{u>t}$; unfreeze parameters of Net_t

  for epoch $= 1, \dots, N_t^{(\text{ep})}$ **do**

    if *epoch* $= 1$ **or** *epoch* mod `REFRESH_EVERY`(*epoch*) $= 0$ **then**

      $A(t) \leftarrow$ SAMPLEBUFFERATTIME(t, B, ρ) ▷ pre-clearing assets at time t, as per Equation 4.18

      $(P^{\text{post}}(t), \widetilde{Y}(t), W(t)) \leftarrow$ TD0BATCHTARGETS$(t, Net_t, Net_{t+1}, A(t), L, b, \rho, \sigma, \Delta t, a_{\text{dead}}, k_{\text{MC}}^{\text{train}})$ ▷ simulation refresh only at scheduled epochs, using $k_{\text{MC}}^{\text{train}}$ replicas

    end if

    $\ell_{\text{CE}}(t) \leftarrow \text{loss}_{\text{CE}}\big(\widetilde{Y}(t), P^{\text{post}}(t), W(t)\big)$ ▷ binary cross–entropy, weighted by $W(t) = S^+(t)$

    Update parameters of Net_t by one gradient step on $\ell_{\text{CE}}(t)$ (e.g. Adam with step size $\eta(\text{epoch})$)

    Periodically: compute validation $\text{MSE}_t = \frac{1}{|\mathcal{V}|}\sum_{(i,\omega)\in\mathcal{V}}\big(P_{i,\omega}^{\text{post}}(t) - \widetilde{Y}_{i,\omega}(t)\big)^2$ *on a held-out buffer using* $k_{\text{MC}}^{\text{val}}$ *replicas, and use the best* MSE_t *for early stopping.*

  end for

end for

return Net_0 and its post-clearing map $P^{\text{post}}(0, \cdot)$

Algorithm 3 TD0BATCHTARGETS - one TD batch with clearing at t and $t{+}1$

Require: time index $t < T$; current network Net_t; frozen continuation model Net_{t+1}; mini-batch of pre-clearing assets $A(t)$; interbank matrix L, liabilities b; asset parameters $(\rho, \sigma, \Delta t, a_{\text{dead}})$; training Monte Carlo replicas $k_{\text{MC}}^{\text{train}}$

Ensure: post-clearing prediction $P^{\text{post}}(t)$, TD target $\widetilde{Y}(t)$, risk-set weights $W(t)$

$P^{\text{raw}}(t) \leftarrow Net_t(A(t))$ ▷ raw to-maturity default probabilities at time t
$(S^{+}(t), P^{\text{post}}(t)) \leftarrow \text{CLEARANDPROJECTPD}(A(t), P^{\text{raw}}(t), L, b, a_{\text{dead}})$
$W(t) \leftarrow S^{+}(t)$ ▷ only post-clearing survivors contribute to continuation loss
$A(t{+}1) \leftarrow \text{PROPAGATEASSETS}(A(t), S^{+}(t), \rho, \sigma, \Delta t, a_{\text{dead}})$ ▷ one-step ABM shock with absorbing defaults
$P^{\text{raw}}(t{+}1) \leftarrow Net_{t+1}(A(t{+}1))$
$(S^{+}(t{+}1), P^{\text{post}}(t{+}1)) \leftarrow \text{CLEARANDPROJECTPD}(A(t{+}1), P^{\text{raw}}(t{+}1), L, b, a_{\text{dead}})$
$\widetilde{Y}(t) \leftarrow (1 - S^{+}(t)) + S^{+}(t) \cdot P^{\text{post}}(t{+}1)$ ▷ Equation (4.11), componentwise; if $k_{\text{MC}}^{\text{train}} > 1$ the calls above are implicitly repeated and averaged over $k_{\text{MC}}^{\text{train}}$ Monte Carlo replicas.
return $(P^{\text{post}}(t), \widetilde{Y}(t), W(t))$

Algorithm 4 ClearAndProjectPD – clearing with fixed valuations and post-clearing PDs at time t (index dropped everywhere)

Require: pre-clearing assets $A \in \mathbb{R}^{B\times N}$; raw to-maturity default probabilities $P^{\text{raw}} \in [0,1]^{B\times N}$; interbank matrix $L \in \mathbb{R}^{N\times N}$, liabilities $b \in \mathbb{R}^{N}$; sentinel asset level $a_{\text{dead}} \le 0$
Ensure: post-clearing survivorship $S^{+} \in \{0,1\}^{B\times N}$, post-clearing probabilities $P^{\text{post}} \in [0,1]^{B\times N}$

$S^{-} \leftarrow \mathbb{1}\{A > a_{\text{dead}}\}$ ▷ pre-clearing survivorship inferred from assets
$F \leftarrow 1 - P^{\text{raw}}$ ▷ payment probabilities for interbank liabilities
$surv \leftarrow S^{-}$
for $r = 1,\dots,N$ **do** ▷ at most one new default per round
    $pay \leftarrow surv \cdot F$ ▷ componentwise product: apply haircuts
    receive $\leftarrow pay\, L + A$ ▷ matrix product along the counterparty index
    survivors_step $\leftarrow \mathbb{1}\{\text{receive} > b\}$ ▷ limited liability and absolute priority
    new_surv $\leftarrow \min\{surv, \text{survivors\_step}\}$ ▷ monotone: no resurrection
    if new_surv $= surv$ **then**
        break ▷ clearing fixed point reached
    else
        $surv \leftarrow$ new_surv
    end if
end for
$S^{+} \leftarrow surv$
$P^{\text{post}} \leftarrow (1 - S^{+}) + S^{+} \cdot P^{\text{raw}}$ ▷ defaults at t get PD $=1$, survivors keep P^{raw}
return (S^{+}, P^{post})

Appendix D

DeepSets parametrization for the symmetric experiments

The TD construction in Section 4.6 does not require a symmetric neural network. In a general liability network, the predictor at time t may be any sufficiently flexible map from the labelled pre-clearing state to N raw default-probability logits. In some of the numerical experiments of this chapter, however, Assumption 4.17 makes bank labels economically arbitrary: the liability matrix is homogeneous and complete, initial conditions are identical, and the shock law is exchangeable across agents. It is therefore natural to restrict the approximation class to permutation-equivariant maps. If a permutation π is applied to the vector of asset states, the vector of predicted logits should be permuted in the same way.

DeepSets provide a simple way to impose this restriction by parameter sharing. The basic representation theorem of Zaheer et al. (2017) shows that permutation-invariant set functions can be written, under suitable conditions, in the form $\rho(\sum_i \phi(x_i))$. For outputs indexed by the same agents, the analogous equivariant construction combines a shared transformation of the element itself with a symmetric summary of the other elements; see also the parameter-sharing perspective of Ravanbakhsh et al. (2017). This should still be interpreted as a finite-capacity approximation class: Wagstaff et al. (2019) show that the latent dimension matters for universal representation of set functions.

In the code, the node feature is one-dimensional,

$$x_i(t) = A_i(t),$$

where already defaulted agents are represented by the sentinel value a_{dead}. There is no separate survivorship feature and no engineered breach-proxy feature. For a batch of scenarios the input tensor is therefore $x(t) \in \mathbb{R}^{B\times N\times 1}$. The implemented DeepSets predictor has the leave-one-out form

$$e_i(t) = \phi_t\big(x_i(t)\big), \tag{D.1}$$

$$\bar{e}_{-i}(t) = \frac{1}{N-1}\sum_{j\neq i} e_j(t), \tag{D.2}$$

$$z_i(t) = g_t\big(x_i(t), \bar{e}_{-i}(t)\big), \tag{D.3}$$

$$P_i^{\mathrm{raw}}(t) = \frac{1}{1+\exp(-z_i(t))}, \tag{D.4}$$

with the convention $\bar{e}_{-i} = 0$ when $N = 1$. Both ϕ_t and g_t are shared across banks. In the baseline implementation, ϕ_t is a two-layer encoder with embedding dimension 64, and g_t is a two-hidden-layer ReLU decoder with hidden width 128 and scalar output. The final decoder bias is initialized at $\log(0.02/0.98)$, so the untrained raw PD is close to two percent before clearing.

This architecture is permutation-equivariant by construction. If the banks are relabelled by π, then the encoded vectors e_i and the leave-one-out averages $\bar{e}_{-i}$ are relabelled in the same way. Since the same decoder g_t is used for every bank, the logits satisfy

$$z(\pi x) = \pi z(x). \tag{D.5}$$

The model therefore cannot attach meaning to an arbitrary bank index. It can still assign different PDs to different banks in the same scenario whenever their realized asset values differ. For the same reason, the training target is not averaged across banks. Exchangeability is a distributional symmetry, not a statement that all banks in a heterogeneous realization have the same conditional default probability.

Only the neural parametrization is changed. The DeepSets module returns the same object as the flattened MLP: one raw logit per agent and scenario. These logits are transformed into raw probabilities, passed through the same interim clearing map, and trained with the same TD target

$$\widetilde{Y}_i(t) = \big(1 - S_i^+(t)\big) + S_i^+(t)P_i^{\mathrm{post}}(t+1).$$

Thus the loss, backward training schedule, target refresh, zero-warmstart and deterministic terminal clearing module are unchanged.

DeepSets should be understood as a stabilizing and parameter-sharing approximation tailored to the homogeneous fully connected experiments; it is not required for the definition of strong self-consistency. For heterogeneous balance sheets, non-symmetric liabilities, or agent-specific covariates, one should either return to a labelled predictor or use a richer architecture that explicitly includes the relevant node and edge information.

Bibliography

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015a). Networks, shocks, and systemic risk. Technical report, National Bureau of Economic Research. https://doi.org/10.3386/w20931.

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015b). Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608. https://doi.org/10.1257/aer.20130456.

Acerbi, C. and Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7):1487–1503. https://doi.org/10.1016/S0378-4266(02)00283-2.

Acharya, V. V. (2009). A theory of systemic risk and design of prudential bank regulation. *Journal of financial stability*, 5(3):224–255. https://doi.org/10.1016/j.jfs.2009.05.001.

Acharya, V. V. and Yorulmazer, T. (2007). Too many to fail—an analysis of time inconsistency in bank closure policies. *Journal of Financial Intermediation*, 16(1):1–31. https://doi.org/10.1016/j.jfi.2006.06.001.

Acharya, V. V., Pedersen, L. H., Philippon, T., and Richardson, M. (2017). Measuring systemic risk. *The Review of Financial Studies*, 30(1):2–47. https://doi.org/10.1093/rfs/hhw088.

Adrian, T. and Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7):1705–1741. https://doi.org/10.1257/aer.20120555.

Ágoston, K. C. (2010). CVaR számítás SRA algoritmussal. *SZIGMA Matematikai-közgazdasági folyóirat*, 41(1-2):61–73.

Ahn, D.-h. and Kim, E. (2019). Optimal intervention under stress scenarios: A case of the korean financial system. *Operations Research Letters*, 47(4): 257–263. https://doi.org/10.1016/j.orl.2019.03.013.

Allen, F. and Babus, A. (2009). Networks in finance. *The network challenge: strategy, profit, and risk in an interlinked world*, 367. URL `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1094883`.

Allen, F. and Carletti, E. (2008). Mark-to-market accounting and liquidity pricing. *Journal of Accounting and Economics*, 45(2-3):358–378. https://doi.org/10.1016/j.jacceco.2007.02.005.

Allen, F. and Gale, D. (2000). Financial contagion. *Journal of political economy*, 108(1):1–33. https://doi.org/10.1086/262109.

Amini, H., Cont, R., and Minca, A. (2016). Resilience to contagion in financial networks. *Mathematical Finance*, 26(2):329–365. https://doi.org/10.1111/mafi.12051.

Amini, H., Minca, A., and Sulem, A. (2017). Optimal equity infusions in interbank networks. *Journal of Financial Stability*, 31:1–17. https://doi.org/10.1016/j.jfs.2017.05.005.

Amini, H., Minca, A., and Sulem, A. (2022). A dynamic contagion risk model with recovery features. *Mathematics of Operations Research*, 47(2):1412–1442. https://doi.org/10.1287/moor.2021.1174.

Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228. https://doi.org/10.1111/1467-9965.00068.

Balog, D., Bátyi, T. L., Csóka, P., and Pintér, M. (2011). Tőkeallokációs módszerek és tulajdonságaik a gyakorlatban (methods of capital allocation and their characteristics in practice). *Közgazdasági Szemle*, 58(7-8):619–632.

Balog, D., Bátyi, T. L., Csóka, P., and Pintér, M. (2017). Properties and comparison of risk capital allocation methods. *European Journal of Operational Research*, 259(2):614–625. https://doi.org/10.1016/j.ejor.2016.10.052.

Banerjee, T. and Feinstein, Z. (2022). Pricing of debt and equity in a financial network with comonotonic endowments. *Operations Research*. https://doi.org/10.1287/opre.2022.2404.

Banerjee, T., Bernstein, A., and Feinstein, Z. (2025). Dynamic clearing and contagion in financial networks. *European Journal of Operational Research*, 321(2):664–675. https://doi.org/10.1016/j.ejor.2024.09.046.

Bardoscia, M., Battiston, S., Caccioli, F., and Caldarelli, G. (201506). Debtrank: A microscopic foundation for shock propagation. *PLOS ONE*, 10(6): 1–13. https://doi.org/10.1371/journal.pone.0130406.

Bardoscia, M., Barucca, P., Brinley Codd, A., and Hill, J. (2017). The decline of solvency contagion risk. https://doi.org/10.2139/ssrn.2996689.

Bardoscia, M., Barucca, P., Codd, A. B., and Hill, J. (2019). Forward-looking solvency contagion. *Journal of Economic Dynamics and Control*, 108:103755. https://doi.org/10.1016/j.jedc.2019.103755.

Barucca, P., Bardoscia, M., Caccioli, F., D'Errico, M., Visentin, G., Caldarelli, G., and Battiston, S. (2020). Network valuation in financial systems. *Mathematical Finance*, 30(4):1181–1204. https://doi.org/10.1111/mafi.12272.

Basel Committee on Banking Supervision. (2010). Guidance for national authorities operating the countercyclical capital buffer. Technical report, Bank for International Settlements, Basel. URL `https://www.bis.org/publ/bcbs187.htm`.

Basel Committee on Banking Supervision. (2011). Basel III: A global regulatory framework for more resilient banks and banking systems. Technical report, Bank for International Settlements, Basel. URL `https://www.bis.org/publ/bcbs189.htm`.

Basel Committee on Banking Supervision. (2012). A framework for dealing with domestic systemically important banks. Technical report, Bank for International Settlements, Basel. URL `https://www.bis.org/publ/bcbs233.htm`.

Basel Committee on Banking Supervision. (2013). Global systemically important banks: Updated assessment methodology and the higher loss absorbency requirement. Technical report, Bank for International Settlements, Basel. URL `https://www.bis.org/publ/bcbs255.htm`.

Basel Committee on Banking Supervision and International Organization of Securities Commissions. (2015). Margin requirements for non-centrally cleared derivatives. Technical report, Bank for International Settlements and International Organization of Securities Commissions. URL `https://www.bis.org/bcbs/publ/d317.htm`.

Battiston, S., Gatti, D. D., Gallegati, M., Greenwald, B., and Stiglitz, J. E. (2012a). Default cascades: When does risk diversification increase stability? *Journal of Financial Stability*, 8(3):138–149. https://doi.org/10.1016/j.jfs.2012.01.002.

Battiston, S., Puliga, M., Kaushik, R., Tasca, P., and Caldarelli, G. (2012b). Debtrank: Too central to fail? financial networks, the fed and systemic risk. *Scientific Reports*, 2. https://doi.org/10.1038/srep00541.

Battiston, S., Caldarelli, G., D'Errico, M., and Gurciullo, S. (2016). Leveraging the network: a stress-test framework based on debtrank. *Statistics & Risk Modeling*, 33(3-4):117–138. https://doi.org/10.1515/strm-2015-0005.

Berlinger, E., Michaletzky, M., and Szenes, M. (2011). A fedezetlen bankközi forintpiac hálózati dinamikájának vizsgálata a likviditási válság előtt és után (examination of the network dynamics of the uncovered interbank forint markel before the liquidity crisis and after). *Közgazdasági Szemle*, 58(3): 229–252.

Berlinger, E., Daróczi, G., Dömötör, B., and Vadász, T. (2017). Pénzügyi hálózatok mag-periféria szerkezete: A magyar bankközi fedezetlen hitelpiac vizsgálata 2003-2012. *Közgazdasági Szemle*, 44:1160–1185. https://doi.org/10.18414/KSZ.2017.11.1160.

Bertsekas, D. P. and Tsitsiklis, J. N. (1995). Neuro-dynamic programming: an overview. In *Proceedings of 1995 34th IEEE conference on decision and control*, volume 1, pages 560–564. IEEE. https://doi.org/10.1109/CDC.1995.478953.

Billio, M., Getmansky, M., Lo, A. W., and Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559. https://doi.org/10.1016/j.jfineco.2011.12.010.

Birch, A. and Aste, T. (2014). Systemic losses due to counterparty risk in a stylized banking system. *Journal of Statistical Physics*, 156(5):998–1024. https://doi.org/10.1007/s10955-014-1040-4.

Biró, P., van de Klundert, J., Manlove, D., et al. (2021). Modelling and optimisation in European kidney exchange programmes. *European Journal of Operational Research*, 291(2):447–456. https://doi.org/10.1016/j.ejor.2019.09.006.

Bluhm, M., Faia, E., and Krahnen, J. P. (2014). Endogenous banks' networks, cascades and systemic risk. Technical report, SAFE Working Paper. https://doi.org/10.2139/ssrn.2235520.

Borsos, A., Mero, B., et al. (2020). Shock propagation in the banking system with real economy feedback. Technical report, Magyar Nemzeti Bank (Central Bank of Hungary). URL `https://www.mnb.hu/letoltes/mnb-wp-2020-6-final-1.pdf`.

Boss, M., Elsinger, H., Summer, M., and Thurner 4, S. (2004). Network topology of the interbank market. *Quantitative Finance*, 4(6):677–684. https://doi.org/10.1080/14697680400020325.

Bricco, J. and Xu, T. (2019). Interconnectedness and contagion analysis: A practical framework. IMF Working Paper WP/19/220, International Monetary Fund. URL `https://ssrn.com/abstract=3491243`.

Brownlees, C. and Engle, R. F. (2017). SRISK: A conditional capital shortfall measure of systemic risk. *The Review of Financial Studies*, 30(1):48–79. https://doi.org/10.1093/rfs/hhw060.

Brunnermeier, M. K. and Cheridito, P. (2019). Measuring and allocating systemic risk. *Risks*, 7(2):46. https://doi.org/10.3390/risks7020046.

Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *The Review of Financial Studies*, 22(6):2201–2238. https://doi.org/10.1093/rfs/hhn098.

Caccioli, F., Shrestha, M., Moore, C., and Farmer, J. D. (2014). Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking & Finance*, 46:233–245. https://doi.org/10.1016/j.jbankfin.2014.05.021.

Caccioli, F., Farmer, J. D., Foti, N., and Rockmore, D. (2015). Overlapping portfolios, contagion, and financial stability. *Journal of Economic Dynamics and Control*, 51:50–63. https://doi.org/10.1016/j.jedc.2014.09.041.

Cai, J., Eidam, F., Saunders, A., and Steffen, S. (2018). Syndication, interconnectedness, and systemic risk. *Journal of Financial Stability*, 34:105–120. https://doi.org/10.1016/j.jfs.2017.11.013.

Calafiore, G. C., Fracastoro, G., and Proskurnikov, A. V. (2023). Clearing payments in dynamic financial networks. *Automatica*, 158:111299. https://doi.org/10.1016/j.automatica.2023.111299.

Capponi, A. and Chen, P.-C. (2015). Systemic risk mitigation in financial networks. *Journal of Economic Dynamics and Control*, 58:152–166. https://doi.org/10.1016/j.jedc.2015.06.004.

Carmona, R., Fouque, J., and Sun, L. (2015). Mean field games and systemic risk. *Communications in Mathematical Sciences*. https://doi.org/10.4310/CMS.2015.v13.n4.a4.

Chen, Y. and Goldberg, D. (2018). Beating the curse of dimensionality in options pricing and optimal stopping. *arXiv preprint arXiv:1807.02227*. URL `https://people.orie.cornell.edu/dag369/ftp/Beating_Curse_Dim_Opt_Stop_Options_Pricing_Goldberg_ARXIV_8_17_2018.pdf`.

Chong, C. and Kluppelberg, C. (2018). Contagion in financial systems: A bayesian network approach. *SIAM Journal on Financial Mathematics*, 9(1): 28–53. https://doi.org/10.1137/17M1116659.

Cifuentes, R., Ferrucci, G., and Shin, H. S. (2005). Liquidity risk and contagion. *Journal of the European Economic Association*, 3(2-3):556–566. https://doi.org/10.1111/j.1542-4774.2005.tb00986.x.

Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions. (2012). Principles for financial market infrastructures. Technical report, Bank for International Settlements and International Organization of Securities Commissions. URL `https://www.bis.org/cpmi/publ/d101a.htm`.

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

Craig, B. and von Peter, G. (2014). Interbank tiering and money center banks. *Journal of Financial Intermediation*, 23(3):322–347. https://doi.org/10.1016/j.jfi.2014.02.003.

Csóka, P. (2017). Az arányos csődszabály karakterizációja körbetartozások esetén. *Közgazdasági Szemle*, 64(9):930–942. https://doi.org/10.18414/KSZ.2017.9.930.

Csóka, P. and Herings, P. J.-J. (2021). An axiomatization of the proportional rule in financial networks. *Management Science*, 67(5):2799–2812. https://doi.org/10.1287/mnsc.2020.3700.

Csóka, P. and Herings, P. J.-J. (2024). Uniqueness of clearing payment matrices in financial networks. *Mathematics of Operations Research*, 49(1):232–250. https://doi.org/10.1287/moor.2023.1354.

Csóka, P. and Kondor, G. (2020). Csődszabályok pénzügyi hálózatokban. *Alkalmazott Matematikai Lapok*, 37(2):1–13.

Csóka, P. and Sass, Z. (2023). Rendszerkockázati jelentőség mérése minimálisan szükséges tőkeinjekció segítségével. *SZIGMA Matematikai-közgazdasági folyóirat*, 54(1):1–20.

Csóka, P., Herings, P. J.-J., and Kóczy, L. Á. (2009). Stable allocations of risk. *Games and Economic Behavior*, 67(1):266–276. https://doi.org/10.1016/j.geb.2008.11.001.

Csóka, P. and Herings, P. J.-J. (2018). Decentralized clearing in financial networks. *Management Science*, 64(10):4681–4699. https://doi.org/10.1287/mnsc.2017.2847.

Davis, P. J. and Rabinowitz, P. (2007). *Methods of numerical integration*. Courier Corporation. ISBN 9780486453392.

Dawid, A. P. and Musio, M. (2014). Theory and applications of proper scoring rules. *Metron*, 72(2):169–183. https://doi.org/10.1007/s40300-014-0039-y.

Delpini, D., Battiston, S., Caldarelli, G., and Riccaboni, M. (2019). Systemic risk from investment similarities. *PloS one*, 14(5):e0217141. https://doi.org/10.1371/journal.pone.0217141.

Demange, G. (2016). Contagion in financial networks: a threat index. *Management Science*. https://doi.org/10.1287/mnsc.2016.2592.

Demange, G. (2018). Contagion in financial networks: a threat index. *Management Science*, 64(2):955–970. https://doi.org/10.1287/mnsc.2016.2592.

Denault, M. (2001). Coherent allocation of risk capital. *Journal of Risk*, 4: 1–34. https://doi.org/10.21314/JOR.2001.053.

D'Errico, M. and Roukny, T. (2021). Compressing over-the-counter markets. *Operations Research*. https://doi.org/10.1287/opre.2020.2035.

Dhaene, J., Denuit, M., Goovaerts, M. J., Kaas, R., and Vyncke, D. (2002). The concept of comonotonicity in actuarial science and finance: applications. *Insurance: Mathematics and Economics*, 31(2):133–161. https://doi.org/10.1016/S0167-6687(02)00135-X.

Diebold, F. X. and Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1):119–134. https://doi.org/10.1016/j.jeconom.2014.04.012.

Drehmann, M. and Tarashev, N. (2013). Measuring the systemic importance of interconnected banks. *Journal of Financial Intermediation*, 22(4):586–607. https://doi.org/10.1016/j.jfi.2013.05.004.

Eisenberg, L. and Noe, T. H. (2001). Systemic risk in financial systems. *Management Science*, 47(2):236–249. https://doi.org/10.1287/mnsc.47.2.236.9835.

Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review*, 104(10):3115–53. https://doi.org/10.1257/aer.104.10.3115.

Embrechts, P., McNeil, A., and Straumann, D. (2002). Correlation and dependence in risk management: Properties and pitfalls. In Dempster, M. A. H., editor, *Risk Management: Value at Risk and Beyond*, pages 176–223. Cambridge University Press. https://doi.org/10.1017/CBO9780511615337.008.

European Banking Authority. (2014). Guidelines on the criteria to determine the conditions of application of article 131(3) of directive 2013/36/EU

(CRD) in relation to the assessment of other systemically important institutions (O-SIIs). EBA Guidelines EBA/GL/2014/10, European Banking Authority. URL https://www.eba.europa.eu/activities/single-rulebook/regulatory-activities/own-funds/guidelines-criteria-assess-other-systemically-important-institutions-o-siis.

Feinstein, Z. and Sojmark, A. (2022). Endogenous distress contagion in a dynamic interbank model: how possible future losses may spell doom today. *arXiv preprint arXiv:2211.15431*. URL https://arxiv.org/abs/2211.15431.

Feinstein, Z., Rudloff, B., and Weber, S. (2017). Measures of systemic risk. *SIAM Journal on Financial Mathematics*, 8(1):672–708. https://doi.org/10.1137/16M1056085.

Fernández-Villaverde, J., Nuño, G., and Perla, J. (2024). Taming the curse of dimensionality: quantitative economics with deep learning. Technical report, National Bureau of Economic Research. https://doi.org/10.3386/w33117.

Financial Stability Board. (2009). Improving financial regulation: Report of the financial stability board to G20 leaders. Technical report, Financial Stability Board, Basel. URL https://www.fsb.org/publications/r_090925b.pdf.

Fischer, T. (2014). No-arbitrage pricing under systemic risk: Accounting for cross-ownership. *Mathematical Finance*, 24(1):97–124. https://doi.org/10.1111/j.1467-9965.2012.00526.x.

Forte, F., Elosegui, P., and Montes-Rojas, G. (2022). Network structure and fragmentation of the argentinean interbank markets. https://doi.org/10.1016/j.latcb.2022.100066.

Frey, R. and Hledik, J. (2014). Correlation and contagion as sources of systemic risk. *Available at SSRN 2541733*. https://doi.org/10.2139/ssrn.2541733.

Frey, R. and McNeil, A. J. (2003). Dependent defaults in models of portfolio credit risk. *Journal of Risk*, 6:59–92. https://doi.org/10.21314/JOR.2003.089.

Friedli, S. and Velenik, Y. (2017). *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction*. Cambridge University Press, Cambridge. ISBN 9781107184824. https://doi.org/10.1017/9781316882603.

Furfine, C. H. (2003). Interbank exposures: Quantifying the risk of contagion. *Journal of Money, Credit and Banking*, 35(1):111–128. https://doi.org/10.1353/mcb.2003.0004.

Gai, P. and Kapadia, S. (2010). Contagion in financial networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2120):2401–2423. https://doi.org/10.1098/rspa.2009.0410.

Gai, P., Haldane, A., and Kapadia, S. (2011). Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5):453–470. https://doi.org/10.1016/j.jmoneco.2011.05.005.

Ghosh, S., Khanra, P., Kundu, P., Ji, P., Ghosh, D., and Hens, C. (2023). Dimension reduction in higher-order contagious phenomena. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 33(5). https://doi.org/10.1063/5.0152959.

Glasserman, P. and Young, H. P. (2015). How likely is contagion in financial networks? *Journal of Banking & Finance*, 50:383–399. https://doi.org/10.1016/j.jbankfin.2014.02.006.

Glasserman, P. and Young, H. P. (2016). Contagion in financial networks. *Journal of Economic Literature*, 54(3):779–831. https://doi.org/10.1257/jel.20151228.

Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102 (477):359–378. https://doi.org/10.1198/016214506000001437.

Gordy, M. B. (2003). A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation*, 12(3):199–232. https://doi.org/10.1016/S1042-9573(03)00040-8.

Gourieroux, C., Héam, J.-C., and Monfort, A. (2012). Bilateral exposures and systemic solvency risk. *Canadian Journal of Economics/Revue canadienne d'économique*, 45(4):1273–1309. https://doi.org/10.1111/j.1540-5982.2012.01750.x.

Grassi, A., Kosiahn, M., Lelli, C., Martín, M. L., Moers, M., Sydow, M., Vincent, M., and Wiersema, G. (2025). Integrating contagion risk into the 2025 EU-wide stress test: a system-wide analysis with amplification effects between banks and non-banks. *Macroprudential Bulletin*, 32. URL `https://ideas.repec.org/a/ecb/ecbmbu/202500323.html`.

Greenwood, R., Landier, A., and Thesmar, D. (2015). Vulnerable banks. *Journal of Financial Economics*, 115(3):471–485. https://doi.org/10.1016/j.jfineco.2014.11.006.

Gualdi, S., Cimini, G., Primicerio, K., Di Clemente, R., and Challet, D. (2016). Statistically validated network of portfolio overlaps and systemic risk. *Scientific Reports*, 6(1):39467. https://doi.org/10.1038/srep39467.

Haldane, A. G. and May, R. M. (2011). Systemic risk in banking ecosystems. *Nature*, 469(7330):351–355. https://doi.org/10.1038/nature09659.

Hurd, T. R. (2016). *Contagion! Systemic Risk in Financial Networks*. Springer. https://doi.org/10.1007/978-3-319-33930-6.

Hurd, T. R., Cellai, D., Melnik, S., and Shao, Q. (2014). Illiquidity and insolvency: a double cascade model of financial crises. *Available at SSRN 2424877*. https://doi.org/10.2139/ssrn.2424877.

Irons, B. M. and Tuck, R. C. (1969). A version of the aitken accelerator for computer iteration. *International Journal for Numerical Methods in Engineering*, 1(3):275–277. https://doi.org/10.1002/nme.1620010306.

Jackson, M. O. and Pernoud, A. (2024). Credit freezes, equilibrium multiplicity, and optimal bailouts in financial networks. *The Review of Financial Studies*, 37(7):2017–2062. https://doi.org/10.1093/rfs/hhae003.

Karatzas, I. and Shreve, S. (2014). *Brownian motion and stochastic calculus*. springer. https://doi.org/10.1007/978-1-4612-0949-2.

Kusnetsov, M. and Veraart, L. A. (2019). Interbank clearing in financial networks with multiple maturities. *SIAM Journal on Financial Mathematics*, 10(1):37–67. https://doi.org/10.1137/18M1196236.

Laux, C. and Leuz, C. (2009). The crisis of fair-value accounting: Making sense of the recent debate. *Accounting, Organizations and Society*, 34(6-7): 826–834. https://doi.org/10.1016/S0361-3682(09)00043-9.

Lee, C.-C., Chen, P.-F., and Zeng, J.-H. (2020). Bank income diversification, asset correlation and systemic risk. *South African Journal of Economics*, 88 (1):71–89. https://doi.org/10.1111/saje.12235.

Lehar, A. (2005). Measuring systemic risk: A risk management approach. *Journal of Banking & Finance*, 29(10):2577–2603. https://doi.org/10.1016/j.jbankfin.2004.09.007.

Li, D. X. (2000). On default correlation: A copula function approach. *The Journal of Fixed Income*, 9(4):43–54. https://doi.org/10.3905/jfi.2000.319253.

Li, S. and Wang, C. (2021). Network structure, portfolio diversification and systemic risk. *Journal of Management Science and Engineering*, 6(2):235–245. https://doi.org/10.1016/j.jmse.2021.06.006.

Lorenz, J., Battiston, S., and Schweitzer, F. (2009). Systemic risk in a unifying framework for cascading processes on networks. *The European Physical Journal B*, 71(4):441–460. https://doi.org/10.1140/epjb/e2009-00347-4.

Magyar Nemzeti Bank. (2025). Macroprudential report 2025. URL `https://www.mnb.hu/en/publications/reports/macroprudential-report/macroprudential-report-2025`.

Magyar Nemzeti Bank. (2026). Capital buffer for other systemically important institutions (O-SII). URL `https://www.mnb.hu/en/financial-stability/macroprudential-policy/the-macroprudential-toolkit/capital-buffer-for-other-systemically-important-institutions-o-sii`.

Nier, E., Yang, J., Yorulmazer, T., and Alentorn, A. (2007). Network models and financial stability. *Journal of Economic Dynamics and Control*, 31(6): 2033–2060. https://doi.org/10.1016/j.jedc.2007.01.014.

Nishijima, T. (2021). Universal approximation theorem for neural networks. *arXiv preprint arXiv:2102.10993*. https://doi.org/10.48550/arXiv.2102.10993.

Plantin, G., Sapra, H., and Shin, H. S. (2008). Marking-to-market: Panacea or pandora's box? *Journal of Accounting Research*, 46(2):435–460. https://doi.org/10.1111/j.1475-679X.2008.00281.x.

Poledna, S., Martínez-Jaramillo, S., Caccioli, F., and Thurner, S. (2021). Quantification of systemic risk from overlapping portfolios in the financial system. *Journal of Financial Stability*, 52:100808. https://doi.org/10.1016/j.jfs.2020.100808.

Raffestin, L. (2014). Diversification and systemic risk. *Journal of Banking & Finance*, 46:85–106. https://doi.org/10.1016/j.jbankfin.2014.06.018.

Ravanbakhsh, S., Schneider, J., and Póczos, B. (2017). Equivariance through parameter-sharing. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2892–2901. PMLR. URL `https://proceedings.mlr.press/v70/ravanbakhsh17a.html`.

Reinhart, C. M. and Rogoff, K. S. (2009). *This Time is Different: Eight Centuries of Financial Folly.* Princeton University Press, Princeton, NJ. https://doi.org/10.1515/9781400831722.

Rogers, L. C. and Veraart, L. A. (2013). Failure and rescue in an interbank network. *Management Science*, 59(4):882–898. https://doi.org/10.1287/mnsc.1120.1569.

Rombach, M. P., Porter, M. A., Fowler, J. H., and Mucha, P. J. (2014). Core-periphery structure in networks. *SIAM Journal on Applied mathematics*, 74 (1):167–190. https://doi.org/10.1137/120881683.

Roukny, T., Battiston, S., and Stiglitz, J. E. (2018). Interconnectedness as a source of uncertainty in systemic risk. *Journal of Financial Stability*, 35: 93–106. https://doi.org/10.1016/j.jfs.2016.12.003.

Schuldenzucker, S. and Seuken, S. (2020). Portfolio compression in financial networks: Incentives and systemic risk. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 79–79. https://doi.org/10.1145/3391403.3399530.

Schuldenzucker, S., Seuken, S., and Battiston, S. (2020). Default ambiguity: Credit default swaps create new systemic risks in financial networks. *Management Science*, 66(5):1981–1998. https://doi.org/10.1287/mnsc.2019.3304.

Shapley, L. S. (1953). A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317. https://doi.org/10.1515/9781400881970-018.

Sonin, I. M. and Sonin, K. (2025). A continuous-time model of financial clearing. *Annals of Operations Research*, pages 1–29. https://doi.org/10.1007/s10479-024-06389-0.

Soramäki, K., Bech, M. L., Arnold, J., Glass, R. J., and Beyeler, W. E. (2007). The topology of interbank payment flows. *Physica A: Statistical Mechanics and its Applications*, 379(1):317–333. https://doi.org/10.1016/j.physa.2006.11.093.

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement Learning: An Introduction.* MIT Press, 2 edition. URL `http://incompleteideas.net/book/the-book-2nd.html`.

Sznitman, A.-S. (1991). Topics in propagation of chaos. In Hennequin, P.-L., editor, *École d'Été de Probabilités de Saint-Flour XIX—1989*, volume 1464 of *Lecture Notes in Mathematics*, pages 165–251. Springer, Berlin, Heidelberg. https://doi.org/10.1007/BFb0085169.

Tarashev, N., Tsatsaronis, K., and Borio, C. (2016). Risk attribution using the shapley value: Methodology and policy applications. *Review of Finance*, 20(3):1189–1213. https://doi.org/10.1093/rof/rfv028.

Tarski, A. (1955). A lattice-theoretical fixpoint theorem and its applications. *Pacific journal of Mathematics*, 5(2):285–309. https://doi.org/10.2140/pjm.1955.5.285.

Tasca, P., Mavrodiev, P., and Schweitzer, F. (2014). Quantifying the impact of leveraging and diversification on systemic risk. *Journal of Financial Stability*, 15:43–52. https://doi.org/10.1016/j.jfs.2014.08.006.

Tirole, J. (2011). Illiquidity and all its friends. *Journal of Economic Literature*, 49(2):287–325. https://doi.org/10.1257/jel.49.2.287.

Upper, C. (2011). Simulation methods to assess the danger of contagion in interbank markets. *Journal of Financial Stability*, 7(3):111–125. https://doi.org/10.1016/j.jfs.2010.12.001.

Vasicek, O. (2002). The distribution of loan portfolio value. *Risk*, 15(12). URL `https://www.bancoinvest.pt/docs/default-source/fundamentais-risco-credito/distributionloanportfoliovalue.pdf`.

Veraart, L. A. M. (2020). Distress and default contagion in financial networks. *Mathematical Finance*, 30(3):705–737. https://doi.org/10.1111/mafi.12247.

Veraart, L. A. M. (2022). When does portfolio compression reduce systemic risk? *Mathematical Finance*, 32(3):727–778. https://doi.org/10.1111/mafi.12346.

Visentin, G., Battiston, S., and D'Errico, M. (2016). Rethinking financial contagion. https://doi.org/10.48550/arXiv.1608.07831. URL `https://arxiv.org/pdf/1608.07831.pdf`. Working paper.

Wagner, W. (2010). Diversification at financial institutions and systemic crises. *Journal of Financial Intermediation*, 19(3):373–386. https://doi.org/10.1016/j.jfi.2009.07.002.

Wagstaff, E., Fuchs, F., Engelcke, M., Posner, I., and Osborne, M. A. (2019). On the limitations of representing functions on sets. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6487–6494. PMLR. URL `https://proceedings.mlr.press/v97/wagstaff19a.html`.

Walter, G. (2002). Var-limitrendszer melletti hozammaximalizálás: a kaszinóhatás. *Közgazdasági Szemle*, 3:212–234. URL `http://efolyoirat.oszk.hu/00000/00017/00080/pdf/walter.pdf`.

Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R., and Smola, A. J. (2017). Deep sets. In *Advances in Neural Information Processing Systems*, volume 30, pages 3391–3401. URL `https://mlanthology.org/neurips/2017/zaheer2017neurips-deep/`.