Reinforcement Learning for Global-Local Control of Dynamic Networks with Applications to Systemic Risk

Nils Detering
Heinrich Heine University Düsseldorf

together with Florian Grell (Heinrich Heine University Düsseldorf) and Thilo Meyer-Brandis (University of Munich)

Swissquote Workshop on Graph Learning in Financial Networks

Contents

Motivation: Systemic Risk

- Before the financial crisis of 2007/2008, banks were mostly considered in isolation.
- Regulation primarily aimed at limiting the probability of default of individual institutions.

Problem:

- Financial systems are highly interconnected:
  - loans between banks,
  - common asset holdings,
  - derivatives exposures,
  - fire-sale externalities,
  - ...
- Shocks can therefore propagate through the network.
- Local problems can lead to a systemic crises.

Motivation: Systemic Risk

- Before the financial crisis of 2007/2008, banks were mostly considered in isolation.
- Regulation primarily aimed at limiting the probability of default of individual institutions.

Problem:

- Financial systems are highly interconnected:
  - loans between banks,
  - common asset holdings,
  - derivatives exposures,
  - fire-sale externalities,
  - ...
- Shocks can therefore propagate through the network.
- Local problems can lead to a systemic crises.

Stability of individual banks does not automatically imply stability of the system.

Literature: Systemic Risk in Financial Networks

Motivation and basic model

- Financial systems exhibit complex, heterogeneous dependency structures, naturally modeled as weighted graphs (nodes: institutions; edges: monetary exposures,....).
- Default contagion first studied explicitly via clearing equilibrium payments Eisenberg and Noe (2001)

More advanced contagion channels and extensions

- Beyond defaults, important contagion channels include fire sales, common asset holdings,.....
- Extensions incorporate multiple contagion channels, bankruptcy costs, and systemic dynamics Rogers and Veraart (2013); Weber and Weske (2017a); Capponi et al. (2016); Gouriéroux et al. (2012); Cifuentes et al. (2005); Nier et al. (2007); Bichuch and Chen (2020); Bichuch and Feinstein (2022),....

Random networks and asymptotic analysis

- Random network approaches (Gai and Kapadia (2010); Amini et al. (2016); Detering et al. (2019)) analyze systemic damage using asymptotic methods as networks grow large.
- Modeling of local interactions by preserving empirical network statistics (e.g., degree heterogeneity).

Literature: Systemic Risk in Financial Networks

Fewer works on network control related to systemic risk: Endogenous network formation, optimal bailouts, ... and dynamic evolution turn out challenging to investigate analytically.
Notable exceptions (incomplete list):

- Amini et al. (2017) and Minca and Sulem (2014) study optimal equity infusions into a financial network. Problem reduces to a combinatorial optimization problem.
- Capponi and Chen (2015) analyses a lender of last resort and shows that time sequence of clearing payments is unique when a liquidity loan allocation strategy is specified. Capponi et al. (2022) consider optimal bailouts and the feedback effects from increased sovereign's default risk.
- Veraart (2022) considers portfolio compression and states it as an optimization problem.
- Papachristou et al. (2023) studies dynamic interventions in an Eisenberg-Noe setting and shows that the problem can be turned into a sequence of linear programs.
- Carmona et al. (2013) and follow up articles Fouque and Ichiba (2013); Carmona et al. (2016); Garnier et al. (2013) study mean field setting for inter-bank borrowing and lending by diffusion processes coupled through their drifts.

Reinforcement Learning on networks

Technical challenges with numerical procedures for network control:

- Financial networks are heterogeneous and sparse.
- It is well known that heterogeneous, sparse networks are most difficult to control (see Liu et al. (2011)).
- Moreover, control policies often need to be transparent and fair which adds constraints.
- Controls are often of *Global-Local* nature.

Our contribution:

- We address *Global-Local* control of financial networks using reinforcement learning on graphs.
- Policy gradient methods with Graph Neural Networks (GNN),
- Our approach is scalable and leverages relational information, symmetry structures, permutation invariance,....
- We can ensure that certain fairness and transparency requirements hold.
- **Case study:** How do regulators intervene optimally in case the financial system is under stress and needs to be supported to prevent a damaging cascade?

Contents

Model Setup: MDP on a dynamic graph

Graph Structure:

- Graph with fixed set of n nodes
- Let V, W, U be finite sets.
- Each node has an m-dimensional feature vector
- Graph described by:
  - Node features $\in V^m$,
  - weighted adjacency matrix $\in W^{n \times n}$,
  - a global feature $\in U^k$

State Space:

$$\mathcal{S} := V^{n \times m} \times W^{n \times n} \times U^k$$

Model Setup: MDP on a dynamic graph

Markov Decision Process (MDP):

- Markov Decision Process $(\mathcal{S}, \mathcal{A}, \mathcal{R}, p)$:
  - consisting of the state space $\mathcal{S}$ as defined above,
  - a finite action space $\mathcal{A} := \bigcup_{s\in\mathcal{S}} \mathcal{A}_s$, where $\mathcal{A}_s$ is set of possible actions in state s,
  - a set of possible rewards $\mathcal{R} \subset \mathbb{R}$.
  - Transition dynamics described by Markov kernel:

$$p : \mathcal{P}(\mathcal{S}) \otimes \mathcal{P}(\mathcal{R}) \times \mathcal{S} \times \mathcal{A} \to [0,1] \tag{1}$$

- Usually factorization of kernel:

$$p(\{s', r\}; s, a) = h(\{s'\}; s, a) \cdot q(\{r\}; s', s, a)$$

  - h: State transition kernel
  - q: Reward kernel (often deterministic)
- Given distribution μ over $\mathcal{S}$ and a policy π, there exists stochastic process $(S_t, A_t, R_t)_{t\in\mathbb{N}_0}$ with values in $(\mathcal{S} \times \mathcal{A} \times \mathcal{R})$ such that conditional prob. are given by (1).

Model Setup: Learning Objective

Learning Objective, infinite horizon:

$$\max_{\pi} \mathbb{E}_s^{\pi}\left[\sum_{t=0}^{\infty} \gamma^t R_t\right], \quad \gamma \in (0,1)$$

Optimal policy π^* exists, stationary, deterministic:

$$\pi_t^*(\cdot; s_0, \ldots, s_t) = \delta_{f(s_t)}(\cdot), \quad f : \mathcal{S} \to \mathcal{A}.$$

Contents

Intervention in a Dynamic Financial Network

Financial network setup:

- Discrete multi-period model of a network with n banks.
- At each time step $t = 0, 1, \ldots,$ a shock may trigger a default cascade.
- A regulator can inject capital (prior to the shock) to prevent or mitigate the cascade.

System state:

$$(X(t), L(t), S(t))_{t=0,\ldots,T} \in \mathcal{S} = V^{n\times 2} \times W^{(n+1)\times(n+1)} \times U^k$$

- **Assets:** $X(t) \in V^{n\times 2}$, $X_{i,1}(t)$ risk-free, $X_{i,2}(t)$ risky.
- **Liabilities:** $L(t) \in W^{(n+1)\times(n+1)}$, $L_{i,j}(t)$: debt of bank i to bank $j > 0$, $L_{i,1}(t)$: to external parties.

$$\text{Total liabilities } \sum_j L_{i,j}(t), \quad \text{Claims } \sum_j L_{j,i}(t)$$

- **Market price:** $S(t) \in U$, value of the risky asset at time t.

Objective: Dynamic capital injections to limit systemic risk.

Intervention in a Dynamic Financial Network

- **Assets:** Of bank i at time t

$$A_i(t) = \sum_{j=1}^{n} L_{j,i}(t) + X_{i,1}(t) + X_{i,2}(t)\, S(t).$$

- **Liabilities:** $\sum_{j=1}^{n} L_{i,j}(t)$, and the **equity** $C_i(t) = A_i(t) - \sum_{j=1}^{n} L_{i,j}(t)$.
- **Contagion effect:** If $C_i(t) < 0$ for at least one bank i, clearing is performed according to the fixed-point model of Weber and Weske (2017b), taking into account bankruptcy costs, cross-holdings, and fire sales.
- The network **after clearing** at time $t+1$ is given by

$$(X(t+1), L(t+1), S(t+1)) = \Lambda\big((X(t), L(t), S(t) + N(t)\big),$$

where $\Lambda : V^{n\times 2} \times W^{(n+1)\times(n+1)} \times U^k \to V^{n\times 2} \times W^{(n+1)\times(n+1)} \times U^k$ denotes the clearing operation and $N(t)$ describes non fire-sales related stochastic fluctuations of the risky assets.

Illustration: Clearing in a Financial Network

Idea: If a bank defaults, its repayments to creditors are reduced. These losses propagate along the edges of the network until a new equilibrium (*clearing state*) is reached.

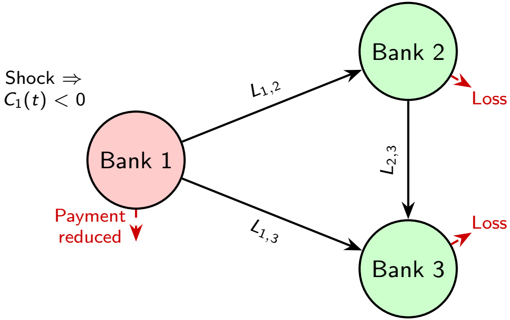

Intervention in a Dynamic Financial Network

Capital injections by the regulator:

- At each time step t, a regulator (e.g., lender of last resort) can choose a capital injection from an action space:

$$A_s \subset \left\{ a \in \mathbb{R}_+^n \;\middle|\; \sum_i |a_i| \leq c \right\}$$

  where $c > 0$ denotes the maximum available total budget.

- The allocated amount a_i for state $s \in \mathcal{S}$ is added to the risk-free assets of bank i:

$$X_{i,1}(t) \leftarrow X_{i,1}(t) + a_i, \qquad i = 1, \ldots, n.$$

- The regulator's objective is to maximize $\max_\pi \mathbb{E}_s^\pi \left[\sum_{t=0}^\infty \gamma^t R_t \right]$, with periodic reward (tradeoff between capital usage and stabilization)

$$R_t = -\bar{C} \sum_{i=1}^n a_i + \sum_{i=1}^n \left(C_i(t+1) - C_i(t) \right),$$

  for a fixed weight $\bar{C} > 1$.

Contents

Policy gradient with Graph Neural Networks

- Standard approach to finding such an optimal policy is via solving the Bellman optimality equation.
- However, often $p : \mathcal{P}(\mathcal{S}) \otimes \mathcal{P}(\mathcal{R}) \times \mathcal{S} \times \mathcal{A} \to [0,1]$ is not know.
- Reinforcement learning offers model independent, sample based numerical procedure.
- Q learning not feasible because $\mathcal{S} = V^{n\times m} \times W^{n\times n} \times U^k$ is huge.
- Therefore focus on a family of parametrised differentiable policies $\{\pi_\theta : \theta \in \Theta\}$ with a parameter space $\Theta \subset \mathbb{R}^d$ and maximize $J_s(\theta) := \mathbb{E}_s^{\pi_\theta}\left[\sum_{t=0}^{\infty} \gamma^t R_t\right]$ over Θ
  - Gradient can be written as

$$\nabla_\theta J_s(\theta) = \mathbb{E}_s^{\pi_\theta}\left[\sum_{t=0}^{\infty} \gamma^t \nabla_\theta \left(\log(\pi_\theta(\{A_t\}; S_t))\right) \Psi_t\right],$$

and we can use gradient descent with $\Psi_t = \sum_{k=t}^{\infty} \gamma^{k-1} R_k$ or modifications.

Policy Representation with Graph Neural Networks

Motivation: Our state space is a *network* - nodes represent financial institutions, edges represent monetary liabilities. To ensure the policy π_θ can exploit the network structure, we parametrize it using a **Message Passing Graph Neural Network (MP-GNN)**.

Idea of Message Passing:

- Each node gathers information from its neighborhood.
- These *messages* are aggregated and incorporated into the node state.
- After several layers, each node contains a representation of its local environment.

Policy Representation with Graph Neural Network

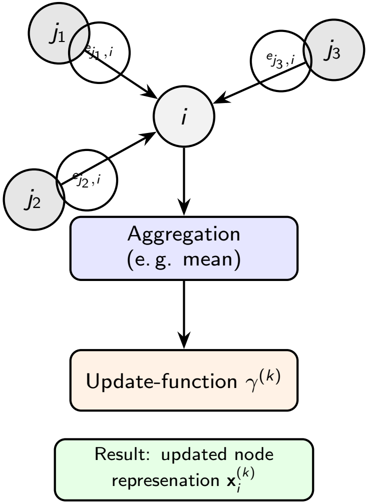

Message Passing Network GNN (MP-GNN): Formal Definition

Let K fixed number of layers and σ activation function. Let further $m_k \in \mathbb{N}$ for $k = 1, \ldots K$ with $m_1 = m$ (feature dimension). Define a map

$$f : \mathbb{R}^m \to \mathbb{R}^{m_K}, \mathbf{x}_i^{(0)} \mapsto \mathbf{x}_i^{(K)}$$

recursively by:

$$\mathbf{x}_i^{(k)} = \gamma^{(k)} \left[\mathbf{x}_i^{(k-1)}, \bigoplus_{j \in \mathcal{N}^-(i)} \phi\left(\mathbf{x}_j^{(k-1)}, \mathbf{e}_{j,i}\right), \bigoplus_{j \in \mathcal{N}^+(i)} \phi\left(\mathbf{x}_j^{(k-1)}, \mathbf{e}_{j,i}\right) \right],$$

where $\bigoplus_{j \in \mathcal{N}^{+-}(i)}$ represents mean aggregation over the in- and out-neighborhood $\mathcal{N}^-(i)$ and $\mathcal{N}^+(i)$ of i, and

$$\phi\left(\mathbf{x}_j^{(k-1)}, \mathbf{e}_{j,i}\right) = e_{j,i} x_j^{(k-1)}, \qquad \gamma^{(k)}(z) = \sigma\left(\mathbf{W}^{(k)} \cdot z + b^{(k)}\right).$$

Here, $\mathbf{W}^{(k)} \in \mathbb{R}^{m_k \times 3m_{k-1}}$, $b^{(k)} \in \mathbb{R}^{m_k}$, $k = 1 \ldots K$ are a trainable weight matrices and bias vectors.

Result: Each node obtains an embedding that takes information from its local environment into account - suitable for parametrizing policies π_θ on networks.

Contents

Desirable graph function properties

- Functions on networks might need to fulfill some fairness constraints.
  - Reshuffling of node indices or positions should not change the function value.
  - Depending on context, node with *same or similar* neighborhood should be treated equally.

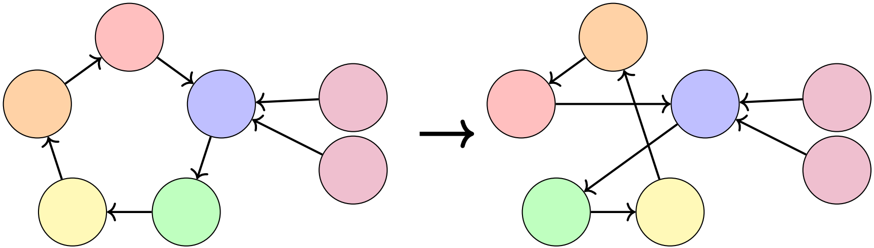

- Denote by $\mathcal{F}_n^{\mathrm{inv}}$ and $\mathcal{F}_n^{\mathrm{eq}}$ the set of permutation invariant (for graph labels) and equivariant (node labels) functions.

k-hop indistinguishability and k-hop-fair maps

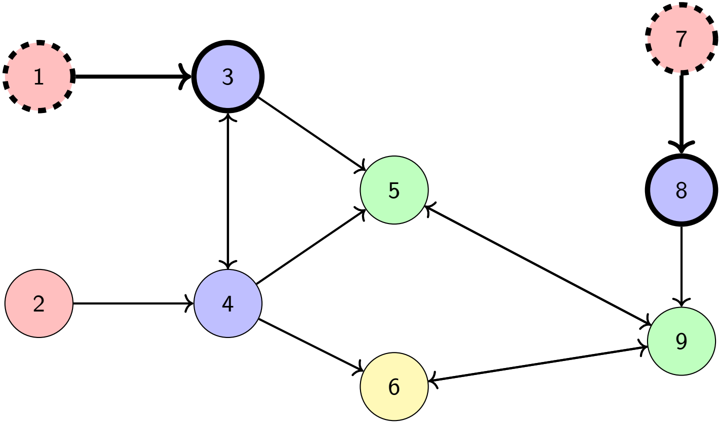

Node 1 and 7 have isomorphic 1-hop neighborhoods.

k-hop indistinguishability and k-hop-fair maps

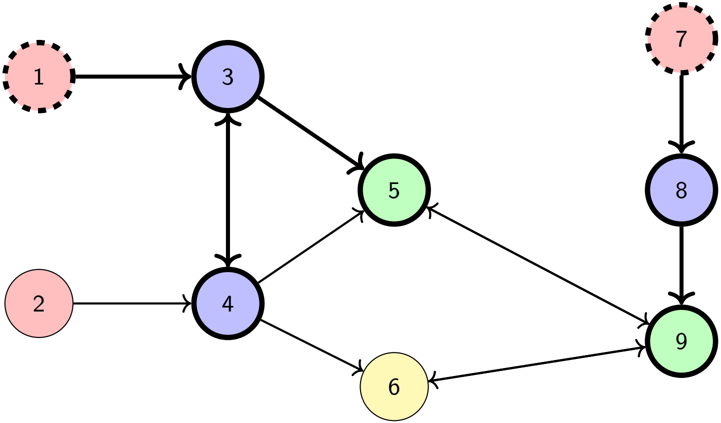

Node 1 and 7 have non-isomorphic 2-hop neighborhoods. We say that 1 and 7 are 1-hop indistinguishable

k-hop indistinguishability and k-hop-fair maps

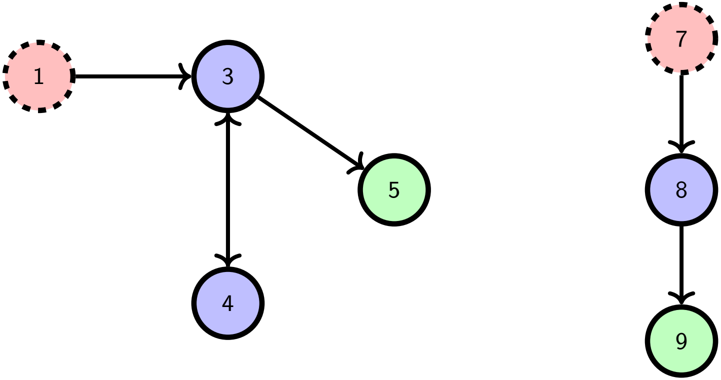

Node 1 and 7 have non-isomorphic 2-hop neighborhoods.

Invariant Outputs and Universality for MP-GNN

To obtain a global (invariant) network output $f(G) \in \mathbb{R}^d$, a *permutation-invariant readout* is applied, e.g.:

$$f(G) = \rho\left(\bigoplus_{i=1}^{n} \gamma_i^{(K)}\right), \qquad \bigoplus \in \{\text{sum, mean, maximum}\}.$$

Universal approximation

- Standard **k-layer-GNNs** (message passing) can represent only a subset of $\mathcal{F}_n^{\mathrm{inv}}$ and $\mathcal{F}_n^{\mathrm{eq}}$, limited by the **1-Weisfeiler–Lehman (1-WL)** isomorphism test.
- As our results reveal, partly this is due to the additional property of **k-layer-GNNs** being k-hop fair.

Literature: Maron et al. (2019), Keriven and Peyré (2019); Xu et al. (2019).

Contents

Architecture for Actor-critic learning

Specification of the GNN-MPN model:

- Two-layer message-passing network (MPN) used as a feature extractor with a 5-dimensional feature representation. The feature extractor is shared by actor and the critic.
- For the **critic** $\hat{V}_\phi \in \mathcal{F}_n^{\mathrm{inv}}$: an additional feedforward layer with output dimension 1 on top of the MPN representation.
- For the **actor** $\pi_\theta \in \mathcal{F}_n^{\mathrm{eq}}$: two parallel feedforward layers, one for the injected capital, one for the node weights (which share each node receives). Node weights are 2-hop fair.
- Action selection is performed via a **softmax** function.
- The risky asset is modeled by an **Ornstein-Uhlenbeck process**.

Architektur

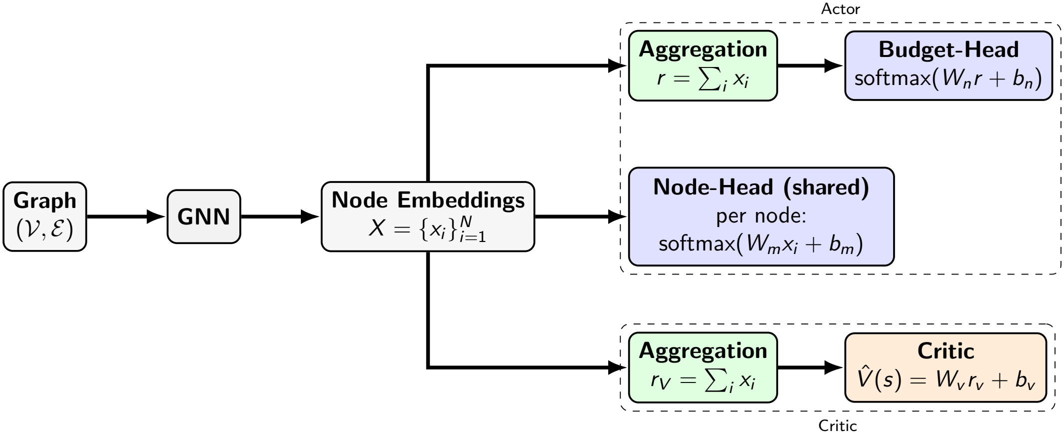

Architecture Dependent Lipschitz Bounds

Robustness of GNN Policies

- The GNN encoder, policy, and policy-gradient estimator are Lipschitz continuous with respect to: network parameters, node features, graph structure.
- Explicit bounds are derived in terms of: graph connectivity, aggregation operator, network depth.
- Deeper GNNs lead to exponentially larger Lipschitz constants.

Convergence guarantees

Consider the projected batch REINFORCE update on a compact hyperrectangle Θ :

$$\theta_{k+1} = \Pi_\Theta\big(\theta_k + \frac{\alpha_k}{m}\sum_{l=1}^{m}\nabla\widehat{J}_l(\theta_k)\big).$$

Theorem 1 (Almost sure convergence)

Assume that the step sizes satisfy

$$\alpha_k > 0, \qquad \alpha_k \xrightarrow[k\to\infty]{} 0, \qquad \sum_{k=0}^{\infty}\alpha_k = \infty, \qquad \sum_{k=0}^{\infty}\alpha_k^2 < \infty.$$

If

$$S_\Theta := \Big\{\theta \in \Theta : \Pi_{T_\Theta(\theta)}\big(\nabla J(\theta)\big) = 0\Big\}$$

has empty interior, then

$$\operatorname{dist}(\theta_k, S_\Theta) \xrightarrow[k\to\infty]{} 0 \qquad \textit{almost surely.}$$

Convergence guarantees

Consider the projected batch REINFORCE update on a compact hyperrectangle Θ :

$$\theta_{k+1} = \Pi_\Theta\big(\theta_k + \frac{\alpha_k}{m}\sum_{l=1}^{m}\nabla\widehat{J}_l(\theta_k)\big).$$

Theorem 2 (Finite-time result for N steps)

Assume $0 < \alpha_k < \frac{2}{L_J}$ *and define for* $k = 0, \ldots, N-1$

$$w_k := \alpha_k - \frac{L_J}{2}\alpha_k^2, \qquad W_N := \sum_{k=0}^{N-1} w_k, \quad \mathbb{P}(\tau = k) = \frac{w_k}{W_N}.$$

Then the expected projected gradient norm at the random iterate $\theta_{\tau+1}$ *satisfies*

$$\Big[\big\|\Pi_{T_\Theta(\theta_{\tau+1})}\big(\nabla J(\theta_{\tau+1})\big)\big\|^2\Big] \leq \mathcal{O}\left(\frac{1}{\sum_{k=0}^{N-1}\alpha_k}\right) + \mathcal{O}\left(\frac{1}{m}\right),$$

where the expectation is taken with respect to τ *and the rollout randomness used to generate the iterates.*

Training

Training setup:

- Training over 40,000 epochs with 100 time steps each using a **Proximal Policy Optimization (PPO)** algorithm with a shared feature extractor for actor and critic.
- After every 19,840 time steps, 10 update steps for the policy and critic are performed.

Some numerical results

A circle with one weak node.

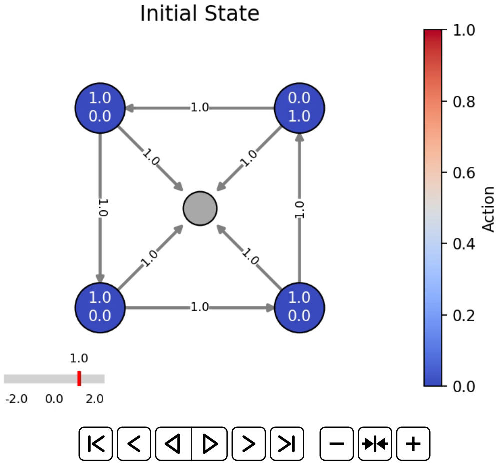

Some numerical results

A star shaped network with one a weak core.

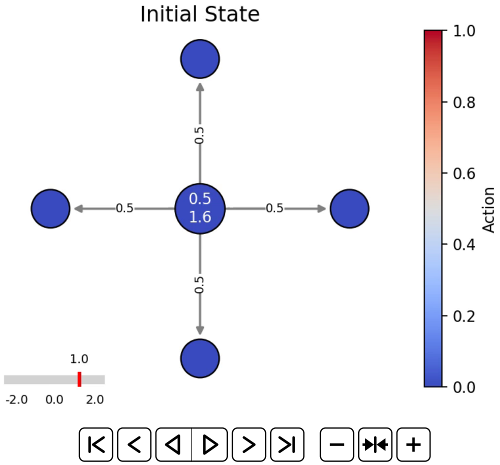

Some numerical results

A network with two symmetric circles.

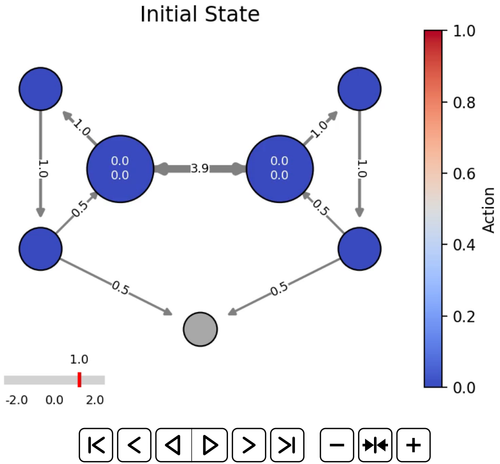

Some numerical results

A network with a stressed community.

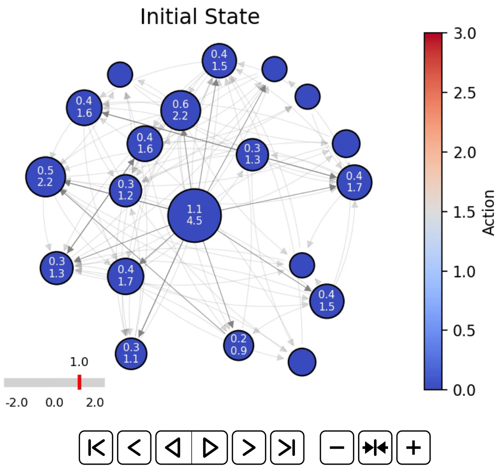

Thank you!

Nils Detering
Department of Mathematics
Heinrich-Heine Universität, Düsseldorf
Germany

nils.detering@hhu.de

References I

Hamed Amini, Rama Cont, and Andreea Minca. Resilience to Contagion in Financial Networks. *Mathematical Finance*, 26(2):329–365, 2016.

Hamed Amini, Andreea Minca, and Agnès Sulem. Optimal equity infusions in interbank networks. *Journal of Financial Stability*, 31:1–17, 2017. ISSN 1572-3089. doi: https://doi.org/10.1016/j.jfs.2017.05.008. URL `https://www.sciencedirect.com/science/article/pii/S1572308917303406`.

Maxim Bichuch and Ke Chen. Systemic risk: the effect of market confidence. *International Journal of Theoretical and Applied Finance*, 23(07):2050043, 2020.

Maxim Bichuch and Zachary Feinstein. A repo model of fire sales with vwap and lob pricing mechanisms. *European Journal of Operational Research*, 296(1):353–367, 2022.

Agostino Capponi and Peng-Chu Chen. Systemic risk mitigation in financial networks. *Journal of Economic Dynamics and Control*, 58:152–166, 2015. ISSN 0165-1889. doi: https://doi.org/10.1016/j.jedc.2015.06.008. URL `https://www.sciencedirect.com/science/article/pii/S0165188915001177`.

Agostino Capponi, Peng-Chu Chen, and David D. Yao. Liability concentration and systemic losses in financial networks. *Operations Research*, 64(5):1121–1134, 2016.

Agostino Capponi, Felix Corell, and Joseph E. Stiglitz. Optimal bailouts and the doom loop with a financial network. *Journal of Monetary Economics*, 128:35–50, 2022. ISSN 0304-3932. doi: https://doi.org/10.1016/j.jmoneco.2022.03.004. URL `https://www.sciencedirect.com/science/article/pii/S0304393222000356`.

References II

Rene Carmona, Jean-Pierre Fouque, and Li-Hsien Sun. Mean field games and systemic risk, 2013. URL https://arxiv.org/abs/1308.2172.

René Carmona, Jean-Pierre Fouque, Seyyed Mostafa Mousavi, and Li-Hsien Sun. Systemic risk and stochastic games with delay. 2016. Working Paper.

Rodrigo Cifuentes, Hyun Song Shin, and Gianluigi Ferrucci. Liquidity risk and contagion. *Journal of the European Economic Association*, 3(2-3):556–566, 2005.

Nils Detering, Thilo Meyer-Brandis, Konstantinos Panagiotou, and Daniel Ritter. Managing default contagion in inhomogeneous financial networks. *SIAM Journal on Financial Mathematics*, 10(2):578–614, 2019.

Larry Eisenberg and Thomas H. Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001.

Jean-Pierre Fouque and Tomoyuki Ichiba. Stability in a model of interbank lending. *SIAM Journal on Financial Mathematics*, 4:784–803, 2013.

P. Gai and S. Kapadia. Contagion in Financial Networks. *Proc. R. Soc. A*, pages 2401–2423, 2010.

Josselin Garnier, George Papanicolaou, and Tzu-Wei Yang. Diversification in financial networks may increase systemic risk. In *Handbook on Systemic Risk*, pages 432–443. Cambridge University Press, 2013.

Christian Gouriéroux, J.-C. Héam, and Alain Monfort. Bilateral exposures and systemic solvency risk. *Canadian Journal of Economics*, 45(4):1273–1309, 2012.

References III

Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/ea9268cb43f55d1d12380fb6ea5bf572-Paper.pdf.

Yang-Yu Liu, Jean-Jacques Slotine, and Albert-László Barabási. Controllability of complex networks. *Nature*, 473 (7346):167–173, 2011. doi: 10.1038/nature10011. URL https://doi.org/10.1038/nature10011.

Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/bb04af0f7ecaee4aae62035497da1387-Paper.pdf.

Andreea Minca and Agnès Sulem. *Statistics & Risk Modeling*, 31(1):23–48, 2014. doi: doi:10.1515/strm-2013-1165. URL https://doi.org/10.1515/strm-2013-1165.

Erland Nier, Jing Yang, Tanju Yorulmazer, and Amadeo Alentorn. Network models and financial stability. *Journal of Economic Dynamics and Control*, 31(6):2033–2060, 2007.

Marios Papachristou, Siddhartha Banerjee, and Jon Kleinberg. Dynamic interventions for networked contagions. In *Proceedings of the ACM Web Conference 2023*, WWW '23, pages 3519–3529, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394161. doi: 10.1145/3543507.3583470. URL https://doi.org/10.1145/3543507.3583470.

References IV

Leonard C.G. Rogers and Luitgard A.M. Veraart. Failure and rescue in an interbank network. *Management Science*, 59(4):882–898, 2013.

Luitgard Anna Maria Veraart. When does portfolio compression reduce systemic risk? *Mathematical Finance*, 32 (3):727–778, 2022. doi: https://doi.org/10.1111/mafi.12346. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/mafi.12346.

Stefan Weber and Kerstin Weske. The joint impact of bankruptcy costs, fire sales and cross-holdings on systemic risk in financial networks. *Probability, Uncertainty and Quantitative Risk*, 2(1):9, June 2017a.

Stefan Weber and Kerstin Weske. The joint impact of bankruptcy costs, fire sales and cross-holdings on systemic risk in financial networks. *Probability, Uncertainty and Quantitative Risk*, 2(1):9, 2017b. doi: 10.1186/s41546-017-0020-9. URL https://doi.org/10.1186/s41546-017-0020-9.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=ryGs6iA5Km.