informs
https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Articles in Advance, pp. 1–16
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Contextual Areas

Incentive-Aware Models of Financial Networks

Akhil Jalan,[a,]* Deepayan Chakrabarti,[b] Purnamrita Sarkar[c]

[a] Department of Computer Science, University of Texas at Austin, Austin, Texas 78712; [b] McCombs School of Business, University of Texas at Austin, Austin, Texas 78712; [c] Department of Statistics and Data Science, University of Texas at Austin, Austin, Texas 78712

*Corresponding author

Contact: akhiljalan@utexas.edu, https://orcid.org/0000-0003-0148-082X (AJ); deepayan.chakrabarti@mccombs.utexas.edu, https://orcid.org/0000-0002-3863-4928 (DC); purna.sarkar@austin.utexas.edu, https://orcid.org/0000-0002-6542-1215 (PS)

Received: December 22, 2022
Revised: October 14, 2023; January 28, 2024; April 3, 2024; May 26, 2024
Accepted: June 8, 2024
Published Online in Articles in Advance: August 12, 2024

Area of Review: Financial Engineering

https://doi.org/10.1287/opre.2022.0678

Abstract. Financial networks help firms manage risk but also enable financial shocks to spread. Despite their importance, existing models of financial networks have several limitations. Prior works often consider a static network with a simple structure (e.g., a ring) or a model that assumes conditional independence between edges. We propose a new model where the network emerges from interactions between heterogeneous utility-maximizing firms. Edges correspond to contract agreements between pairs of firms, with the contract size being the edge weight. We show that, almost always, there is a unique "stable network." All edge weights in this stable network depend on all firms' beliefs. Furthermore, firms can find the stable network via iterative pairwise negotiations. When beliefs change, the stable network changes. We show that under realistic settings, a regulator cannot pin down the changed beliefs that caused the network changes. Also, each firm can use its view of the network to inform its beliefs. For instance, it can detect outlier firms whose beliefs deviate from their peers. However, it cannot identify the deviant belief: Increased risk-seeking is indistinguishable from increased expected profits. Seemingly minor news may settle the dilemma, triggering significant changes in the network.

Funding: This work was supported by the National Science Foundation [Grants 2217069, 2019844, and DMS 2109155], McCombs Research Excellence Grants, and a Dell Faculty Award.

Supplemental Material: The computer code and data that support the findings of this study are available within this article's supplemental material at https://doi.org/10.1287/opre.2022.0678.

Keywords: financial networks • utility maximization • heterogeneous agents • dynamic games

1. Introduction

The financial crisis of 2008 showed the need for mitigating systemic risks in the financial system. There has been much recent work on categorizing such risks (Elliott et al. 2014; Glasserman and Young 2015, 2016; Birge 2021; Jackson and Pernoud 2021). Although the causes of systemic risk are varied, they often share one feature. This shared feature is the network of interconnections between firms via which problems at one firm spread to others. One example is the weighted directed network of debt between firms. If one firm defaults on its debt, its creditors suffer losses. Some creditors may be forced into default, triggering a default cascade (Eisenberg and Noe 2001). Another example is the implicit network between firms holding similar assets. Sales by one firm can lead to mark-to-market valuation losses at other firms. These can snowball into fire sales (Caballero and Simsek 2013, Cont and Minca 2016, Feinstein 2020, Feinstein and Søjmark 2021).

The structure of interfirm networks plays a vital role in the financial system. Small changes in network structure can lead to jumps in credit spreads in over-the-counter (OTC) markets (Eisfeldt et al. 2021). Network density, diversification, and interfirm cross-holdings can affect how robust the networks are to shocks and how such shocks propagate (Elliott et al. 2014, Acemoglu et al. 2015). The network structure also affects the design of regulatory interventions (Amini et al. 2015, Galeotti et al. 2020, Calafiore et al. 2022, Papachristou and Kleinberg 2022).

Despite its importance, many prior works use simplistic descriptions of the network structure. For instance, they often assume that the network is fixed and observable. However, only regulators may have access to the entire network. Furthermore, shocks or regulatory interventions can change the network. Others assume that the network belongs to a general class. For instance, Caballero and Simsek (2013) assume a ring network between banks. Amini et al. (2015) derive tractable optimal interventions for core-periphery networks. However, financial networks can exhibit complex structure (Peltonen et al. 2014, Eisfeldt et al. 2021). Leverage levels, size heterogeneity, and other factors can affect the network topology (Glasserman and Young 2016). Hence, there is a need for models to help reason about financial networks.

In this paper, we design a model for a weighted network of contracts between agents, such as firms, countries, or individuals. The contracts can be arbitrary, and the edge weights denote contract sizes. In designing the model, we have two main desiderata. First, the model must account for heterogeneity between firms. This follows from empirical observations that differences in dealer characteristics lead to different trade risk exposures in OTC markets (Eisfeldt et al. 2021). Second, each firm seeks to maximize its utility and selects its contract sizes accordingly. In effect, each firm tries to optimize its portfolio of contracts (Markowitz 1952). The model must reflect this behavior. From this starting point, we ask the following questions:

1. How does a network emerge from interactions between heterogeneous utility-maximizing firms?
2. How does the network respond to regulatory interventions?
3. How can the network structure inform the beliefs that firms hold about each other?

Next, we review the relevant literature.

1.1. Imputing Financial Networks

We often have only partial information about the structure of a financial network. For example, we may know the total liability of each bank in a network. From this, we want to reconstruct all the interbank liabilities (Squartini et al. 2018). One approach is to pick the network that minimizes the Kullback-Leibler divergence from a given input matrix (Upper and Worms 2004). Mastromatteo et al. (2012) use message-passing algorithms, whereas Gandy and Veraart (2017) use a Bayesian approach. However, such random graph models often do not reflect the sparsity and power-law degree distributions of financial networks (Upper 2011). Furthermore, these models do not account for the utility-maximizing behavior of firms.

1.2. General-Purpose Network Models

The simplest and most well-explored network model is the random graph model (Erdös and Rényi 1959, Gilbert 1959). Here, every pair of nodes is linked independently with probability p. Generalizations of this model allow for different degree distributions and edge directionality (Aiello et al. 2000). Exponential random graph models remove the need for independence, but parameter estimation is costly (Frank and Strauss 1986, Wasserman and Pattison 1996, Hunter and Handcock 2006, Caimo and Friel 2011). Several models add node-specific latent variables to model the heterogeneity of nodes. For example, in the stochastic blockmodel and its variants, nodes are members of various latent communities. The community affiliations of two nodes determine their probability of linkage (Holland et al. 1983, Chakrabarti et al. 2004, Airoldi et al. 2008, Mao et al. 2018). Instead of latent communities, Hoff et al. (2002) assign a latent location to each node. Here, the probability of an edge depends on the distance between their locations.

All the latent variable models assume conditional independence of edges given the latent variables. However, in financial networks, contracts between firms are not independent. Two firms will sign a contract only if the marginal benefit of the new contract is higher than the cost. This cost/benefit tradeoff depends on all other contracts signed with other firms. Unlike our model, existing general-purpose models do not account for such utility-maximization behavior.

1.3. Network Games

Here, the payoffs of nodes are dependent on the actions of their neighbors (Tardos 2004). One well-studied class of network games is linear-quadratic games, with linear dynamics and quadratic payoff functions. Prior work has explored the stability of Nash equilibria (Guo and De Persis 2021) and algorithms to learn the agents' payoff functions (Leng et al. 2020). However, our model does not yield a linear-quadratic game except in exceptional cases. Instead, our process involves nonlinear rational functions of the beliefs of firms. Thus, our setting differs from linear-quadratic games. Recently, network games have been extended to settings where the number of players tends to infinity (Carmona et al. 2022). However, we only consider finite networks.

1.4. Games to Form Networks

Several works study the stability of networks. In a pairwise-stable network, no pair of agents want to form or sever edges. This may be achieved via side-payments between agents, which our model also uses (Jackson and Wolinsky 2003). Pairwise stability has been extended to strong stability for networks (Jackson and Van den Nouweland 2005), and to weighted networks with edge weights in $[0, 1]$ (Bich and Morhaim 2020, Bich and Teteryatnikova 2023). We introduce an analogous notion called higher-order Nash stability against any deviating coalition. However, the weights in our network are not bounded in $[0, 1]$ and can be negative. Furthermore, our edge weights denote contract size, requiring agreement from both parties. In contrast, prior works typically interpret edge weights as the engagement level in an ongoing interaction.

Sadler and Golub (2021) study a network game with endogenous network formation, whose stable points are both pairwise stable and Nash equilibria. We show similar results for our model. However, they consider unweighted networks and focus on the case of separable games. In our setting, this corresponds to the case where all firms are uncorrelated. However, in financial networks, correlations are widespread and help firms diversify their contracts.

Several authors study the effect of exogenous inputs on production networks (Herskovic 2018, Elliott et al. 2022). Acemoglu and Azar (2020) also model endogenous network formation but differ from our approach. Prices in their model equal the minimum unit cost of production. For us, prices are determined by pairwise negotiations between firms. Also, each firm in their model only considers a discrete set of choices among possible suppliers. In our model, firms can choose both their counterparties and the contract sizes.

1.5. Risk-Sharing and Exchange Economies

The pricing of risk is a well-studied problem (Arrow and Debreu 1954; Bühlmann 1980, 1984; Tsanakas and Christofides 2006; Banerjee and Feinstein 2022). Most models typically price risk via a global market. However, in our model, all contracts are pairwise, and the contract terms and payments between a buyer and seller are bespoke. There is no global contract or global market price. Because contracts are pairwise, each firm under our model must consider counterparty risks and the correlations between them. A firm i may make large payments and accept a *negative* reward for a contract with firm j to diversify the risk from contracts with other firms. Finally, in our model, agents can hedge their risk by betting against one another. In contrast, Bühlmann equilibria always result in comonotonic endowments, which firms cannot use as hedges for each other (Yaari 1987, Banerjee and Feinstein 2022).

1.6. Network Valuation Adjustment

Some recent works price the risk due to exposure to the entire financial network (Banerjee and Feinstein 2022, Feinstein and Søjmark 2022). The network is usually treated as exogenous and fully known to all firms. In contrast, we consider endogenous network formation resulting from pairwise interactions between firms. The network valuation algorithm of Barucca et al. (2020) works with incomplete information but is not designed for network formation, and it needs firms to share information not required to form their contracts.

1.7. Properties of Equilibria

Another line of work considers the efficiency or social welfare of equilibria (Jackson and Pernoud 2021, Elliott and Golub 2022). Galeotti et al. (2020) show that welfare-maximizing interventions rely mainly on the top or bottom eigenvectors of the network. Elliott et al. (2022) show an efficiency-stability tradeoff for their model of supply network formation. Like prior work, we show that stable equilibria exist and are nondominated. However, our emphasis is on potentially valuable insights for regulators and firms. For instance, we show a negative result about the ability of regulators to infer the causes of changes to the network structure. The linkage between firms' utilities and their beliefs and its effect on stability is not considered in prior work.

1.8. Our Contributions

We develop a new network model of contracts between heterogeneous agents, such as firms, countries, or individuals. Each agent aims to maximize a mean-variance utility parametrized by its beliefs. But for two agents to sign a contract, both must agree to the contract size. For a stable network, all agents must agree to all their contracts. We show that such constraints are solvable by allowing agents to pay each other. By choosing prices appropriately, every agent maximizes its utility in a stable network.

1.8.1. Characterization of Stable Networks (Section 2). We show that unique stable networks exist for almost all choices of agents' beliefs. These networks are robust against actions by cartels, a condition that we call higher-order Nash stability. The agents can also converge to the stable network via iterative pairwise negotiations. The convergence is exponential in the number of iterations. Hence, the stable network can be found quickly. Finally, we show how to infer the agents' beliefs by observing network snapshots over time, under certain conditions.

1.8.2. Limits of Regulation (Section 3). A financial regulator can observe the entire network but not the agents' beliefs. Suppose firm i changes its beliefs about firm j. Then the contract size between i and j will change. Indirectly, other contracts will change too. We show empirically that in realistic settings, the indirect effects can be as significant as the direct effects. In such cases, the regulator cannot infer the underlying cause of changes in the network. Similarly, suppose the regulator intervenes with one firm, affecting its beliefs. The resulting network changes need not be localized to that firm's neighborhood in the network. Thus, targeted interventions can have strong ripple effects. Broad-based interventions aimed at increasing stability can also have adverse effects. For instance, increasing margin requirements on contracts may even increase some contract sizes.

1.8.3. Outlier Detection by Firms (Section 4). A firm i can observe its contracts with counterparties but not the entire network. Suppose another firm j (say, a real estate firm) has beliefs that are very different from its peers. Then, we prove that under certain conditions, j's contract size with i is also an outlier compared with other real-estate firms. Therefore, firm i can use the network to detect outliers and update its beliefs. However, suppose all real estate firms change their beliefs. This changes all their contract sizes without creating outliers. We show that i cannot determine the cause of this change. For example, firm i would observe the same change whether all real estate firms had become more risk seeking or profitable. However, firm i may want to increase its exposure if they are more profitable but reduce exposure if they are more risk seeking. Because the data cannot identify the proper action, firm i remains uncertain. Exogenous, seemingly insignificant information may persuade firm i one way or another. Thus, minor news may trigger drastic changes in the network.

1.9. Notation

We use lowercase letters, with or without subscripts, to denote scalars (e.g., c, γ_i). Lowercase bold letters denote vectors ($\boldsymbol{\mu}_i, \boldsymbol{w}$), and uppercase letters denote matrices (W, P, Σ_i). We use $\boldsymbol{\mu}_{i;j}$ to refer to the jth component of the vector $\boldsymbol{\mu}_i$, and $\Sigma_{i;jk}$ for the (j, k) cell of matrix Σ_i. We use $\boldsymbol{v}^T$ to denote the transpose of a vector $\boldsymbol{v}$, and $\|\cdot\|_p$ to denote the ℓ_p norm of a vector or matrix. We say $A \succeq 0$ if A is positive semidefinite, $A \succ 0$ if it is positive definite, and $A \succeq B$ if $A - B \succeq 0$. The vectors $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n$ denote the standard basis in $\mathbb{R}^n$, and I_n is the $n \times n$ identity matrix. If $A \in \mathbb{R}^{m\times n}, B \in \mathbb{R}^{p\times q}$ then $A \otimes B \in \mathbb{R}^{mp\times nq}$ denotes their tensor product: $(A \otimes B)_{ij,k\ell} = A_{ik}B_{j\ell}$. For an appropriate matrix M, $\mathrm{tr}(M)$ calculates its trace, $\mathrm{vec}(M)$ vectorizes M by stacking its columns into a single vector, and $\mathrm{uvec}(M)$ vectorizes the upper-triangular off-diagonal entries of M. For an integer $r \geq 1$, we use $[r]$ to denote the set of integers $[r] := \{1, 2, \ldots, r\}$.

2. Proposed Model

We consider a *weighted* network $W \in \mathbb{R}^{n\times n}$ between n agents (such as firms, countries, or individuals). The element W_{ij} represents the size of a contract between agents i and j. We make no assumptions about the content of the contract. For instance, the contract could be an interest rate swap, a stock swap, or an insurance contract. We assume that each pair of firms can form a contract of a standard type and negotiate only on the contract size and price (discussed below). Because contracts need mutual agreement, $W_{ij} = W_{ji}$. We take W_{ii} to represent i's investment in itself. A negative contract ($W_{ij} = W_{ji} < 0$) is a valid contract that reverses the content of a positive contract. For example, if a positive contract is a derivative trade between two firms, the negative contract swaps the roles of the two firms.

Let $\boldsymbol{w}_i$ denote the ith column of W (i.e., $\boldsymbol{w}_{i;j} = W_{ji}$ for all j). Each agent i would prefer to set its contract sizes $\boldsymbol{w}_i$ to maximize its utility. However, other agents will typically have different preferences. Therefore, to achieve an agreement about the contract size W_{ij}, agents i and j can agree on a price for the contract. For example, i may agree to pay j an amount $P_{ji} \cdot W_{ji}$ in cash at the beginning of the contract. Because payments are zero-sum and $W_{ji} = W_{ij}$, we must have $P_{ji} = -P_{ij}$. We do not model how firms raise funds to pay the price.

Each contract yields a stochastic payout, and agents have beliefs about these payouts. We represent agent i's beliefs by a vector $\boldsymbol{\mu}_i$ of expected returns and a covariance matrix $\Sigma_i \succ 0$. Thus, Σ_i represents firm i's perceived risk of trading with other firms, and includes both contract-specific risk and counterparty risk. We do *not* assume that the contracts are zero-sum or that the beliefs are correct, even approximately. Thus, the overall expected return from all contracts of i is $\boldsymbol{w}_i^T(\boldsymbol{\mu}_i - P\boldsymbol{e}_i)$, and the variance of the overall return is $\boldsymbol{w}_i^T \Sigma_i \boldsymbol{w}_i$. We assume that each agent has a mean-variance utility (Markowitz 1952):

$$\text{agent } i\text{'s utility} \quad g_i(W, P) := \boldsymbol{w}_i^T(\boldsymbol{\mu}_i - P\boldsymbol{e}_i) - \gamma_i \cdot \boldsymbol{w}_i^T \Sigma_i \boldsymbol{w}_i, \tag{1}$$

where $\gamma_i > 0$ is a risk-aversion parameter. In practice, we expect the set $\{\gamma_i\}_{i\in[n]}$ to be not too heterogeneous (Metrick 1995, Kimball et al. 2008, Ang 2014, Paravisini et al. 2017). Equation (1) ignores costs for contract formation; we will consider these in Section 3.1. Also, we assume that P_{ji} does not change the perceived risk.

Example 1 (Insurance Contract). Suppose firm i buys fire insurance from insurer j. Then, $\boldsymbol{\mu}_{i;j}$ is the buyer's expected insurance payout minus the insurance premium. The expected payout depends on the probability of a fire, for which the buyer and insurer may have different estimates. Also, the insurance contract is negatively correlated with the buyer's other contracts (reflected in Σ_i). This is because the buyer gains a payout from the insurer in case of a fire but incurs losses on other contracts. Hence, the buyer i may be willing to accept a contract with negative expected reward, and even pay a higher-than-usual premium P_{ji} per contract.

Example 2 (Interest Rate Swap Contract). Suppose firm i makes fixed-rate payments to firm j, and receives floating-rate payments in return. Then, $\boldsymbol{\mu}_{i;j}$ is the expected net present value of these payments for i from a standard unit-sized contract. This value depends on i's forecast of future interest rates and need for floating-rate income, for example, to match future liabilities. Hence, it may be quite different from $\boldsymbol{\mu}_{j;i}$. Also, the firms agree to a price $P_{ij} = -P_{ji}$ per contract. If $P_{ij} > 0$, then firm j must pay firm i the price $P_{ij} \cdot W_{ij}$; if $P_{ij} < 0$, then firm i makes the payment.

Example 3 (Loan Contract). Suppose borrower i takes a loan of size W_{ij} from lender j. Then, $\boldsymbol{\mu}_{j;i} \cdot W_{ij}$ represents the lender j's expected value for this loan. The expected value depends on the repayment schedule, the collateral, j's estimate of the probability of default, the recovery rate in case of default, etc. The borrower's expected value $\boldsymbol{\mu}_{i;j} \cdot W_{ij}$ depends on the planned use of this loan. For example, if the borrower wants the loan to purchase equipment, $\boldsymbol{\mu}_{i;j}$ is the net present value of expected extra profits due to that equipment. Hence, $\boldsymbol{\mu}_{i;j}$ may not be a function of $\boldsymbol{\mu}_{j;i}$. Now, the borrower and lender must settle on a contract price to reach an agreement on the contract size. If the standard loan contract requires the lender to give cash to the borrower at the beginning of the contract, this loan amount can be adjusted for the price. Otherwise, if the borrower firm needs to pay the price, it must arrange a separate bridge loan.

The model above allows contracts between all pairs of agents. However, some edges may be prohibited due to logistical or legal reasons. For each agent i, let $J_i \subseteq [n]$ denote the ordered set of agents with whom i can form an edge. Therefore, if $k \notin J_i$ (and hence $i \notin J_k$), we have $W_{ik} = W_{ki} = P_{ik} = P_{ki} = 0$. Similarly, if $i \notin J_i$, then self-loops are prohibited ($W_{ii} = P_{ii} = 0$). We will encode these constraints in the binary matrix $\Psi_i \in \mathbb{R}^{|J_i| \times n}$ where $\Psi_{i;jk} = 1$ if k is the jth element of J_i, and $\Psi_{i;jk} = 0$ otherwise. In other words, Ψ_i is obtained from I_n by deleting the rows corresponding to the prohibited counterparties of i. Thus, for any $v \in \mathbb{R}^n$, $\Psi_i v$ selects the elements of v corresponding to J_i. If all edges are allowed, we have $\Psi_i = I_n$ for all i.

Definition 1 (Network Setting). A *network setting* $(\boldsymbol{\mu}_i, \gamma_i, \Sigma_i, \Psi_i)_{i \in [n]}$ captures the beliefs and constraints of n agents. When there are no constraints (i.e., all edges are allowed), we drop the $\Psi_i = I_n$ terms to simplify the exposition. Finally, we will use $M \in \mathbb{R}^{n \times n}$ to denote a matrix whose ith column is $\boldsymbol{\mu}_i$, and Γ to denote a diagonal matrix with $\Gamma_{ii} = \gamma_i$.

2.1. Characterizing Stable Points

In the above model, every agent tries to optimize its own utility (Equation (1)). We now characterize the conditions under which selfish utility-maximization leads to a stable network.

Definition 2 (Feasibility). A tuple (W, P) is feasible if $W = W^T$, $P = -P^T$, and W and P obey the constraints encoded in $(\Psi_i)_{i \in [n]}$.

Definition 3 (Stable Point). A feasible (W, P) is stable if each agent achieves its maximum possible utility given prices P:

$$g_i(W, P) = \max_{\text{feasible}(W', P) \text{ under } \{\Psi_i\}} g_i(W', P) \quad \forall i \in [n].$$

Example 4. Suppose we only have two firms with the following setting:

$$\text{mean beliefs } M = \begin{bmatrix} 0 & 3 \\ 1 & 4 \end{bmatrix}$$

$$\text{covariance } \Sigma_1 = \Sigma_2 = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$$

$$\text{risk aversion } \gamma_1 = \gamma_2 = 1.$$

Therefore, both firms perceive a benefit from trading ($M_{12} > 0, M_{21} > 0$). If trading is disallowed, the optimum W is diagonal with $W_{11} = 0$ and $W_{22} = 1$ (and P is the zero matrix). The corresponding utilities are zero for firm 1 and 2 for firm 2. Suppose we allow trading but do not allow pricing (Figure 1(a)). Then, the two firms can each improve their utility by trading, but achieve their optimum utilities at different contract sizes. Hence, they may be unable to agree to a contract. In Figure 1(b), firm 2 pays firm 1 a specially chosen price of 5/3 per unit contract. At this price, both firms achieve their optimum utilities at the same contract size $W_{12} = W_{21} = 2/3$. Hence, they can agree to a contract. By paying the price, firm 2 shares some of its utility with firm 1 to achieve agreement on the contract. This choice of W and P is a stable point (Figure 1(c)). The following results show that this is the *only* stable point. □

Define $Q_i = \Psi_i^T (2\gamma_i \Psi_i \Sigma_i \Psi_i^T)^{-1} \Psi_i$. When all edges are allowed, $\Psi_i = I_n$ and $Q_i = (2\gamma_i \Sigma_i)^{-1}$. Let $F = \{(i,j) : 1 \le i < j \le n, \Psi_i \boldsymbol{e}_j \ne \boldsymbol{0}\}$ denote the ordered pairs $i < j$ where P_{ij} is allowed to be nonzero. Note that $|F| \le n(n-1)/2$. For any $n \times n$ matrix X, let $\text{uvec}(X)_F \in \mathbb{R}^{|F|}$ be a vector whose entries are the ordered set $\{X_{ij} | (i,j) \in F\}$.

Theorem 1 (Existence and Uniqueness of Stable Point). *Define $n \times n$ matrices A, $B_{(i,j)}$, and $C_{(i,j)}$ as follows:*

$$A_{ij} = \boldsymbol{e}_i^T Q_j M \boldsymbol{e}_j, B_{(i,j)} = \boldsymbol{e}_i \boldsymbol{e}_j^T Q_i,$$

$$C_{(i,j)} = (B_{(i,j)} - B_{(j,i)}) - (B_{(i,j)} - B_{(j,i)})^T.$$

Let Z_F be the $|F| \times |F|$ matrix whose rows are the ordered sets $\{\text{uvec}(C_{(i,j)})_F | (i,j) \in F\}$. Then, we have the following:

1. *A stable point (W, P) under $\{\Psi_i\}$ exists if and only if $\text{uvec}(A - A^T)_F$ lies in the column space of Z_F.*
2. *If a stable point (W, P) exists, then $Z_F \text{uvec}(P)_F = \text{uvec}(A - A^T)_F$.*
3. *A unique stable point always exists if Z_F is full rank.*

Theorem 1 is proved in the appendix, Section A.1. When the Σ_i are random variables, we give a simple

Figure 1. Example of a Stable Point for a Borrower (Firm 1) and a Lender (Firm 2)

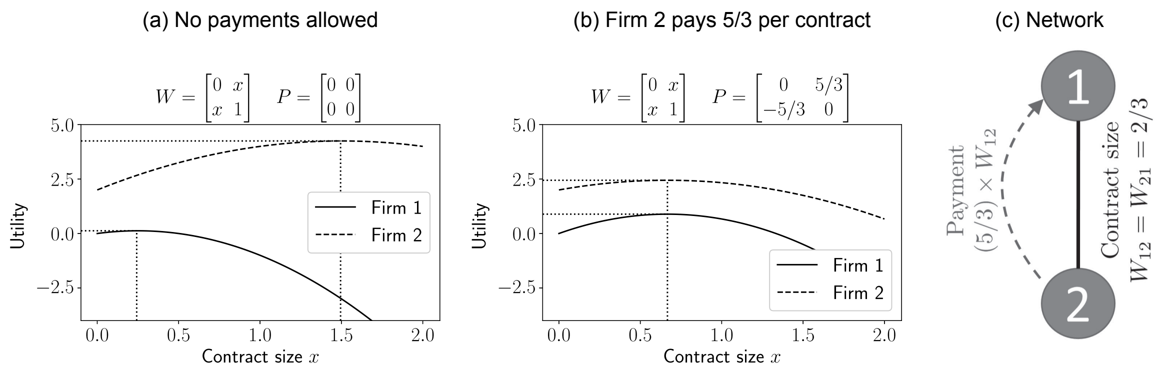

Notes. (a) When the borrower cannot pay the lender an additional payment, the firms may be unable to agree to a contract, even if trading improves their utilities. (b) By allowing for contract-specific payments, both firms can agree on a contract size. In effect, the borrower (Firm 2) shares its utility with the lender (Firm 1) to achieve agreement. (c) The stable network is shown.

sufficient condition that a stable point exists and is unique with probability 1 (see Sections A.1 and A.2 in the supplemental material). Also, the appendix, Section 8.2, provides closed-form formulas for the stable point when all agents have the same covariance ($\Sigma_i = \Sigma$ for all $i \in [n]$). This occurs when the risk of a contract is primarily counterparty risk (so $\Sigma_{i;jk}$ depends on j and k, not i), and there is reliable public data on such risks (say, via credit rating agencies).

Next, we consider some properties of the stable point. For two feasible tuples (W_1, P_1) and (W_2, P_2), let (W_2, P_2) *dominate* (W_1, P_1) if for all $i \in [n], g_i(W_1, P_1) \leq g_i(W_2, P_2)$, with at least one inequality being strict.

Theorem 2 (Stable Points Cannot Be Dominated). *Suppose a stable point (W, P) exists. Then, there is no feasible* (W', P') *that dominates (W, P).*

The proofs of Theorem 2 and all subsequent claims are provided in the supplemental material.

The stable point obeys a strong form of robustness that we call *higher-order Nash stability*. This strengthens the notions of *pairwise stability* (Hellmann 2013) and *pairwise Nash* (Calvó-Armengol and Ilkiliç 2009, Sadler and Golub 2021) by allowing for agent coalitions, instead of just considering pairs of agents. It is also closely related to the concept of *strong Nash equilibrium*, which strengthens Nash equilibrium by requiring that no subset of agents can deviate at equilibrium without at least one agent being worse off (Mazalov and Chirkova 2019).

Definition 4 (Agent Action). At a given feasible point (W, P), an "action" by agent i is the ordered set $(w'_{i,j}, p'_{i,j})_{j \in J_i}$, where $J_i \subseteq [n]$ is the set of permissible edges for agent i. The action represents a set of proposed changes to i's existing contracts. Each agent $j \in J_i$ responds as follows:

1. If the new (w'_{ij}, p'_{ij}) raises j's utility, then j agrees to the revised contract and price.

2. Otherwise, i must either keep the existing contract or cancel it ($w_{ij} = p_{ij} = 0$). We assume that i cancels the contract if and only if this strictly increases i's utility.

We call the shifted (W', P') the *resulting network*.

Definition 5 (Higher-Order Nash Stability). A feasible (W, P) is higher-order Nash stable if:

1. *Nash equilibrium*: No agent i has an action such that the resulting network (W', P') is strictly better for i.

2. *Cartel robustness*: For any proper subset $S \subset [n]$ of agents, there is no feasible point (W', P') that differs from (W, P) only for indices $\{i, j\}$ with $i \in S, j \in S$ such that all agents in S have higher utility under (W', P') than (W, P).

Theorem 3 (Higher-Order Nash Stability). *Any stable point (W, P) is higher-order Nash stable.*

2.2. Finding the Stable Point via Pairwise Negotiations

To compute the stable point in Theorem 1, we must know the beliefs of all agents. However, in practice, contracts are set iteratively by negotiations among pairs of agents. We will now formalize the process of pairwise negotiations and characterize the conditions under which such negotiations can converge to the stable point.

We propose a multiround pairwise negotiation process. In round $t+1$, every pair of agents i and j update the price $P_{ij}(t)$ to $P_{ij}(t+1)$ (and hence $P_{ji}(t)$ to $P_{ji}(t+1)$) as follows. First, they agree to a price P'_{ij} between themselves, *assuming optimal contract sizes with all other agents at the current prices P(t).* In other words, we assume that the other agents will accept the prices in $P(t)$ and the contract sizes preferred by i and j. Under this condition, P'_{ij} is the price at which i's optimal contract size with j is also j's optimal size with i. We provide an explicit formula for P'_{ij} in Section A.5 of the supplemental material. All pairs of agents calculate these prices *simultaneously*.

We create a new price matrix P' from these prices. Then, we set $P(t+1) = (1-\eta)P(t) + \eta P'$, where $\eta \in (0,1)$ is a dampening factor chosen to achieve convergence. Algorithm 1 shows the details.

Algorithm 1 (Pairwise Negotiations)

procedure PAIRWISE($\eta \in (0,1)$)
  $t \leftarrow 0$
  $P(0) \leftarrow$ any skew-symmetric matrix
  while $P(t)$ has not converged **do**
    $\forall i,j \in [n], P'_{ij} \leftarrow$ pairwise-negotiated price for (i,j) (Section A.5 in the supplemental material)
    $P(t+1) \leftarrow (1-\eta)P(t) + \eta P'$
    $t \leftarrow t+1$
  end while
end procedure

Example 5 (Pairwise Negotiations for Loan Contracts). Consider a three-firm loans network containing a national bank (firm 1), local bank (firm 2), and local firm (firm 3). Suppose that the local firm cannot access the national bank, so the edge between firms 1 and 3 is prohibited. The other parameters are

$$\Sigma_1 = \Sigma_2 = \Sigma_3 = \begin{bmatrix} 1 & 0.25 & 0.75 \\ 0.25 & 1 & 0.6 \\ 0.75 & 0.6 & 1 \end{bmatrix},$$

$$M = \begin{bmatrix} 0 & 0.9 & 0.9 \\ 0.75 & 0 & 0.95 \\ 0.5 & 0.8 & 0 \end{bmatrix}, \gamma_1 = \gamma_2 = \gamma_3 = 1.$$

Figure 2 shows how pairwise negotiations via Algorithm 1 converge to the stable network.

Now, we will show that Algorithm 1 converges. First, we define *global asymptotic stability* (following Callier and Desoer 1994).

Definition 6 (Global Asymptotic Stability). The pairwise negotiation process is globally asymptotically stable for a given network setting and dampening factor η if, for any initial price matrix $P(0)$, there exists a matrix $P^\star$ such that the sequence of price matrices $P(t)$ converges to $P^\star$ in Frobenius norm: $\lim_{t\to\infty} \|P(t) - P^\star\|_F = 0$.

When pairwise negotiations are globally asymptotically stable, the limiting matrix $P^\star$ must be skew-symmetric because each $P(t)$ is skew-symmetric. Also, because prices are updated whenever two agents disagree on the size of the contract between them, all agents agree on their contract sizes at $P^\star$. Hence, $P^\star$ must be a stable point for the given network setting.

Now, we show that for a range of η, pairwise negotiations are globally asymptotically stable (Section A.7 in the supplemental material presents an example).

Theorem 4 (Convergence Conditions and Rate). *Let Q_i be defined as in Theorem 1. Define the following $n^2 \times n^2$ matrices:*

$$K := \sum_{r=1}^{n} e_r e_r^T \otimes Q_r + Q_r \otimes e_r e_r^T$$

$$L_{(i-1)n+j,(i-1)n+j} = Q_{i;j,j} + Q_{j;i,i} \quad \forall i,j \in [n]$$

$$(L \text{ is diagonal}).$$

Let $L^\dagger$ denote the pseudoinverse of L, and $(L^\dagger K)|_R$ denote the principal submatrix of $L^\dagger K$ containing the rows/columns $(i-1)n+j$ such that the edge (i, j) is not prohibited.

Figure 2. Pairwise Negotiations for the Setting of Example 5

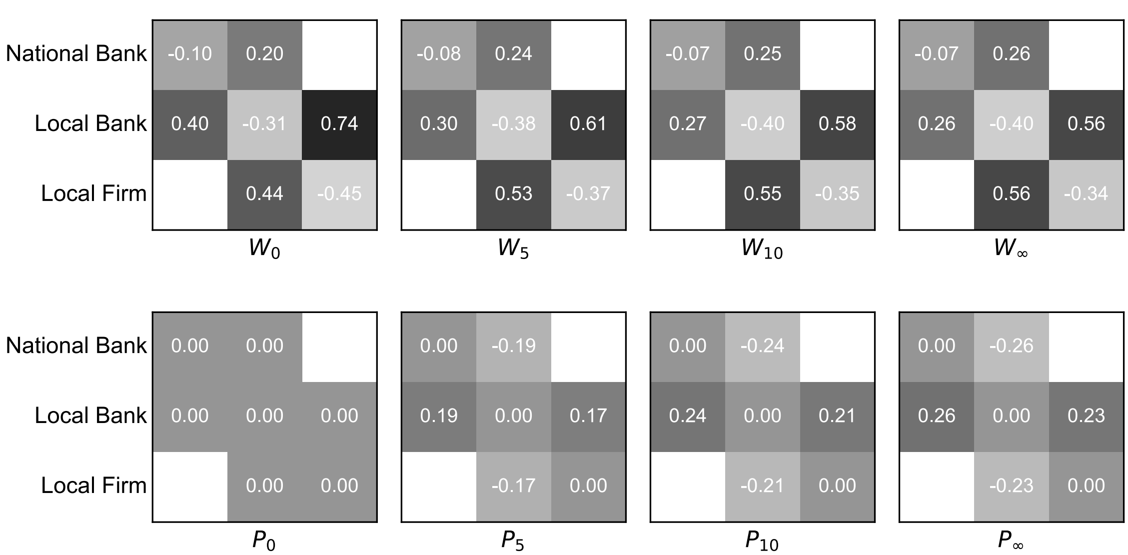

Notes. The contracts matrix W_t and payments matrix P_t after t = 0, 5, 10 steps of Algorithm 1 ($\eta = 0.5$) converge to the stable point $(W, P) = (W_\infty, P_\infty)$. Cells corresponding to forbidden edges are empty.

Let $\lambda_{\max}, \lambda_{\min}$ be the largest and smallest eigenvalues of the matrix $(L^{\dagger}K)|_R$ respectively. Let $\eta^* = \frac{2}{\lambda_{\max}}$. Then, we have:

1. *For all $\eta \in (0, \eta^*)$, pairwise negotiations with η are globally asymptotically stable.*
2. *For such an η, the convergence is exponential in the number of rounds t:*

$$\|P(t) - P^{\star}\|_F \le \frac{\alpha^t}{1-\alpha} \cdot \|P(1) - P(0)\|_F,$$

$$\text{where } \alpha = \max\{|1 - \eta\lambda_{\min}|, |1 - \eta\lambda_{\max}|\}.$$

Here, $P^{\star}$ is the stable point to which the negotiation converges.

Remark 1. For clarity of exposition, we restrict $\eta \in (0,1)$ in Algorithm 1. However, Theorem 4 shows that we only need $\eta < \eta^*$ for convergence to the stable point.

2.3. Pairwise Negotiations Under Random Covariances

Thus far, we made no assumptions about agents' beliefs. In this section, we analyze the convergence of pairwise negotiations for "data-driven" agents. Specifically, each agent i now *estimates* its covariance matrix. For this section only, we will call the covariance matrix $\hat{\Sigma}_i$ instead of Σ_i to emphasize that it is an estimated quantity.

Suppose each agent i observes m independent data samples. Each sample is a vector of the returns of unit contracts with all n agents. The samples for agent i are collected in a matrix $X_i \in \mathbb{R}^{n\times m}$, with one column per sample. The sample covariance of these data is $\hat{\Sigma}_i$.

We assume that all agents observe samples from the same return distribution, which has covariance Σ. Under a wide range of conditions, $\|\hat{\Sigma}_i - \Sigma\| \to 0$ in probability (Vershynin 2018). Hence, at convergence, the maximum allowed dampening rate $\eta^{\star}$ in Theorem 4 would be a function of Σ. However, for finite sample sizes, each agent's $\hat{\Sigma}_i$ can be different. Hence, the maximum dampening $\hat{\eta}^{\star}$ may be less than $\eta^{\star}$. The smaller the $\hat{\eta}^{\star}$, the worse the rate of convergence of pairwise negotiations. However, even with a few samples, $\hat{\eta}^{\star}$ is close to $\eta^{\star}$, as the next theorem shows.

Theorem 5 (Small Sample Sizes Are Sufficient for Fast Convergence). *Suppose that $\|\Sigma\|, \|\Sigma^{-1}\|, \|\Gamma\|$, and $\|\Gamma^{-1}\|$ are O(1) with respect to n and all edges are allowed. Also, suppose that each sample column of X_i is drawn independently from a $\mathcal{N}(\mathbf{0}, \Sigma)$ distribution, and let $\hat{\boldsymbol{\mu}} = \frac{1}{m}\sum_i X_i$ and $\hat{\Sigma}_i := \frac{1}{m-1}\sum_i (X_i - \hat{\boldsymbol{\mu}})(X_i - \hat{\boldsymbol{\mu}})^T$. Let $\hat{\eta}^{\star}$ be the maximum dampening factor using $(\hat{\Sigma}_i)_{i\in[n]}$ as defined in Theorem 4. Let $\eta^{\star}$ be the dampening factor if $\hat{\Sigma}_i$ were replaced by Σ for all i. If $m = \lceil n \log n \rceil$, then for large enough n, $\hat{\eta}^{\star} \ge (1 - o(1))\eta^{\star}$ with probability at least $1 - \exp(-\Omega(n))$.*

Theorem 5 shows that data-driven agents using a broad range of dampening factors are still likely to find the stable point via pairwise negotiations. Furthermore, the amount of data they need is comparable to the number of agents (up to a logarithmic factor). We note that if firms use data sets of fixed sizes $m_1, \dots, m_n$, then the conclusion of Theorem 5 still holds, as long as $\min_i m_i \ge \lceil n \log n \rceil$. For example, firms might use different look-back periods for covariance estimation.

2.4. Inferring Beliefs from the Network Structure

Suppose we are given a network that lies at a unique stable point as defined in Theorem 1. How can we infer the beliefs of the agents?

2.4.1. Nonidentifiability of Beliefs. Suppose we are given a network W that is generated using a single covariance $\Sigma_i = \Sigma \succ 0$. We want to infer the agents' beliefs (M, Γ, Σ). By Corollary A.1,

$$\frac{1}{2}\mathrm{vec}(M + M^T) = (\Gamma \otimes \Sigma + \Sigma \otimes \Gamma)\mathrm{vec}(W).$$

Clearly, the agents' beliefs can only be specified up to an appropriate scaling of M, Γ, and Σ. But even if we specify a scale (e.g., $\mathrm{tr}[\Gamma] = \mathrm{tr}[\Sigma] = 1$), for any valid choice of Γ and Σ we can find a corresponding M. Thus, even in the simple setting of identical covariance and fixed scale, the network W cannot be used to select a unique combination of the parameters (M, Γ, Σ). By a similar argument, we cannot identify the underlying beliefs even if we observe multiple networks generated using the same Σ and Γ (but different M). Thus, we need further assumptions in order to infer beliefs.

Assumption 1. *Consider a sequence of networks $W(t)$ over timesteps $t \in [T]$. We assume that* (a) $\Gamma(t) = I$ *and* $\Sigma_i(t) = \Sigma$ *for all $t \in [T]$,* (b) *for all $i, j \in [n]$, $M_{ij}(t)$ varies independently according to a Brownian motion with the same parameters for all (i, j), and* (c) $\mathrm{tr}\Sigma = 1$.

The first assumption is motivated by the observations in portfolio theory that errors in mean estimation are far more significant than covariance estimation errors (Chopra and Ziemba 1993). Therefore, accounting for variations in Σ may be less important than variations in M (but see Remark 2). The homogeneity of risk aversion was noted in Section 2, and this justifies setting $\Gamma = I$. The second assumption is common in the literature on pricing models (Geman et al. 2001, Bianchi et al. 2013). The third assumption fixes the scale, as discussed above.

Proposition 1. *Finding the maximum likelihood estimator of Σ under Assumption 1 is equivalent to the following semidefinite program (SDP):*

$$\min_{\Sigma} \sum_{t=1}^{T-1} \|\Sigma(W(t+1) - W(t)) + (W(t+1) - W(t))\Sigma\|_F^2$$

$$\text{s.t. } \Sigma \succeq 0, \mathrm{tr}(\Sigma) = 1.$$

Remark 2 (Generalization to Time-Varying $\boldsymbol{\Sigma}$). Instead of a constant covariance Σ, the time range may be split into intervals, with covariance $\Sigma_{(j)}$ in interval j. Then, we can add a regularizer $\nu \cdot \sum_j \|\Sigma_{(j+1)} - \Sigma_{(j)}\|$ for some $\nu > 0$ to the objective of the SDP to penalize differences between successive covariances. This allows the covariance to evolve while keeping the objective convex. The time intervals can be tuned based on heuristics or prior information.

3. Insights for Regulators

A financial regulator can observe the network but does not know the firms' beliefs. The regulator may ask the following. What changes in beliefs caused recently observed changes in the network? What are the side effects of different regulatory interventions? To answer these questions, we need to know how changes in firms' beliefs or utility functions affect the network. That is the subject of this section.

3.1. Effect of Friction in Contract Formation

Our model imposes no costs for contract formation. This is reasonable for large firms where the fixed costs associated with contract negotiations may be small relative to the contract sizes. However, in an overheating market, a regulator may impose frictions by penalizing large contracts, for example by increasing margin requirements.

We model contract costs via an adding a penalty term $F_i(\boldsymbol{w}_i)$ to the utility of agent i in Equation (1):

agent i's utility

$$g_i(W, P) := \boldsymbol{w}_i^T(\boldsymbol{\mu}_i - P\boldsymbol{e}_i) - \gamma_i \cdot \boldsymbol{w}_i^T \Sigma_i \boldsymbol{w}_i - F_i(\boldsymbol{w}_i). \tag{2}$$

Theorem 6. *Consider a network setting where $\Sigma_i = \Sigma$ and all edges are allowed. Suppose that for each firm $i \in [n]$, the function $F_i : \mathbb{R}^n \longrightarrow \mathbb{R}$ is twice differentiable, and there exist strictly increasing functions $f_{ji} : \mathbb{R} \longrightarrow \mathbb{R}$ such that for all $\boldsymbol{x} \in \mathbb{R}^n$, $\nabla F_i(\boldsymbol{x}) = [f_{1i}(x_1), \ldots, f_{ni}(x_n)]^T$. Then, there exists a unique stable point.*

Example 6. By imposing frictions, the regulator may *increase* the sizes of certain contracts. For example, let $F_i(\boldsymbol{w}_i) = \epsilon \cdot w_{i;i}^2 + \lambda \cdot \sum_{j \neq i} w_{i;j}^2$ for some $\lambda > \epsilon > 0$. Thus, the cost of interfirm trades scales with the square of the contract size (we assume $\epsilon \approx 0$). Consider a network setting with three firms, with $\gamma_i = 1$, $\Sigma_i = \Sigma = \begin{bmatrix} 0.1 & 0.1 & 0.1 \\ 0.1 & 1 & 0.5 \\ 0.1 & 0.5 & 1 \end{bmatrix}$, and $M = \begin{bmatrix} 0 & 1000 & 111.233 \\ 1000 & 1 & 0.1 \\ 1000 & 0.1 & 1 \end{bmatrix}$. Then, $W_{23} = W_{32} \approx 0$ without frictions (when $F_i(\boldsymbol{w}_i) = 0$) but $|W_{23}| > 0$ for $\lambda > 0$.

3.2. Effect of Changes in Firms' Beliefs

Regulatory actions can change the risk and expected return perceptions of firms. The next theorem shows the effect of such belief changes on the stable point.

Theorem 7. *Suppose $\Sigma_i = \Sigma$ for all firms and let M be the matrix of expected returns.*

1. **Change in beliefs about expected returns:** *Let Σ have the eigendecomposition $\Sigma = V\Lambda V^T$. Then for $i, j, k, \ell \in [n]$,*

$$\frac{\partial W_{ij}}{\partial M_{k\ell}} = \frac{1}{2\sqrt{\gamma_i \gamma_j \gamma_k \gamma_\ell}} \cdot \sum_{s,t \in [n]} \frac{V_{is}V_{ks}V_{jt}V_{\ell t} + V_{is}V_{\ell s}V_{jt}V_{kt}}{\lambda_s + \lambda_t}. \tag{3}$$

In particular, W_{ij} is monotonically increasing with respect to M_{ij}.

2. **Risk scaling:** *If the covariance Σ changes to $c\Sigma$ ($c > 0$), then W changes to $(1/c)W$.*

3. **Increase in perceived risk:** *Suppose $\gamma_i = \gamma$ for all i, and the covariance Σ increases to $\Sigma' > \Sigma$. Let W and W' be the stable points under Σ and Σ' respectively. Then,* $\mathrm{tr}(M^T(W' - W)) < 0$.

This shows that, in general, an increase in risk leads to a decrease in the weighted average of the contract sizes. The weights are given by the expected return beliefs of the firms. However, individual contracts between firms can increase, as can the norm $\|W\|_F$. This is because increases in the covariance Σ may also increase correlations, which can offer better hedging opportunities. By hedging some risks, larger contract sizes can be supported.

Theorem 7 also shows that a change in the perceived expected return $M_{k\ell}$ affects all contracts W_{ij}. Can we trace the changes in W back to the underlying changes in M? For instance, consider the following problem.

Definition 7 (Source Detection Problem). Suppose that a financial regulator observes two networks W and W', with the only difference being a small change in a single entry of M (say, M_{ij}). Can the regulator identify the pair (i, j)?

One approach is to try to infer all beliefs of all firms, and then identify the changed belief. However, as discussed in Section 2.4, the beliefs are only identifiable under extra assumptions and more data. An alternative approach for the source detection problem is to find the entry (i, j) with the largest change $|W_{ij} - W'_{ij}|$. The intuition is that a change in M_{ij} has a direct effect on W_{ij} and (hopefully weaker) indirect effects on other contracts. Thus, the source detection problem is closely tied to the following:

Definition 8 (Targeted Intervention Problem). Can a regulator induce a small change in a single entry of M (say, M_{ij}) such that the change in W_{ij} is significantly larger than changes in other entries of W?

When all eigenvalues of Σ are equal (that is, $\Sigma \propto I_n$), a change in $M_{k\ell}$ only affects $W_{k\ell} (= W_{\ell k})$, as can be seen from Corollary A.1. However, when the eigenvalues are skewed, the terms in Equation (3) corresponding to the smallest eigenvalues have greater weight. In such circumstances, the indirect effect of a change in $M_{k\ell}$ on other W_{ij} can be significant. The following empirical results show that this is indeed the case.

3.2.1. Empirical Results for the Source Detection Problem (Simulated Data). Here, we set the covariance $\Sigma = D^{1/2}(R + \mathcal{E})D^{1/2}$, where D is a diagonal matrix, R a correlation matrix, and $\mathcal{E}$ a noise matrix. If $\mathcal{E} = 0$, then D_{ii} would be the variance of firm i. We set D_{ii} according to a power law: $D_{ii} = i^{-\alpha}$ for an $\alpha > 0$. Larger values of α correspond to greater skew in the variances. We choose R to be an equi-correlation matrix with one along the diagonal and $\rho \in (0,1)$ everywhere else. We draw the error matrix $\mathcal{E}$ from a scaled Wishart distribution: $\mathcal{E} = \|R\|_2 \cdot \mathcal{W}(\sqrt{\epsilon} \cdot I_n, n)/n$ for some chosen the noise level ϵ. As ϵ increases, the noise $\mathcal{E}$ dominates R.

Figure 3 shows the success rate of source detection over 1,000 experiments for various values of (ϵ, α) for $\rho = 0.1$ and $n = 50$. As α increases, the variances become more skewed, and the source detection can fail even with $\epsilon = 0$ noise. When ϵ grows, the success rate for the source detection problem goes to zero. This suggests that skew combined with noise makes source detection difficult. These trends occur even if we only test whether the source belongs to the 10 most changed contracts (Figure 3(b)), as opposed to single largest change (Figure 3(a)). We observe similar results for real-world choices of Σ, as we show next.

3.2.2. Empirical Results for the Source Detection Problem (Real-World Data). We consider two data sets: (a) a trade network between 46 large economies (OECD 2022) and (b) a simulated network between 96 portfolio managers following various Fama-French strategies (Fama and French 2015). For each data set, we construct a "ground-truth" covariance Σ using all available data (the details are in Section B of the supplemental materials). Then, using m independent samples $\boldsymbol{x}_i \sim \mathcal{N}(0, \Sigma)$, we build a "data-driven" covariance $\hat{\Sigma} = (1/(m-1)) \sum_{i=1}^{m}(\boldsymbol{x}_i - \hat{\boldsymbol{\mu}})(\boldsymbol{x}_i - \hat{\boldsymbol{\mu}})^T$, where $\hat{\boldsymbol{\mu}} = (1/m)\sum_{i=1}^{m} \boldsymbol{x}_i$ is the sample mean. We use this $\hat{\Sigma}$ to construct the financial network.

Figure 4 shows the success rate over 500 experiments for various choices of the sample size m. The success rate increases monotonically with m. The reason for this behavior lies in the spectra of Σ and $\hat{\Sigma}$. We find that in both data sets, the largest and smallest eigenvalues of Σ are separated by several orders of magnitude. This gap becomes even more extreme in the data-driven $\hat{\Sigma}$; the fewer the samples m, the greater the gap (Figure 5). In fact, we observe that the smallest eigenvalue of $\hat{\Sigma}$ is much smaller than the second-smallest eigenvalue: $\lambda_n \ll \lambda_{n-1}$. Zhao et al. (2019) make similar observations.

In summary, the experiments on both simulated and real-world data sets highlight the difficulty of source detection and targeted intervention in realistic networks. The reason is the skew in the eigenvalues coupled with noise, which affects the eigenvectors. Skewed eigenvalues correspond to trade combinations (eigenvectors) that are seemingly low risk. Hence, firms use such trades to diversify. This implies that these eigenvectors have an outsized effect on the network and how it responds to local changes. Intuitively, if these eigenvectors are "random," the effect of a changed belief $M_{k\ell}$ affects the rest of the network randomly. Hence, the direct effects on $W_{k\ell}$ may be less than the indirect effects on other W_{ij}. We explore this theoretically in Section A.12 of the supplemental material.

Figure 3. Source Detection Problem in a Noisy Scaled Equi-Correlation Model of Σ

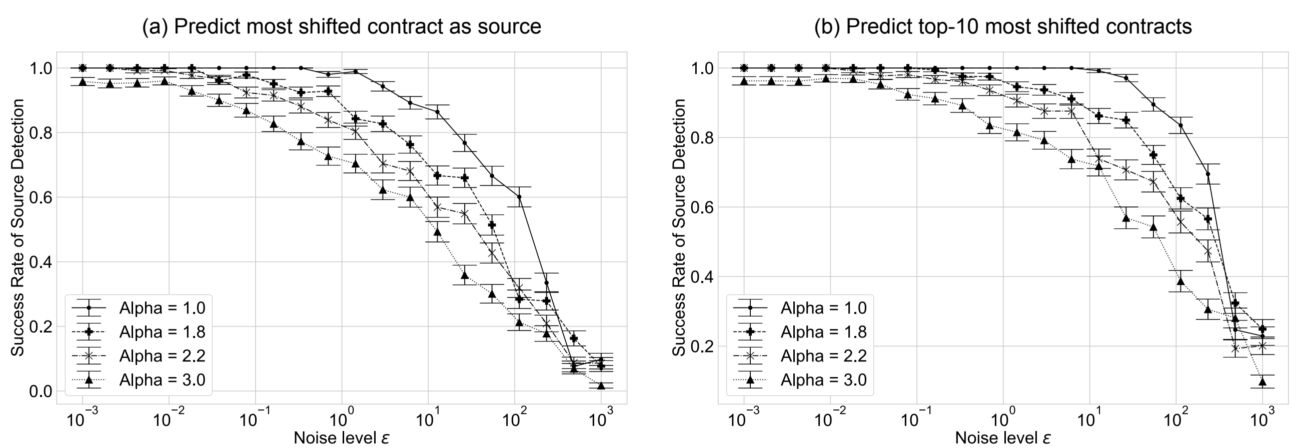

Notes. We rank the entries of W by the magnitude of change induced by a change in one entry of M (M_{ij}). (a) Fraction of times W_{ij} is most-changed entry of W. (b) Fraction of times W_{ij} is among the top 10 most changed entries of W. The success rate goes to zero as α and ϵ increase.

Figure 4. Source Detection Problem on Real-World Data

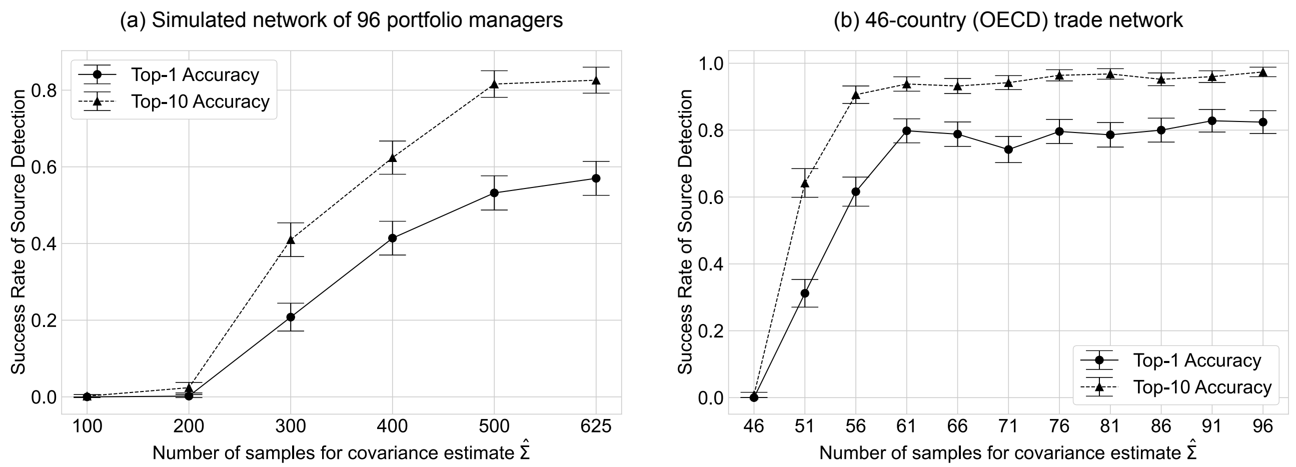

Note. The success rate scales monotonically with the number of samples used to construct the data-driven covariance matrix $\hat{\Sigma}$.

4. Insights for Firms

Until now, we treated the beliefs of firms as fixed and exogenous. In this section, we consider how a firm can use its contracts to gain insights into other firms and update its beliefs.

For instance, suppose a firm j faces a crisis, for example, a looming debt payment that may make it insolvent. The firm may then become risk seeking (i.e., lower its γ_j), hoping that the risks pay off. Another firm i may be unaware of the crisis, so i's risk perceptions (perhaps based on historical data) would be outdated. Can firm i *infer* the lower γ_j, solely from i's contracts $\boldsymbol{w}_i$ with all firms? What if a group of firms become risk-seeking, and not just one firm?

4.1. Detecting Outlier Firms

Intuitively, firm i will try to answer these questions by comparing the behavior of firm j against other similar firms. We formalize this by assuming that each firm j belongs to a community θ_j, for example, banking, real estate, or insurance, and so on. The community of each firm is publicly known. Firms in the same community are perceived to have similar return distributions:

$$M_{ij} = f(\theta_i, \theta_j) + \epsilon'_{\theta_i, j}, \quad \Sigma_{ij} = g(\theta_i, \theta_j),$$
$$\gamma_i = h(\theta_i) + \epsilon_i \tag{4}$$

for some unknown deterministic functions $f(.)$, $g(.)$, and $h(.)$ and random error terms ϵ_i and $\epsilon'_{\theta_i, j}$. We also assume that all firms use the same covariance Σ.

Now, suppose one firm j is an outlier, with very different beliefs from other firms in its community. For firm i to detect the outlier firm j, the contract size W_{ij} should deviate from a cluster of contracts $\{W_{ij'} \mid \theta_{j'} = \theta_j\}$ of other firms from the same community as firm j. Now, outlier detection methods often assume independent

Figure 5. Eigenvalues of Estimated Covariance Matrices Are Skewed, and the Degree of Skew Depends on the Number of Samples m

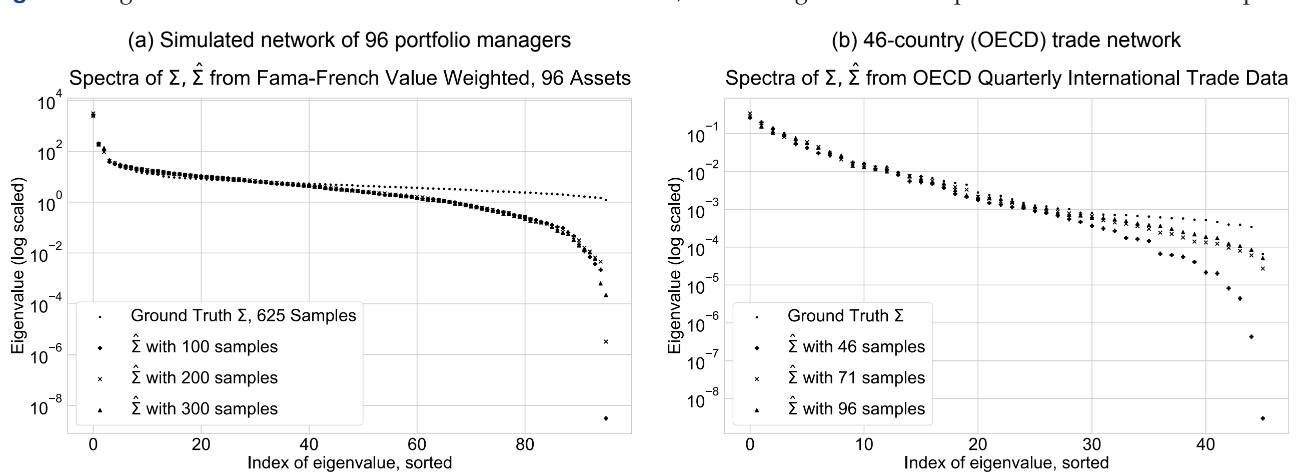

Note. As m decreases, so does the smallest eigenvalue λ_n and the ratio λ_n/λ_{n-1}.

Figure 6. Success Rate for Detecting Outlier Risk-Seeking Firms

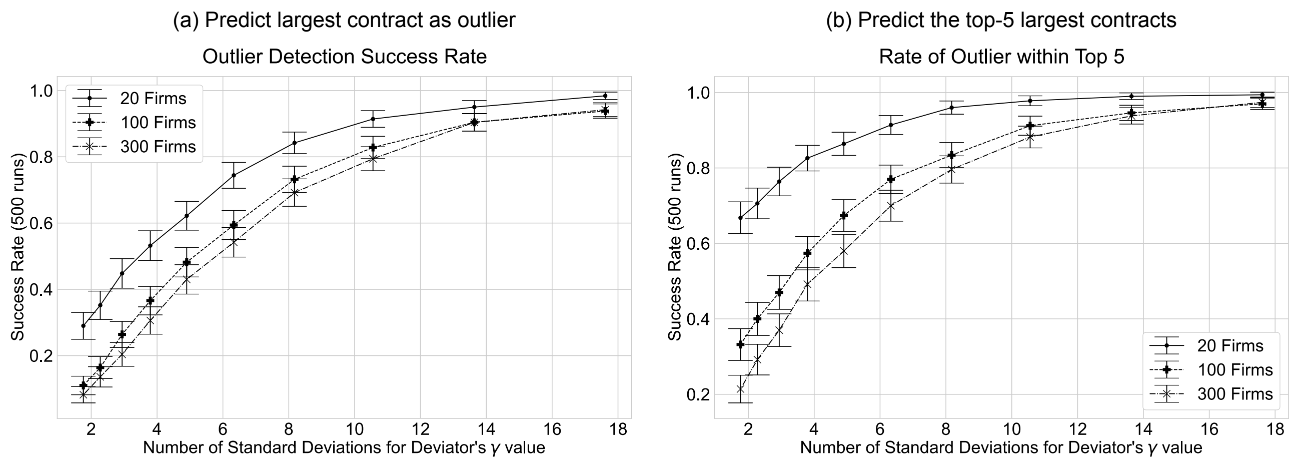

Note. Detection is easier when there are fewer firms and when the risk-seeking firm's γ_{outlier} is more standard deviations away from the γ of the normal firms.

datapoints. In our model, all contracts are dependent. However, we can still do outlier detection if the contracts are appropriately exchangeable. We prove below this is the case.

Definition 9. An intracommunity permutation is a permutation $\pi : [n] \to [n]$ such that $\pi(i) = j$ implies that $\theta_i = \theta_j$.

Proposition 2. *Suppose M, Σ, Γ exhibit community structure (Equation (4)), and all the error terms $(\epsilon_i)_{i\in[n]}$ and $(\epsilon'_{\theta_i,j})_{i,j\in[n]}$ are independent and identically distributed. Let $\pi : [n] \to [n]$ be any intracommunity permutation, and let $\Pi : \mathbb{R}^n \to \mathbb{R}^n$ be the corresponding column-permutation matrix: $\Pi(\boldsymbol{e}_i) = \boldsymbol{e}_{\pi(i)}$. Then, W and $\Pi^T W \Pi$ are identically distributed.*

Corollary 1. *Let $j_1, \ldots, j_m \in [n]$ belong to the same community: $\theta_{j_1} = \cdots = \theta_{j_m}$. Suppose the conditions of Proposition 2 hold. Then, for any $i \in [n]$, the joint distribution of $(W_{i,j_1}, \ldots, W_{i,j_m})$ is exchangeable.*

4.1.1. Empirical Results for Outlier Detection.

We generate community-based networks (Equation (4)) such that $\gamma_i \sim N(1, \sigma^2)$ truncated to $[0.5, 1.5]$. The smaller the σ, the more closely the γ_i values cluster around one. For the outlier risk-seeking firm, we set $\gamma_{\text{outlier}} = 0.5$. For clarity of exposition, we set $\epsilon' = 0$ everywhere.

To detect outliers under exchangeability (Corollary 1), we can use methods based on conformal prediction (Guan and Tibshirani 2022). Here, we use a simpler approach: pick the firm j with the largest contract size as the outlier; $\hat{j} := \arg\max_{j\in\{j_1,\ldots,j_m\}} |W_{i,j}|$. To test sensitivity to false negatives, we also test whether the outlier is among the five largest contracts in $\{|W_{i,j}| : j = j_1, \ldots, j_m\}$. We run 500 experiments for each choice of σ and count the frequency with which the outlier firm is detected via its contract size. Further details are presented in Section B.3 of the supplemental material.

Figure 6 shows the results. We characterize the degree of outlierness by how many standard deviations away γ_{outlier} is from the baseline of one. The smaller the σ, the more the outlierness. The success rate increases with increasing outlierness, as expected. It also increases when the number of firms n is reduced. This is because contract sizes depend on the γ values of all firms; fewer firms reduces the chances of any one firm attaining large contract sizes due to randomness.

4.2. Risk Aversion vs. Expected Returns

The discussion above shows that a firm can detect outlier counterparties. However, the firm cannot determine *why* the counterparty is an outlier, as the following theorem shows.

Theorem 8 (Nonidentifiability of Risk Aversion Versus Expected Returns). *Consider two network settings $S = (\boldsymbol{\mu}_i, \Sigma, \gamma_i)_{i\in[n]}$ and $S' = (\boldsymbol{\mu}_i, \Sigma, \gamma_i{}')_{i\in[n]}$ that differ only in the risk aversions of firms $J = \{j \mid \gamma_j \neq \gamma'_j\} \subseteq [n]$. Then, there exists a setting $S^\dagger = (\boldsymbol{\mu}_i^\dagger, \Sigma, \gamma_i)_{i\in[n]}$ such that $\boldsymbol{\mu}_i = \boldsymbol{\mu}_i^\dagger$ for all $i \notin J$ and the stable networks under $S^\dagger$ and S' are identical.*

Thus, one cannot determine whether an outlier is more risk seeking than its community or expects higher profits. However, risk-seeking behavior may be indicative of stress while higher profits than similar firms are unlikely. Hence, in either case, the firm detecting the outlier may choose to reduce its exposure to the outlier. However, this approach fails if an entire community shifts its behavior. The following example illustrates the problem.

Example 7. Consider two communities numbered 1 and 2, with n_1 and n_2 firms, respectively. Let the setting

S of Theorem 8 correspond to

$$M_{ij} = \begin{cases} a & \text{if } \theta_i = \theta_j = 1 \\ b & \text{if } \theta_i = \theta_j = 2 \\ c/2 & \text{otherwise} \end{cases}$$

$$\Sigma_{ij} = \begin{cases} 1 & \text{if } \theta_i = \theta_j = 1 \\ 1 & \text{if } \theta_i = \theta_j = 2 \\ 0 & \text{otherwise} \end{cases} \quad \gamma_i = 1.$$

Now, suppose that under setting S', $\gamma_i \longmapsto \gamma_i + \delta$ for some small δ for all nodes i in community 1. The change in the network would be the same if we had updated the columns corresponding to community 1 in the M matrix instead (setting $S^{\dagger}$):

$$M^{\dagger}_{ij} = M_{ij} + \Delta(\theta_i, \theta_j)\Delta(\theta_i, \theta_j) + O(\delta^2)$$
$$= \begin{cases} -\delta a/2 & \text{if } \theta_i = \theta_j = 1 \\ -\delta b \cdot n_2/(n_1 + n_2) & \text{if } \theta_i = 2, \theta_j = 1 \\ 0 & \text{if } \theta_j = 2. \end{cases}$$

Thus, a firm from community 2 cannot determine whether the network change was due to a change in $(\gamma_i)_{\theta_i=1}$ or $(\boldsymbol{\mu}_i)_{\theta_i=1}$. For instance, when $b > 0$, an increase in risk-seeking ($\delta < 0$) looks the same as an increase in trading benefits ($\Delta(1,2) > 0$). In the former case, firms in community 2 should *reduce* their exposure to community 1 firms. However, in the latter case, they should *increase* exposure. Because the data cannot be used to choose the appropriate action, the behaviors of firms may be guided by their prior beliefs or inertia. When such beliefs change due to external events (e.g., due to news about one firm in community 1), the resulting change in the network may be drastic.

5. Conclusions

We proposed a model of a weighted undirected financial network of contracts. The network emerges from the beliefs of the participant firms. The link between the two is utility maximization coupled with pricing. For almost all belief settings, our approach yields a unique network. This network satisfies a strong higher-order Nash stability property. Furthermore, the firms can converge to this stable network via iterative pairwise negotiations.

The model yields two insights. First, a regulator is unable to reliably identify the causes of a change in network structure, or engage in targeted interventions. The reason is that firms seek to diversify risk by exploiting correlations. We find that in realistic settings, there are often combinations of trades that offer seemingly low risk. Hence, all firms aim to use such trades. The overdependence on a few such combinations leads to a pattern of connections between firms that thwarts targeted regulatory interventions.

The second insight is that firms can use the network to update their beliefs. For instance, they can identify counterparties that behave very differently from their peers. However, the cause of the outlierness remains hidden. If all firms in one line of business become more risk-seeking, the result is indistinguishable from that business becoming more profitable. Innocuous events (such as a news story) may cause beliefs to change suddenly, leading to drastic changes in the network. In addition to identifying risky counterparties, firms may use the network to update their mean and covariance beliefs. For example, a firm that suffers significant losses on its current trades may be judged by others to be a riskier counterparty for future trades. We leave this for future work.

Our work focuses on mean-variance utility, but some of our results are applicable in other settings too. A second-order Taylor approximation of a twice-differentiable concave utility matches the form of a mean-variance utility. Hence, results based on mean-variance utility can be useful guides for small perturbations around a stable point. Some of our results for pairwise negotiations and targeted interventions are based on such perturbation arguments.

Finally, contract formation under budget constraints is an important direction for future work. In Theorem 6, we only consider contract frictions that depend on a firm's contract sizes. To model budget constraints, we must also consider the contract prices. These require different techniques than our approach, which is based on results from Sandberg and Willson (1972) (see Section A.15 in the supplemental material).

Acknowledgments

The authors thank Paul Glasserman (area editor), the associate editor, and both anonymous referees for invaluable review comments and Stathis Tompaidis, Marios Papachristou, Kshitij Kulkarni, and David Fridovich-Keil for valuable discussions.

Appendix

A.1. Proof of Theorem 1

Recall that $Q_i = \Psi_i^T(2\gamma_i\Psi_i\Sigma_i\Psi_i^T)^{-1}\Psi_i$, $F = \{(i,j) : 1 \le i < j \le n, \Psi_i \boldsymbol{e}_j \ne \boldsymbol{0}\}$, and $\text{uvec}(X)_F \in \mathbb{R}^{|F|}$ is a vector whose entries are the ordered set $\{X_{ij} | (i,j) \in F\}$. Note that $\Psi_i\Sigma_i\Psi_i^T$ is positive definite because it is a principal submatrix of the positive definite matrix Σ_i.

Proof of Theorem 1. For clarity of exposition, we first prove the result when all edges are allowed, and then consider the case of disallowed edges.

(1) **All edges allowed.** Here, $E = \{i,j | 1 \le i < j \le n\}$, and we use uvec(.) and Z to refer to $\text{uvec}(.)_E$ and Z_E in the theorem statement. For any price matrix P with $P = -P^T$, consider

the matrix W whose jth column has the utility-maximizing contract sizes for agent j:

$$W_{ij} = \boldsymbol{e}_i^T \Psi_j^T (2\gamma_j \Psi_j \Sigma_j \Psi_j^T)^{-1} \Psi_j (M-P)\boldsymbol{e}_j = \boldsymbol{e}_i^T Q_j (M-P)\boldsymbol{e}_j.$$

The tuple (W, P) is stable if $W = W^T$. Therefore, for all $i < j$, we require

$$W_{ij} = W_{ji} \tag{A.1}$$

$$\Longleftrightarrow \boldsymbol{e}_i^T Q_j (M-P)\boldsymbol{e}_j = \boldsymbol{e}_j^T Q_i (M-P)\boldsymbol{e}_i$$

$$\Longleftrightarrow \boldsymbol{e}_i^T Q_j M \boldsymbol{e}_j - \boldsymbol{e}_j^T Q_i M \boldsymbol{e}_i = \boldsymbol{e}_i^T Q_j P \boldsymbol{e}_j - \boldsymbol{e}_j^T Q_i P \boldsymbol{e}_i$$

$$\Longleftrightarrow \boldsymbol{e}_i^T (A - A^T)\boldsymbol{e}_j = \boldsymbol{e}_i^T (Q_j P - (Q_i P)^T)\boldsymbol{e}_j. \tag{A.2}$$

Because $P = -P^T$, we must have $P = R - R^T$, where R is upper-triangular with zero on the diagonal. Hence, using $Q_i = Q_i^T$, we have

$$\begin{aligned}
\boldsymbol{e}_i^T (Q_j P - (Q_i P)^T)\boldsymbol{e}_j &= \boldsymbol{e}_i^T (Q_j P + P Q_i)\boldsymbol{e}_j \\
&= tr P(\boldsymbol{e}_j \boldsymbol{e}_i^T Q_j + Q_i \boldsymbol{e}_j \boldsymbol{e}_i^T) \\
&= tr(R - R^T)(B_{(j,i)} + B_{(i,j)}^T) \\
&= tr R^T C_{(i,j)} \\
&= \text{uvec}(R)^T \text{uvec}(C_{(i,j)}),
\end{aligned}$$

where we used the upper-triangular nature of R in the last step. Plugging into Equation (A.2), a stable point exists if and only if there is an appropriate vector $\boldsymbol{p} := \text{uvec}(R) \in \mathbb{R}^{n(n-1)/2}$ such that for all $1 \le i < j \le n$, $\boldsymbol{e}_i^T (A - A^T)\boldsymbol{e}_j = \text{uvec}(C_{(i,j)})^T \boldsymbol{p}$. This is equivalent to $\text{uvec}(A - A^T) = Z\boldsymbol{p}$. If such a solution vector $\boldsymbol{p}$ exists, then by definition, it corresponds to a matrix $P = -P^T$ via $P = R - R^T$ and $\boldsymbol{p} = \text{uvec}(R)$.

(2) **Disallowed edges.** If $\{i, j\}$ is a prohibited edge then $\Psi_i \boldsymbol{e}_j = \Psi_j \boldsymbol{e}_i = \boldsymbol{0}$, so $B_{(i,j)} = B_{(j,i)} = 0$, so $\boldsymbol{e}_{ij}^T Z = \boldsymbol{0}^T$. Also, $A_{ij} = A_{ji} = 0$ so $\text{uvec}(A - A^T)_{ij} = 0$. Therefore, the equality $\boldsymbol{e}_i^T (A - A^T)\boldsymbol{e}_j = \text{uvec}(C_{(i,j)})^T \boldsymbol{x}$ is achieved for any solution vector $\boldsymbol{x}$ if $\{i, j\}$ is a prohibited edge. We can therefore reduce the linear system $Z\boldsymbol{p} = \text{uvec}(A - A^T)$ from part (1) by deleting rows of Z corresponding to prohibited edges.

Similarly, because the system is constrained by $\boldsymbol{p}_{ij} = 0$ for prohibited edges $\{i, j\}$, the columns of Z corresponding to such edges have no effect on the solution set.

We conclude that the linear system in (1) is equivalent to the (unconstrained) reduced system $Z_F \boldsymbol{p}_F = \text{uvec}(A - A^T)_F$. Each solution $\boldsymbol{p}_F$ corresponds to a skew-symmetric P by construction. Finally, if Z_F has full rank then the unique reduced solution is $\boldsymbol{p}_F = Z_F^{-1} \text{uvec}(A - A^T)_F$. □

A.2. Stable Network for the Shared Covariance Case

In the case of a shared covariance matrix for all agents, we can give a closed form expression for the stable network.

Corollary A.1 (Shared $\boldsymbol{\Sigma}$, All Edges Allowed). *Suppose $\Sigma_i = \Sigma$ and $\Psi_i = I_n$ for all $i \in [n]$. Let $(\lambda_i, \boldsymbol{v}_i)$ denote the ith eigenvalue and eigenvector of $\Gamma^{-1/2}\Sigma\Gamma^{-1/2}$. Then, the network W can be written in two equivalent ways:*

$$\text{vec}(W) = \frac{1}{2}(\Gamma \otimes \Sigma + \Sigma \otimes \Gamma)^{-1} \text{vec}(M + M^T),$$

$$W = \Gamma^{-1/2} \left(\sum_{i=1}^n \sum_{j=1}^n \frac{\boldsymbol{v}_i^T \Gamma^{-1/2}}{2(\lambda_i + \lambda_j)} (M + M^T)\Gamma^{-1/2} \boldsymbol{v}_j \boldsymbol{v}_i \boldsymbol{v}_j^T \right) \Gamma^{-1/2}.$$

The prices can be written as

$$\text{vec}(P) = (\Gamma^{-1} \otimes \Sigma^{-1} + \Sigma^{-1} \otimes \Gamma^{-1})^{-1} \text{vec}(\Sigma^{-1} M \Gamma^{-1} - \Gamma^{-1} M^T \Sigma^{-1})$$

$$P = \Gamma^{1/2} \left(\sum_{i=1}^n \sum_{j=1}^n \frac{\boldsymbol{v}_i^T \Gamma^{1/2}}{\lambda_i^{-1} + \lambda_j^{-1}} (\Sigma^{-1} M \Gamma^{-1} - \Gamma^{-1} M^T \Sigma^{-1})\Gamma^{1/2} \boldsymbol{v}_j \boldsymbol{v}_i \boldsymbol{v}_j^T \right) \Gamma^{1/2}.$$

Proof. We first prove the identity with $\text{vec}(W)$.

For each agent i the optimal set of contracts is given as $\boldsymbol{w}_i = (2\gamma_i \Sigma_i)^{-1}(M - P)\boldsymbol{e}_i$. Because $\Sigma_i = \Sigma$ for all i, we obtain $W = \frac{1}{2}\Sigma^{-1}(M - P)\Gamma^{-1}$. Hence $M - P = 2\Sigma W \Gamma$. Using $W = W^T$ and $P^T = -P$ for a stable feasible point (W, P), we obtain $\Sigma W \Gamma + \Gamma W \Sigma = \frac{1}{2}(M + M^T)$.

Vectorization implies $(\Gamma \otimes \Sigma + \Sigma \otimes \Gamma)\text{vec}(W) = \frac{1}{2}\text{vec}(M + M^T)$. It remains to show that $(\Gamma \otimes \Sigma + \Sigma \otimes \Gamma)$ is invertible.

Let $K := (\Gamma \otimes \Sigma + \Sigma \otimes \Gamma)$ for shorthand. Notice $K = (\Gamma^{1/2} \otimes \Gamma^{1/2})(I \otimes \Gamma^{-1/2}\Sigma\ \Gamma^{-1/2} + \Gamma^{-1/2}\Sigma\Gamma^{-1/2} \otimes I)(\Gamma^{1/2} \otimes \Gamma^{1/2})$. Let $K' = (I \otimes \Gamma^{-1/2}\Sigma\Gamma^{-1/2} + \Gamma^{-1/2}\Sigma\Gamma^{-1/2} \otimes I)$. Because $(\Gamma^{1/2} \otimes \Gamma^{1/2})$ is invertible, it suffices to show K' is invertible.

Properties of Kronecker products imply that if a matrix $A \in \mathbb{R}^{n \times n}$ has strictly positive eigenvalues, then $\sigma(I \otimes A + A \otimes I) = \{\lambda + \mu : \lambda, \mu \in \sigma(A)\}$ counting mutiplicities (Horn and Johnson 1994). Let $\boldsymbol{v} \neq \boldsymbol{0}$. Then, because $\Sigma > 0$ and $\Gamma^{-1/2} > 0$, we obtain $\boldsymbol{v}^T \Gamma^{-1/2}\Sigma\Gamma^{-1/2}\boldsymbol{v} = (\Gamma^{-1/2}\boldsymbol{v})^T \Sigma (\Gamma^{-1/2}\boldsymbol{v}) > 0$. Hence, $\Gamma^{-1/2}\Sigma\Gamma^{-1/2} > 0$, so K' is invertible, and hence K is invertible. This proves the first identity.

Next, we prove the second identity. Properties of Kronecker products imply that $(K')^{-1}$ has eigendecomposition $(K')^{-1} = \sum_{i=1}^n \sum_{j=1}^n \frac{1}{\lambda_i + \lambda_j}(\boldsymbol{v}_i \otimes \boldsymbol{v}_j)(\boldsymbol{v}_i \otimes \boldsymbol{v}_j)^T$.

Therefore, because $(\Gamma^{1/2} \otimes \Gamma^{1/2})^{-1} = (\Gamma^{-1/2} \otimes \Gamma^{-1/2})$, we obtain

$$\begin{aligned}
\text{vec}(W) &= (\Gamma^{-1/2} \otimes \Gamma^{-1/2}) \sum_{i=1}^n \sum_{j=1}^n \frac{1}{\lambda_i + \lambda_j} (\boldsymbol{v}_i \otimes \boldsymbol{v}_j)(\boldsymbol{v}_i \otimes \boldsymbol{v}_j)^T (\Gamma^{-1/2} \otimes \Gamma^{-1/2}) \text{vec}\left(\frac{M + M^T}{2}\right) \\
&= (\Gamma^{-1/2} \otimes \Gamma^{-1/2}) \sum_{i=1}^n \sum_{j=1}^n \frac{1}{2(\lambda_i + \lambda_j)} \text{vec}(\Gamma^{-1/2}(M + M^T)\Gamma^{-1/2}) \\
&= (\Gamma^{-1/2} \otimes \Gamma^{-1/2}) \text{vec}\left(\sum_{i=1}^n \sum_{j=1}^n \frac{\boldsymbol{v}_i^T \Gamma^{-1/2}}{2(\lambda_i + \lambda_j)} (M + M^T)\Gamma^{-1/2} \boldsymbol{v}_j \boldsymbol{v}_i \boldsymbol{v}_j^T \right)
\end{aligned}$$

$$W = \Gamma^{-1/2} \left(\sum_{i=1}^n \sum_{j=1}^n \frac{\boldsymbol{v}_i^T \Gamma^{-1/2}}{2(\lambda_i + \lambda_j)} (M + M^T)\Gamma^{-1/2} \boldsymbol{v}_j \boldsymbol{v}_i \boldsymbol{v}_j^T \right) \Gamma^{-1/2}.$$

Finally, the formulas for vec(P) and P follow from similar reasoning, using $W = W^T$ and $W = \frac{1}{2}\Sigma^{-1}(M - P)\Gamma^{-1}$. □

References

Acemoglu D, Azar PD (2020) Endogenous production networks. *Econometrica* 88(1):33–82.
Acemoglu D, Ozdaglar A, Tahbaz-Salehi A (2015) Systemic risk and stability in financial networks. *Amer. Econom. Rev.* 105(2):564–608.
Aiello W, Chung F, Lu L (2000) A random graph model for massive graphs. *Proc. 32nd Annual ACM Sympos. Theory Comput. (STOC '00)* (Association for Computing Machinery, New York), 171–180.
Airoldi EM, Blei DM, Fienberg SE, Xing EP (2008) Mixed membership stochastic blockmodels. *J. Machine Learn. Res.* 9:1981–2014.
Amini H, Minca A, Sulem A (2015) Control of interbank contagion under partial information. *SIAM J. Financial Math.* 6(1):1195–1219.
Ang A (2014) *Asset Management: A Systematic Approach to Factor Investing*, online edn. (Oxford Academic, New York)
Arrow KJ, Debreu G (1954) Existence of an equilibrium for a competitive economy. *Econometrica* 22(3):265–290.
Banerjee T, Feinstein Z (2022) Pricing of debt and equity in a financial network with comonotonic endowments. *Oper. Res.* 70(4): 2085–2100.
Barucca P, Bardoscia M, Caccioli F, D'Errico M, Visentin G, Caldarelli G, Battiston S (2020) Network valuation in financial systems. *Math. Finance* 30(4):1181–1204.
Bianchi S, Pantanella A, Pianese A (2013) Modeling stock prices by multifractional Brownian motion: An improved estimation of the pointwise regularity. *Quant. Finance* 13(8):1317–1330.
Bich P, Morhaim L (2020) On the existence of pairwise stable weighted networks. *Math. Oper. Res.* 45(4):1393–1404.
Bich P, Teteryatnikova M (2023) On perfect pairwise stable networks. *J. Econom. Theory* 207:105577.
Birge JR (2021) Modeling investment behavior and risk propagation in financial networks. Preprint, submitted May 18, https://dx.doi.org/10.2139/ssrn.3847443.
Bühlmann H (1980) An economic premium principle. *ASTIN Bull.* 11(1):52–60.
Bühlmann H (1984) The general economic premium principle. *ASTIN Bull.* 14(1):13–21.
Caballero RJ, Simsek A (2013) Fire sales in a model of complexity. *J. Finance* 68(6):2549–2587.
Caimo A, Friel N (2011) Bayesian inference for exponential random graph models. *Soc. Networks* 33(1):41–55.
Calafiore GC, Fracastoro G, Proskurnikov AV (2022) Control of dynamic financial networks. *IEEE Control Systems Lett.* 6:3206–3211.
Callier FM, Desoer CA (1994) *Linear System Theory* (Springer Verlag, Berlin).
Calvó-Armengol A, Ilkiliç R (2009) Pairwise-stability and Nash equilibria in network formation. *Internat. J. Game Theory* 38(1):51–79.
Carmona R, Cooney DB, Graves CV, Lauriere M (2022) Stochastic graphon games: I. The static case. *Math. Oper. Res.* 47(1):750–778.
Chakrabarti D, Zhan Y, Faloutsos C (2004) R-MAT: A recursive model for graph mining. *Proc. 4th SIAM Internat. Conf. Data Mining* (SIAM, Philadelphia), 442–446.
Chopra VK, Ziemba WT (1993) The effect of errors in means, variances, and covariances on optimal portfolio choice. *J. Portfolio Management* 19(2):6–11.
Cont R, Minca A (2016) Credit default swaps and systemic risk. *Ann. Oper. Res.* 247(7):523–547.
Eisenberg L, Noe TH (2001) Systemic risk in financial systems. *Management Sci.* 47(2):236–249.
Eisfeldt AL, Herskovic B, Rajan S, Siriwardane E (2021) OTC intermediaries. Research Paper 18-05, Office of Financial Research, United States Department of the Treasury, Washington, DC.
Elliott M, Golub B (2022) Networks and economic fragility. *Annu. Rev. Econom.* 14:665–696.
Elliott M, Golub B, Jackson MO (2014) Financial networks and contagion. *Amer. Econom. Rev.* 104(10):3115–3153.
Elliott M, Golub B, Leduc MV (2022) Supply network formation and fragility. *Amer. Econom. Rev.* 112(8):2701–2747.
Erdös P, Rényi A (1959) On random graphs I. *Publicationes Math.* 6:290–297.
Fama EF, French KR (2015) A five-factor asset pricing model. *J. Financ. Econom.* 116(1):1–22.
Feinstein Z (2020) Capital regulation under price impacts and dynamic financial contagion. *Eur. J. Oper. Res.* 281(2):449–463.
Feinstein Z, Søjmark A (2021) Dynamic default contagion in heterogeneous interbank systems. *SIAM J. Financial Math.* 12(4):SC83–SC97.
Feinstein Z, Søjmark A (2022) Endogenous network valuation adjustment and the systemic term structure in a dynamic interbank model. Preprint, submitted November 28, https://arxiv.org/abs/2211.15431.
Frank O, Strauss D (1986) Markov graphs. *J. Amer. Statist. Assoc.* 81(395):832–842.
Galeotti A, Golub B, Goyal S (2020) Targeting interventions in networks. *Econometrica* 88(6):2445–2471.
Gandy A, Veraart LA (2017) A Bayesian methodology for systemic risk assessment in financial networks. *Management Sci.* 63(12):4428–4446.
Geman H, Madan DB, Yor M (2001) Asset prices are Brownian motion: Only in business time. Avellaneda M, ed. *Quantitative Analysis in Financial Markets: Collected Papers of the New York University Mathematical Finance Seminar*, vol. II (World Scientific, Singapore), 103–146.
Gilbert EN (1959) Random graphs. *Ann. Math. Statist.* 30(4):1141–1144.
Glasserman P, Young HP (2015) How likely is contagion in financial networks? *J. Bank. Finance* 50:383–399.
Glasserman P, Young HP (2016) Contagion in financial networks. *J. Econom. Literature* 54(3):779–831.
Guan L, Tibshirani R (2022) Prediction and outlier detection in classification problems. *J. Roy. Statist. Soc. Ser. B Statist. Methodology* 84(2):524.
Guo M, De Persis C (2021) Linear quadratic network games with dynamic players: Stabilization and output convergence to Nash equilibrium. *Automatica J. IFAC* 130:109711.
Hellmann T (2013) On the existence and uniqueness of pairwise stable networks. *Internat. J. Game Theory* 42(1):211–237.
Herskovic B (2018) Networks in production: Asset pricing implications. *J. Finance* 73(4):1785–1818.
Hoff PD, Raftery AE, Handcock MS (2002) Latent space approaches to social network analysis. *J. Amer. Statist. Assoc.* 97(460):1090–1098.
Holland PW, Laskey KB, Leinhardt S (1983) Stochastic blockmodels: First steps. *Soc. Networks* 5(2):109–137.
Horn RA, Johnson CR (1994) *Topics in Matrix Analysis* (Cambridge University Press, Cambridge, UK).
Hunter DR, Handcock MS (2006) Inference in curved exponential family models for networks. *J. Comput. Graph. Statist.* 15(3):565–583.
Jackson MO, Pernoud A (2021) Systemic risk in financial networks: A survey. *Annu. Rev. Econom.* 13:171–202.
Jackson MO, Van den Nouweland A (2005) Strongly stable networks. *Games Econom. Behav.* 51(2):420–444.
Jackson MO, Wolinsky A (2003) A strategic model of social and economic networks. Dutta B, Jackson MO, eds. *Networks and Groups* (Springer Verlag, Heidelberg), 23–49.
Kimball MS, Sahm CR, Shapiro MD (2008) Imputing risk tolerance from survey responses. *J. Amer. Statist. Assoc.* 103(483):1028–1038.
Leng Y, Dong X, Wu J, Pentland A (2020) Learning quadratic games on networks. *Proc. Internat. Conf. Machine Learn.* (PMLR, New York), 5820–5830.
Mao X, Sarkar P, Chakrabarti D (2018) Overlapping clustering models, and one (class) SVM to bind them all. Bengio S, Wallach H, Larochelle H, Grauman K, N CB, R G, eds. *Adv. Neural Inform.*

Processing Systems, vol. 31 (Neural Information Processing Systems Foundation, Inc., La Jolla, CA), 2126–2136.

Markowitz H (1952) Portfolio selection. *J. Finance* 7(1):77–91.

Mastromatteo I, Zarinelli E, Marsili M (2012) Reconstruction of financial networks for robust estimation of systemic risk. *J. Statist. Mechanisms Theory Experiment* 2012(3):1–15.

Mazalov V, Chirkova JV (2019) *Networking Games: Network Forming Games and Games on Networks* (Academic Press, Cambridge, MA).

Metrick A (1995) A natural experiment in "Jeopardy!". *Amer. Econom. Rev.* 85(1):240–253.

OECD (2022) OECD statistics. Accessed September 21, 2022, https://stats.oecd.org/.

Papachristou M, Kleinberg J (2022) Allocating stimulus checks in times of crisis. Laforest F, Troncy R, Medini L, Herman I, eds. *Proc. ACM Web Conf.* (Association for Computing Machinery, New York), 16–26.

Paravisini D, Rappoport V, Ravina E (2017) Risk aversion and wealth: Evidence from person-to-person lending portfolios. *Management Sci.* 63(2):279–297.

Peltonen TA, Scheicher M, Vuillemey G (2014) The network structure of the CDS market and its determinants. *J. Financial Stability* 13:118–133.

Sadler E, Golub B (2021) Games on endogenous networks. Preprint, submitted February 2, https://arxiv.org/abs/2102.01587.

Sandberg I, Willson A Jr (1972) Existence and uniqueness of solutions for the equations of nonlinear DC networks. *SIAM J. Appl. Math.* 22(2):173–186.

Squartini T, Caldarelli G, Cimini G, Gabrielli A, Garlaschelli D (2018) Reconstruction methods for networks: The case of economic and financial systems. *Phys. Rep.* 757:1–47.

Tardos E (2004) Network games. *Proc. 36th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 341–342.

Tsanakas A, Christofides N (2006) Risk exchange with distorted probabilities. *ASTIN Bull.* 36(1):219–243.

Upper C (2011) Simulation methods to assess the danger of contagion in interbank markets. *J. Financial Stability* 7(3): 111–125.

Upper C, Worms A (2004) Estimating bilateral exposures in the German interbank market: Is there a danger of contagion? *Eur. Econom. Rev.* 48:827–849.

Vershynin R (2018) *High-Dimensional Probability: An Introduction with Applications in Data Science*, vol. 47 (Cambridge University Press, Cambridge, UK).

Wasserman S, Pattison P (1996) Logit models and logistic regressions for social networks: I. An introduction to Markov graphs and p^*. *Psychometrika* 61(3):401–425.

Yaari ME (1987) The dual theory of choice under risk. *Econometrica* 55(1):95–115.

Zhao L, Chakrabarti D, Muthuraman K (2019) Portfolio construction by mitigating error amplification: The bounded-noise portfolio. *Oper. Res.* 67(4):965–983.

Akhil Jalan is a PhD student in the Department of Computer Science, University of Texas at Austin. His research is on theoretical and applied aspects of network science and statistics.

Deepayan Chakrabarti is an associate professor in the Department of Information, Risk, and Operations Management at McCombs School of Business, University of Texas at Austin. His research includes statistical modeling of networks and decision making under uncertainty.

Purnamrita Sarkar is an associate professor of statistics at the University of Texas at Austin. She is interested in foundational problems arising in networks, resampling, streaming algorithms, and general unsupervised learning.