

FDIC Center for Financial Research

Working Paper

No. 2011-03

SystemicRiskComponentsandDepositInsurancePremia

December 2010



Federal Deposit Insurance Corporation • Center for Financial Research

Systemic Risk Components and Deposit Insurance Premia*

Ming Liu
Morgan Stanley
1585 Broadway
New York, NY 10036-8293, USA

Jeremy Staum[†]
Dept. of IEMS
McCormick School of Engineering
Northwestern University
2145 Sheridan Road
Evanston, IL 60208-3119, USA

December 10, 2010

Abstract

In light of recent events, there have been proposals to establish a theory of financial system risk management analogous to portfolio risk management. One important aspect of portfolio risk management is risk attribution, the process of decomposing a risk measure into components that are attributed to individual assets or activities. The theory of portfolio risk attribution has limited applicability to systemic risk because systems can have richer structure than portfolios. We take a first step towards a theory of systemic risk attribution and illuminate the design process for systemic risk attribution by developing some schemes for attributing systemic risk in an application to deposit insurance.

1 Introduction

Risk attribution can play an important role in systemic risk management, as it does in portfolio risk management. In portfolio risk management, risk attribution means decomposing the risk of a portfolio or enterprise into risk components that are attributed to components of the portfolio. One application is allocating risk capital to sub-portfolios or businesses (Denault, 2001). Another is computing risk-adjusted performance, which can be used to provide managers with incentives to consider the risk profile of the whole portfolio or enterprise when taking risks (Tasche, 2004). In systemic risk management, Gauthier et al. (2010) and Tarashev et al. (2010) use portfolio risk components as tools for setting capital requirements for banks.

Our primary contribution is to the theory and methodology of systemic risk attribution. We take a first step towards a theory of systemic risk attribution that is substantively different from the theory of portfolio risk attribution and is appropriate for use with systemic risk models that include interactions between components of the system. We concur with Tarashev et al. (2010, pp. 6, 9) that the theory of portfolio risk attribution, which treats a portfolio as a sum of assets that do not interact, does not fit models of systemic risk that include interactions. Using simple models of the financial system and a simple measure of systemic risk, we construct systemic risk attribution schemes that handle the interactions between financial institutions. In contrast to the literature on portfolio risk components, we see many reasonable ways to allocate systemic risk in models with interactions. We make a contribution to the design of systemic risk allocation schemes by discussing design principles and giving examples.

*The authors thank discussant Myron Kwast, Rosalind Bennett, Ken Jones, Paul Kupiec, Amiyatosh Purnanandam, and other participants in the 2009 FDIC Center for Financial Research Workshop for their comments. The views expressed are those of the authors. The second author gratefully acknowledges support from the FDIC Center for Financial Research and an IBM Faculty Award.

[†]Corresponding author. Email: j-staum@northwestern.edu

Our second contribution is to apply systemic risk components to deposit insurance. Acharya et al. (2010) use a model of bank resolution costs to analyze the implications of systemic risk for deposit insurance premia. We develop systemic risk attribution schemes in a model inspired by theirs and in a network model of the financial system, and discuss the use of systemic risk components as deposit insurance premia.

The remaining sections of the introduction discuss systemic risk and systemic risk attribution. Deposit insurance and responsibility for externalities in deposit insurance are the topic of Section 2. Section 3 develops systemic risk attribution schemes in a model of bank resolution costs. It contains introductions to the tools we use for creating risk components, the Shapley value (Section 3.2) and Aumann-Shapley value (Section 3.3). Section 4 develops systemic risk attribution schemes in a network model. Section 5 concludes and discusses future research directions.

1.1 Systemic Risk

Systemic risk management includes three challenges: modeling systemic risk, defining a measure of systemic risk, and operationalizing systemic risk measurement by gathering the necessary data. We review some relevant literature and explain how our methods address these issues.

Systemic risk involves risk that arises because of the structure of the financial system and interactions between financial institutions. Systemic risk is not the same as *systematic risk*, which is risk explained by factors that influence the economy as a whole. Systemic risk includes systematic risk and also risks arising from phenomena such as *contagion*, the transmission of losses or distress from one institution to another. Contagion can take several forms. In *asset price contagion*, an institution's sale of assets into an illiquid market can cause a decline in asset prices and thus losses to others. For example, a *fire sale* in which a distressed financial institution is compelled to sell assets can cause losses to other institutions that own similar assets (Brunnermeier and Pedersen, 2009; Cifuentes et al., 2005; Diamond and Rajan, 2005; Krishnamurthy, 2010; Shin, 2008). The models of bank resolution costs in Section 3, inspired by fire sales and by Acharya et al. (2010), include effects that a bank has on the value of other banks' assets due to market illiquidity, as in asset price contagion. In *counterparty contagion*, loss transmission occurs when an institution is unable to make promised payments to others (Eisenberg and Noe, 2001). In Section 4, we use a model of counterparty contagion that extends the model of Eisenberg and Noe (2001) to include banks' depositors, and is a special case of the model of Elsinger (2007).

The systemic risk measure is often taken to be a risk measure, such as value at risk (VaR) or expected shortfall, applied to the distribution of the aggregate losses in the financial system (Adrian and Brunnermeier, 2009; Gauthier et al., 2010; Tarashev et al., 2010). One alternative measure is the fraction of the banking system that enters bankruptcy (Lehar, 2005, §5.1). A different approach is to measure systemic risk as a market price or value. Huang et al. (2009, 2010) measure systemic risk as the price of insurance against losses in excess of 15% to a portfolio of twelve large banks' bonds. In measuring systemic risk from the perspective of a deposit insurer, Lehar (2005, § 5.3) considers the price of a portfolio of put options, each written on the assets of one insured bank, and the volatility of this price. This price is the expected insured loss where the probabilities in the expectation come from a *pricing measure* (see, e.g., Föllmer and Schied, 2004). Under a pricing measure, the expectation of an asset's discounted future value equals the asset's price. This reflects systematic risk and investors' preferences. Our systemic risk measure is the deposit insurer's expected cost. The expectation may be computed either under the pricing measure or with real-world probabilities.

Systemic risk models involving real-world probabilities require a great deal of data to model systematic risk and connect it to outcomes for the system. Modeling systematic risk requires a macroeconomic model and data on each bank's exposure to risk factors. For example, BankCaR (Frye and Pelz, 2008) models the joint distribution of U.S. banks' loan charge-offs using the call report data on banks' assets. Systematic risk is modeled in a similar manner in RAMSI (Aikman et al., 2009), which also incorporates fire sales and counterparty contagion. Modeling fire sales requires a model of how an asset's price depends on the quantity that is sold. It would be valuable to have further study of market liquidity when a substantial portion of the stock of an entire class of assets is sold. Our methods in Section 3 could be operationalized using BankCaR and a model of illiquidity costs in bank resolution. Modeling counterparty contagion requires data

on bilateral links among financial institutions that is only available, in imperfect form, to supervisors in some countries. Elsinger et al. (2006) model systemic risk including counterparty contagion using Austrian supervisory data on banks’ assets and interbank loans. Using Canadian supervisory data that includes also cross-shareholdings and derivative securities traded among banks, Gauthier et al. (2010) model systemic risk including fire sales and counterparty contagion. Our methods in Section 4 could be operationalized using these supervisory data sets.

Public data on market prices, e.g. of equity, equity options, and credit default swaps, suffices to operationalize some methods that involve only a pricing measure (Huang et al., 2009, 2010; Lehar, 2005). This approach is relatively easy to implement, but it does not discriminate between the sources of systemic risk such as systematic risk and various forms of contagion. To distinguish among these sources of risk in models that use a pricing measure, one could build a model with real-world probabilities, as described above, and then choose an equivalent pricing measure (see, e.g., Föllmer and Schied, 2004).

1.2 Systemic Risk Components

To date, schemes for attributing systemic risk to individual institutions have been drawn directly from the theory of portfolio risk components, with the exception of ΔCoVaR (Adrian and Brunnermeier, 2009). Gauthier et al. (2010) is the only study to apply these schemes to a model featuring contagion; others have done systemic risk attribution for models that include only systematic risk. Because the financial system is not simply a portfolio of institutions, we see a need to develop a new theory of systemic risk attribution, different from the theory of portfolio risk attribution.

Within the topic of risk attribution, we focus on *risk components*, which have the property that their sum equals the risk measure of the portfolio or system. In our case, the systemic risk measure is the deposit insurer’s expected cost, the aggregate premium required to insure the system. Summing to the required aggregate premium is a desirable property for systemic risk components to have if they are to be used as deposit insurance premia.

Several methods have been used to attribute systemic risk to institutions, but not all of them yield risk components. The ΔCoVaR measure of an institution’s contribution to systemic risk (Adrian and Brunnermeier, 2009) is not a risk component: the sum of CoVaR or ΔCoVaR over all institutions does not equal VaR of the whole system (Tarashev et al., 2010, p. 4). Likewise, incremental VaR is not a risk component because it does not sum to VaR of the whole system (Gauthier et al., 2010, §2.2). Component VaR and Shapley value have been applied to systemic risk and do yield risk components. Component VaR was applied to systemic risk by Gauthier et al. (2010) in setting capital requirements and by Lehar (2005, §5.3) in deposit insurance. Its name is misleading because component VaR is not related to VaR . It is proportional to beta, the covariance between an asset and the portfolio. Component VaR is the risk component derived from marginal contribution to risk in the special case that the risk measure is variance (see, e.g., Goldberg et al., 2009). In general, this method of risk attribution is also known as gradient allocation, Euler allocation, or the Aumann-Shapley value. Tarashev et al. (2010, p. 1) propose to apply the Shapley value to attribute systemic risk for “any measure of risk that treats the system as a portfolio of institutions.” Denault (2001) proposed the Shapley and Aumann-Shapley values as portfolio risk components.

We show how to apply the Shapley and Aumann-Shapley values to models that include interactions between institutions and do not treat the system as a portfolio of institutions. Tarashev et al. (2010, p. 6) state that their risk attribution schemes, which we call *portfolio Shapley values*, do not handle counterparty contagion. However, Gauthier et al. (2010, §2.4) use a portfolio Shapley value with a model that includes counterparty contagion. The resulting risk components should be interpreted carefully. In the setting of systemic risk, the usual interpretation of Shapley value involves the impact on systemic risk of an institution’s participation in the system, assessed by comparing the risk of a system in which the institution participates to the risk of a system in which it does not (Section 3.2). Portfolio Shapley values yield portfolio risk components and assess the impact of an institution’s inclusion in the portfolio by comparing the risk of a portfolio that includes the institution’s profit or loss to the risk of a portfolio that excludes it. They do not compare different systems. To interpret portfolio Shapley values as yielding systemic risk components

that assess the impact of an institution’s participation in the system, one must assume that participating institutions behave the same way regardless of others’ participation. As Tarashev et al. (2010, p. 9) point out, this assumption does not fit a model that includes interactions. Gauthier et al. (2010, footnote 4) recognize the difficulty of assessing the impact of an institution’s participation in the system when modeling counterparty contagion, stating “removing a bank [from the system] would leave holes in the remaining banks’ balance sheets.” We address this problem of holes in balance sheets in Section 4.2 and apply the Shapley and Aumann-Shapley values to a model that includes counterparty contagion in Sections 4.3 and 4.4. We apply them to a simpler model with interactions in Sections 3.2 and 3.3.

Because our systemic risk measure is an expectation, which is linear, our analysis omits an important feature of the theory of portfolio risk attribution: the non-linearity of the risk measure as a function of the portfolio value considered as a random variable. We leave for future research the topic of systemic risk attribution when the systemic risk measure is non-linear. Because of the linearity of the Shapley and Aumann-Shapley values, our systemic risk components are simply expectations of allocations of the deposit insurer’s cost in each scenario. We focus on cost allocation in each scenario. Cost allocation via the Shapley value (Shubik, 1962) and Aumann-Shapley value (Billera and Heath, 1982; Mirman and Tauman, 1982) is well-established. The cost allocation problem is interesting when cost is a non-linear function of the system parameters, as is the case in systemic risk: due to interactions, the cost in insuring a system containing two banks is not necessarily the sum of the costs if each of the two banks existed in isolation. This is the phenomenon that is absent from portfolio risk theory because, in each scenario, portfolio value is linear in the portfolio weights.

Cost allocation is also interesting in systemic risk because there is more than one reasonable specification of cost as a function of system parameters. Different assumptions about responsibility or counterfactuals, i.e., what the cost would be if the system parameters were different, lead to different cost allocations. For example: Who is responsible for a bad loan, the lender, the borrower, or both? How would the system be different if one institution were absent? This freedom to think differently about responsibility and counterfactuals creates flexibility for a systemic risk analyst to design a systemic risk allocation scheme to have desired properties.

2 Deposit Insurance and Responsibility for Contagion

When systemic risk is measured as the deposit insurer’s expected cost, systemic risk components could be used in two ways. They provide insight into the causes of systemic risk, quantifying such things as how much systemic risk is due to one particular bank or to loans to a certain industry. For this purpose, the assumptions about responsibility or counterfactuals are key. For example, one can design a systemic risk allocation that assigns responsibility to defaulting borrowers, or one that assigns responsibility to leverage at banks. A belief about what would happen if a particular bank were absent can ground an assessment of how much systemic risk it causes. A deposit insurer could also use systemic risk components as premia, but this poses a challenging design problem. It requires an understanding of the incentives that such a deposit insurance premium scheme would create for participants in the financial system, and the consequences of those incentives. The potential benefit of such a scheme is that it could provide incentives to act in ways that lower systemic risk.

In this regard, it could do for systemic risk something similar to what the *fair-market* deposit insurance premium scheme (Duffie et al., 2003; Pennachi, 2006) did for systematic risk. In deposit insurance, systematic risk has to do primarily with the effect of dependence among banks’ assets on the distribution of insured loss (Bennett et al., 2005; Jarrow et al., 2003; Kuritzkes et al., 2005; Lehar, 2005). The *actuarially fair* deposit insurance premium for a bank is the expectation, using real-world probabilities, of the loss in insuring that bank’s depositors. The *fair-market* premium is the same except that it uses probabilities that incorporate the market price of systematic risk due to investors’ risk-aversion. Because insured losses tend to occur in bad economic scenarios, in which there is high marginal utility of wealth, fair-market premia tend to be higher than actuarially fair premia. The fair-market premium for a bank with high exposure to systematic risk is

higher than the premium for a similar bank with low exposure to systematic risk. Pennachi (2006) finds that in the absence of fair market deposit insurance premia, banks have an incentive to amass a systematically risky asset portfolio, which increases the variability of insured loss.

It also concentrates insolvencies in bad economic scenarios, in which it can be difficult to resolve insolvent banks by selling them. Acharya et al. (2010) model resolution costs due to limited liquidity of insolvent banks' assets. They also show that the possibility that insolvent banks would be bailed out instead of closed during a systemic crisis creates an incentive for banks to herd and increase systematic risk. To remove this incentive, they propose deposit insurance premia that increase with systematic risk. They avoid the difficulties of cost allocation because their model includes only two banks. We explore the allocation of resolution costs in a model inspired by theirs, but including many banks. In the model of resolution costs in Section 3.1.1, the picture is fairly clear: each insolvent bank causes and has sole responsibility for its demand for liquidity, which depletes the common pool of liquidity. The result is a simple cost allocation problem with negative externalities. The simplicity arises because each bank's solvency and demand for liquidity can be assessed in isolation: only the cost of resolution depends on the entire system.

In models of contagion with more complicated interactions among financial institutions, events can have multiple causes, allowing for multiple perspectives on responsibility. Contagion implies an externality: one institution's behavior has a significant impact on others' survival or value, and it could be held responsible for that impact (Acharya et al., 2010, §1). However, contagion takes two: one to transmit and one to receive contagion. The recipient could be held responsible for behavior that created vulnerability to contagion.

Consider fire sales triggered by capital requirements. Suppose that bank A, in distress, conducted a fire sale, causing the price of an illiquid security to decline, which caused bank B to become insolvent and experience an insured loss. Who is responsible for the insured loss, bank A, bank B, or both? Bank B's behavior is also a cause of the insured loss, which would not have occurred if bank B had had lower leverage or less exposure to the illiquid security. Is bank A responsible for the negative externality of its fire sale, or is bank B responsible for its vulnerability to the externality? Is market liquidity a common good, like financial stability, and are institutions entitled to some normal level of it? In this view, an institution that contributes to systemic risk is like a polluter, held responsible for the harm that others suffer from its pollution (Acharya et al., 2010, §6). Or is market liquidity a resource that anyone is free to exploit? Or is it neither a good nor a resource, but rather a quality describing the competition between sellers to find buyers? In these views, banks are simply competing businesses, each responsible for ensuring that it has sufficient liquidity, or for the consequences if it does not.

Consider counterparty contagion. Suppose that bank B lent to bank A, and then bank A invested in sub-prime mortgages. When too many of the mortgages defaulted, bank A defaulted, and the loss on the loan to bank A caused Bank B to become insolvent and experience an insured loss. Who is responsible for it? Among the causes of the insured loss are bank B's leverage, its decision to lend to bank A, bank A's investment decisions, and the housing market. The effect of bank A's investment decisions on bank B illustrates the externality in counterparty risk: after an obligation is undertaken, the obligor's behavior affects the creditor. Indeed, a limited-liability borrower that has already obtained a loan has an incentive to pursue higher-risk strategies: it would capture all the upside, whereas its lender would bear the extreme losses. Debt covenants are a mechanism for limiting such behavior, which has an attraction for decision-makers only because some of the negative consequences are borne by others, and which would not be preferred if everyone's interests were aligned. Collateral is a mechanism for limiting the damage to the creditor in the event of default.

One might argue that, because creditors can use such mechanisms to protect themselves from the externalities of their obligors' behavior, they should be held responsible for doing so. Likewise, one might argue that, because banks can avoid suffering contagion from fire sales by not investing in illiquid securities, they should be held responsible for doing so. Applied to a model in which the deposit insurer's cost is the sum of insured losses at all banks, these arguments justify the actuarially fair or fair-market premium: each bank would be responsible for the loss in insuring its own deposits and nothing else.

An objection to this argument is that, considering the complexity of modern finance, the adequacy of the available mechanisms is questionable. For example, consider over-the-counter derivatives trades. It is not easy to assess the ability of a large complex financial institution to fulfill its obligation, much less design a

contract with features that control this ability, similar to what debt covenants do. Data such as the obligor’s credit ratings and credit default swap rate provide information about its unconditional creditworthiness, not its ability to pay in those scenarios in which it has an obligation to make a payment on a particular derivative security. The demand to collateralize derivatives exposures is so great that the supply of high-quality collateral is insufficient; reliance on lower-quality collateral is itself a source of systemic risk (Gorton, 2009).

When externalities can not be dealt with by other means, one may consider using regulation in an attempt to reduce their negative consequences. Consider using systemic risk components as deposit insurance premia that are sensitive to effects of an institution’s behavior on the cost to insure all deposits, whether they are its deposits or not. At present, we do not have an attractive solution for dealing with contagion in this way. There are symmetrical objections to deposit insurance premia that hold banks responsible for the loss in insuring their own deposits and to deposit insurance premia that hold banks responsible for the impact that their actions have on insured losses anywhere in the system. The former leave banks exposed to externalities of the behavior of their obligors or of banks that might conduct fire sales; the latter leave banks exposed to externalities of the behavior of their creditors or of banks that might be harmed by fire sales. Under the latter scheme, a bank would be charged a higher deposit insurance premium if its creditors or banks holding the same illiquid assets that it does were to become weaker and their deposits subject to greater insured losses in case of its own default or distress. Furthermore, the latter scheme has the disadvantage that it might well allocate positive systemic risk components to non-banks that are sources of contagion. Even though these institutions might be regarded as responsible for costs to the deposit insurer, charging them deposit insurance premia seems impractical. Nonetheless, models of contagion could still be useful in assessing expected costs, and systemic risk components would be useful in analyzing the causes of systemic risk, even if they were not to be used as deposit insurance premia.

We concur with Acharya et al. (2010) that deposit insurance premia should reflect bank resolution costs. Using systemic risk components as deposit insurance premia would be a way for the deposit insurer to recover total expected costs and to give banks an incentive to reduce systemic risk by making investments that are collectively less likely to result in clustering of bank insolvencies. However, like portfolio risk components, systemic risk components can be negative. Depending on one’s perspective on responsibility for bank resolution costs, this might be appropriate in some sense, but negative deposit insurance premia would entail practical disadvantages. We show how to design schemes that yield non-negative systemic risk components.

3 Resolution Costs

In this section, we treat models of the deposit insurer’s cost to resolve insolvent banks, inspired by Acharya et al. (2010). The key feature of the models is that the deposit insurer incurs an extra cost, due to illiquidity of banks’ assets, in scenarios in which the aggregate assets of insolvent banks are too large. Thus, there is an interaction among banks: the insolvency of one bank can decrease the value of the assets of another insolvent bank.

3.1 Models of Bank Resolution Costs

The system contains n banks. Bank i has d_i in deposits and the rest of its funding is equity. Its assets have book value a_i . Systematic risk enters the model as the randomness of the vector \mathbf{a} . The net book value of bank i is $u_i = a_i - d_i$. Define $U_i = \text{sgn}(u_i)$, U_i^+ indicating whether bank i is solvent, and U_i^- indicating whether bank i is insolvent. The aggregate book values of solvent banks’ assets and insolvent banks’ assets are $\sum_{i=1}^n U_i^+ a_i$ and $\sum_{i=1}^n U_i^- a_i$, respectively. The deposit insurer’s cost is the difference between the total deposits of insolvent banks and the amount for which their assets are sold, $\sum_{i=1}^n U_i^- d_i - s(\sum_{i=1}^n U_i^- a_i, \sum_{i=1}^n U_i^+ a_i)$. For any y , $s(0, y) = 0$ and the function $s(\cdot, y)$ is concave because the average price per dollar of book value for the assets of insolvent banks decreases as the supply of these assets increases.

There is a negative externality of one bank's insolvency: it can make other insolvent banks worth less. We decompose the cost to the deposit insurer as

$$\sum_{i=1}^n U_i^- d_i - s\left(\sum_{i=1}^n U_i^- a_i, \sum_{i=1}^n U_i^+ a_i\right) = \sum_{i=1}^n \ell_i + \tilde{c}\left(\sum_{i=1}^n U_i^- a_i, \sum_{i=1}^n U_i^+ a_i\right) \quad (1)$$

where $\ell_i = u_i^-$ is the book value of insured losses at bank i and the extra cost due to illiquidity is $\tilde{c}(\sum_{i=1}^n U_i^- a_i, \sum_{i=1}^n U_i^+ a_i) = \sum_{i=1}^n a_i - s(\sum_{i=1}^n U_i^- a_i, \sum_{i=1}^n U_i^+ a_i)$. For any y , the function $\tilde{c}(\cdot, y)$ is convex and $\tilde{c}(0, y) = 0$.

We will apply the Shapley and Aumann-Shapley values to allocate the deposit insurer's cost to banks. Both cost allocation procedures are *additive* (see, e.g., Moulin and Sprumont, 2007), which implies that applying the procedure to the left side of Equation (1) yields the same allocation as applying the procedure separately to the two terms on the right side and summing the results. The first term on the right side, the book value of insured losses $\sum_{i=1}^n \ell_i$, is separable: it is merely a sum of costs due to each bank and involves no interactions among them. Both procedures satisfy the principle of *separable costs* (see, e.g., Sudhölter, 1998), which implies that, when applied to $\sum_{i=1}^n \ell_i$, they must allocate ℓ_i to bank i . Overall, the cost allocated to bank i is the sum of ℓ_i and its component of the illiquidity cost $\tilde{c}(\sum_{i=1}^n U_i^- a_i, \sum_{i=1}^n U_i^+ a_i)$.

3.1.1 Fire Sale Model.

In this model, the average price per dollar of book value of assets sold is $\exp(-\alpha x)$ when the book value of assets sold is x dollars. This is the functional form used by Cifuentes et al. (2005) in their model of fire sales. The insolvent banks' assets are sold for $s(x, y) = x \exp(-\alpha x)$ and the illiquidity cost is $\tilde{c}(x, y) = x(1 - \exp(-\alpha x))$. It depends only on the assets of insolvent banks, x , and not on assets of solvent banks, y , so we define a cost function \hat{c} by $\hat{c}(x) = \tilde{c}(x, y)$ for any y . Now we have a simple, classic, well-studied cost allocation problem: allocating the cost required to meet the aggregate demand for a single good to the individuals who make demands (see, e.g., Moulin and Sprumont, 2007; Sudhölter, 1998). We interpret $z_i = U_i^- a_i$ as the demand for liquidity imposed by bank i , $z = \sum_{i=1}^n z_i$ as the aggregate liquidity demanded, and $\hat{c}(z) = z(1 - \exp(-\alpha z))$ as the cost of providing it.

3.1.2 Acquisition Model.

This model is more closely related to that of Acharya et al. (2010). Solvent banks offer the highest price for the assets of insolvent banks, compared to other potential buyers, because they can make the best use of those assets. If the book value x of insolvent banks' assets is small enough, the deposit insurer can sell all of them to solvent banks. If x is too large, some of the assets must be sold to non-bank buyers for a lower price. Our model is intended merely as an illustration, not as a realistic model of the bargaining process that occurs when the deposit insurer sells insolvent banks or their assets. We assume that

- solvent banks pay book value for insolvent banks' assets,
- a solvent bank can at most double its size by acquisition, i.e. its purchases can not exceed its own book value before it made the purchases,
- an insolvent bank can be divided up and its assets sold to different buyers, and
- non-bank buyers pay $1 - \beta$ times book value for insolvent banks' assets, where $0 < \beta < 1$.

If the assets of insolvent banks, $x = \sum_{i=1}^n U_i^- a_i$, are less than the assets of solvent banks, $y = \sum_{i=1}^n U_i^+ a_i$, then solvent banks buy all the assets of insolvent banks for $s(x, y) = x$. Otherwise, the deposit insurer first sells assets with book value y to solvent banks for y and then sells the remaining assets with book value $x - y$ to non-bank buyers for $(1 - \beta)(x - y)$. The total proceeds are $s(x, y) = y + (1 - \beta)(x - y) = (1 - \beta)x + \beta y$. In addition to the negative externality, that the insolvency of one bank makes other insolvent banks worth less,

there is also a positive externality: the existence of a solvent bank makes insolvent banks worth more. The illiquidity cost is $\tilde{c}(x, y) = \beta \max\{x - y, 0\}$. The acquisition model also fits into the univariate cost allocation framework. The net demand for liquidity imposed by bank i is $z_i = -U_i a_i$. A solvent bank has negative demand for liquidity, meaning that it supplies liquidity to the system: if bank i is solvent then it supplies the deposit insurer with the capacity to sell up to a_i of insolvent banks' assets without incurring a illiquidity cost. Let $z = \sum_{i=1}^n z_i$ be the aggregate net demand for liquidity. The illiquidity cost is $\hat{c}(z) = \beta \max\{z, 0\}$.

3.2 The Shapley Value

The Shapley value creates one kind of fairness by using incremental costs and averaging over all possible orderings of participants in the system. Consider the problem of allocating the cost $\hat{c}(\sum_{i=1}^n z_i)$ to the participants' non-negative demands z_1, \dots, z_n . One unfair procedure allocates the cost $\hat{c}(z_1)$ to participant 1 and the incremental cost $\hat{c}(\sum_{j=1}^i z_j) - \hat{c}(\sum_{j=1}^{i-1} z_j)$ to participant i for $i = 2, \dots, n$. If the cost function \hat{c} is strictly convex, the marginal cost is increasing in the aggregate demand. This procedure gives an unfair advantage to participant 1, who is first in line to pay the lowest price per unit, and avoids paying for any of the negative externalities. The Shapley value gives an equal weight to the incremental cost due to participant i under any ordering. Let π be a permutation of $1, \dots, n$ and $\pi(i)$ be its i th element, i.e. the identity of the participant who comes i th in this ordering. Then $\pi^{-1}(i)$ is the position of participant i in this ordering. The Shapley value allocates to demand z_i the cost

$$\frac{1}{n!} \sum_{\pi} \left(\hat{c} \left(\sum_{h=1}^{\pi^{-1}(i)} z_{\pi(h)} \right) - \hat{c} \left(\sum_{h=1}^{\pi^{-1}(i)-1} z_{\pi(h)} \right) \right)$$

where the summation \sum_{π} is over all $n!$ permutations of $1, \dots, n$. For any subset S of $\{1, \dots, n\}$ that contains i , there are $(n - |S|)! (|S| - 1)!$ permutations of $1, \dots, n$ in which i is the $|S|$ th participant and the first $|S|$ participants are in the set S . Therefore the cost allocation is also equal to

$$\frac{1}{n!} \sum_{S \ni i} (n - |S|)! (|S| - 1)! \left(\hat{c} \left(\sum_{j \in S} z_j \right) - \hat{c} \left(\sum_{j \in S \setminus \{i\}} z_j \right) \right). \quad (2)$$

These two formulations of the Shapley value can be found in Moulin and Sprumont (2007) and Denault (2001), along with further theoretical discussion. Because the summation in Equation (2) is over 2^{n-1} sets, it may be computationally infeasible when n is large, but the Shapley value can be approximated by Monte Carlo (David et al., 2005).

The Shapley value applies to more general cost allocation problems. Call the n participants *players*. (The Shapley and Aumann-Shapley values have their roots in cooperative game theory.) Suppose that if only players in the subset $S \subseteq \{1, \dots, n\}$ were participating in the system, the cost would be $c(S)$. It is standard to assume that $c(\emptyset) = 0$, i.e., there is zero cost for a system with no participants. We represent the cost function c as a vector \mathbf{c} whose 2^n elements are the costs associated with all subsets of $\{1, \dots, n\}$. The Shapley value involves a comparison of the actual system to $2^n - 1$ counterfactual systems in which not every player is participating. The Shapley value allocates to participant i the cost

$$(\mathbf{S}\mathbf{c})_i = \frac{1}{n!} \sum_{S \ni i} (n - |S|)! (|S| - 1)! (c(S) - c(S \setminus \{i\})) \quad (3)$$

where \mathbf{S} is a matrix of size $n \times 2^n$ whose coefficients are defined by Equation (3); $\mathbf{S}\mathbf{c}$ is the vector of costs allocated to all n participants.

Because the Shapley value is a linear operator represented by the matrix \mathbf{S} , it commutes with expectation: $\mathbf{S}\mathbf{E}[\mathbf{c}] = \mathbf{E}[\mathbf{S}\mathbf{c}]$. That is, the Shapley value of expected cost is the expectation of the Shapley value of cost. Although it could be more efficient to compute $\mathbf{S}\mathbf{E}[\mathbf{c}]$ than $\mathbf{E}[\mathbf{S}\mathbf{c}]$, we focus on cost allocation in each scenario. The interaction among banks is more transparent within a single scenario.

Designing a cost allocation scheme based on the Shapley value means choosing the players and the cost function c so that the cost allocation has the desired interpretation or properties. An interpretation of what it means for a player to participate determines the cost function c and thus the incremental costs for which the player is responsible. Anything in the system not affected by players' participation is viewed as immutable and not responsible for costs. The special case of Equation (3) in which $c(S) = \hat{c}(\sum_{j \in S} z_j)$ is Equation (2). This is the natural cost function to use when cost depends only on the sum of non-negative demands from the system's participants, as in the fire sale model. In this interpretation, the players are the participants in the system, and for player i to participate means that he imposes his demand z_i on the system; for player i not to participate means that he imposes zero demand. The acquisition model admits this interpretation, but also others, leading to different Shapley values. We also consider Shapley values for which the players are the insolvent banks, or the solvent banks, or banks' leverage.

3.2.1 All Banks

Let all banks be players and the cost of the system in which players in the set S participate be $c(S) = \hat{c}(\sum_{i \in S} z_i) = \tilde{c}(\sum_{i \in S} U_i^- a_i, \sum_{i \in S} U_i^+ a_i) = \beta \max \{ \sum_{i \in S} U_i a_i, 0 \}$. It is as though a bank that does not participate does not exist. This scheme holds an insolvent bank responsible for contributing to the deposit insurer's cost by requiring liquidation and holds a solvent bank responsible for reducing cost by providing liquidity. It is non-negative for an insolvent bank and non-positive for a solvent bank.

3.2.2 Insolvent Banks

One might wish to design a scheme yielding a non-negative cost allocation, especially if its expectation is to be used as a deposit insurance premium. One way to do so is to let the players be only the insolvent banks and the cost be $c(S) = \tilde{c}(\sum_{i \in S} U_i^- a_i, \sum_{i=1}^n U_i^+ a_i)$. There is no allocation of cost to solvent banks. This scheme considers the supply of liquidity to be fixed; it considers only systems in which all the solvent banks exist and supply liquidity. All it does is hold insolvent banks responsible for contributing to the cost.

3.2.3 Solvent Banks

The preceding scheme only takes account of negative externalities. To design a scheme that only takes account of positive externalities, we separately allocate to insolvent banks the cost in the absence of positive externalities and allocate to solvent banks the benefit due to positive externalities. The cost in the absence of positive externalities is $\tilde{c}(\sum_{i=1}^n U_i^- a_i, 0) = \sum_{i=1}^n \beta U_i^- a_i$, because all the insolvent banks' assets would have to be sold to non-banks. By the principle of separable costs, the allocation of this cost to bank i is $\beta U_i^- a_i$, which is its stand-alone cost, the illiquidity cost of resolving it in the absence of any other banks. Next we use the Shapley value to allocate the benefit due to positive externalities, taking only solvent banks as players. The benefit generated by the participation of the set S of solvent banks is $b(S) = \tilde{c}(\sum_{i=1}^n U_i^- a_i, 0) - \tilde{c}(\sum_{i=1}^n U_i^- a_i, \sum_{i \in S} U_i^+ a_i)$. The cost allocation to a solvent bank i is $-(\mathbf{Sb})_i$.

3.2.4 Leverage

The preceding schemes regarded non-participation as absence from the system, but we can imagine other kinds of changes to the system. This scheme considers changes in which some banks' leverage is eliminated by replacing their deposits with equity. In a system in which player i participates, bank i 's deposits are d_i . In a counterfactual system in which player i does not participate, bank i 's deposits are zero and therefore it is solvent. The cost function is $c(S) = \tilde{c}(\sum_{i \in S} U_i^- a_i, \sum_{i=1}^n U_i^+ a_i + \sum_{i \in S} U_i^- a_i)$. The cost allocation is non-negative and holds insolvent banks responsible for the impact on illiquidity cost of their being insolvent and needing liquidation instead of being solvent and providing liquidity.

3.3 The Aumann-Shapley Value

The setup for the Aumann-Shapley value is similar to that for the Shapley value except that we replace the binary distinction between participating or not ($i \in S$ or $i \notin S$) with a continuous variable, the participation level $\lambda_i \in [0, 1]$ of player i . A participation level of 0 means not participating, and a participation level of 1 means participating at the same level as in the actual system. Let $c(\boldsymbol{\lambda})$ represent the cost of the system in which the participation level of player i is λ_i for $\lambda = 1, \dots, n$. Again, it is standard to assume that $c(\mathbf{0}) = 0$: there is zero cost in a system with no participation. To get a unique Aumann-Shapley value, we also assume that there is a set $\mathcal{D} \subseteq [0, 1]$ of Lebesgue measure one such that c is differentiable at $\gamma\mathbf{1}$ for all $\gamma \in \mathcal{D}$, and that the function that maps $\gamma \in [0, 1]$ to $c(\gamma\mathbf{1})$ is absolutely continuous. That is, the first assumption is that c is sufficiently smooth around the line segment $\{\gamma\mathbf{1} : \gamma \in (0, 1)\}$, which we call the *diagonal*, so that its gradient exists almost everywhere along the diagonal; the second is that the total change in cost along the diagonal, $c(\mathbf{1}) - c(\mathbf{0})$, which equals the cost of the actual system, can be recovered by integrating its rate of change along the diagonal. Then the Aumann-Shapley value yields the cost allocation

$$\mathbf{A}c = \int_0^1 \nabla c(\gamma\mathbf{1}) d\gamma. \quad (4)$$

The differentiability assumption does not hold for every c of interest. If not, one may consider a set of cost allocations, each replacing the gradient in Equation (4) with a subgradient of the cost function (Denault, 2001; Tsanakas, 2009). A specific example appears in Section 3.3.1. See also Buch and Dorfleitner (2008) and Cherny and Orlov (2011) on non-differentiability in the theory of risk components.

If the cost function c is positively homogeneous, the gradient is the same for all $\gamma \in (0, 1]$, so $\mathbf{A}c = \nabla c(\mathbf{1})$. That is, the cost allocation to a player equals the sensitivity of cost to the player's participation level. If c is not positively homogeneous, $(\mathbf{A}c)_i$ can be interpreted as the average sensitivity of cost to the participation level of player i as all players' participation levels go from 0 to 1 in lock step.

Because the Aumann-Shapley value in Equation (4) is a linear operator, under appropriate conditions, it commutes with expectation: $\mathbf{A}E[c] = E[\mathbf{A}c]$. That is, the Aumann-Shapley value of expected cost is the expectation of the Aumann-Shapley value of cost. Let $c(\boldsymbol{\lambda}, \omega)$ represent the cost of system $\boldsymbol{\lambda}$ in scenario ω and P be the probability measure.

Theorem 1. *If*

1. *for almost every $\gamma \in (0, 1)$, there exists a neighborhood N_γ of $\gamma\mathbf{1}$ and an integrable random variable L_γ such that, with probability 1, $c(\cdot, \omega)$ is differentiable at $\gamma\mathbf{1}$ and $|c(\boldsymbol{\lambda}', \omega) - c(\boldsymbol{\lambda}'', \omega)| \leq L_\gamma \|\boldsymbol{\lambda}' - \boldsymbol{\lambda}''\|$ for any $\boldsymbol{\lambda}', \boldsymbol{\lambda}'' \in N_\gamma$, and*
2. *with probability 1, the positive and negative parts of $\nabla c(\cdot, \omega)$ have finite integrals along the diagonal,*

then $\mathbf{A}E[c] = E[\mathbf{A}c]$.

Proof. Condition 2 implies that $E[\mathbf{A}c] = \int_\Omega \int_0^1 \nabla c(\gamma\mathbf{1}, \omega) d\gamma dP(\omega) = \int_0^1 \int_\Omega \nabla c(\gamma\mathbf{1}, \omega) dP(\omega) d\gamma$ by Fubini's theorem. It follows from the more general argument in Glasserman (2004, §7.2.2) that Condition 1 implies that for almost every $\gamma \in (0, 1)$, $\int_\Omega \nabla c(\gamma\mathbf{1}, \omega) dP(\omega) = \nabla(\int_\Omega c(\gamma\mathbf{1}, \omega) dP(\omega)) = \nabla E[c(\gamma\mathbf{1})]$. \square

The Aumann-Shapley value provides a different kind of fairness than the Shapley value. Consider again the problem of allocating the cost $\hat{c}(z)$ to the demands z_1, \dots, z_n , where $z = \sum_{i=1}^n z_i$. With the interpretation that the participation level λ_i of player i is the fraction of his actual demand z_i that he demands in the counterfactual system specified by $\boldsymbol{\lambda}$, we get $c(\boldsymbol{\lambda}) = \hat{c}(\sum_{i=1}^n \lambda_i z_i)$. The Aumann-Shapley value allocates to participant i the cost $\int_0^1 z_i \hat{c}'(\gamma z) d\gamma = (z_i/z) \int_0^z \hat{c}'(u) du = z_i(\hat{c}(z)/z)$. This allocation is *average-cost pricing* because the price all participants pay per unit of demand is the average cost $\hat{c}(z)/z$. If the cost function is positively homogeneous, this allocation is also *marginal pricing* because $\hat{c}(z) = \hat{c}'(z)z$, so the price equals the marginal cost $\hat{c}'(z)$.

Table 1: Prices associated with bank assets in Aumann-Shapley allocations of illiquidity cost in the acquisition model.

| Allocation Scheme | Scenario | Price p_- for Insolvent Banks | Price p_+ for Solvent Banks |
|-------------------|------------|--|-------------------------------|
| All Banks | $x < y$ | 0 | 0 |
| | $x = y$ | equal to each other, between 0 and β | |
| | $x > y$ | β | β |
| Insolvent Banks | $x \leq y$ | 0 | 0 |
| | $x \geq y$ | $\beta(1 - y/x)$ | 0 |
| Solvent Banks | $x \leq y$ | β | $\beta x/y$ |
| | $x \geq y$ | β | β |

Now suppose that the demands are random. If $\mathbf{AE}[c] = \mathbf{E}[\mathbf{Ac}]$, the Aumann-Shapley value allocates the expected cost $\mathbf{E}[(\hat{c}(z)/z)z_i]$ to participant i . The average cost $\hat{c}(z)/z$ serves as a state price for each scenario. For example, in the fire sale model, the average illiquidity cost is $1 - \exp(-\alpha z)$, where z is the aggregate book value of insolvent banks' assets. Thus, a bank gets a higher cost allocation if it has a greater tendency to become insolvent in scenarios in which the aggregate assets of insolvent banks are larger. This provides an incentive for banks to hold assets with lower systematic risk.

Designing a cost allocation scheme based on the Aumann-Shapley value means identifying players who bear the responsibility for costs and building a cost function c whose response to the participation level λ_i of player i identifies the cost impact for which player i is responsible. Average-cost pricing is a natural choice in the fire sale model and the acquisition model. In the acquisition model, we might choose to distinguish between the positive and negative demands, i.e. insolvent and solvent banks, as we did in Section 3.2. This leads to three schemes, analogous to those in Sections 3.2.1–3.2.3. They all yield a cost allocation to bank i of the form $(p_- U_i^- - p_+ U_i^+) a_i$, where p_- and p_+ are prices of assets for insolvent and solvent banks, respectively. Table 1 shows these prices, which depend on whether the aggregate assets x in insolvent banks exceed the aggregate assets y in solvent banks, because the illiquidity cost is zero if $x \leq y$ and is $\beta(x - y)$ if $x \geq y$.

There is no Aumann-Shapley value scheme analogous to the Shapley value scheme in which leverage is set to zero (Section 3.2.4). The cost in the acquisition model is discontinuous with respect to the deposits \mathbf{d} at a point where $d_i = a_i$ for any i , meaning that bank i is on the border between solvency and insolvency. Consequently, we can not find a suitably differentiable cost function c that shows how substituting deposits for equity is responsible for the illiquidity cost.

3.3.1 All Banks

Let all banks be players and let the participation level λ_i multiply the size of bank i 's deposits and assets. Participation levels have no effect on solvency. Then $c(\boldsymbol{\lambda}) = \tilde{c}(\sum_{i=1}^n \lambda_i U_i^- a_i, \sum_{i=1}^n \lambda_i U_i^+ a_i) = \hat{c}(\sum_{i=1}^n \lambda_i U_i a_i) = \beta \max\{\sum_{i=1}^n \lambda_i U_i a_i, 0\}$. If the aggregate net demand $\sum_{i=1}^n U_i a_i$ is non-zero, the Aumann-Shapley value is given by average-cost pricing. Because c is positively homogeneous, this is also marginal pricing. The price is β if the assets of insolvent banks exceed the assets of solvent banks, and 0 if the reverse is true. This scheme holds banks responsible for demanding or supplying liquidity, but only in scenarios in which a marginal change in a bank's size affects the illiquidity cost. If the assets $\sum_{i=1}^n U_i^- a_i$ of insolvent banks and $\sum_{i=1}^n U_i^+ a_i$ of solvent banks are non-zero and equal to each other, then c is not differentiable at $\gamma \mathbf{1}$ for any $\gamma \in (0, 1)$. The set of subgradients at $\gamma \mathbf{1}$, for any $\gamma \in (0, 1)$, is $\{-p\mathbf{U} : 0 \leq p \leq \beta\}$, corresponding to a price p anywhere between 0 and β . When aggregate net demand is zero, the concept of average cost breaks down, and the concept of marginal cost becomes ill-specified: the marginal cost is β for increases in aggregate net demand but 0 for decreases in aggregate net demand.

3.3.2 Insolvent Banks

To get a non-negative cost allocation, we modify the preceding scheme by letting only insolvent banks be players. Then $c(\boldsymbol{\lambda}) = \tilde{c}(\sum_{i=1}^n \lambda_i U_i^- a_i, \sum_{i=1}^n U_i^+ a_i) = \beta \max \{ \sum_{i=1}^n \lambda_i U_i^- a_i - \sum_{i=1}^n U_i^+ a_i, 0 \}$. Again, the Aumann-Shapley value is given by average-cost pricing, but here demand is measured as the assets of insolvent banks only. The price is $\beta \max \{ \sum_{i=1}^n U_i^- a_i - \sum_{i=1}^n U_i^+ a_i, 0 \} / \sum_{i=1}^n U_i^- a_i$. This scheme holds insolvent banks responsible for demanding liquidity, but only in scenarios in which a marginal change in a bank's size affects the illiquidity cost.

3.3.3 Solvent Banks

This scheme is similar to that in Section 3.2.3. To apply the Aumann-Shapley value to the benefit of positive externalities to solvent banks, we express the benefit as a function of their participation levels: $b(\boldsymbol{\lambda}) = \tilde{c}(\sum_{i=1}^n U_i^- a_i, 0) - \tilde{c}(\sum_{i=1}^n U_i^- a_i, \sum_{i=1}^n \lambda_i U_i^+ a_i)$. The non-negative allocation \mathbf{Ab} of the benefit is given by average-cost pricing in which the price for solvent banks' assets is $\beta \max \{ \sum_{i=1}^n U_i^- a_i / \sum_{i=1}^n U_i^+ a_i, 1 \}$. The cost allocation to a solvent bank i is $-(\mathbf{Ab})_i$. As in Section 3.2.3, the cost allocation to an insolvent bank is its stand-alone cost, which corresponds to the price β for insolvent banks' assets.

3.4 Incentives: Monotonicity and Mergers

There is no perfect cost allocation principle: none satisfies all of the axioms that one might like (Moulin and Sprumont, 2007, §5). The Shapley value has a monotonicity property that the Aumann-Shapley value lacks; the Aumann-Shapley value has a property of invariance to mergers and splits that the Shapley value lacks.

Suppose that the system arises from each player's choice of its participation level. *Monotonicity* means that a player can not improve its cost allocation by making a choice that increases the system's cost. Lack of monotonicity implies that the cost allocation principle creates an incentive for players to do something bad for the system as a whole. For example, in the acquisition model, consider a scenario in which $\sum_{i=2}^n U_i^+ a_i < \sum_{i=1}^n U_i^- a_i < \sum_{i=1}^n U_i^+ a_i$. The illiquidity cost is zero, so the all-banks Aumann-Shapley scheme (Section 3.3.1) gives bank 1, which is solvent, a zero cost allocation. Consider system $\boldsymbol{\lambda}$, where $\lambda_j = 1$ for all $j = 2, \dots, n$. This represents a system changed by player 1 alone. In system $\boldsymbol{\lambda}$, this scheme allocates to bank 1 a cost of $-\beta \lambda_1 a_i$ if $\lambda_1 < (\sum_{i=1}^n U_i^- a_i - \sum_{i=2}^n U_i^+ a_i) / a_i$, or a cost of zero if λ_1 exceeds that threshold. This shows that player 1 could have chosen a participation level λ_1 , different from the participation level $\lambda_1 = 1$ that it chose in the actual system, that would have resulted in a lower cost allocation to player 1 but a greater cost to the deposit insurer. The discontinuity in this example also illustrates how the Aumann-Shapley value can be quite sensitive to changes in the system.

Suppose that it is possible for players to merge or split. The simplest way to think about this, in the context of our model of bank resolution costs, is *not* to think of mergers between independent banks, which have a substantive effect on the system. For example, a merger between a solvent bank and an insolvent bank changes the aggregate assets of solvent banks. Instead, we think of mergers or splits of players with no effect on the system, only an effect on how players are counted and on the set or space of counterfactual systems that we imagine when allocating costs. For example, a split of player i into new players i and $n+1$ replaces the cost function c_n on $[0, 1]^n$ with the cost function c_{n+1} on $[0, 1]^{n+1}$ given by $c_{n+1}(\boldsymbol{\lambda}) = c_n(\boldsymbol{\lambda}')$ such that $\lambda'_i = \lambda_i + \lambda_{n+1}$ and $\lambda'_j = \lambda_j$ for $j \neq i$. *Invariance to mergers and splits* means that, whether two players merge into one or one player splits into two, the sum of the costs allocated to the two parts equals the cost allocated to the combined player. In the model of bank resolution costs, with banks as players, the split of one bank into two banks that are smaller but otherwise identical would count as splitting a player. Charging deposit insurance premia to bank holding companies instead of to banks would be a matter of merging players. A cost allocation principle that lacks invariance to mergers and splits could create undesirable incentives for merging and splitting or arguments about the level of consolidation at which deposit insurance premia should be charged. For example, in the acquisition model, consider a scenario in which there are three banks: bank A is solvent with assets of \$200 million, banks B and C are insolvent with assets of \$150 million and \$100 million. The illiquidity cost is \$10 million, based on $\beta = 0.2$. The all-banks Shapley value

Table 2: Probabilities of bank solvency (+) and insolvency (-) in examples of the acquisition model.

| Bank | A | + | + | + | + | - | - | - | - |
|---------|---------------------|-------|------|------|-------|------|-------|-------|--------|
| | B | + | + | - | - | + | + | - | - |
| | C | + | - | + | - | + | - | + | - |
| Example | Independent | 93.5% | 2.1% | 2.1% | 0.05% | 2.1% | 0.05% | 0.05% | 0.001% |
| | Equally Correlated | 95.0% | 1.2% | 1.2% | 0.4% | 1.2% | 0.4% | 0.4% | 0.2% |
| | B,C More Correlated | 95.8% | 0.4% | 0.4% | 1.2% | 1.5% | 0.1% | 0.1% | 0.5% |

Table 3: Allocations of expected illiquidity cost in examples of the acquisition model.

| Example | | Independent | | | Equally Correlated | | | B,C More Correlated | | |
|---------------------------------|----------------|-------------|------|------|--------------------|-----|----|---------------------|-----|-----|
| Expected Illiquidity Cost | | 25.6 | | | 354 | | | 413 | | |
| Allocation Scheme and Principle | | A | B | C | A | B | C | A | B | C |
| All Banks | Shapley | 119 | -57 | -36 | 245 | 69 | 40 | 131 | 156 | 126 |
| | Aumann-Shapley | 30.6 | -3.0 | -2.0 | 320 | 20 | 14 | 224 | 113 | 75 |
| Insolvent Banks | Shapley | 14.9 | 6.7 | 4.0 | 187 | 103 | 64 | 191 | 134 | 88 |
| | Aumann-Shapley | 15.3 | 7.4 | 2.9 | 191 | 109 | 55 | 192 | 135 | 86 |
| Solvent Banks | Shapley | 228 | -100 | -103 | 284 | 55 | 16 | 137 | 148 | 127 |
| | Aumann-Shapley | 156 | -35 | -95 | 242 | 91 | 20 | 123 | 182 | 107 |
| Leverage | Shapley | 12.8 | 8.6 | 4.2 | 169 | 119 | 66 | 186 | 138 | 89 |

(Section 3.2.1) allocates costs of about -\$21.7 million to bank A, \$18.3 million to bank B, and \$13.3 million to bank C. If banks B and C were merged, their cost allocation would be \$30 million, less than the \$31.7 million total for these two banks when they are separate. This is an incentive for them to merge, which is not desirable from a systemic perspective.

3.5 Numerical Examples

To illustrate the seven systemic risk component schemes and the differences among them, we use some simple examples of the acquisition model. The examples all feature three banks called A, B, and C, and the examples differ only in the correlation among the banks' asset returns. The initial book values of the assets of banks A, B, and C are 200, 150, and 100 million dollars, respectively. Each bank is financed by deposits equal to 95% of its assets and equity equal to 5% of its assets. The interest rate on deposits is zero and the return on a bank's assets is 3% with probability 97.8% and is -20% with probability 2.2%. In these examples, a bank is solvent if its asset return is positive and insolvent if its asset return is negative. The illiquidity cost ratio $\beta = 0.2$. In the first example, the banks' asset returns are independent. In the second example, they are correlated in a symmetric manner. The conditional probability that a bank has a loss is 1.2% given that the other two do not, is 25% given that exactly one of the other two does, and is 33% given that both the others do. The third example is like the second, but with greater correlation between banks B and C. Table 2 gives the probabilities for each of the eight scenarios. Table 3 reports systemic risk components, which are allocations of expected illiquidity cost to banks A, B, and C. All numbers are in thousands of dollars. The allocations (rows of three numbers) sum to the expected illiquidity cost (number at the head of the column), up to rounding error. They may be compared to the expected book value of insured losses at banks A, B, and C, which are 660, 495, and 330 thousand dollars, respectively.

The example of independent banks exhibits dramatic differences among schemes. The all-banks and solvent-banks schemes, which give credit to banks for supplying liquidity, provide extreme allocations: the systemic risk components for banks B and C are negative, whereas bank A's exceeds the systemic risk—by more than a factor of 8, when using the solvent-banks Shapley value. The reason that bank A fares poorly

under these schemes is that it is so big that not all of its assets are useful in providing liquidity when it is solvent and other banks fail. Thus, when computing the Shapley value, the incremental benefit of bank A's solvency tends to be a smaller fraction of its size than for other banks; when computing the Aumann-Shapley value, the conditional expectation of the price of solvent banks' assets given that bank A is solvent is smaller than for other banks. The insolvent-banks and deposits schemes exhibit a modest penalty for A's excessive size: they allocate 50-60% of systemic risk to bank A, which has 44% of the assets in the system. In the examples with correlation, there are no negative systemic risk components because banks are less likely to be solvent when other banks are insolvent, and thus earn less credit for supplying liquidity. Still, the all-banks and solvent-banks schemes are least favorable to bank A in the equally-correlated example.

The third example illustrates that having higher-than-average systematic risk, as well as being larger than average, can cause a bank's systemic risk component to be disproportionately large compared to its size. Most importantly, comparing the three examples, we see that increasing the correlation between the assets of banks increases their systemic risk components: they have the correct sensitivity to actions that increase systemic risk. This illustrates how using the systemic risk components to set deposit insurance premia would create an incentive not to herd in asset selection.

Comparing the second and third examples, we see that when banks B and C become more correlated, bank A's systemic risk component decreases sharply under the all-banks and solvent-banks schemes, but increases slightly under the other schemes. The former effect occurs because there is an increased probability of scenario $+ - -$, in which the liquidity that bank A supplies is most valuable. The latter effect has to do with an increased probability of scenario $- - -$, in which bank A's insolvency costs the most, because there is no solvent bank available to purchase insolvent banks' assets from the deposit insurer.

4 Counterparty Contagion

In this section, we treat a model of systemic risk with counterparty contagion. It is a generalization of the model of Eisenberg and Noe (2001) and a special case of the model of Elsinger (2007). There is a network whose edges represent unsecured liabilities that nodes owe one another. The amount that a node is able to pay depends on the amount that it receives from other nodes. Thus, there is contagion in the sense that the default of one node on its liability to a second node can cause the second node to default. If the second node is a bank, contagion can cause a loss in insuring its depositors.

4.1 A Model with Counterparty Contagion

The financial system contains n nodes. Node i receives a non-negative cashflow of e_i . Systematic risk is present in the model in the randomness of the vector \mathbf{e} . If node i is a bank, it has deposits $d_i > 0$, all insured; otherwise, $d_i = 0$. Node i also has liabilities to other nodes. They are junior to deposits. Let L_{ij} represent the liability of node i to node j for all $i \neq j$, and $L_{ii} = 0$ for all i . The liabilities of node i sum to its promised payment $\bar{p}_i = \sum_{j \neq i} L_{ij}$. Some nodes may be unable to pay their liabilities in full. Let p_i be the total amount that node i pays to other nodes. Because of the equal priority of liabilities to each node, it is helpful to define $\Pi_{ij} = L_{ij}/\bar{p}_i$, the fraction of the liabilities of node i that are owed to node j . (If $\bar{p}_i = 0$, let $\Pi_{ij} = 0$ for all j .) The amount that node i pays to node j is $p_i \Pi_{ij}$.

We have a model of network flows: the matrix of flows between nodes is $\text{diag}(\mathbf{p})\mathbf{\Pi}$ where \mathbf{p} is the *payment vector*. The vector of promised payments is $\bar{\mathbf{p}} = \mathbf{L}\mathbf{1}$. The vector of total outflows from each node is \mathbf{p} and the vector of total inflows into each node is $\mathbf{\Pi}^\top \mathbf{p}$, so $(\mathbf{\Pi}^\top - \mathbf{I})\mathbf{p}$ is the vector of net flow into each node. Define $\mathbf{w} = \mathbf{e} - \mathbf{d}$, the vector of *primary value* at each node, before payments are made on liabilities between nodes. Where \mathbf{u} represents *terminal value* at each node, the *balance equation* is

$$\mathbf{u} = \mathbf{w} + (\mathbf{\Pi}^\top - \mathbf{I})\mathbf{p}. \quad (5)$$

At node i , the terminal equity value is $v_i = u_i^+$ and the insured loss is $\ell_i = u_i^-$. The flows must also satisfy the *capacity constraints* $\mathbf{0} \leq \mathbf{p} \leq \bar{\mathbf{p}}$, $\mathbf{0} \leq \ell \leq \mathbf{d}$, and $\mathbf{v} \geq \mathbf{0}$ and the *priority constraints* $v_i > 0 \Rightarrow p_i = \bar{p}_i$ and $p_i > 0 \Rightarrow \ell_i = 0$ for all $i = 1, \dots, n$.

Eisenberg and Noe (2001) and Elsinger (2007) discuss the existence, uniqueness, and computation of a *clearing payment vector*, which satisfies the balance equation and constraints. We assume that there is a unique clearing payment vector \mathbf{p}^* . Let \mathbf{u}^* , \mathbf{v}^* , and $\boldsymbol{\ell}^*$ represent terminal value, terminal equity, and insured loss when $\mathbf{p} = \mathbf{p}^*$.

The systemic risk measure is the expected insured loss $E[\mathbf{1}^\top \boldsymbol{\ell}^*]$. We use the Shapley and Aumann-Shapley values to create systemic risk components. The Shapley value commutes with expectation, and so does the Aumann-Shapley value, under the conditions of Theorem 1. Therefore we focus on allocating the cost $\mathbf{1}^\top \boldsymbol{\ell}^*$ within a single scenario.

4.2 Bilateral Deals and Financial System Structure

To apply the Shapley or Aumann-Shapley values, we must choose a set of m players to whom cost will be allocated, and create a cost function that responds to their participation levels. The natural correspondence between subsets of $\{1, \dots, m\}$ and vertices of $[0, 1]^m$ allows us to use this cost function for both Shapley and Aumann-Shapley values. To do this in a way that yields useful interpretations, we create a framework depicted in (6). In this section, we will describe one way in which the financial network, specified by the data $(\mathbf{e}, \mathbf{d}, \mathbf{L})$, can arise as a function Ψ of fundamental data on individual nodes and bilateral deals between nodes. In Sections 4.3–4.4, we create cost allocation schemes by expressing the fundamental data as a function Φ of players' participation levels, thus assigning to players responsibility for the system's structure. The composition $\Psi \circ \Phi$ maps a vector $\boldsymbol{\lambda} \in [0, 1]^m$ of participation levels to network data $\mathbf{e}(\boldsymbol{\lambda})$, $\mathbf{d}(\boldsymbol{\lambda})$, and $\mathbf{L}(\boldsymbol{\lambda})$. Finally, the cost $c(\boldsymbol{\lambda}) = \mathbf{1}^\top \boldsymbol{\ell}^*(\Psi(\Phi(\boldsymbol{\lambda})))$, where $\boldsymbol{\ell}^*(\mathbf{e}, \mathbf{d}, \mathbf{L})$ is the vector of insured losses in the financial network whose data is $(\mathbf{e}, \mathbf{d}, \mathbf{L})$. The structure of the framework is

$$[0, 1]^m \xrightarrow{\Phi} \text{fundamental data} \xrightarrow{\Psi} \underbrace{\mathbb{R}_+^n \times \mathbb{R}_+^n \times \mathbb{R}_+^{n^2}}_{\text{network data}} \xrightarrow{\mathbf{1}^\top \boldsymbol{\ell}^*} \mathbb{R}_+. \quad (6)$$

We next provide a specification of Ψ using two kinds of deals: loans and swaps.

4.2.1 Loans and Swaps

So far, we have discussed only what happens at a single instant when payment on liabilities is due. However, loans affect the financial system not only when they mature (time 1), but also when they are made (time 0). Let the loan from node i to node j have principal D_{ij} and per-period interest rate \bar{r}_{ij} . Node j 's time-0 liabilities are $\sum_{i \neq j} D_{ij}$ of borrowing from other nodes, s_j of equity, and d_j of deposits. (For simplicity, we assume zero interest on deposits.) Node j 's time-0 assets are $\sum_{k \neq j} D_{jk}$ of lending to other nodes and t_j of its baseline asset. We distinguish two kinds of nodes, *financial* and *industrial*, by their baseline assets. For financial nodes, the baseline asset is cash. For simplicity, we assume that cash earns a zero rate of return. Each industrial node invests entirely in its own baseline asset, which is a productive technology with lifespan of one period. Industrial nodes do not make loans.¹ The accounting equation states that assets equal liabilities: where a_j is the time-0 size of node j 's balance sheet,

$$\mathbf{a} = \mathbf{t} + \mathbf{D}\mathbf{1} = \mathbf{s} + \mathbf{d} + \mathbf{D}^\top \mathbf{1}. \quad (7)$$

Like loans, each node's baseline asset matures at time 1. Let r_i be the rate of return on the baseline asset of node i . Then the time-1 cashflows generated by baseline assets are

$$\mathbf{e} = (\mathbf{I} + \text{diag}(\mathbf{r}))\mathbf{t}. \quad (8)$$

In this model, the fundamental data is the payments \mathbf{X} owed on swaps, the principal \mathbf{D} and interest rates $\bar{\mathbf{r}}$ on loans between nodes, the rates of return \mathbf{r} on baseline assets, and time-0 equity \mathbf{s} and deposits \mathbf{d} . We

¹This inessential assumption leads to a simpler presentation. Counterparty contagion between industrial firms, involving phenomena such as credit issued by suppliers, has been studied by Kiyotaki and Moore (1997) and many authors since.

assume that all of this data is non-negative, except for \mathbf{r} . We assume that there is no netting agreement or collateral for swaps. Then

$$\mathbf{L} = \mathbf{X} + \mathbf{D} \circ (\mathbf{I} + \bar{\mathbf{r}}), \quad (9)$$

where \circ represents the element-wise product. Equations (7)–(9) specify the function Ψ . Our model assumes that the equity is held by shareholders outside the financial system and that swap payments, interest payments, and principal repayments have equal seniority. These assumptions can be dropped using the results of Elsinger (2007).

We need to imagine systems in which a loan is absent or smaller while avoiding the problem of “holes in... banks’ balance sheets” (Gauthier et al., 2010) in a counterfactual system in which a loan is smaller or absent. We do this by plugging the holes, adjusting each node’s time-0 balance sheet to maintain equality of its time-0 assets and liabilities when a loan shrinks. There is more than one way to do this, and, as we shall see, the scheme for adjusting balance sheets has a profound impact on cost allocations. We consider two schemes. In both, the time-0 size of the lender’s balance sheet is fixed, and a loan substitutes for the baseline asset on the lender’s balance sheet: if the loan from node i to node j is eliminated, the investment of node i in its baseline asset increases to $t_i + D_{ij}$. In the *debt/equity* scheme, suitable for use with the Shapley or Aumann-Shapley value, the time-0 size of the borrower’s balance sheet is also fixed, and a loan substitutes for equity on the balance sheet of the borrower: if the loan from node i to node j is eliminated, the equity of node j increases to $s_j + D_{ij}$. In the *debt/baseline* scheme, a loan affects the size of the borrower’s balance sheet: if the loan from node i to node j is eliminated, the investment of node j in its baseline asset decreases to $t_j - D_{ij}$. There is a difficulty in using this model with the Shapley value: $t_j - D_{ij}$ might be negative, representing an infeasible system. In Section 4.4.7 we will see that this need not be an obstacle to using the debt/baseline scheme with the Aumann-Shapley value.

4.2.2 A Simpler System of Loans Only

The cost allocation problem is interesting even in a simple special case of this framework. In this special case, we can obtain simpler, more illuminating formulae for cost allocations. There are also schemes that yield non-negative cost allocations in this special case.

Definition 1. *A classic lending system has no swaps and no loans between financial nodes, only loans from financial nodes to industrial nodes.*

There are no cycles in the graph of a classic lending system. The debt/baseline scheme can be used with the Shapley value in a classic lending system because any borrower is an industrial node whose investment in its baseline asset exceeds its debt.

Definition 2. *In a system with competitive rates, interest rates depend only on the borrower: there is a vector $\bar{\mathbf{r}}$ such that $\bar{r}_{ij} = \bar{r}_j$ for all i and j .*

In a system with competitive rates, the recovery rate on principal for any loan to node j is $\rho_j = (1 + \bar{r}_j)p_j^*/\bar{p}_j$. In a classic lending system with competitive rates, ρ_j is $1 + \bar{r}_j$ if node j pays its obligations in full (i.e., $p_j^* = \bar{p}_j$) and $e_j/\sum_{i \neq j} D_{ij}$ if it does not, in which case $p_j^* = e_j$. The insured loss at bank i is

$$\ell_i^* = \sum_{j \neq i} D_{ij}(1 - \rho_j)^+ - \sum_{j \neq i} D_{ij}(\rho_j - 1)^+ - s_i. \quad (10)$$

The terms are node i ’s loan losses, profits on loans, and time-0 equity. The first of these can be interpreted as a demand for money at bank i , the latter two as a supply of it. Some of the same issues are present in the cost allocation problem in the counterparty contagion model as in the acquisition model of bank resolution costs. Which of these three sources of demand and supply do we hold responsible for the insured losses? How do we handle the indirect interactions among nodes due to their interactions with a common entity? In the model of bank resolution costs, all banks access a common pool of liquidity; in the counterparty contagion model, some nodes have creditors or obligors in common.

4.3 Cost Allocations from the Shapley Value

Who are the players, and what happens to the fundamental data when a player does not participate? An answer to this question specifies the function Φ in (6).

The Shapley value is based on a fair division among players of the costs due to their interactions. It would be questionable to seek fairness between, e.g., a swap and a bank's equity. We consider only schemes in which all players are nodes or all players are deals. One choice is which type of node or deal is a player, e.g., all nodes, only banks, or only industrial nodes. For example, one could design a systemic risk attribution scheme that identifies the industries that contribute most to systemic risk by using only industrial nodes as players.

We must also specify what participation means. Based on Equation (9), a swap that does not participate simply disappears. According to the debt/equity scheme, a loan that does not participate disappears, leaving in its place cash on the lender's balance sheet and equity on the borrower's balance sheet. We consider three alternatives for nodes' participation. First, if a bank does not participate in the deposit insurance program, replace its insured deposits with uninsured deposits. The cost function is simply $c(S) = \sum_{i \in S} \ell_i^*$ because the deposit insurer pays claims only at banks in S . The cost allocated to bank i is its own depositors' insured losses ℓ_i^* . Second, if a bank does not participate in the deposit-taking function of the banking system, replace its deposits with equity. Each bank is responsible for the effect its deposits have on insured losses anywhere in the system: because of the seniority of deposits to liabilities to other nodes, substituting deposits for equity at a bank can cause insured losses at its creditors. Third, if a node does not participate in the financial system, eliminate it and any deal in which it is involved. Each node is responsible for the effects anywhere in the system of its deposits, equity, and the deals to which it is a party.

Another choice is whether players include all nodes or deals of the chosen type or only a class of bad nodes or deals. We can design schemes that yield non-negative systemic risk components by choosing only players that have non-negative incremental costs of participation in Equation (3), while ensuring that the cost of the system in which none of the players participate is zero. A bank with insured losses has a positive incremental cost of participation: it adds to the total insured losses, and due to the seniority of deposits, it makes no payments to other nodes. A *loss-making loan* is one whose principal exceeds the time-1 payment that the borrower makes. Its incremental cost of participation is non-negative: deleting it, replacing it with cash on the lender's balance sheet and equity on the borrower's balance sheet as prescribed in Section 4.2, makes both lender and borrower better able to fulfill their obligations. Analogously, we define a *loss-making industrial node* as one whose baseline asset experiences a loss rate exceeding its leverage ratio (time-0 equity to assets), i.e., for industrial node j to be loss-making means $r_j < -s_j/a_j$. Such a node imposes a loss on at least one of its lenders. However, because of the equal priority of swap payments, interest payments, and principal repayments, some of its counterparties may profit from their deals with it, and some lenders to an industrial node that is not loss-making may lose from their lending to it. In a classic lending system with competitive rates, an industrial node is loss-making if and only if its recovery rate is less than 1. All counterparties of loss-making industrial nodes experience losses in their deals with those nodes, and all counterparties of the other industrial nodes profit from their deals with those nodes. In such a system, loss-making industrial nodes have non-negative incremental costs of participation. The system from which they have been deleted has zero cost because there are no swaps, and all of the assets have non-negative returns. Therefore the scheme in which the players are the loss-making industrial nodes yields a non-negative allocation.

Combining the considerations above, we describe ten Shapley value schemes in terms of the changes to the system that they contemplate, i.e., the consequences of players' non-participation: (1) replace any bank's insured deposits with uninsured deposits, (2) replace any bank's deposits with equity, (3) eliminate any bank that has insured losses, (4) eliminate any loss-making loan, (5) eliminate any loss-making industrial node, (6) eliminate any bank, (7) eliminate any loan, (8) eliminate any deal, (9) eliminate any industrial node, and (10) eliminate any node. Schemes 1–4 yield non-negative cost allocations because they entail non-negative incremental costs of participation. In a classic lending system with competitive rates, scheme 5 also yields a non-negative cost allocation. Any scheme that allocates cost to deals, such as schemes 4, 7, and 8, can be transformed into a scheme that allocates cost to nodes by re-allocating the cost allocated to a deal to the

Table 4: Aumann-Shapley allocations of insured loss in the contagion model.

| Effect of loan on borrower's balance sheet | Prices | Funding of node j | Swap payment X_{ij} from node i to node j | Loan of D_{ij} at rate \bar{r}_{ij} from node i to node j |
|--|----------|-------------------------|---|---|
| replace equity | marginal | $-\zeta_j q_j$ | $-\theta_{ij} X_{ij}$ | $(\zeta_i - \theta_{ji}(1 + \bar{r}_{ij}))D_{ij}$ |
| replace equity | average | $-\bar{\zeta}_j q_j$ | $-\widehat{\theta}_{ij} X_{ij}$ | $(\widehat{\zeta}_i - \widehat{\theta}_{ji}(1 + \bar{r}_{ij}))D_{ij}$ |
| invest in baseline | marginal | $-\zeta_j s_j(1 + r_j)$ | $-\theta_{ij} X_{ij}$ | $(\zeta_i - \theta_{ji}(1 + \bar{r}_{ij}) - \zeta_j(1 + r_j))D_{ij}$ |

nodes that are parties to the deal.

4.4 Cost Allocations from the Aumann-Shapley Value

The same design issues discussed in Section 4.3 apply to schemes based on the Aumann-Shapley value. We limit our exposition to a few schemes that illustrate the issues particular to applying the Aumann-Shapley value to the model of counterparty contagion. The scheme that substitutes uninsured deposits for insured deposits is simple: based on the cost function $c(\boldsymbol{\lambda}) = \boldsymbol{\lambda}^\top \boldsymbol{\ell}^*$, the Aumann-Shapley value allocates to bank i its own depositors' insured losses ℓ_i^* , the same as the corresponding Shapley value scheme. The other Aumann-Shapley value schemes are not so simple: they involve linear programming (LP) sensitivity analysis. These schemes can all be seen as allocating costs to nodes and to deals and then re-allocating the cost allocated to a deal to the nodes that are parties to the deal. In the obligor-responsibility scheme (Section 4.4.2), the cost allocated to a deal is re-allocated to the obligor, i.e., the borrower in a loan or the party to a swap who is obligated to make a payment. Based on the obligor-responsibility scheme, in Section 4.4.3 we construct a scheme analogous to the Shapley value scheme that eliminates loss-making industrial nodes. This scheme yields a non-negative allocation in a classic lending system with competitive rates. In the creditor-responsibility scheme (Section 4.4.4), the cost allocated to a deal is re-allocated to the creditor, i.e., the lender in a loan or the recipient of a swap payment. It involves a more complicated LP sensitivity analysis than the preceding schemes. The obligor- and creditor-responsibility schemes are extreme in the way they assign responsibility for cost. The shared-responsibility scheme (Section 4.4.5) re-allocates to each party to a deal half the cost allocated to the deal: it yields a cost allocation that is the average of those from the obligor- and creditor-responsibility schemes. Nonetheless, it can still yield an allocation that seems too extreme. The obligor-, creditor-, and shared-responsibility schemes all use marginal prices (see Section 3.3) and the debt/equity scheme. In Section 4.4.6 we develop a scheme using average prices instead of marginal prices; in Section 4.4.7 we use the debt/baseline scheme instead of the debt/equity scheme.

Table 4 shows how these schemes allocate costs to nodes and deals. The marginal prices ζ_i of money at node i and θ_{ij} of the promised payment from node i to node j are given by Equations (13) and (17). The average-cost scheme uses average prices given in Equations (22)–(24). In the debt/equity scheme, deals have no impact on balance sheet sizes, so in the absence of deals, each node i would generate the cashflow $a_i(1 + r_i)$. The primary value (cashflow minus deposits) attributable to node j is $q_j = a_j(1 + r_j) - d_j$. If node j is industrial, $q_j = (s_j + \sum_{i \neq j} D_{ij})(1 + r_j) = e_j$, its actual time-1 cashflow. If node j is financial, $q_j = a_j - d_j = w_j + \sum_{k \neq j} D_{jk} = s_j + \sum_{i \neq j} D_{ij}$, the portion of its financing that is junior to deposits. To summarize, the amount of money each node would have had available at time 1 to pay its creditors if it had not made any loans is

$$\mathbf{q} = \text{diag}(1 + \mathbf{r})\mathbf{a} - \mathbf{d} = \text{diag}(1 + \mathbf{r})(\mathbf{s} + \mathbf{D}^\top \mathbf{1}) = \mathbf{w} + \mathbf{D}\mathbf{1}. \quad (11)$$

In the debt/baseline scheme, the cashflow $D_{ij}(1 + r_j)$ at node j is attributable to the loan from node i to node j , so the primary value attributable to node j is $(a_j - \sum_{i \neq j} D_{ij})(1 + r_j) - d_j = s_j(1 + r_j)$, the time-1 value of investing its equity in its baseline asset.

4.4.1 Technical Issues

We require different properties of the cost function $c = \mathbf{1}\ell^* \circ \Psi \circ \Phi$ for the Aumann-Shapley and Shapley values. For the Shapley value, we need c to be defined on the vertices of a unit hypercube. For the Aumann-Shapley value, we need c to be defined on an open set containing the diagonal $\{\gamma\mathbf{1} : \gamma \in (0, 1)\}$. We are also concerned about differentiability properties of c , for the sake of getting a unique Aumann-Shapley value.

It is permissible to specify the fundamental data as a function Φ of participation levels in such a way that system $\Phi(\boldsymbol{\lambda})$ is infeasible for some $\boldsymbol{\lambda}$ in the unit hypercube, as long as it is feasible for every $\boldsymbol{\lambda}$ in an open set containing the diagonal. For example, some values of $\boldsymbol{\lambda}$ may specify an infeasible system, violating constraints such as non-negativity of time-0 equity $\mathbf{s}(\boldsymbol{\lambda})$ and investment $\mathbf{t}(\boldsymbol{\lambda})$ in the baseline asset. In the schemes we investigate, these constraints hold on some open set containing the diagonal if the time-0 equity and investment in the baseline asset are strictly positive in the actual system.

In the settings we study, getting a unique Aumann-Shapley value depends on differentiability properties of $\mathbf{1}\ell^*$, and positive homogeneity of Ψ is relevant. The cost $\mathbf{1}\ell^*$ is positively homogeneous. We assume that Ψ is differentiable and positively homogeneous, as is true for Ψ specified in Section 4.2. Then positive homogeneity of Φ implies positive homogeneity of c . If c is positively homogeneous, then differentiability of c at $\mathbf{1}$ is sufficient to get a unique Aumann-Shapley value $\mathbf{Ac} = \nabla c(\mathbf{1})$. If c is not positively homogeneous, then it is sufficient that c be differentiable at $\gamma\mathbf{1}$ for all $\gamma \in \mathcal{D}$, where $\mathcal{D} \subseteq [0, 1]$ has Lebesgue measure one. We assume that Φ is differentiable, as is true for all the choices we consider below. Then it is differentiability properties of $\mathbf{1}\ell^*$ on which differentiability of c depends. In Appendix B, we use LP sensitivity analysis to compute the gradient of $\mathbf{1}\ell^*$ where it is differentiable. It is differentiable where the network does not contain a borderline node.

Definition 3. *Node i is borderline if $p_i^* = \bar{p}_i^* > 0$ and $v_i^* = 0$ (the borderline of default on liabilities to other nodes) or $p_i^* = 0$, $\ell_i^* = 0$, and $d_i > 0$ (the borderline of insured loss).*

If there is no borderline node, then the LP has a unique dual-optimal solution, which provides the gradient of $\mathbf{1}\ell^*$; if there is a borderline node, then the dual-optimal solutions provide the subgradients of $\mathbf{1}\ell^*$. If c is positively homogeneous, then there is a unique Aumann-Shapley value $\mathbf{Ac} = \nabla c(\mathbf{1})$ if there is no borderline node; if there is a borderline node, any subgradient of $\mathbf{1}\ell^*$ at the actual system $\Psi(\Phi(\mathbf{1}))$ yields a cost allocation. If c is not positively homogeneous and for almost every $\gamma \in (0, 1)$ the network $\Psi(\Phi(\gamma\mathbf{1}))$ has no borderline nodes, then there is a unique Aumann-Shapley value. Henceforth, we focus on the case of a unique Aumann-Shapley value.

In the model of bank resolution costs, we gave an example in which the Aumann-Shapley value was discontinuous (Section 3.4). At a point $\boldsymbol{\lambda}$ at which the cost function was non-differentiable, there was a discontinuity in its gradient. The same phenomenon occurs in the counterparty contagion model: there can be a discontinuity in the gradient of the cost function at a point $\boldsymbol{\lambda}$ at which the network $\Psi(\Phi(\boldsymbol{\lambda}))$ contains a borderline node.

4.4.2 Obligor Responsibility

This debt/equity scheme gives responsibility for deals to the obligor. The participation level λ_j of node j multiplies the swap payments it has promised to make and everything on the liability side of its time-0 balance sheet: deposits $\mathbf{d}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{d}$, borrowing $\mathbf{D}(\boldsymbol{\lambda}) = \mathbf{D}\text{diag}(\boldsymbol{\lambda})$, and equity $\mathbf{s}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{s}$. The investment in baseline assets $\mathbf{t}(\boldsymbol{\lambda}) = \mathbf{a}(\boldsymbol{\lambda}) - \mathbf{D}(\boldsymbol{\lambda})\mathbf{1} = \text{diag}(\boldsymbol{\lambda})\mathbf{a} - \mathbf{D}\boldsymbol{\lambda}$. The liability matrix $\mathbf{L}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{L}$, so the promised payments $\bar{\mathbf{p}}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\bar{\mathbf{p}}$ and the matrix of payment proportions $\mathbf{\Pi}(\boldsymbol{\lambda}) = \mathbf{\Pi}$. The primary value $\mathbf{w}(\boldsymbol{\lambda}) = (\mathbf{I} + \text{diag}(\mathbf{r}))\mathbf{t}(\boldsymbol{\lambda}) - \mathbf{d}(\boldsymbol{\lambda}) = (\mathbf{I} + \text{diag}(\mathbf{r}))(\text{diag}(\boldsymbol{\lambda})\mathbf{a} - \mathbf{D}\boldsymbol{\lambda}) - \text{diag}(\boldsymbol{\lambda})\mathbf{d}$. The function Φ is positively homogeneous.

Our analysis of this scheme is based on a formulation of the cost allocation problem as an LP game in the formulation of Samet and Zemel (1984).² They and previous authors showed that cost allocation problems

²Despite the network flows interpretation, this LP game is not a flow game in the formulation of Kalai and Zemel (1982), because nodes must respect equal priorities in their payments to other nodes.

Table 5: Node behavior in the counterparty contagion model and marginal prices of resources in the associated linear programming game.

| Behavior of Node i at Time 1 | | | | Marginal Prices of Resources | |
|--------------------------------|----------------|-------------------------|-------------------------|------------------------------|------------------------------|
| Color | Equity v_i^* | Payment p_i^* | Insured Loss ℓ_i^* | Primary Value w_i | Promised Payment \bar{p}_i |
| green | $v_i^* > 0$ | $p_i^* = \bar{p}_i$ | $\ell_i^* = 0$ | $\zeta_i = 0$ | $0 \leq \eta_i \leq 1$ |
| yellow | $v_i^* = 0$ | $0 < p_i^* < \bar{p}_i$ | $\ell_i^* = 0$ | $0 \leq \zeta_i \leq 1$ | $\eta_i = 0$ |
| red | $v_i^* = 0$ | $p_i^* = 0$ | $\ell_i^* > 0$ | $\zeta_i = 1$ | $\eta_i = 0$ |

of this type are fruitfully studied in terms of dual-optimal solutions to the LP. The LP we consider can be formulated as

$$\min_{\ell, \mathbf{p}} \mathbf{1}^\top \ell \quad \text{subject to} \quad (\mathbf{I} - \mathbf{\Pi}^\top) \mathbf{p} - \ell \leq \mathbf{w}, \quad \mathbf{p} \leq \bar{\mathbf{p}}, \quad \ell \geq \mathbf{0}, \quad \mathbf{p} \geq \mathbf{0}. \quad (12)$$

In Appendix A we show that the clearing payment vector \mathbf{p}^* provides an optimal solution to this LP because the seniority of deposits makes the clearing payment vector minimize insured losses.

In an LP game, non-negative marginal prices are assigned to each of the *resources* on the right side of constraints. In the LP (12), the resources are the primary value \mathbf{w} , whose marginal prices are $\boldsymbol{\zeta}$, and the promised payments $\bar{\mathbf{p}}$, whose marginal prices are $\boldsymbol{\eta}$. Suppose that there are no borderline nodes. Each marginal price is the rate of decrease in the optimal value of the objective, $\mathbf{1}^\top \ell^*$, as the amount of the corresponding resource increases. The marginal prices, given in Equation (13), quantify how valuable the resources are in mitigation of insured losses, i.e., in the “production” of payments to depositors or the “transportation” of money to depositors: increasing the money available at some node, or the capacity for flow out of a node, may increase the amount of money that can be transported to depositors. Table 5 summarizes how the marginal prices of a node’s primary value and promised payments depend on its behavior. We use colors to denote nodes with different behaviors.

Definition 4. *If the terminal equity value v_i^* and payment made p_i^* are 0 and the insured loss ℓ_i^* is positive, then node i is red. If the terminal equity value $v_i^* = 0$, the payment made p_i^* is positive but less than the promised payment \bar{p}_i , and the insured loss $\ell_i^* = 0$, then node i is yellow. If the terminal equity value v_i^* is positive, the payment made p_i^* equals the promised payment \bar{p}_i , and the insured loss $\ell_i^* = 0$, then node i is green.*

Let \mathcal{R} , \mathcal{Y} , and \mathcal{G} indicate the sets of red, yellow, and green nodes, respectively. When a set is used as a subscript of a vector or matrix, the result is a vector or matrix formed by selecting the rows or columns whose indices are in the set. Proposition 1, proved in Appendix B, says that the marginal prices are

$$\boldsymbol{\zeta}_{\mathcal{R}} = \mathbf{1}, \quad \boldsymbol{\zeta}_{\mathcal{Y}} = (\mathbf{I} - \mathbf{\Pi}_{\mathcal{Y}\mathcal{Y}})^{-1} \mathbf{\Pi}_{\mathcal{Y}\mathcal{R}} \mathbf{1}, \quad \boldsymbol{\zeta}_{\mathcal{G}} = \mathbf{0}, \quad \boldsymbol{\eta}_{\mathcal{R}} = \boldsymbol{\eta}_{\mathcal{Y}} = \mathbf{0}, \quad \text{and} \quad \boldsymbol{\eta}_{\mathcal{G}} = \mathbf{\Pi}_{\mathcal{G}} \boldsymbol{\zeta}. \quad (13)$$

Proposition 1. *If there are no borderline nodes, then for all $i = 1, \dots, n$, $\zeta_i = -\partial \mathbf{1}^\top \ell^* / \partial w_i$ and $\eta_i = -\partial \mathbf{1}^\top \ell^* / \partial \bar{p}_i$.*

The marginal price ζ_j of money at node j is the fundamental quantity in our sensitivity analysis. It is the value, for reducing total insured losses, of an extra dollar available at node j . It is helpful to interpret $\boldsymbol{\zeta}$ as a value function for a discrete-time Markov chain (see, e.g., Nelson, 2002, Ch. 6) in which the states are nodes and each transition represents a dollar moving from one node to another. Green and red nodes are absorbing states, because a perturbation of the money available there has no impact on the network flows. Adding a dollar at a red node reduces its insured loss, so $\boldsymbol{\zeta}_{\mathcal{R}} = \mathbf{1}$. Adding a dollar at a green node increases its terminal equity value and has no effect on insured losses, so $\boldsymbol{\zeta}_{\mathcal{G}} = \mathbf{0}$. Adding a dollar at a yellow node increases its outflow, the total payment it makes to other nodes. Therefore we use $\mathbf{\Pi}_{\mathcal{Y}}$ as the matrix of transition probabilities from yellow nodes. If node i is yellow, then Π_{ij} is the fraction of its outflow p_i^* that goes to node j , and it is the probability that a randomly selected dollar at node i goes to node j . The marginal price of money ζ_i at node i is the fraction of its outflow that is ultimately absorbed by red nodes,

where it reduces insured losses. Equivalently, ζ_i is an absorption probability, the probability that a randomly selected dollar at node i is ultimately absorbed by a red node. It is a continuation value: $\zeta_i = \sum_{j=1}^n \Pi_{ij} \zeta_j$ is the conditional expectation of the value at the next step in the Markov chain given that the state is currently i . The formula for $\zeta_{\mathcal{Y}}$ in Equation (13) is the solution to the system of equations $\zeta_{\mathcal{Y}} = \mathbf{\Pi}_{\mathcal{Y}} \zeta$, $\zeta_{\mathcal{R}} = \mathbf{1}$, and $\zeta_{\mathcal{G}} = \mathbf{0}$. The marginal prices of promised payments from green nodes are also continuation values: $\eta_{\mathcal{G}} = \mathbf{\Pi}_{\mathcal{G}} \zeta$ because adding a dollar to the promised payment \bar{p}_i of a green node i adds a dollar to its outflow p_i^* .

The last step in LP game cost allocation is to use marginal prices to allocate costs to the game's players, who each supply some of the resources. In a typical LP game, players supply non-negative amounts of the resources. Our LP game is unusual in that some players supply negative amounts of some resources, i.e., they demand positive amounts, and the net supply of some resources is negative. In particular, the net supply of primary value at every red bank is negative: its deposits exceed its cash. Supplying a resource results in a non-positive term in the player's cost allocation, whereas demanding a resource results in a non-negative term. This scheme assigns responsibility to obligors, so each node is the supplier of its promised payment. Because the debt/equity scheme holds each node responsible for reducing its lenders' cash, node j demands D_{ij} of the primary value at node i for all $i \neq j$. Because the net supply of primary value at node j is w_j , the amount supplied by node j is $q_j = w_j + \sum_{k \neq j} D_{jk}$ (see the discussion of Table 4). Another way to derive the resource holdings is $\partial \bar{p}_j / \partial \lambda_j = \bar{p}_j$, $\partial w_j / \partial \lambda_j = (1 + r_j) a_j - d_j = q_j$, and for $i \neq j$, $\partial \bar{p}_i / \partial \lambda_j = 0$ and $\partial w_i / \partial \lambda_j = -(1 + r_i) D_{ij}$. The cost allocation to node j is

$$\sum_{i \neq j} \zeta_i D_{ij} - \eta_j \bar{p}_j - \zeta_j q_j. \quad (14)$$

Each node is responsible for the principal it borrowed and for the payments it promised to make.

To see more explicitly how this scheme assigns responsibility, consider the special case of a classic lending system. The cost allocated to financial node i is 0 if it is green and $-s_i$ if it is red, because a bank's initial equity serves as a cushion for its depositors, but the cushion has zero marginal value unless the bank has insured losses. There is a striking violation of monotonicity (Section 3.4) if we consider changing leverage (which does not change with participation levels in this scheme): an increase in the initial equity of a bank could transform it from red to green, which would increase its cost allocation but decrease the system's cost. If industrial node j is green, its cost allocation is $-\sum_{i \in \mathcal{R}} D_{ij} \bar{r}_{ij}$, because the loan from a red bank i to a green node j results in a profit of $D_{ij} \bar{r}_{ij}$ for bank i , which reduces its insured losses. Suppose further that the system has competitive rates. The cost allocation to a yellow industrial node j is $\sum_{i \in \mathcal{R}} D_{ij} - (\sum_{i \in \mathcal{R}} D_{ij} / \sum_{h \neq j} D_{hj}) e_j = (1 - \rho_j) \sum_{i \in \mathcal{R}} D_{ij}$, which is the sum of the losses it causes to red banks. This is negative if $1 < \rho_j < 1 + \bar{r}_j$. Equation (10) implies that the sum of the losses on bank i 's loss-making loans, for which the corresponding borrowers are held responsible, is the sum of the profits on its other loans, its insured loss ℓ_i^* , and its initial equity s_i .

4.4.3 Loss-Making Industrial Nodes

To design a scheme similar to the obligor-responsibility scheme, but yielding non-negative allocations in a classic lending system with competitive rates, we make the players the loss-making industrial nodes. The function Φ is not positively homogeneous: the data is given by the same formulae as in Section 4.4.2, but with λ_i replaced by 1 if node i is not a player. Suppose that there are borderline nodes in the network $\Psi(\Phi(\gamma \mathbf{1}))$ for only finitely many values of $\gamma \in (0, 1)$, and let these values be arranged in the increasing sequence $\gamma_0, \dots, \gamma_m$, in which $\gamma_0 = 0$ and $\gamma_m = 1$. For any other value of $\gamma \in [0, 1]$, let $\zeta(\gamma)$ and $\eta(\gamma)$ represent the marginal prices from Proposition 1, but with the data for system $\Phi(\gamma \mathbf{1})$ plugged into the LP (12). The marginal prices are piecewise constant in γ , with their points of discontinuity contained in the set $\{\gamma_1, \dots, \gamma_{m-1}\}$, and can be found via parametric linear programming (see, e.g., Gass, 2003, Ch. 8). Define $\delta_h = \gamma_h - \gamma_{h-1}$, the width of the h th interval, and $\mu_h = (\mu_{h-1} + \mu_h)/2$, its midpoint. The cost

allocation to node j is 0 if node j is not a player. If node j is a player, its cost allocation is

$$\sum_{h=1}^m \delta_h \left(\sum_{i \neq j} D_{ij} \zeta_i(\mu_h) - \bar{p}_j \eta_j(\mu_h) - q_j \zeta_j(\mu_h) \right), \quad (15)$$

which follows by replacing Equation (14) with its integral as γ ranges from 0 to 1. The quantities $\sum_{h=1}^m \delta_h \zeta(\mu_h)$ and $\sum_{h=1}^m \delta_h \boldsymbol{\eta}(\mu_h)$ can be interpreted as average prices as the loss-making industrial nodes are scaled down to zero size.

Consider the special case of a classic lending system with competitive rates. For a green financial node i , let $\tilde{\gamma}_i = 1$. For a red bank i , let $\tilde{\gamma}_i = (s_i + \sum_{j \neq i} D_{ij}(\rho_j - 1)^+) / \sum_{j \neq i} D_{ij}(1 - \rho_j)^+$. Using Equation (10), for any financial node i and $\gamma \in [0, 1]$, node i is red in system $\Phi(\gamma \mathbf{1})$ if $\gamma > \tilde{\gamma}_i$ and is green in system $\Phi(\gamma \mathbf{1})$ if $\gamma < \tilde{\gamma}_i$, and $\int_0^1 \zeta_i(\gamma) d\gamma = 1 - \tilde{\gamma}_i = \ell_i^* / \sum_{j \neq i} D_{ij}(1 - \rho_j)^+$, which is the average cost of loan losses at node i . If node j is a player, then it is yellow in system $\Phi(\gamma \mathbf{1})$ for all $\gamma \in [0, 1]$, so $\eta_j(\gamma) = 0$ and $\zeta_j(\gamma) = \sum_{i \neq j} \Pi_{ji} \zeta_i(\gamma) = \sum_{i \neq j} D_{ij} \zeta_i(\gamma) / \sum_{i \neq j} D_{ij}$. The non-negative cost allocation to node j is $\sum_{i \neq j} \ell_i^* D_{ij}(1 - \rho_j)^+ / \sum_{k \neq i} D_{ik}(1 - \rho_k)^+$. That is, a borrower is responsible for a fraction of the insured losses of each of its lenders, and this fraction is the fraction of the lender's loan losses due to lending to this borrower.

4.4.4 Creditor Responsibility

This debt/equity scheme assigns responsibility for deals to the creditor. The participation level λ_i of node i multiplies the swap payments promised to node i and everything on the asset side of its time-0 balance sheet: lending $\mathbf{D}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{D}$ and the investment in the baseline asset $\mathbf{t}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{t}$. Therefore the time-0 balance sheet size $\mathbf{a}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{a}$ and the liability matrix $\mathbf{L}(\boldsymbol{\lambda}) = \mathbf{L}\text{diag}(\boldsymbol{\lambda})$. As in the obligor-responsibility scheme, λ_i also multiplies node i 's deposits $\mathbf{d}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{d}$. Therefore the primary value $\mathbf{w}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{w}$ and the function Φ is positively homogeneous.

The payment proportions $\boldsymbol{\Pi}(\boldsymbol{\lambda})$ vary with $\boldsymbol{\lambda}$ in a complicated way, which makes the LP (12) inconvenient for analyzing this scheme. We reformulate the LP by defining f_i to be the fraction of node i 's liabilities that it pays. If $\bar{p}_i > 0$, then $f_i = p_i / \bar{p}_i$. If node i has no promised payments, then the payment fraction f_i is indeterminate in a sense, but we choose the payment fraction to be consistent with the priority constraints $v_i > 0 \Rightarrow f_i = 1$ and $f_i > 0 \Rightarrow \ell_i = 0$, as follows. Let \mathbf{f}^* denote the vector of payment fractions corresponding to the clearing payment vector \mathbf{p}^* . If node i is green, $f_i^* = 1$; if node i is yellow, $\bar{p}_i > 0$ and $0 < f_i^* < 1$; and if node i is red, $f_i^* = 0$. The reformulated LP is

$$\min_{\boldsymbol{\ell}, \mathbf{f}} \mathbf{1}^\top \boldsymbol{\ell} \quad \text{subject to} \quad (\text{diag}(\mathbf{L}\mathbf{1}) - \mathbf{L}^\top) \mathbf{f} - \boldsymbol{\ell} \leq \mathbf{w}, \quad \mathbf{f} \leq \mathbf{1}, \quad \boldsymbol{\ell} \geq \mathbf{0}, \quad \mathbf{f} \geq \mathbf{0}. \quad (16)$$

Suppose there are no borderline nodes. Then Proposition 1 gives the marginal prices $\boldsymbol{\zeta}$ for \mathbf{w} , and Proposition 2, proved in Appendix B, says that the marginal prices $\boldsymbol{\Theta}$ for \mathbf{L} are given by

$$\theta_{ij} = f_i^* (\zeta_j - \zeta_i). \quad (17)$$

The interpretation of this formula is that node i pays the fraction f_i^* of its liabilities, and $\zeta_i - \zeta_j$ is the marginal cost of moving money from node i to node j . The relationship between the marginal prices θ_{ij} for promised payments from node i to other nodes and the marginal price η_i for its total promised payments is

$$\eta_i = \sum_{j \neq i} \Pi_{ij} \theta_{ij}. \quad (18)$$

Proposition 2. *If there are no borderline nodes, then for all $i \neq j$, $\theta_{ij} = -\partial \mathbf{1}^\top \boldsymbol{\ell}^* / \partial L_{ij}$.*

For all j and $i \neq j$, $\partial w_j / \partial \lambda_j = w_j$, $\partial L_{ij} / \partial \lambda_j = L_{ij}$, and $\partial w_j / \partial \lambda_i = \partial L_{ij} / \partial \lambda_i = 0$. Therefore the cost

allocation to node j is

$$-\sum_{i \neq j} \theta_{ij} L_{ij} - \zeta_j w_j = \zeta_j \sum_{k \neq j} D_{jk} - \sum_{i \neq j} \theta_{ij} L_{ij} - \zeta_j q_j \quad (19)$$

$$= \sum_{i \neq j} \zeta_i f_i^* L_{ij} - \zeta_j \left(w_j + \sum_{i \neq j} f_i^* L_{ij} \right) = \sum_{i \neq j} \zeta_i p_i^* \Pi_{ij} - \zeta_j \left(w_j + \sum_{i \neq j} p_i^* \Pi_{ij} \right). \quad (20)$$

Equation (19) follows from Equation (11) and shows that each node is responsible for the principal it has lent and for the payments it is supposed to receive. Equation (20) follows by using Equation (17) and observing that $f_i^* L_{ij} = p_i^* \Pi_{ij}$ is the amount paid by node i to node j . The amount of money available at node j after receiving payments from other nodes is $w_j + \sum_{i \neq j} p_i^* \Pi_{ij}$, and its marginal price is ζ_j . In Equation (20), the term $-\zeta_j (w_j + \sum_{i \neq j} p_i^* \Pi_{ij})$ is the insured loss ℓ_j^* if node j is red, 0 if node j is green, and $-\zeta_j p_j^*$ if node j is yellow. The term $\sum_{i \neq j} \zeta_i p_i^* \Pi_{ij} = \sum_{i \in \mathcal{Y}} \zeta_i p_i^* \Pi_{ij}$ because $\zeta_i p_i^* = 0$ unless node i is yellow. It is a cost allocated to node j as a creditor: the scheme holds it responsible for taking money from obligors who can pay something but can not pay in full. The sum of the costs allocated to all nodes for receiving payments from yellow nodes equals the sum of the negative cost allocations to all yellow nodes for the payments they make: $\sum_{j=1}^n \sum_{i \in \mathcal{Y}} \zeta_i p_i^* \Pi_{ij} = \sum_{i \in \mathcal{Y}} \zeta_i p_i^*$.

Consider the special case of a classic lending system. The cost allocation to green industrial nodes is 0. If industrial node i is yellow, its cost allocation is $-\zeta_i p_i^* = -p_i^* \sum_{j \in \mathcal{R}} \Pi_{ij}$, the negative of the sum of the payments it makes to red banks. There is a striking violation of monotonicity: an increase in the initial equity of a borrower could transform it from yellow to green, which would increase its cost allocation but decrease the system's cost. Cf. the similar violation of monotonicity in increasing the initial equity of a lender in Section 4.4.2.

4.4.5 Shared Responsibility

This debt/equity scheme shares the responsibility for deals equally between the parties to the deal. Any deal between nodes i and j is multiplied by the geometric average $\sqrt{\lambda_i \lambda_j}$ of their participation levels. We choose the function that maps (λ_1, λ_2) to $\sqrt{\lambda_1 \lambda_2}$ because it is positively homogeneous, symmetric, and, like the Shapley value scheme that eliminates a deal whenever either party to the deal is not participating, it maps $(1, 1)$ to 1 but $(\lambda, 0)$ and $(0, \lambda)$ to 0 for any λ . The use of the geometric average leads to an arithmetic average in Equation (21) that entails equal sharing of responsibility between creditor and obligor. We have $\mathbf{D}(\boldsymbol{\lambda}) = \sqrt{\text{diag}(\boldsymbol{\lambda})} \mathbf{D} \sqrt{\text{diag}(\boldsymbol{\lambda})}$ and $\mathbf{L}(\boldsymbol{\lambda}) = \sqrt{\text{diag}(\boldsymbol{\lambda})} \mathbf{L} \sqrt{\text{diag}(\boldsymbol{\lambda})}$. As in the obligor- and creditor-responsibility schemes, $\mathbf{a}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda}) \mathbf{a}$ and $\mathbf{d}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda}) \mathbf{d}$. The investment in baseline assets $\mathbf{t}(\boldsymbol{\lambda}) = \mathbf{a}(\boldsymbol{\lambda}) - \mathbf{D}(\boldsymbol{\lambda}) \mathbf{1}$ and the cashflow $\mathbf{e}(\boldsymbol{\lambda}) = (\mathbf{I} + \text{diag}(\mathbf{r})) \mathbf{t}(\boldsymbol{\lambda})$. Because industrial nodes do not lend or take deposits and financial nodes' baseline asset is cash, which has a zero rate of return, the primary value $\mathbf{w}(\boldsymbol{\lambda}) = (\mathbf{I} + \text{diag}(\mathbf{r})) (\mathbf{a}(\boldsymbol{\lambda}) - \mathbf{D}(\boldsymbol{\lambda}) \mathbf{1}) - \mathbf{d}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda}) \mathbf{w} + \text{diag}(\boldsymbol{\lambda}) \mathbf{D} \mathbf{1} - \mathbf{D}(\boldsymbol{\lambda}) \mathbf{1}$. The function Φ is positively homogeneous.

For all j and $i \neq j$, $\partial L_{ij} / \partial \lambda_i = \partial L_{ij} / \partial \lambda_j = L_{ij} / 2$, $\partial w_i / \partial \lambda_j = -D_{ij} / 2$, and $\partial w_j / \partial \lambda_j = w_j + \sum_{k \neq j} D_{jk} / 2$. Equation (11) defined $q_j = w_j + \sum_{k \neq j} D_{jk}$, the primary value that node j would have had if there were no lending. If there are no borderline nodes, the cost allocated to node j is

$$\frac{1}{2} \left(\sum_{i \neq j} (\zeta_i D_{ij} - \theta_{ij} L_{ij}) + \sum_{k \neq j} (\zeta_j D_{jk} - \theta_{jk} L_{jk}) \right) - \zeta_j q_j. \quad (21)$$

From Equations (14), (18), and (19), it follows that this is the average of the costs allocated to node j by the obligor- and creditor-responsibility schemes. Another interpretation is that the shared-responsibility scheme equally shares costs allocated to each deal between the deal's creditor and obligor. The cost allocated to a swap that results in a promised payment of X_{ij} from node i to node j is $-\theta_{ij} X_{ij}$. The cost allocated to a loan of D_{ij} at rate \bar{r}_{ij} from node i to node j is $\zeta_i D_{ij} - \theta_{ji} D_{ij} (1 + \bar{r}_{ij})$, where the first term is for removing

the principal from node i and the second term is for the promised repayment of principal with interest by node j to node i . Because this is a debt/equity scheme, there is no term for moving the principal to node j : the assumption is that, in the absence of this loan, node j would have replaced the debt with equity.

4.4.6 An Average-Cost Scheme

Because the obligor-, creditor-, and shared-responsibility schemes all use the same marginal prices, which produce allocations that are discontinuous as functions of the system data and may be extreme, we develop a debt/equity scheme based on shared responsibility but using average prices. As in previous schemes, the participation level λ_i of node i multiplies its deposits and time-0 balance sheet size: $\mathbf{d}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{d}$ and $\mathbf{a}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{a}$. Any deal between nodes i and j is multiplied by the product $\lambda_i\lambda_j$ of their participation levels: $\mathbf{D}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{D}\text{diag}(\boldsymbol{\lambda})$ and $\mathbf{L}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{L}\text{diag}(\boldsymbol{\lambda})$. Although we use sensitivities to perturbations that change the matrix $\boldsymbol{\Pi}(\boldsymbol{\lambda})$ of payment proportions, along the diagonal, this matrix is constant: $\mathbf{L}(\gamma\mathbf{1}) = \gamma^2\mathbf{L}$, so $\boldsymbol{\Pi}(\gamma\mathbf{1}) = \boldsymbol{\Pi}$. The investment in baseline assets $\mathbf{t}(\boldsymbol{\lambda}) = \mathbf{a}(\boldsymbol{\lambda}) - \mathbf{D}(\boldsymbol{\lambda})\mathbf{1} = \text{diag}(\boldsymbol{\lambda})(\mathbf{t} + \mathbf{D}(\mathbf{1} - \boldsymbol{\lambda}))$ and the cashflow $\mathbf{e}(\boldsymbol{\lambda}) = (\mathbf{I} + \text{diag}(\mathbf{r}))\mathbf{t}(\boldsymbol{\lambda})$. The primary value $\mathbf{w}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})((\mathbf{I} + \text{diag}(\mathbf{r}))(\mathbf{t} + \mathbf{D}(\mathbf{1} - \boldsymbol{\lambda})) - \mathbf{d}) = \text{diag}(\boldsymbol{\lambda})(\mathbf{w} + \mathbf{D}(\mathbf{1} - \boldsymbol{\lambda}))$ because industrial nodes do not lend or take deposits and financial nodes' baseline asset is cash, which has a zero rate of return. The function that maps (λ_1, λ_2) to $\lambda_1\lambda_2$ is symmetric and maps $(1, 1)$ to 1 and $(\lambda, 0)$ and $(0, \lambda)$ to 0 for any λ , but it is not positively homogeneous; accordingly, Φ is not positively homogeneous. Loosely speaking, it becomes easier for nodes to fulfill their obligations in system $\Phi(\gamma\mathbf{1})$ as γ decreases, as can be seen by comparing the primary value $\mathbf{w}(\gamma\mathbf{1}) = \gamma(\mathbf{w} + (1 - \gamma)\mathbf{D}\mathbf{1})$ to $\mathbf{0}$, the boundary of insured loss, and to $\bar{\mathbf{p}}(\gamma\mathbf{1}) = \gamma^2\bar{\mathbf{p}}$, the boundary of default. Alternatively, one may observe that the leverage ratio in system $\Phi(\gamma\mathbf{1})$, $\text{diag}(\mathbf{a}(\gamma\mathbf{1}))^{-1}\mathbf{s}(\gamma\mathbf{1}) = \text{diag}(\mathbf{a})^{-1}(\mathbf{s} + (1 - \gamma)\mathbf{D}^\top\mathbf{1})$, increases as γ decreases. The intent of the scheme's design is to make the Aumann-Shapley value use prices averaged over a range of systems in which there are fewer or milder defaults and insured losses. This happens in a classic lending system: see Proposition 3.

As in Section 4.4.3, suppose $\{\gamma_1, \dots, \gamma_{m-1}\}$ are the only values of $\gamma \in (0, 1)$ such that there are borderline nodes in the network $\Psi(\Phi(\gamma\mathbf{1}))$, let $\gamma_0 = 0$ and $\gamma_m = 1$, and let δ_h be the width and μ_h be the midpoint of (γ_{h-1}, γ_h) . For any value of $\gamma \in [0, 1] \setminus \{\gamma_0, \dots, \gamma_m\}$, the marginal prices $\zeta(\gamma)$ and $\Theta(\gamma)$ are as described in Propositions 1 and 2, but with $\mathbf{w}(\gamma\mathbf{1})$ and $\mathbf{L}(\gamma\mathbf{1})$ plugged into the LP (16). Let $\mathcal{H}_R(i)$ be the set of all h such that node i is red in system $\Phi(\mu_h\mathbf{1})$, and define $\mathcal{H}_Y(i)$ and $\mathcal{H}_G(i)$ similarly for yellow and green. Appendix B.3 shows that, for all i and j , the average prices

$$\int_0^1 \zeta_i(\gamma) d\gamma = \bar{\zeta}_i = \sum_{h \in \mathcal{H}_R(i)} \delta_h + \sum_{h \in \mathcal{H}_Y(i)} \delta_h \zeta_i(\mu_h), \quad (22)$$

$$2 \int_0^1 \gamma \zeta_i(\gamma) d\gamma = \hat{\zeta}_i = \sum_{h \in \mathcal{H}_R(i)} (\gamma_h^2 - \gamma_{h-1}^2) + \sum_{h \in \mathcal{H}_Y(i)} (\gamma_h^2 - \gamma_{h-1}^2) \zeta_i(\mu_h), \quad \text{and} \quad (23)$$

$$2 \int_0^1 \gamma \theta_{ij}(\gamma) d\gamma = \hat{\theta}_{ij} = \sum_{h \in \mathcal{H}_G(i)} (\gamma_h^2 - \gamma_{h-1}^2) \zeta_j(\mu_h) + \sum_{h \in \mathcal{H}_Y(i)} (\gamma_h^2 - \gamma_{h-1}^2) \theta_{ij}(\mu_h). \quad (24)$$

The non-zero sensitivities of the LP data to participation level λ_j are, for any $\gamma \in (0, 1)$ and any $i, k \neq j$,

$$\frac{\partial L_{ij}}{\partial \lambda_j}(\gamma\mathbf{1}) = \gamma L_{ij}, \quad \frac{\partial L_{jk}}{\partial \lambda_j}(\gamma\mathbf{1}) = \gamma L_{jk}, \quad \frac{\partial w_j}{\partial \lambda_j}(\gamma\mathbf{1}) = w_j + (1 - \gamma) \sum_{k \neq j} D_{jk}, \quad \text{and} \quad \frac{\partial w_i}{\partial \lambda_j}(\gamma\mathbf{1}) = -\gamma D_{ij}.$$

The cost allocated to node j is

$$\frac{1}{2} \left(\sum_{i \neq j} (\hat{\zeta}_i D_{ij} - \hat{\theta}_{ij} L_{ij}) + \sum_{k \neq j} (\hat{\zeta}_j D_{jk} - \hat{\theta}_{jk} L_{jk}) \right) - \bar{\zeta}_j q_j, \quad (25)$$

where $q_j = w_j + \sum_{k \neq j} D_{jk}$ as in Equation (11). Like the shared-responsibility scheme of Section 4.4.5, this scheme equally shares costs allocated to each deal between the deal's creditor and obligor. As illustrated in Table 4, the cost allocations to deals are similar, but with the average prices $\widehat{\zeta}$ and $\widehat{\theta}$ replacing the marginal prices ζ and θ .

Consider the special case of a classic lending system. In system $\Phi(\gamma \mathbf{1})$, an industrial node j has promised to pay $\gamma^2 \bar{p}_j$ and has γe_j available to pay it, so it pays the fraction $f_j^*(\gamma) = \min\{1, e_j/\gamma \bar{p}_j\}$ of its obligations. This fraction is non-decreasing as a function of $\gamma \in [0, 1]$. Let $\tilde{\gamma}_j = \min\{1, e_j/\bar{p}_j\}$ for an industrial node j . In system $\Phi(\gamma \mathbf{1})$, industrial node j is yellow if $\gamma > \tilde{\gamma}_j$ and is green if $\gamma < \tilde{\gamma}_j$. For a green financial node i , let $\tilde{\gamma}_i = 1$. For a red bank i , $\tilde{\gamma}_i$ is described by Lemma 1, whose proof shows that node i 's initial equity equals the loss on its loans in system $\Phi(\tilde{\gamma}_i \mathbf{1})$.

Lemma 1. *In a classic lending system, if node i is a red bank with time-0 equity $s_i > 0$, then there is a unique value $\tilde{\gamma}_i$ of $\gamma \in (0, 1)$ that solves*

$$s_i = \gamma \sum_{j \neq i} D_{ij} (1 - (1 + \bar{r}_{ij}) f_j^*(\gamma)). \quad (26)$$

In system $\Phi(\gamma \mathbf{1})$, node i is red if $\gamma > \tilde{\gamma}_i$ and green if $\gamma < \tilde{\gamma}_i$.

Proof. Denote the right side of Equation (26) as $v(\gamma)$. In system $\Phi(\gamma \mathbf{1})$, node i 's loss on its loans is $\gamma v(\gamma)$ and its initial equity is γs_i , so it is red if $v(\gamma) > s_i$ and green if $v(\gamma) < s_i$.

The function v is continuous in γ , $v(0) = 0 < s_i$, and $v(1) = \sum_{j \neq i} (D_{ij} - p_j^* \Pi_{ji})$. Because node i is a red bank, $w_i + \sum_{j \neq i} p_j^* \Pi_{ji} = -\ell_i^* < 0$. From this and $\sum_{j \neq i} D_{ij} = a_i - t_i = (s_i + d_i) - e_i = s_i - w_i$, it follows that $v(1) > s_i$. Therefore there exists a solution $\tilde{\gamma}_i$ in $(0, 1)$ to $v(\gamma) = s_i$.

For all $\gamma \in [0, 1] \setminus \{\gamma_0, \dots, \gamma_m\}$, we have $v(\gamma) = \sum_{j \neq i} 1_{\{\gamma > \tilde{\gamma}_j\}} D_{ij} (\gamma - (1 + \bar{r}_{ij}) e_j / \bar{p}_j) - \gamma \sum_{j \neq i} 1_{\{\gamma < \tilde{\gamma}_j\}} D_{ij} \bar{r}_{ij}$ and $v'(\gamma) = \sum_{j \neq i} D_{ij} (1_{\{\gamma > \tilde{\gamma}_j\}} - \bar{r}_{ij} 1_{\{\gamma < \tilde{\gamma}_j\}})$. The function v' is non-decreasing. From $\int_0^{\tilde{\gamma}_i} v'(\gamma) d\gamma = v(\tilde{\gamma}_i) - v(0) = s_i > 0$, it follows that $v'(\tilde{\gamma}_i) > 0$. Therefore $v(\gamma) > s_i$ if $\gamma > \tilde{\gamma}_i$. Because v is convex, it achieves its maximum on $[0, \tilde{\gamma}_i]$ at a single point, 0 or $\tilde{\gamma}_i$; because $v(\tilde{\gamma}_i) = s_i > 0 = v(0)$, the maximum is s_i achieved at $\tilde{\gamma}_i$. Therefore $v(\gamma) < s_i$ if $\gamma < \tilde{\gamma}_i$. \square

By Lemma 1, for a classic lending system whose red banks all have strictly positive time-0 equity, the average prices in Equations (22)–(24) are as follows. If node i is financial, then $\bar{\zeta}_i = 1 - \tilde{\gamma}_i$ and $\widehat{\zeta}_i = 1 - \tilde{\gamma}_i^2$. Define $\tilde{\gamma}_{ij} = \max\{\tilde{\gamma}_i, \tilde{\gamma}_j\}$. If node i is industrial,

$$\begin{aligned} \bar{\zeta}_i &= \int_{\tilde{\gamma}_i}^1 \sum_{j \neq i} \Pi_{ij} \zeta_j(\gamma) d\gamma = \sum_{j \neq i} \Pi_{ij} (1 - \tilde{\gamma}_{ij}), \\ \widehat{\zeta}_i &= 2 \int_{\tilde{\gamma}_i}^1 \gamma \sum_{j \neq i} \Pi_{ij} \zeta_j(\gamma) d\gamma = \sum_{j \neq i} \Pi_{ij} (1 - \tilde{\gamma}_{ij}^2), \quad \text{and} \\ \widehat{\theta}_{ij} &= 2 \int_0^{\tilde{\gamma}_i} \gamma \zeta_j(\gamma) d\gamma + 2 \int_{\tilde{\gamma}_i}^1 \gamma f_i^*(\gamma) (\zeta_j(\gamma) - \zeta_i(\gamma)) d\gamma = \tilde{\gamma}_{ij}^2 - \tilde{\gamma}_j^2 + 2 \frac{e_i}{\bar{p}_i} (1 - \tilde{\gamma}_{ij} - \bar{\zeta}_i) \end{aligned}$$

because $\gamma f_i^*(\gamma) = e_i / \bar{p}_i$ for $\gamma > \tilde{\gamma}_i$. Furthermore, the average prices $\bar{\zeta}$ and $\widehat{\zeta}$ for money at each node are less than the corresponding marginal prices ζ in the actual system $\Phi(\mathbf{1})$.

Proposition 3. *In a classic lending system whose red banks all have strictly positive time-0 equity, $\mathbf{0} \leq \bar{\zeta} \leq \zeta$, $\mathbf{0} \leq \widehat{\zeta} \leq \zeta$, and the marginal prices $\zeta(\gamma)$ are non-decreasing as a function of γ on $[0, 1]$.*

Proof. By Lemma 1, the set of red financial nodes in system $\Phi(\gamma \mathbf{1})$ is non-decreasing as a function of γ . By the analysis preceding Lemma 1, the set of yellow industrial nodes in system $\Phi(\gamma \mathbf{1})$ is non-decreasing as a function of γ . From Equation (13), $\zeta_i(\gamma) = \sum_{j \neq i} \Pi_{ij}(\gamma \mathbf{1}) \zeta_j(\gamma)$ for all $i \in \mathcal{Y}$, and $\mathbf{\Pi}(\gamma \mathbf{1}) = \mathbf{\Pi}$ for all $\gamma \in (0, 1]$, it follows that $\zeta(\gamma)$ is a non-negative, non-decreasing function of γ . Because $\bar{\zeta}$ and $\widehat{\zeta}$ are weighted averages of $\zeta(\gamma)$, they are non-negative and do not exceed $\zeta = \zeta(1)$. \square

To see more explicitly how this scheme assigns responsibility, consider a classic lending system with competitive rates. In system $\Phi(\gamma\mathbf{1})$, the recovery rate on lending to node j is $\rho_j(\gamma) = (1 + \bar{r}_j)f_j^*(\gamma) = \min\{1 + \bar{r}_j, \rho_j/\gamma\}$. The cost allocated to an industrial node j is $\int_0^1 (\gamma \sum_{i \neq j} D_{ij}(\zeta_i(\gamma) - \theta_{ji}L_{ji}) - e_j \zeta_j(\gamma)) d\gamma$. Because $L_{ji} = D_{ij}(1 + \bar{r}_j)$ and $\rho_j(\gamma) = (1 + \bar{r}_j)f_j^*(\gamma)$, from Equation (17) it follows that $\theta_{ji}(\gamma)L_{ji} = \rho_j(\gamma)D_{ij}(\zeta_i(\gamma) - \zeta_j(\gamma))$. Because $\zeta_j(\gamma) = 0$ for $\gamma \in (0, \tilde{\gamma}_j)$ and $\sum_{i \neq j} D_{ij}\gamma\rho_j(\gamma) = \sum_{i \neq j} D_{ij}\rho_j = e_j$ for $\gamma \in (\tilde{\gamma}_j, 1]$, $\zeta_j(\gamma)(\sum_{i \neq j} D_{ij}\gamma\rho_j(\gamma) - e_j) = 0$ for almost every $\gamma \in [0, 1]$. The cost allocation is

$$\begin{aligned} \int_0^1 \left(\sum_{i \neq j} D_{ij}\gamma(\zeta_i(\gamma)(1 - \rho_j(\gamma)) + \zeta_j(\gamma)\rho_j(\gamma)) - e_j \zeta_j(\gamma) \right) d\gamma &= \sum_{i \neq j} D_{ij} \int_0^1 \gamma \zeta_i(\gamma)(1 - \rho_j(\gamma)) d\gamma \\ &= \frac{1}{2} \sum_{i \neq j} D_{ij} \alpha_{ij}, \end{aligned}$$

where

$$\begin{aligned} \alpha_{ij} &= 2 \int_0^1 \gamma \zeta_i(\gamma)(1 - \rho_j(\gamma)) d\gamma = 2 \int_{\tilde{\gamma}_i}^1 \gamma(1 - \rho_j(\gamma)) d\gamma \\ &= 2 \int_{\tilde{\gamma}_i}^{\tilde{\gamma}_{ij}} (-\gamma \bar{r}_j) d\gamma + 2 \int_{\tilde{\gamma}_{ij}}^1 (\gamma - \rho_j) d\gamma = -(\tilde{\gamma}_{ij}^2 - \tilde{\gamma}_i^2) \bar{r}_j + (1 - \tilde{\gamma}_{ij}^2) - 2(1 - \tilde{\gamma}_{ij})\rho_j \\ &= -\bar{r}_j(\tilde{\gamma}_{ij}^2 - \tilde{\gamma}_i^2) + (1 - \tilde{\gamma}_{ij})(1 + \tilde{\gamma}_{ij} - 2\rho_j). \end{aligned}$$

The terms relate to the profit or loss on the loan. The first term is for systems $\Phi(\gamma\mathbf{1})$ in which node j is green and the second is for systems $\Phi(\gamma\mathbf{1})$ in which node j is yellow. For bank i , $q_i = s_i$, and the cost allocation is

$$\int_0^1 \left(\sum_{j \neq i} D_{ij}\gamma(\zeta_i(\gamma)(1 - \rho_j(\gamma)) + \zeta_j(\gamma)\rho_j(\gamma)) - s_i \zeta_i(\gamma) \right) d\gamma = \sum_{j \neq i} D_{ij} \left(\frac{1}{2} \alpha_{ij} + \beta_j \right) - s_i(1 - \tilde{\gamma}_i)$$

where

$$\beta_j = \int_0^1 \gamma \zeta_j(\gamma)\rho_j(\gamma) d\gamma = \rho_j \int_{\tilde{\gamma}_j}^1 \zeta_j(\gamma) d\gamma = \rho_j \int_{\tilde{\gamma}_j}^1 \frac{\sum_{i \neq j} D_{ij}\zeta_i(\gamma)}{\sum_{i \neq j} D_{ij}} d\gamma = \rho_j \frac{\sum_{i \neq j} D_{ij}(1 - \tilde{\gamma}_{ij})}{\sum_{i \neq j} D_{ij}}$$

follows from $\zeta_j(\gamma) = \sum_{i \neq j} \Pi_{ji}\zeta_i(\gamma)$ and $\Pi_{ji} = D_{ij}/\sum_{h \neq j} D_{hj}$. The cost β_j is allocated to a creditor for the payments it collects from its obligor. Like the obligor-responsibility scheme, this scheme allocates a negative cost to red banks' equity and a cost to a red bank's loss on its loans (or negative cost to its profit), but this scheme splits this cost between lender and borrower instead of allocating it all to the borrower. Like the creditor-responsibility scheme, this scheme allocates a cost to payments received from yellow nodes, but it does not give a corresponding negative cost allocation to the yellow nodes.

4.4.7 A Debt/Baseline Scheme

This debt/baseline scheme is based on shared responsibility and marginal prices, but it uses the marginal prices to allocate costs to loans differently than the shared-responsibility scheme of Section 4.4.5. As in Section 4.4.5, $\mathbf{D}(\boldsymbol{\lambda}) = \sqrt{\text{diag}(\boldsymbol{\lambda})}\mathbf{D}\sqrt{\text{diag}(\boldsymbol{\lambda})}$, $\mathbf{L}(\boldsymbol{\lambda}) = \sqrt{\text{diag}(\boldsymbol{\lambda})}\mathbf{L}\sqrt{\text{diag}(\boldsymbol{\lambda})}$, and $\mathbf{d}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{d}$. Using the debt/baseline scheme, $\mathbf{a}(\boldsymbol{\lambda}) = \mathbf{s}(\boldsymbol{\lambda}) + \mathbf{d}(\boldsymbol{\lambda}) + \mathbf{D}(\boldsymbol{\lambda})^\top \mathbf{1}$, $\mathbf{s}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})\mathbf{s}$, and $\mathbf{t}(\boldsymbol{\lambda}) = \mathbf{a}(\boldsymbol{\lambda}) - \mathbf{D}(\boldsymbol{\lambda})\mathbf{1}$. The cashflow $\mathbf{e}(\boldsymbol{\lambda}) = (\mathbf{I} + \text{diag}(\mathbf{r}))\mathbf{t}(\boldsymbol{\lambda})$. Because industrial nodes do not lend or take deposits and financial nodes' baseline asset is cash, which has a zero rate of return, the primary value $\mathbf{w}(\boldsymbol{\lambda}) = (\mathbf{I} + \text{diag}(\mathbf{r}))\mathbf{t}(\boldsymbol{\lambda}) - \mathbf{d}(\boldsymbol{\lambda}) = \text{diag}(\boldsymbol{\lambda})(\mathbf{I} + \text{diag}(\mathbf{r}))\mathbf{s} + \sqrt{\text{diag}(\boldsymbol{\lambda})}((\mathbf{I} + \text{diag}(\mathbf{r}))\mathbf{D}^\top - \mathbf{D})\sqrt{\text{diag}(\boldsymbol{\lambda})}\mathbf{1}$. The function Φ is positively homogeneous.

For all j and $i \neq j$, $\partial L_{ij}/\partial \lambda_i = \partial L_{ij}/\partial \lambda_j = L_{ij}/2$, $\partial w_i/\partial \lambda_j = ((1 + r_i)D_{ji} - D_{ij})/2$, and $\partial w_j/\partial \lambda_j = s_j + ((1 + r_j)\sum_{i \neq j} D_{ij} - \sum_{k \neq j} D_{jk})/2$. If there are no borderline nodes, the cost allocated to node j is

$$\frac{1}{2} \left(\sum_{i \neq j} ((\zeta_i - \zeta_j(1 + r_j))D_{ij} - \theta_{ij}L_{ij}) + \sum_{k \neq j} ((\zeta_j - \zeta_k(1 + r_k))D_{jk} - \theta_{jk}L_{jk}) \right) - \zeta_j(1 + r_j)s_j. \quad (27)$$

The scheme equally shares costs allocated to each deal between the creditor and obligor in the deal. The cost allocated to a swap that results in a promised payment of X_{ij} from node i to node j is $-\theta_{ij}X_{ij}$, just as in Section 4.4.5. The difference is in the cost allocated to a loan of D_{ij} at rate \bar{r}_{ij} from node i to node j , which is

$$(\zeta_i - \zeta_j(1 + r_j) - \theta_{ji}(1 + \bar{r}_{ij}))D_{ij} = \zeta_i(1 - f_j^*(1 + \bar{r}_{ij}))D_{ij} + \zeta_j(f_j^*(1 + \bar{r}_{ij}) - (1 + r_j))D_{ij}. \quad (28)$$

Compared to Section 4.4.5, the left side of Equation (28) has an extra term, $-\zeta_j(1 + r_j)D_{ij}$. In the debt/baseline scheme, the loan makes node j 's time-0 balance sheet larger by D_{ij} , and $(1 + r_j)D_{ij}$ is the increase in node j 's time-1 cashflow. To compensate for this difference from the debt/equity scheme, in the debt/baseline scheme, node j receives a negative cost allocation only for the time-1 value of the investment of its equity in its baseline asset, $(1 + r_j)s_j$, not for $q_j = (1 + r_j)a_j - d_j$. On the right side of Equation (28), the factor multiplying ζ_i is node i 's loss on the loan and the factor multiplying ζ_j is the loss the loan causes for node j : its payment to node i minus the time-1 value of its time-0 investment of D_{ij} in its baseline asset. Both factors may be positive or negative. The loan's cost allocation is 0 if nodes i and j are both green or both red. If there is no default on the loan (i.e., node j is green), the cost allocation is $-\zeta_i\bar{r}_{ij}D_{ij}$, the benefit of the profit earned by node i . If the borrower, node j , is a red bank, the cost allocation is $-(1 - \zeta_i)D_{ij}$, the benefit of moving the principal D_{ij} from node i to node j , where it protects node j 's depositors.

Consider the special case of a classic lending system with competitive rates. The cost allocated to the loan from bank i to industrial node j is $\zeta_i(1 - \rho_j)D_{ij} + \zeta_j(\rho_j - (1 + r_j))D_{ij}$. The first term in the loan's cost allocation, $\zeta_i(1 - \rho_j)D_{ij}$, is a cost for the loss on the loan, which is negative if the loan is profitable. The sum of the second term over all loans to node j is $\sum_{i \neq j} \zeta_j(\rho_j - (1 + r_j))D_{ij} = \zeta_j(\rho_j - (1 + r_j)) \sum_{i \neq j} D_{ij}$, which is 0 if node j is green and is $\zeta_j s_j(1 + r_j)$ if node j is yellow, because $\rho_j \sum_{i \neq j} D_{ij} = a_j(1 + r_j) = (s_j + \sum_{i \neq j} D_{ij})(1 + r_j)$. Thus, the second term in the loan's cost allocation is $\zeta_j s_j(1 + r_j)D_{ij} / \sum_{h \neq j} D_{hj}$. If node j is yellow, this term is the cost of using up a fraction of node j 's equity, and this fraction is the fraction of node j 's debt that comes from bank i .

4.5 Numerical Examples

We first examine the behavior of several schemes in a single scenario. In this example, banks 1 and 2 have deposits $d_1 = d_2 = 480$ and equity $s_1 = 64$ and $s_2 = 32$, all figures being quoted in millions of dollars. There are four industrial nodes, $j = 3, 4, 5,$ and 6 , each with equity $s_j = 80$ and $D_{1j} + D_{2j} = 320$ in bank debt, on which they pay $\bar{r}_j = 5\%$ interest. The bank debt is $D_{13} = D_{26} = 320$ and $D_{14} = D_{24} = D_{15} = D_{25} = 160$. Consider the scenario in which nodes 3 and 4 pay their obligations in full, while nodes 5 and 6 experience a return of $r_5 = r_6 = -30\%$ on their assets and can only repay $p_5^* = p_6^* = 280$. Bank 1 makes a profit and pays its depositors in full. Bank 2 earns 4 on its loan to node 4 but loses 10 and 20 on its loans to nodes 5 and 6, resulting in an insured loss of $\ell_2^* = 20$. Table 6 shows several schemes' cost allocations. The allocations (rows) sum to the total insured loss, which is 20.

In this simple example, with no direct links between banks, several schemes (the Shapley and Aumann-Shapley values for schemes that replace insured deposits with uninsured deposits, or replace deposits with equity, or eliminate banks with insured losses) yield the "standard" allocation, in which bank 2 is responsible for its own insured loss $\ell_2^* = 20$. Some schemes allocate a positive cost to bank 1, because they assign it responsibility for the bad effect of its loan to node 5 on bank 2: bank 1's claim on node 5 impairs node 5's ability to repay bank 2. Some schemes allocate a positive cost to bank 2 because it has an insured loss, but the obligor-responsibility scheme allocates it a negative loss because its equity reduces its insured loss, and the obligor-responsibility scheme assigns responsibility for its loan losses entirely to the borrowers. All the schemes allocate zero cost to node 3, whose presence has no impact on insured losses. Some schemes allocate a negative cost to node 4 because the profit that bank 2 earns on the loan to node 4 reduces bank 2's insured loss. Some schemes allocate a positive cost to the defaulting nodes 5 and 6 because of their responsibility for bank 2's loan losses, but others allocate a negative cost to them because their equity or the payments they make to bank 2 reduce its insured loss.

Table 6: Allocations of total insured loss in one scenario to nodes.

| Allocation Scheme and Principle | | Banks | | Industrial Nodes | | | |
|---------------------------------|----------------|-------|------|------------------|------|------|------|
| | | 1 | 2 | 3 | 4 | 5 | 6 |
| Standard | both | 0 | 20 | 0 | 0 | 0 | 0 |
| Banks | Shapley | 10 | 10 | 0 | 0 | 0 | 0 |
| All Nodes | Shapley | 5.4 | 6.1 | 0 | -2.3 | 4.7 | 6.1 |
| Obligor Responsibility | Aumann-Shapley | 0 | -32 | 0 | -8 | 20 | 40 |
| Creditor Responsibility | Aumann-Shapley | 70 | 370 | 0 | 0 | -140 | -280 |
| Shared Responsibility | Aumann-Shapley | 35 | 169 | 0 | -4 | -60 | -120 |
| Average Cost | Aumann-Shapley | 3.0 | 15.3 | 0 | -0.3 | 0.7 | 1.4 |
| Debt/Baseline | Aumann-Shapley | 7 | 29 | 0 | -4 | -4 | -8 |
| Industrial Nodes | Shapley | 0 | 0 | 0 | -4 | 10 | 14 |
| Loss-Making Industrial Nodes | Shapley | 0 | 0 | 0 | 0 | 10 | 10 |
| | Aumann-Shapley | 0 | 0 | 0 | 0 | 6.7 | 13.3 |

Table 7: Aumann-Shapley allocations of total insured loss in one scenario to nodes and loans.

| Scheme for Loans | Prices | Funding of node j | | | Loan from node i to node j | | | |
|------------------|----------|---------------------|------|-------|--------------------------------|--------|--------|--------|
| | | 2 | 5 | 6 | 1 to 5 | 2 to 4 | 2 to 5 | 2 to 6 |
| debt/equity | marginal | -32 | -140 | -280 | 70 | -8 | 90 | 320 |
| debt/equity | average | -1.4 | -5.9 | -11.9 | 5.9 | -0.7 | 7.3 | 26.5 |
| debt/baseline | marginal | -32 | -28 | -56 | 14 | -8 | 34 | 96 |

The Shapley value demonstrates its propensity for fairness as equality: in the schemes that allocate costs only to banks or only to loss-making industrial nodes, it allocates equal costs to the two banks or the two loss-making industrial nodes because there is zero insured loss unless both are present. In contrast, in all the schemes here, the ratio of the Aumann-Shapley value's allocations to nodes 6 and 5 equals the 2:1 ratio of bank 2's loss on the loan to node 6 to that on the loan to node 5.

The obligor- and creditor-responsibility schemes are extreme, with multiple nodes receiving cost allocations whose absolute value exceeds the total insured loss. What they accomplish is to show where scarce resources are produced and consumed: money at banks with insured losses, in the obligor-responsibility scheme, and money at defaulting industrial nodes, in the creditor-responsibility scheme. Although it is an average between schemes with opposite behavior, the shared-responsibility scheme is also extreme. The average-cost and debt/baseline schemes are more moderate.

Table 7 shows how the Aumann-Shapley value schemes allocate costs to nodes' funding and loans. It presents the non-zero numerical values in this scenario of the entries of Table 4. The obligor-, creditor-, and shared-responsibility schemes all allocate to funding and loans the costs given on the first row of the table; they differ only in how they re-allocate the costs of loans to the lender and borrower. This allocation of costs to funding and loans is extreme because the marginal prices are large ($\zeta_2 = 1$ for money at bank 2, $\zeta_5 = \Pi_{52}\zeta_2 = 1/2$ for money at node 5, and $\zeta_6 = \Pi_{62}\zeta_2 = 1$ for money at node 6) and the debt/equity scheme assigns responsibility to loans for removing principal from the lender but not for the cashflow generated by the borrower's investment. The average-cost scheme generates a more moderate allocation by using average prices, which are much smaller than the marginal prices because bank 2 is close to the boundary of insured loss. The debt/baseline scheme generates a more moderate allocation by assigning loans responsibility for increasing the borrower's cashflow and holding each node responsible only for its equity funding.

Next we consider some related examples which have four scenarios because each of two industrial nodes independently has a return on assets of 10% with probability 90% or -50% with probability 10%. The scenarios are denoted ++, +-, -+, and --, where the first and second symbols indicate whether the asset

Table 8: Allocations of expected total insured loss in examples of the counterparty contagion model.

| Expected Total Insured Loss | Example | No Swap | | Node 2 Owes Node 1 | | | | Node 3 Owes Node 4 | | | |
|---------------------------------|---------|---------|------|--------------------|-----|------|------|--------------------|-----|------|------|
| | | 1&2 | 3&4 | 76 | | | | 139 | | | |
| Allocation Scheme and Principle | | 1 | 2 | 3 | 4 | 1 | 2 | 3 | 4 | | |
| Insured vs. Uninsured | both | 56 | 0 | 20 | 56 | 0 | 0 | 86 | 53 | 0 | 0 |
| Deposits vs. Equity | S | 56 | 0 | 14 | 62 | 0 | 0 | 86 | 53 | 0 | 0 |
| Banks With Insured Losses | S | 56 | 0 | 20 | 56 | 0 | 0 | 87 | 52 | 0 | 0 |
| Banks | S | 56 | 0 | 38 | 38 | 0 | 0 | 87 | 52 | 0 | 0 |
| All Nodes | S | 28 | 28 | 17 | 17 | 17 | 25 | 38 | 27 | 37 | 37 |
| Obligor Responsibility | A-S | -24 | 80 | -24 | -46 | 80 | 66 | -24 | -24 | 107 | 80 |
| Creditor Responsibility | A-S | 296 | -240 | 296 | 260 | -240 | -240 | 272 | 296 | -213 | -216 |
| Shared Responsibility | A-S | 136 | -80 | 136 | 107 | -80 | -87 | 124 | 136 | -53 | -68 |
| Average Cost | A-S | 47 | 9 | 18 | 45 | 4 | 8 | 59 | 45 | 18 | 18 |
| Debt/Baseline | A-S | 56 | 0 | 56 | 27 | 0 | -7 | 53 | 56 | 18 | 12 |
| Industrial Nodes | S | 0 | 56 | 0 | 0 | 26 | 50 | 0 | 0 | 70 | 70 |
| Loss-Making Industrial Nodes | S | 0 | 56 | 0 | 0 | 19 | 57 | 0 | 0 | 83 | 56 |

returns r_3 and r_4 , respectively, are positive or negative. Bank 1 has deposits $d_1 = 3200$ and equity $s_1 = 240$, and has lent $D_{13} = 3200$ to industrial node 3, which has equity $s_3 = 1600$, all figures being quoted in millions of dollars. Bank 2 and industrial node 4 are identical. The interest rates are 5%. There are no other deals in the first example. Each bank has a loan loss of 800 and an insured loss of 560 when its borrower defaults due to a return on assets of -50%. In the second example, there is also a swap resulting in a promised payment $L_{21} = 480$ from bank 2 to bank 1 in scenarios $-+$ and $--$. This swap reduces total insured loss. In scenario $-+$, bank 2 defaults and pays to bank 1 only $p_2^* = 400$, of which 240 is its own equity and 160 is the profit earned on its loan to node 4. This leaves bank 1 with an insured loss of 160. In scenario $--$, bank 2 pays nothing and both banks' insured losses are 560. In the third example, instead there is a swap resulting in a promised payment $L_{34} = 480$ from node 3 to node 4 in scenarios $-+$ and $--$. This swap increases the total insured loss because it has equal priority with loan repayments: in scenarios $-+$ and $--$, node 3's total payment $p_3^* = 2400$ is split in the proportions $\Pi_{31} = 0.875$ and $\Pi_{34} = 0.125$, so node 1's insured loss is $\ell_3^* = 860$. In scenario $--$, this increase of 300 in insured loss is counterbalanced by a decrease of 300 in node 2's insured loss, but not in scenario $-+$, in which case $\ell_4^* = 0$ whether or not this swap exists. Table 8 presents systemic risk components from several schemes applied to these examples.

Including the promised swap payment from bank 2 to bank 1 in moving from the first to second example decreases systemic risk, i.e., expected total insured loss, by 36. In the schemes that replace insured with uninsured deposits or delete banks with insured losses, this decrease is felt entirely in the risk allocated to bank 1, which experiences the reduction in expected insured loss. The scheme that deletes any bank shares the reduction equally between the two banks' risk components, because the participation of both banks is necessary for the benefit to occur. According to the scheme that substitutes equity for deposits, the presence of the swap increases the risk component of bank 2, whose deposits prevent it from making its promised payment to bank 1 and thus are responsible for some of bank 1's insured loss. For a similar reason, when the swap is present, node 4 has a slightly higher risk component under the scheme that allocates risk to loss-making industrial nodes. Node 4's risk component decreases under schemes that give it credit for the profit bank 2 earns in lending to it, which enables bank 2 to pay more to bank 1 on the swap. In the obligor-responsibility scheme, the decrease in systemic risk is split (unequally) between bank 2 and node 4, because they are the obligors whose additional payments reduce the insured loss at bank 1. The same thing happens in the debt/baseline scheme for different reasons: bank 2's equity is more valuable because it also mitigates insured loss at bank 1, and the loan from bank 2 to node 4 is more valuable because the profits it generates flow through bank 2 to reduce bank 1's insured loss in scenario $-+$. Even though the

swap decreases systemic risk, it has zero marginal value in the debt/baseline scheme, because the promised payment is already so large that bank 2 can never make it. For this reason, bank 2’s risk component is the only one to decrease under the creditor-responsibility scheme: in scenario $-+$, bank 2 is yellow and receives credit for the payments it makes to bank 1, which is red. Under the all-nodes scheme, every risk component decreases, but node 4’s decreases the least. This is because it contributes less money than bank 2 to mitigating the insured loss at bank 1 and because its participation has no value without the participation of bank 2: without bank 2, node 4 is disconnected from bank 1. Under the average-cost scheme, every risk component decreases, but those of nodes 2 and 4 do not decrease much. In scenario $-+$, nodes 2 and 4 receive only a modest negative cost allocation because bank 1 is not far from the boundary of default and this scheme uses average prices, not marginal prices.

Comparing the first and third examples, the promised swap payment from node 3 to node 4 increases systemic risk by 27. The payment node 3 makes to node 4 increases the insured loss of bank 1 in scenarios $-+$ and $--$, and decreases the insured loss of bank 2 in scenario $--$, which is reflected in the risk components generated by the two schemes that replace insured deposits with another source of funding. Compared to these schemes, the risk component of bank 1 is slightly higher under the schemes that eliminate banks or banks with insured losses. The reason is that, in scenario $--$, eliminating bank 1 enables node 3 to make its entire promised payment to node 4, which increases node 4’s payment to bank 2 and decreases bank 2’s insured loss. The obligor-responsibility scheme holds node 3 entirely responsible for the increase in risk due to the swap, which comes about because node 3 pays less to bank 1. The scheme that allocates risk to loss-making industrial nodes does the same, because the increased risk comes entirely from an increase in total insured loss in scenario $-+$. The schemes that allocate risk to industrial nodes or to all nodes hold nodes 3 and 4 equally responsible for the effects of the swap between them. The all-nodes scheme also allocates some of the additional risk to bank 1, which is harmed by the swap, and slightly decreases the risk component of bank 2, which experiences a small reduction in expected insured loss due to the swap. The average-cost scheme behaves similarly, although its risk components for nodes 3 and 4 are not exactly equal. The introduction of the swap reduces the price ζ_3 of money at node 3 in scenario $-+$ from 1 to 0.875. Therefore, in the creditor-responsibility scheme, bank 1 gets a smaller and node 3 gets a larger allocation. Node 4 gets a larger allocation because it receives money from node 3 when the swap is present. Similar effects occur in the debt/baseline scheme, because the lower price of money at node 3 in scenario $-+$ decreases the risk allocation to the loan from bank 1 to node 3, while the swap itself has a positive risk allocation. Even though node 3 never makes the full promised payment on the swap, the swap has a positive marginal cost in scenario $-+$, because increasing the promised swap payment reduces the fraction of the money at node 3 that goes to bank 1.

5 Conclusions and Research Directions

We can draw some lessons about design principles from our exploration of systemic risk attribution schemes. Working with expected cost as our risk measure, we focused on designing a cost function to which to apply the Shapley or Aumann-Shapley values to allocate cost in each scenario. The first step is to decide which entities should receive a cost allocation. The second step is to design a cost function: the way that cost varies with the entities’ levels of participation in the system is the way that the function assigns responsibility for cost to those entities. This requires imagining counterfactual systems in which some of the entities are not participating fully, e.g., are smaller, weaker, or less connected. The counterfactual systems should be feasible. For example, we showed how to generate feasible counterfactual systems in a model that includes lending, without leaving holes in other banks’ balance sheets when one bank shrinks. We also showed how to design schemes that yield non-negative systemic risk components, by attributing systemic risk only to entities that increase systemic risk. A third step in designing an attribution scheme is choosing between the Shapley and Aumann-Shapley values. The Shapley value allocates equally among participants the costs of their interactions, whereas the Aumann-Shapley value can be interpreted in terms of unit prices and equality on a per-unit basis. The Shapley value quantifies global contributions to risk:

it considers large changes in the system that would occur if some entities did not participate in the system at all. The Aumann-Shapley value is based on local contributions to risk: it considers marginal changes in the system that would occur if participation levels were perturbed. Indeed, if the cost function is positively homogeneous, the Aumann-Shapley value consists of sensitivities of cost to perturbations, and we saw that they can be discontinuous as a function of system data. If the cost function is not positively homogeneous, the Aumann-Shapley value consists of sensitivities averaged across counterfactual systems. These average sensitivities were continuous in the examples we considered. Smoothing can also result from taking an expectation, i.e. averaging across scenarios. If participants in the system have incentives to lower their systemic risk components, the designer of a systemic risk attribution scheme should be aware of vulnerabilities to counterproductive strategic behavior: the Shapley value to mergers and splits of participants, and the Aumann-Shapley value to unilateral actions (Section 3.4).

In models of bank resolution costs (Section 3.1), we found schemes for creating systemic risk components that are worthy of consideration in setting deposit insurance premia (Sections 3.2–3.3). In the fire sale model, they are the natural Shapley and Aumann-Shapley schemes. In the acquisition model, they are the insolvent-banks Shapley and Aumann-Shapley schemes and the leverage Shapley scheme.

One task for future research is theoretical investigation of systemic risk components based on non-linear risk measures. Another is the application of systemic risk components to more models of systemic risk, e.g., of fire sales driven by capital requirements, of funding liquidity, or of counterparty contagion with more detailed features, such as multiple levels of seniority, netting arrangements, or collateral. More applied work will need to be done to design systemic risk attribution schemes for various purposes in systemic risk management.

References

- Acharya, V. V., Santos, J. A. C., Yorulmazer, T., August 2010. Systemic risk and deposit insurance premiums. FRBNY Economic Policy Review, 89–99.
- Adrian, T., Brunnermeier, M. K., August 2009. CoVaR. Staff report 348, Federal Reserve Bank of New York.
- Aikman, D., Alessandri, P., Eklund, B., Gai, P., Kapadia, S., Martin, E., Mora, N., Sterne, G., Willison, M., June 2009. Funding liquidity risk in a quantitative model of systemic stability. Working paper 372, Bank of England.
- Bennett, R. L., Nuxoll, D. A., Jarrow, R. A., Fu, M. C., Zhang, H., 2005. A loss default simulation model of the federal bank deposit insurance funds. In: Kuhl, M. E., Steiger, N. M., Armstrong, F. B., Joines, J. A. (Eds.), 2005 Winter Simulation Conference. pp. 1835–1843.
- Billera, L. J., Heath, D. C., 1982. Allocation of shared costs. *Mathematics of Operations Research* 7 (1), 32–39.
- Brunnermeier, M., Pedersen, L. H., 2009. Market liquidity and funding liquidity. *Review of Financial Studies* 22 (6), 2201–2238.
- Buch, A., Dorfleitner, G., 2008. Coherent risk measures, coherent capital allocations and the gradient allocation principle. *Insurance: Mathematics and Economics* 42, 235–242.
- Cherny, A., Orlov, D., 2011. On two approaches to coherent risk contribution. *Mathematical Finance*(forthcoming).
- Cifuentes, R., Ferrucci, G., Shin, H. S., 2005. Liquidity risk and contagion. *Journal of the European Economic Association* 3 (2-3), 556–566.

- David, L., Massol, O., Moison, A., 2005. The Shapley value as joint cost allocation mechanism: is the story definitely over?
 URL citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.110.9810
- Denault, M., 2001. Coherent allocation of risk capital. *Journal of Risk* 4 (1), 1–34.
- Diamond, D. W., Rajan, R. G., 2005. Liquidity shortages and banking crises. *Journal of Finance* 60 (2), 615–647.
- Duffie, D., Jarrow, R., Purnanandam, A., Yang, W., 2003. Market pricing of deposit insurance. *Journal of Financial Services Research* 24 (2-3), 93–119.
- Eisenberg, L., Noe, T. H., 2001. Systemic risk in financial systems. *Management Science* 47 (2), 236–249.
- Elsinger, H., February 2007. Financial networks, cross holdings, and limited liability.
 URL ssrn.com/abstract=916763
- Elsinger, H., Lehar, A., Summer, M., 2006. Risk assessment for banking systems. *Management Science* 52 (9), 1301–1314.
- Föllmer, H., Schied, A., 2004. *Stochastic Finance: An Introduction in Discrete Time*, 2nd Edition. De Gruyter, Berlin.
- Frye, J., Pelz, E., 2008. BankCaR (Bank Capital-at-Risk): a credit risk model for US commercial bank charge-offs. Working paper 2008-03, Federal Reserve Bank of Chicago.
- Gass, S. I., 2003. *Linear Programming: Methods and Applications*, 5th Edition. Dover, New York.
- Gauthier, C., Lehar, A., Souissi, M., January 2010. Macroprudential regulation and systemic capital requirements. Working paper 2010-4, Bank of Canada.
- Glasserman, P., 2004. *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York.
- Goldberg, L. R., Hayes, M. Y., Menchero, J., Mitra, I., February 2009. Extreme risk management. Working paper 2009-4, MSCI Barra.
- Gorton, G., May 2009. Slapped in the face by the invisible hand: banking and the panic of 2007.
 URL ssrn.com/abstract=1401882
- Huang, X., Zhou, H., Zhu, H., 2009. A framework for assessing the systemic risk of major financial institutions. *Journal of Banking and Finance* 33, 2036–2049.
- Huang, X., Zhou, H., Zhu, H., January 2010. Assessing the systemic risk of a heterogeneous portfolio of banks during the recent financial crisis. Working paper 296, Bank for International Settlements.
- Jarrow, R. A., Bennett, R. L., Fu, M. C., Nuxoll, D. A., Zhang, H., November 2003. A general martingale approach to measuring and valuing the risk to the FDIC deposit insurance funds.
 URL www.fdic.gov/bank/analytical/cfr/Jarrow.pdf
- Kalai, E., Zemel, E., 1982. Totally balanced games and games of flow. *Mathematics of Operations Research* 7 (3), 476–478.
- Kiyotaki, N., Moore, J., January 1997. Credit chains.
 URL http://www.mpls.frb.org/Research/events/1997_01-31/Kiyotaki_CreditChains.pdf
- Krishnamurthy, A., 2010. Amplification mechanisms in liquidity crises. *American Economic Journal: Macroeconomics* 2 (3), 1–30.

- Kuritzkes, A., Schuermann, T., Weiner, S. M., 2005. Deposit insurance and risk management of the U.S. financial system: what is the loss distribution faced by the FDIC? *Journal of Financial Services Research* 27 (3), 217–242.
- Lehar, A., 2005. Measuring systemic risk: a risk management approach. *Journal of Banking and Finance* 29 (10), 2577–2603.
- Liu, M., Staum, J., 2010. Sensitivity analysis of the Eisenberg-Noe model of contagion. *Operations Research Letters* 38 (5), 489–491.
- Mirman, L. J., Tauman, Y., 1982. Demand compatible equitable cost sharing prices. *Mathematics of Operations Research* 7 (1), 40–56.
- Moulin, H., Sprumont, Y., 2007. Fair allocation of production externalities: recent results. *Revue d’Économie Politique* 117 (1), 7–36.
- Nelson, B. L., 2002. *Stochastic Models: Analysis and Simulation*. Dover, New York.
- Pennachi, G., 2006. Deposit insurance, banking regulation, and financial system risks. *Journal of Monetary Economics* 53, 1–30.
- Samet, D., Zemel, E., 1984. On the core and dual set of linear programming games. *Mathematics of Operations Research* 9 (2), 309–316.
- Shin, H. S., 2008. Risk and liquidity in a system context. *Journal of Financial Intermediation* 17, 315–329.
- Shubik, M., 1962. Incentives, decentralized control, the assignment of joint costs and internal pricing. *Management Science* 8 (3), 325–343.
- Sudhölter, P., 1998. Axiomatizations of game theoretical solutions for one-output cost sharing problems. *Games and Economic Behavior* 24, 142–171.
- Tarashev, N., Borio, C., Tsatsaronis, K., May 2010. Attributing systemic risk to individual institutions. Working paper 308, Bank of International Settlements.
- Tasche, D., 2004. Allocating portfolio economic capital to sub-portfolios. In: Dev, A. (Ed.), *Economic Capital: A Practitioner Guide*. Risk Books, London, pp. 275–302.
- Tsanakas, A., 2009. To split or not to split: Capital allocation with convex risk measures. *Insurance: Mathematics and Economics* 44, 268–277.

A Linear Programming and the Clearing Payment Vector

In the counterparty contagion model (Section 4.1), we assume that there is a unique clearing payment vector; sufficient conditions are discussed by Eisenberg and Noe (2001) and Elsinger (2007). Eisenberg and Noe show that in their model, which is a special case of ours, the clearing payment vector is the optimal solution of an LP similar to our LP (12). Elsinger, whose model is more general than ours, provides algorithms for computing the clearing payment vector, but does not discuss LPs. We show that in our model of counterparty contagion, the clearing payment vector provides an optimal solution of our LP (12), which we reformulate here as

$$\min \mathbf{1}^\top \boldsymbol{\ell} \quad \text{subject to} \quad \mathbf{v} - \boldsymbol{\ell} + (\mathbf{I} - \boldsymbol{\Pi}^\top) \mathbf{p} = \mathbf{w}, \quad \mathbf{p} \leq \bar{\mathbf{p}}, \quad \mathbf{v} \geq \mathbf{0}, \quad \boldsymbol{\ell} \geq \mathbf{0}, \quad \mathbf{p} \geq \mathbf{0}. \quad (29)$$

By way of doing so, we provide an LP-based method for computing the clearing payment vector.

The clearing payment vector \mathbf{p}^* satisfies, with some \mathbf{v}^* and $\boldsymbol{\ell}^*$,

- the *balance equation* $\mathbf{v}^* - \boldsymbol{\ell}^* = \mathbf{u}^* = \mathbf{w} + (\boldsymbol{\Pi}^\top - \mathbf{I})\mathbf{p}^*$,
- the *capacity constraints* $\mathbf{0} \leq \mathbf{p}^* \leq \bar{\mathbf{p}}$, $\mathbf{0} \leq \boldsymbol{\ell}^* \leq \mathbf{d}$, and $\mathbf{v}^* \geq \mathbf{0}$, and
- the *priority constraints* $v_i^* > 0 \Rightarrow p_i^* = \bar{p}_i$ and $p_i^* > 0 \Rightarrow \ell_i^* = 0$ for all $i = 1, \dots, n$.

It is a consequence of the priority constraints that $v_i^* > 0 \Rightarrow \ell_i^* = 0$ for all $i = 1, \dots, n$. Together with $\mathbf{v}^* - \boldsymbol{\ell}^* = \mathbf{u}^*$ and the capacity constraint $\mathbf{v}^* \geq \mathbf{0}$, this implies that the insured loss ℓ_i^* and terminal equity v_i^* are the negative and positive parts, respectively, of u_i^* ; they can be computed from \mathbf{p}^* . Thus, the clearing payment vector \mathbf{p}^* is in one-to-one correspondence with the solution $(\mathbf{p}^*, \mathbf{v}^*, \boldsymbol{\ell}^*)$ of LP (29).

Because the objective of LP (29) is to minimize $\mathbf{1}^\top \boldsymbol{\ell}$, the upper bound $\boldsymbol{\ell}^* \leq \mathbf{d}$ is not needed. The only difference between the LP (29) and the definition of the clearing payment vector is that (29) lacks the priority constraints, while the definition of the clearing payment vector makes no mention of minimizing insured loss. Therefore the clearing payment vector corresponds to a feasible solution $(\mathbf{p}^*, \mathbf{v}^*, \boldsymbol{\ell}^*)$ of LP (29). It remains to show that this solution is optimal. There may be other optimal solutions that do not correspond to the clearing payment vector. This is not a problem because we only use LP (29) to study the sensitivity of the optimal objective. To show that $(\mathbf{p}^*, \mathbf{v}^*, \boldsymbol{\ell}^*)$ is an optimal solution of (29), we use a three-step procedure for calculating the clearing payment vector:

1. Solve LP (29) and record the optimal objective value as ℓ_- .
2. Solve the LP

$$\min \mathbf{1}^\top \mathbf{p} \text{ s.t. } \mathbf{v} - \boldsymbol{\ell} + (\mathbf{I} - \boldsymbol{\Pi}^\top)\mathbf{p} = \mathbf{w}, \mathbf{0} \leq \mathbf{v}, \mathbf{0} \leq \boldsymbol{\ell}, \mathbf{0} \leq \mathbf{p} \leq \bar{\mathbf{p}}, \mathbf{1}^\top \boldsymbol{\ell} \leq \ell_-. \quad (30)$$

Record an optimal solution as $(\mathbf{p}^{(2)}, \mathbf{v}^{(2)}, \boldsymbol{\ell}^*)$. Let \mathbf{U}^- be the vector whose i th entry is 1 if $\ell_i^* > 0$ and 0 otherwise.

3. Solve the LP

$$\max \mathbf{1}^\top \mathbf{p} \text{ s.t. } \mathbf{v} - \boldsymbol{\ell}^* + (\mathbf{I} - \boldsymbol{\Pi}^\top)\mathbf{p} = \mathbf{w}, \mathbf{0} \leq \mathbf{v}, \mathbf{0} \leq \mathbf{p} \leq \bar{\mathbf{p}}, \text{diag}(\mathbf{U}^-)(\mathbf{p} + \mathbf{v}) = \mathbf{0}. \quad (31)$$

Call the optimal solution $(\mathbf{p}^*, \mathbf{v}^*)$. The clearing payment vector is \mathbf{p}^* .

As a preliminary, we remark that all three LPs have optimal solutions because they are bounded and feasible. LP (29) is bounded by zero and a feasible solution can be generated by taking $\mathbf{p} = \mathbf{0}$. LP (30) is bounded by zero and the optimal solution of LP (29) is a feasible solution for LP (30). LP (31) is bounded by $\mathbf{1}^\top \bar{\mathbf{p}}$ and $(\mathbf{p}^{(2)}, \boldsymbol{\ell}^{(2)})$ is a feasible solution: Lemma 3 shows that the constraint $\text{diag}(\mathbf{U}^-)(\mathbf{p} + \mathbf{v}) = \mathbf{0}$ is satisfied.

The purpose of LP (30) is to find a payment vector that results in the smallest possible insured loss while ensuring that the nodes with insured losses respect the priority constraints: they do not make any payments to other nodes and have zero terminal value. The next two lemmas show that LP (30) accomplishes this.

Lemma 2. *The optimal solutions of LPs (29) and (30) have the same cost: $\mathbf{1}^\top \boldsymbol{\ell}^* = \ell_-$.*

Proof. The optimal solution of LP (30) is a feasible solution of LP (29), so $\ell_- \leq \mathbf{1}^\top \boldsymbol{\ell}^*$. Also, $\mathbf{1}^\top \boldsymbol{\ell}^* \leq \ell_-$ follows from a constraint in (30). \square

Lemma 3. *For all i such that $\ell_i^* > 0$, $v_i^{(2)} = 0$ and $p_i^{(2)} = 0$.*

Proof. Consider a solution $(\mathbf{p}, \mathbf{v}, \boldsymbol{\ell})$ to LP (30) such that $\ell_i > 0$.

First, suppose that $v_i > 0$. Let $\epsilon = \min\{\ell_i, v_i\}$. Define $\ell'_i = \ell_i - \epsilon$ and $v'_i = v_i - \epsilon$, whereas for all $j \neq i$, $\ell'_j = \ell_j$ and $v'_j = v_j$. The solutions $(\mathbf{p}, \mathbf{v}, \boldsymbol{\ell})$ and $(\mathbf{p}, \mathbf{v}', \boldsymbol{\ell}')$ are either both feasible or both infeasible for LP (29). If they are both infeasible for LP (29), then $(\mathbf{p}, \mathbf{v}, \boldsymbol{\ell})$ is infeasible for LP (30). Suppose instead that both are feasible for LP (29). Because $\mathbf{1}^\top \boldsymbol{\ell}' < \mathbf{1}^\top \boldsymbol{\ell}$, while the optimal value of LP (29) is ℓ_- , it follows that $\ell_- \leq \mathbf{1}^\top \boldsymbol{\ell}' < \mathbf{1}^\top \boldsymbol{\ell}$. Thus $(\mathbf{p}, \mathbf{v}, \boldsymbol{\ell})$ is infeasible for LP (30). Therefore the optimal solution of LP (30), being feasible, must satisfy $v_i^{(2)} = 0$.

Next, suppose instead that $p_i > 0$. Let $\epsilon = \min\{\ell_i, p_i\}$. Define $\ell'_i = \ell_i - \epsilon$ and $p'_i = p_i - \epsilon$, whereas for all $j \neq i$, $p'_j = p_j$ and $\ell'_j = \ell_j + \epsilon \Pi_{ij}$. Because $\sum_{j \neq i} \Pi_{ij} = 1$, $\mathbf{1}^\top \ell' = \mathbf{1}^\top \ell$. Therefore the solutions $(\mathbf{p}, \mathbf{v}, \ell)$ and $(\mathbf{p}', \mathbf{v}, \ell')$ are either both feasible or both infeasible for LP (30). Because $\mathbf{1}^\top \mathbf{p}' < \mathbf{1}^\top \mathbf{p}$, $(\mathbf{p}, \mathbf{v}, \ell)$ is not an optimal solution of LP (30). Therefore the optimal solution of LP (30) must satisfy $p_i^{(2)} = 0$. \square

Proposition 4. *Where $(\mathbf{p}^*, \mathbf{v}^*, \ell^*)$ is computed by the three-step procedure, \mathbf{p}^* is the clearing payment vector and $\mathbf{1}^\top \ell^*$ is the optimal objective value of LP (29).*

Proof. It suffices to show that the procedure's output $(\mathbf{p}^*, \mathbf{v}^*, \ell^*)$ satisfies the balance equation, capacity constraints, and priority constraints. Then \mathbf{p}^* is the clearing payment vector and Lemma 2 implies that $\mathbf{1}^\top \ell^*$ is the optimal objective value of the LP (29). Consequently, comparing (29), (30), and (31), we can see that $(\mathbf{p}^*, \mathbf{v}^*, \ell^*)$ is an optimal solution to LP (29). As discussed previously, LP (29) includes the balance equation and capacity constraints, so it remains only to show that $(\mathbf{p}^*, \mathbf{v}^*, \ell^*)$ satisfies the priority constraints. Because of the constraint $\text{diag}(\mathbf{U}^-)(\mathbf{p} + \mathbf{v}) = \mathbf{0}$ in (31), the priority constraints are satisfied for i such that $\ell_i^* > 0$. For i such that $\ell_i^* = 0$, the only priority constraint to consider is $v_i^* > 0 \Rightarrow p_i^* = 1$.

Suppose that (\mathbf{p}, \mathbf{v}) is feasible for LP (31). Then $(\mathbf{p}, \mathbf{v}, \ell^*)$ is feasible for LPs (29) and (30). Further suppose that for some i , $p_i < \bar{p}_i$, $v_i > 0$, and $\ell_i^* = 0$. Let $\epsilon = \min\{\bar{p}_i - p_i, v_i\}$. Define $p'_i = p_i + \epsilon$ and, for all $j \neq i$, $p'_j = p_j$. Define $v'_i = v_i - \epsilon$ and $\ell'_i = 0$. For all $j \neq i$, define $u'_j = v_j - \ell_j + \epsilon \Pi_{ij}$, $v'_j = \max\{u'_j, 0\}$, and $\ell'_j = \max\{-u'_j, 0\}$. Then $(\mathbf{p}', \mathbf{v}', \ell')$ is feasible for LP (29). Because $\ell' \leq \ell^*$, $(\mathbf{p}', \mathbf{v}', \ell')$ is also feasible for LP (30). Then Lemma 2 implies $\mathbf{1}^\top \ell' \geq \mathbf{1}^\top \ell^*$, so $\ell' = \ell^*$. Therefore $(\mathbf{p}', \mathbf{v}')$ is feasible for LP (31). Because $\mathbf{1}^\top \mathbf{p}' > \mathbf{1}^\top \mathbf{p}$, (\mathbf{p}, \mathbf{v}) is not an optimal solution of LP (31). \square

B Linear Programming Sensitivity Analysis

We use LP sensitivity analysis to find the sensitivity of cost, i.e., total insured loss, in the counterparty contagion model (Section 4.1), proving Propositions 1 and 2. For background material on LP sensitivity analysis, see, e.g., Gass (2003, §8.3). As in Appendix A, we assume that there is a unique clearing payment vector \mathbf{p}^* , and let \mathbf{v}^* and ℓ^* be the corresponding vectors of terminal equity and insured losses. Compared to the sensitivity analysis of the Eisenberg and Noe (2001) model in Liu and Staum (2010), here we obtain more explicit formulae for a more general model, but we do not treat the case in which there are borderline nodes (Def. 3). In the absence of borderline nodes, the cost $\mathbf{1}^\top \ell^*$ is a differentiable function of the financial network data at the point being considered (Section 4.2). Then each sensitivity is a partial derivative of cost with respect to one parameter of the financial network. Borderline nodes can be handled with an approach similar to that of Liu and Staum (2010). Left and right derivatives may differ; multiple subgradients may exist. One must choose an appropriate basis to compute the sensitivity of cost to perturbations in each direction. For example, a right derivative with respect to w_j is computed using a basis that remains a basis when w_j is increased.

B.1 Proof of Proposition 1

First we reformulate LP (12) as

$$\min \mathbf{c}^\top \mathbf{x} \quad \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b}, \quad \mathbf{x} \geq \mathbf{0} \quad (32)$$

where

$$\mathbf{x} = \begin{bmatrix} \mathbf{p} \\ \mathbf{q} \\ \mathbf{v} \\ \ell \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{1} \end{bmatrix}, \quad \mathbf{A} = \begin{bmatrix} \mathbf{I} - \mathbf{\Pi}^\top & \mathbf{0} & \mathbf{I} & -\mathbf{I} \\ & \mathbf{I} & \mathbf{I} & \mathbf{0} \end{bmatrix}, \quad \text{and} \quad \mathbf{b} = \begin{bmatrix} \mathbf{w} \\ \bar{\mathbf{p}} \end{bmatrix}.$$

It follows from Appendix A that $\mathbf{x}^* = (\mathbf{p}^*, \bar{\mathbf{p}} - \mathbf{p}^*, \mathbf{v}^*, \ell^*)$ is an optimal solution of LP (32). If there are borderline nodes, \mathbf{x}^* is a degenerate solution and has multiple bases. Suppose there are no borderline nodes.

Then \mathbf{x}^* is non-degenerate and its unique basis \mathcal{B} is constructed as follows, in terms of the classification of nodes as green, yellow, and red (Def. 4).

- If node i is red, the variables q_i and ℓ_i are in the basis.
- If node i is yellow, the variables p_i and q_i are in the basis.
- If node i is green, the variables p_i and v_i are in the basis.

The absence of borderline nodes implies that \mathbf{b} can be perturbed in any direction without changing the basis. The basic matrix $\mathbf{B} = \mathbf{A}_{\mathcal{B}}$ is formed by selecting the $2n$ columns of \mathbf{A} corresponding to \mathcal{B} . The optimal solution is $\mathbf{x}^* = \mathbf{B}^{-1}\mathbf{b}$. The cost $\mathbf{1}^\top \boldsymbol{\ell}^*$ is the optimal objective value $\mathbf{c}_{\mathcal{B}}^\top \mathbf{B}^{-1}\mathbf{b}$, so the dual-optimal solution $\mathbf{y}^* = \mathbf{c}_{\mathcal{B}}^\top \mathbf{B}^{-1}$ provides the sensitivities of cost to \mathbf{b} . We must demonstrate that $[-\boldsymbol{\zeta} - \boldsymbol{\eta}] = \mathbf{y}^*$, equivalently, $[-\boldsymbol{\zeta} - \boldsymbol{\eta}]\mathbf{B} = \mathbf{c}_{\mathcal{B}}^\top$, the row vector whose first $2n - |\mathcal{R}|$ elements are zero and whose last $|\mathcal{R}|$ elements are one, corresponding to the variable ℓ_i for each $i \in \mathcal{R}$. The element corresponding to p_i is $-\zeta_i + \sum_{j \neq i} \Pi_{ij} \zeta_j - \eta_i$. The vector of such elements for green nodes is $-\boldsymbol{\zeta}_{\mathcal{G}} + \mathbf{\Pi}_{\mathcal{G}} \cdot \boldsymbol{\zeta} - \boldsymbol{\eta}_{\mathcal{G}} = \mathbf{0} + \boldsymbol{\eta}_{\mathcal{G}} - \boldsymbol{\eta}_{\mathcal{G}} = \mathbf{0}$. The vector of such elements for yellow nodes is $-\boldsymbol{\zeta}_{\mathcal{Y}} + \mathbf{\Pi}_{\mathcal{Y}} \cdot \boldsymbol{\zeta} - \boldsymbol{\eta}_{\mathcal{Y}} = -\boldsymbol{\zeta}_{\mathcal{Y}} + \mathbf{\Pi}_{\mathcal{Y}\mathcal{R}} \boldsymbol{\zeta}_{\mathcal{R}} + \mathbf{\Pi}_{\mathcal{Y}\mathcal{Y}} \boldsymbol{\zeta}_{\mathcal{Y}} + \mathbf{\Pi}_{\mathcal{Y}\mathcal{G}} \boldsymbol{\zeta}_{\mathcal{G}} - \mathbf{0} = \mathbf{\Pi}_{\mathcal{Y}\mathcal{R}} \mathbf{1} - (\mathbf{I} - \mathbf{\Pi}_{\mathcal{Y}\mathcal{Y}}) \boldsymbol{\zeta}_{\mathcal{Y}} + \mathbf{0} = \mathbf{0}$. The element corresponding to q_i for a yellow or red node is $-\eta_i = 0$. The element corresponding to v_i for a green node is $-\zeta_i = 0$. The element corresponding to ℓ_i for a red node is $\zeta_i = 1$.

B.2 Proof of Proposition 2

To handle changes to the liability matrix \mathbf{L} that change $\mathbf{\Pi}$, we reformulate LP (16) as

$$\min \mathbf{c}^\top \tilde{\mathbf{x}} \quad \text{subject to} \quad \tilde{\mathbf{A}} \tilde{\mathbf{x}} = \tilde{\mathbf{b}}, \quad \tilde{\mathbf{x}} \geq \mathbf{0} \quad (33)$$

where

$$\tilde{\mathbf{x}} = \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \\ \mathbf{v} \\ \boldsymbol{\ell} \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{1} \end{bmatrix}, \quad \tilde{\mathbf{A}} = \begin{bmatrix} \text{diag}(\mathbf{L}\mathbf{1}) - \mathbf{L}^\top & \mathbf{0} & \mathbf{I} & -\mathbf{I} \\ \mathbf{I} & \mathbf{I} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \text{and} \quad \tilde{\mathbf{b}} = \begin{bmatrix} \mathbf{w} \\ \mathbf{1} \end{bmatrix}.$$

The optimal solution to LP (33) has the same basis \mathcal{B} as the optimal solution to LP (32), if we regard f_i and g_i as corresponding to p_i and q_i , respectively, for all $i = 1, \dots, n$. Thus $\tilde{\mathbf{B}} = \tilde{\mathbf{A}}_{\mathcal{B}}$ is the basic matrix of LP (33), the optimal solution $\tilde{\mathbf{x}}^* = \tilde{\mathbf{B}}^{-1}\tilde{\mathbf{b}}$, and the dual-optimal solution $\tilde{\mathbf{y}}^* = \mathbf{c}_{\tilde{\mathcal{B}}}^\top \tilde{\mathbf{B}}^{-1}$ provides the sensitivities of cost to $\tilde{\mathbf{b}}$. Therefore $\tilde{y}_i^* = \partial \mathbf{1}^\top \boldsymbol{\ell}^* / \partial w_i = y_i^* = -\zeta_i$ for all $i = 1, \dots, n$. The sensitivity of cost to the entry \tilde{a}_{kl} of the constraint matrix $\tilde{\mathbf{A}}$ is $-\tilde{y}_k^* f_l^*$. Recall that diagonal entries of \mathbf{L} are zero. For $i \neq j$, the sensitivity of \tilde{a}_{kl} to L_{ij} is 1 if $k = l = i$, -1 if $k = j$ and $l = i$, and 0 otherwise. Therefore the sensitivity of cost to L_{ij} is $\tilde{y}_j^* f_i^* - \tilde{y}_i^* f_j^* = f_i^*(\zeta_i - \zeta_j)$.

B.3 Derivation of Equations (22)–(24)

Here we derive the average prices used in Section 4.4.6. The marginal prices $\boldsymbol{\zeta}(\gamma)$ are piecewise constant in γ , with their points of discontinuity contained in the set $\{\gamma_1, \dots, \gamma_{m-1}\}$. Therefore $\int_0^1 \zeta_i(\gamma) d\gamma = \sum_{h=1}^m \delta_h \zeta_i(\mu_h)$ and $2 \int_0^1 \gamma \zeta_i(\gamma) d\gamma = \sum_{h=1}^m (\gamma_h^2 - \gamma_{h-1}^2) \zeta_i(\mu_h)$. Equations (22)–(23) follow because $\zeta_i(\gamma) = 0$ if node i is green in system $\Phi(\gamma\mathbf{1})$ and $\zeta_i(\gamma) = 1$ if node i is red in system $\Phi(\gamma\mathbf{1})$. Similarly, $\theta_{ij}(\gamma) = \zeta_j(\gamma)$ if node i is green in system $\Phi(\gamma\mathbf{1})$ and $\theta_{ij}(\gamma) = 0$ if node i is red in system $\Phi(\gamma\mathbf{1})$. It remains to show that $2 \int_{\gamma_{h-1}}^{\gamma_h} \gamma \theta_{ij}(\gamma) d\gamma = (\gamma_h^2 - \gamma_{h-1}^2) \theta_{ij}(\mu_h)$ if node i is yellow in system $\Phi(\gamma\mathbf{1})$ for all $\gamma \in (\gamma_{h-1}, \gamma_h)$.

We have parameterized families of LPs given by (32) and (33) with $\mathbf{w}(\gamma\mathbf{1}) = \gamma(\mathbf{w} + (1 - \gamma)\mathbf{D}\mathbf{1})$ and $\mathbf{L}(\gamma\mathbf{1}) = \gamma^2 \mathbf{L}$ plugged in for \mathbf{w} and \mathbf{L} , respectively. For all $h = 1, \dots, m$, the basis $\mathcal{B}(\gamma)$ is the same for any $\gamma \in (\gamma_{h-1}, \gamma_h)$. Also, $\mathbf{\Pi}(\gamma\mathbf{1}) = \mathbf{\Pi}$ for any $\gamma \in (0, 1]$. Consider h such that $G_i < h \leq Y_i$. For all $\gamma \in (\gamma_{h-1}, \gamma_h)$, the basic matrix $\mathbf{B}(\gamma)$ of the optimal solution $\mathbf{x}^*(\gamma) = (\mathbf{p}^*(\gamma), \bar{\mathbf{p}}(\gamma) - \mathbf{p}^*(\gamma), \mathbf{v}^*(\gamma), \boldsymbol{\ell}^*(\gamma))$

is the same. The payment $p_i^*(\gamma)$ made by node i in system $\Phi(\gamma\mathbf{1})$ forms a piecewise linear function of γ . Consider any $\gamma \in (\gamma_{h-1}, \gamma_h)$. In system $\Phi(\gamma\mathbf{1})$, node i 's payment $p_i^*(\gamma) = ((\mathbf{B}(\mu_h))^{-1}\mathbf{b}(\gamma))_i = \boldsymbol{\psi}_i(\mu_h)\mathbf{b}(\gamma)$, where $\boldsymbol{\psi}_i(\mu_h)$ is the i th row of $(\mathbf{B}(\mu_h))^{-1}$. Its promised payment is $\bar{p}_i(\gamma) = \gamma^2\bar{p}_i$ and its payment fraction $f_i^*(\gamma) = p_i^*(\gamma)/\bar{p}_i(\gamma)$. Therefore

$$\theta_{ij}(\gamma) = f_i^*(\gamma)(\zeta_j(\gamma) - \zeta_i(\gamma)) = (\zeta_j(\mu_h) - \zeta_i(\mu_h))\frac{\boldsymbol{\psi}_i(\mu_h)\mathbf{b}(\gamma)}{\bar{p}_i\gamma^2}.$$

The integral

$$\int_{\gamma_{h-1}}^{\gamma_h} \gamma\theta_{ij}(\gamma) d\gamma = (\zeta_j(\mu_h) - \zeta_i(\mu_h))\frac{\boldsymbol{\psi}_i(\mu_h)}{\bar{p}_i} \int_{\gamma_{h-1}}^{\gamma_h} \frac{1}{\gamma}\mathbf{b}(\gamma) d\gamma,$$

and

$$\begin{aligned} \int_{\gamma_{h-1}}^{\gamma_h} \frac{1}{\gamma}\mathbf{b}(\gamma) d\gamma &= \int_{\gamma_{h-1}}^{\gamma_h} \begin{bmatrix} \mathbf{w} + (1-\gamma)\mathbf{D}\mathbf{1} \\ \gamma\bar{\mathbf{p}} \end{bmatrix} d\gamma \\ &= \begin{bmatrix} (\mathbf{w} + \mathbf{D}\mathbf{1})(\gamma_h - \gamma_{h-1}) - \mathbf{D}\mathbf{1}(\gamma_h^2 - \gamma_{h-1}^2)/2 \\ \bar{\mathbf{p}}(\gamma_h^2 - \gamma_{h-1}^2)/2 \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{w} + (1-\mu_h)\mathbf{D}\mathbf{1} \\ \mu_h\bar{\mathbf{p}} \end{bmatrix} \delta_h \\ &= \frac{1}{\mu_h}\mathbf{b}(\mu_h)\delta_h. \end{aligned}$$

Therefore $\boldsymbol{\psi}_i(\mu_h) \int_{\gamma_{h-1}}^{\gamma_h} \frac{1}{\gamma}\mathbf{b}(\gamma) d\gamma = \boldsymbol{\psi}_i(\mu_h)\mathbf{b}(\mu_h)\delta_h/\mu_h = p_i^*(\mu_h)\delta_h/\mu_h$, so

$$\int_{\gamma_{h-1}}^{\gamma_h} \gamma\theta_{ij}(\gamma) d\gamma = (\zeta_j(\mu_h) - \zeta_i(\mu_h))\frac{p_i^*(\mu_h)}{\bar{p}_i\mu_h}\delta_h = \theta_{ij}(\mu_h)\delta_h\mu_h = \frac{1}{2}(\gamma_h^2 - \gamma_{h-1}^2)\theta_{ij}(\mu_h).$$