



**University of
Zurich^{UZH}**

Department of Informatics

A Formal Analysis of Complexity and Systemic Risk in Financial Networks with Derivatives

Dissertation submitted to the
Faculty of Business, Economics and Informatics
of the University of Zurich

to obtain the degree of
Doktor der Wissenschaften, Dr. sc.
(corresponds to Doctor of Science, PhD)

presented by
Steffen Schuldenzucker
from Bonn, Germany

approved in October 2019

at the request of
Prof. Sven Seuken, Ph.D.
Prof. Constantinos Daskalakis, Ph.D.
Prof. Michael Wellman, Ph.D.

The Faculty of Business, Economics and Informatics of the University of Zurich hereby authorizes the printing of this dissertation, without indicating an opinion of the views expressed in the work.

Zurich, October 23, 2019

Chairman of the Doctoral Board: Prof. Thomas Fritz, Ph.D.

Abstract

The 2008 financial crisis has been attributed by policymakers to “excessive complexity” of the *financial network*, especially due to *financial derivatives*. In a financial network, financial institutions (“banks” for short) are connected by financial contracts. As banks depend on payments from contracts with other banks to cover their own obligations, such a situation creates *systemic risk*, i.e., the risk of a financial crisis. Some of the contracts are financial *derivatives*, where an obligation to pay depends on another variable.

In this thesis, I study in what sense derivatives make a financial network fundamentally “more complex” compared to one without derivatives. I capture the notion of “complexity” formally using tools from finance and theoretical computer science. I reveal new kinds of systemic risk that arise in financial networks specifically because of derivatives and I discuss the impact of recent regulatory policy.

I first focus on a type of derivative called a *credit default swap (CDS)*, in which the writer insures the holder of the contract against the default (i.e., bankruptcy) of a third party, the *reference entity*. I show that, when the reference entity is another bank, then such CDSs introduce a new kind of systemic risk arising from what I call *default ambiguity*. Default ambiguity is a situation where it is impossible to decide which banks are in default following a shock (i.e., a loss in banks’ assets). At a technical level, I show that the *clearing problem* may have no solution or multiple incompatible solutions. In contrast, without CDSs, a unique canonical solution always exists.

I then demonstrate that increased “complexity” due to CDSs also manifests as *computational complexity*. More in detail, I show that the clearing problem leads to NP-complete decision and PPAD-complete approximation problems if CDSs are allowed. This implies a fundamental barrier to the computational analysis of these networks, specifically to macroprudential *stress testing*. Without CDSs, the problems are either trivial or in P. I study the impact of different regulatory policies. My main result is that the aforementioned phenomena can be attributed to *naked CDS positions*.

In a final step, I focus on one specific regulatory policy: mandatory *portfolio compression*, which is a post-trade mechanism by which cycles in the financial network are eliminated. While this always reduces individual exposures, I show that, surprisingly, it can worsen the impact of certain shocks. Banks’ incentives to compress may further be misaligned with social welfare. I provide sufficient conditions on the network structure under which these issues are eliminated. Overall, my results in this thesis contribute to a better understanding of systemic risk and the effects of regulatory policy.

Zusammenfassung

Die Finanzkrise des Jahres 2008 wurde von politischen Entscheidungsträgern als Folge von “unverhältnismäßiger Komplexität” des *Finanznetzwerkes*, besonders aufgrund von *Finanzderivaten*, bezeichnet. In einem Finanznetzwerk sind Finanzinstitutionen (kurz “Banken”) durch Finanzverträge verbunden. Da die Banken von Zahlungen aus ihren Verträgen mit anderen Banken abhängen, um ihren eigenen Verpflichtungen nachzukommen, erzeugt eine solche Situation *systemisches Risiko*, d.h., das Risiko einer Finanzkrise. Einige der Verträge sind *Derivate*, bei denen eine Zahlungsverpflichtung von einer anderen Variable abhängt.

In der vorliegenden Arbeit untersuche ich, in welchem Sinne Derivate ein Finanznetzwerk fundamental “komplexer” machen, verglichen mit einem Netzwerk ohne Derivate. Ich verwende Werkzeuge aus den Finanzwissenschaften und der theoretischen Informatik, um den Begriff der “Komplexität” formal zu erfassen. Ich decke neue Arten von systemischem Risiko auf, die in Finanznetzwerken spezifisch aufgrund von Derivaten auftreten, und ich untersuche die Auswirkungen der jüngsten regulatorischen Maßnahmen.

Ich konzentriere mich zunächst auf eine Klasse von Derivaten namens *Credit Default Swaps* (CDSs), in denen der Verkäufer des Vertrages den Käufer gegen den Kreditausfall einer dritten Partei absichert, des *Referenzschuldners*. Ich zeige, dass solche CDSs, in denen der Referenzschuldner eine andere Bank ist, eine neue Art systemischen Risikos einführen, das ich *default ambiguity* (Ambiguität des Kreditausfalls) nenne. Default ambiguity bezeichnet eine Situation, in der es nach einem Schock (d.h. einem Verlust in den Anlagen der Banken) unmöglich ist, zu bestimmen, welche Banken infolgedessen ausfallen. Auf einer technischen Ebene zeige ich, dass es möglich ist, dass das *Clearing-Problem* keine Lösung oder mehrere inkompatible Lösungen hat. Im Gegensatz dazu existiert ohne CDSs immer eine eindeutige kanonische Lösung.

Ich lege dann dar, dass die erhöhte “Komplexität” durch CDSs sich auch als *Rechenkomplexität* offenbart. Im Detail zeige ich, dass das Clearing-Problem zu NP-vollständigen Entscheidungsproblemen sowie PPAD-vollständigen Approximationsproblemen führt, wenn CDSs erlaubt sind. Dies impliziert eine fundamentale Hürde zur computergestützten Analyse dieser Netzwerke, insbesondere im Kontext makroprudensieller *Stress Tests*. Ohne CDSs sind die Probleme entweder trivial oder in P. Ich untersuche die Auswirkungen verschiedener regulatorischer Maßnahmen. Mein Hauptresultat ist, dass die obengenannten Phänomene auf *ungedeckte CDS-Positionen* zurückgeführt werden können.

Zuletzt konzentriere ich mich auf eine bestimmte regulatorische Maßnahme, nämlich obligatorische *Portfoliokomprimierung*. Portfoliokomprimierung ist ein Mechanismus, durch den Kreise im Finanznetzwerk entfernt werden. Während dieser Vorgang immer die individuellen Risikopositionen reduziert, zeige ich, dass er die Auswirkungen von bestimmten Schocks unerwarteterweise verschlimmern kann. Die Anreize der Banken, eine Komprimierung durchzuführen, können außerdem dem sozialen Wohl entgegenstehen.

Ich gebe hinreichende Bedingungen an, unter denen diese Probleme ausgeschlossen sind. Meine Resultate in dieser Arbeit tragen zu einem besseren Verständnis systemischer Risiken und der Auswirkungen regulatorischer Maßnahmen bei.

Acknowledgements

I would like to thank my advisor Sven Seuken for guiding me through the past five years. Sven’s relentless curiosity, scientific enthusiasm, and pragmatism in the best sense of the word have accompanied me throughout my PhD. Our many discussions centering around some variant of “Why? And why is this important?” have truly shaped my way of thinking. I thank Sven for his constant support, dedication, and sheer availability and for teaching me how to communicate and present my work. Sven will continue to be a role model for me throughout my career.

I would like to thank Constantinos Daskalakis and Michael Wellman for serving as external reviewers for this thesis and for helping me improve the final version of this document through their comments. I am honored by having them on my dissertation committee.

This work would not have been possible without the fruitful collaboration with Stefano Battiston. By suggesting to explore the clearing problem with CDSs, Stefano has given me a jump start into my PhD. I am further grateful to Stefano for sharing his expertise at the intersection of finance and network theory with me.

I want to thank Frank Page for serving as an external reader for my PhD proposal and for his valuable comments in various discussions. For countless discussions that have shaped and refined my research I further thank Marco D’Errico, Helmut Elsinger, Gaston Gonnet, Ariah Klages-Mundt, Aviad Rubinstein, Juan Manuel Sánchez-Cartas, Joseph Stiglitz, Martin Summer, and Peter Widmayer. I would like to especially thank Marc Chesney for his helpful feedback on the first version of my model.

I thank the Deutsche Bundesbank for giving me the opportunity to do a research internship with them, where I could work with transaction-level derivatives data. In particular, I thank Puriya Abbassi for his advice during the internship. This empirical work has taught me a lot about the important intricacies of real financial markets.

I would like to thank my colleagues from the Economics & Computation Research Group at the University of Zurich for creating an inspiring environment and for their constructive feedback: Gianluca Brero, Ludwig Dierks, Stefania Ionescu, Dmitrii Moor, Nils Olberg, Mike Shann, and Jakob Weissteiner. Special thanks go to my colleagues Timo Mennle and Vitor Bosshard, whose continuous constructive comments have shaped me and my work. I thank the members of Stefano Battiston’s research group at the University of Zurich for many interesting discussions. I thank Pouyan Rezakhani and Wei Qiu for giving me the opportunity to advise their master’s theses.

I gratefully acknowledge financial support from the European Union’s Horizon 2020 program and the Swiss National Science Foundation.

I thank all members of my family for their love and support and I thank Pia Leimbacher for her friendship. I am grateful to be able to share this moment with my grandmothers Else Leurs and Hildegard Schuldenzucker. I thank my uncle Klaus-Wilfried Leurs, whose

caring attention has contributed to me starting a PhD in the first place. Most of all, I would like to thank my parents, Ansgar and Ulrike Schuldenzucker, for having given this life to me. My scientific curiosity, my aspiration to contribute to society, and my occasional stubbornness: I got it from you, both of you.

Finally, I would like to thank my wife Felicitas. Fee, your persistent support, your merciless cheerfulness, your caring devotion, and all the things we share and the attempt which to put into words would fill a thesis on its own carried me through my dark times during the last five years and made the light times so much more meaningful. I love you.

Dedications are weird.

I would dedicate this thesis to my parents or my wife, the people I love.
But is it something I did *for them*?
How would a thesis on financial networks benefit them?
And wouldn't it be pretentious to claim that I did this for anyone but myself?
Should I celebrate myself and sing myself?

No.

That is not what we mean by a dedication.
It is something we *offer* to them, like a sacrifice is offered on the altar.
It seems, then, that the only one I can dedicate my thesis to is god.
But who would ever do this?! And I'm lacking the god for it, too.

Instead, I remain silent,
knowing
that the one who wrote this thesis
is not me.

Contents

Abstract	v
Zusammenfassung	viii
Acknowledgements	x
1 Motivation and Overview of Results	1
1.1 Financial Derivatives, Networks, and Systemic Risk	1
1.2 Background, Problem Statement, and Research Questions	3
1.3 Publications Contained in This Thesis	9
1.4 Summary of Contributions	10
1.4.1 Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks	10
1.4.2 The Computational Complexity of Financial Networks with CDSs	11
1.4.3 Monotonic and Non-Monotonic Solution Concepts for Generalized Circuits	12
1.4.4 Portfolio Compression in Financial Networks: Incentives and Systemic Risk	13
1.5 Conclusion and Future Work	15
2 Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks	23
3 The Computational Complexity of Financial Networks with Credit Default Swaps	43
4 Monotonic and Non-Monotonic Solution Concepts for Generalized Circuits	93
5 Portfolio Compression in Financial Networks: Incentives and Systemic Risk	121
Curriculum Vitae	163

1 Motivation and Overview of Results¹

If we go down,
we all go down together.
If we go down, down, down,
We all go down together.

Krewella, *We go down*

1.1 Financial Derivatives, Networks, and Systemic Risk

Aristotle tells the story how philosopher Thales of Miletus made a fortune off the first known *financial derivative* around 600 BC. During one winter, Thales had used his understanding of the weather to predict a good olive harvest for the coming summer. Thales had then paid a small amount of money to reserve all the olive presses in his area for use in summer. When Thales eventually found himself correct and demand for olive presses spiked, he was able to rent them out at a high price.² However, not all stories of this kind have the same ending. In 1868, the time Thomas Mann sets his epos *Buddenbrooks*, derivatives are slowly gaining acceptance at the commodities exchanges. Thomas Buddenbrook, son of the family, enters into a risky contract in which he buys a farmer's whole harvest while it is still growing. Like Thales, Thomas had speculated on a good harvest; in this case, to his ruin: when a hailstorm destroys the whole crop, Thomas is left with nothing.³

When we witness the protagonists of Michael Lewis' true story *The Big Short* profit from the 2008 financial crisis,⁴ derivatives have already become an integral part of our financial system. A financial derivative is a financial contract between two parties in which the payment depends on the future value of another variable. In case of Thales and Thomas, this was the harvest later in the year. In the case of *The Big Short*, the derivative is called a *credit default swap (CDS)* and the variable was the default (i.e., bankruptcy) of another firm, the *reference entity*. When this entity is a mortgage fund or even a bank, the holder of a CDS profits off a financial crisis — provided the seller of the CDS does not default as well.

¹I liberally borrow from my own prior work (Schuldenzucker, Seuken and Battiston, 2019b,a, Schuldenzucker and Seuken, 2019a,b) for parts of Chapter 1.

²Aristotle, *Politics*, section 1259a. See also Crawford and Sen (1996, p. 7).

³Thomas Mann, *Buddenbrooks*, part 8

⁴Michael Lewis, *The Big Short: Inside the Doomsday Machine*. W. W. Norton & Company, 2010

What would have been banned as illegal betting in earlier times is today an indispensable tool by which financial institutions trade and reallocate risk. To see how this can be useful to society, consider a European company seeking a loan. US investors may be willing to supply the loan, but there may be disagreement about the currency the loan should be denominated in: if it is denominated in US dollars, the company is exposed to the risk that dollars become more expensive. If it is denominated in euros, investors are exposed to the risk that the euro devalues. An investment bank may be able to help here: it can bundle a euro-denominated loan with a derivative that will pay the dollar–euro exchange rate in the future.⁵ US investors could then buy the bundle without having to worry about exchange rates: any loss due to a devaluation of the euro would be offset by an equivalent gain in the derivative. Another trader, like a hedge fund, would take on the other side of the derivative and thus the exchange rate risk. The investment bank has thus just enabled an investment that otherwise would not have taken place, reducing funding costs for the company and contributing to economic growth.⁶

Of course, in our example, both the investment bank and the investor depend on the hedge fund to meet its obligation to pay. If the investor or the hedge fund enter into further derivatives with other parties, the process can continue over any number of stages, with each party buying, rebundling, and reselling risk. A network of obligations arises: a graph where the nodes are financial institutions (“banks” for short) and the edges are financial contracts. I call this the *financial network*.⁷

The financial network can serve a stabilizing function because losses at one institution can be spread across many different institutions and are thus more easily absorbed. However, the 2008 financial crisis has told us a different story: the financial network can also be a source of *systemic risk*, which endangers the financial system as a whole. Andrew Haldane (2009), then Executive Director of Financial Stability at the Bank of England, described the crisis as a manifestation of “the behaviour under stress of a complex, adaptive network,” in which “financial innovation [had] increased further network dimensionality, complexity, and uncertainty.” The financial network functioned as an “incendiary device” (Haldane, 2009), through which financial distress would travel to new institutions and markets. *Financial contagion* turned significant but local losses in the US housing market into the worst economic crisis since the Great Depression.

Financial over-the-counter (OTC)⁸ derivatives have played a major role during the

⁵This would most likely come in the form of a *cross-currency swap*, one of the most actively traded derivatives.

⁶See Mehrling (2010) and Mehrling et al. (2013) for more examples of how the *shadow banking sector* generates value through the use of derivatives.

⁷One important property of the financial network is that *cycles*, i.e., closed chains of obligations, are overwhelmingly common. See D’Errico et al. (2018), for example, for an empirical study of the network structure of credit default swap markets, which will be discussed below. One may assume that cycles are redundant structures that should be eliminated. Indeed, a post-trade mechanism called *portfolio compression* specifically aims to eliminate cycles. I will discuss portfolio compression in detail in Section 1.2 and in Chapter 5.

⁸OTC derivatives are derivatives that are traded directly with other financial institutions rather than

crisis.⁹ For example, Fender, Frankel and Gyntelberg (2008) described how the default of Lehman Brothers, which was both a major counterparty and reference entity in the CDS market, had significant repercussions in money markets (i.e., short-term loan markets). Further distress could only be averted by the government rescue of AIG, another major CDS trader. Both firms were among the most important institutions in the CDS market, both as counterparties and as reference entities (Fitch Ratings, 2007). It is hence almost certain that they were counterparties in a significant amount of CDSs where the respective other bank was the reference entity. Such a situation makes the consequences of a government intervention hard to foresee in a way I will make precise below.

After the crisis, financial regulators found themselves under great urgency to act against the “excessive systemic risk arising from the complexity and interconnectedness that characterize our financial system” (then Vice Chair of the Federal Reserve Janet Yellen, 2013). These reforms can be grouped into two directions: the first was tighter *regulation* of financial institutions through reforms such as Basel III, EMIR (in Europe), and the Dodd-Frank act (in the US). The second was improved monitoring and *stress testing*. In a stress test, a regulator such as the European Banking Association evaluates the stability of the financial system under an array of adverse economic scenarios. At the same time, researchers found a renewed interest in systemic risk in financial networks to evaluate and support these regulatory measures. The present thesis is part of this research effort.

The above accounts by policymakers attribute the financial crisis to excessive “complexity” of the financial network. The question remains, though, how exactly we should understand the informal term “complexity” here. In particular, it seems intuitive that derivatives lead to a “more complex” financial network, which should therefore be exposed to more systemic risk compared to a network without derivatives. The goal of this thesis is to capture this notion formally using tools from finance and computer science. Therefore, the present thesis is guided by the following overarching research question:

In what sense are financial networks with derivatives “more complex” than those without and what are the implications for systemic risk?

In the next section, I present some background and I operationalize my overarching research question into three specific research questions.

1.2 Background, Problem Statement, and Research Questions

Researchers have studied network-induced systemic risk since around 2000, where they have mostly focused on financial contagion. Researchers have studied two questions through an exchange. In this thesis, I only consider OTC derivatives.

⁹Regarding the role of OTC derivatives, specifically CDSs, in the 2008 crisis, see also: Financial Crisis Inquiry Commission (2011)

in particular: first, what is the impact of network topology on contagion compared to other factors such as correlation between banks' asset portfolios (Allen and Gale, 2000, Elsinger, Lehar and Summer, 2006, Gai, Haldane and Kapadia, 2011, Acemoglu et al., 2012, Glasserman and Young, 2015)? And second, how can an individual bank's contribution to network-induced systemic risk be measured (Battiston et al., 2012, Hu et al., 2012, Acemoglu, Ozdaglar and Tahbaz-Salehi, 2015, Demange, 2016)?

The overwhelming majority of these pieces of work makes the implicit assumption that the contracts in a financial network are all of the same kind: *debt contracts*, which encode a fixed obligation to pay a certain amount from one bank to another bank.¹⁰ This model has many advantages: debt networks can be represented as weighted graphs and this way we gain access to many standard tools from graph theory. The model is also simple to evaluate in a way that will be made precise below. Going back to my original discussion from Section 1.1, however, one should bear in mind that many of the contracts in the real financial network are actually derivatives, where the obligation to pay is not fixed, but depends on another variable. This raises the question if the debt model is appropriate in this case. In other words, are financial networks that contain derivatives merely "a bit more complicated" in the sense that they by and large exhibit the same phenomena as debt networks? Or are they truly *more complex* in the sense that they exhibit entirely new phenomena, and potentially new systemic risks, that are not visible if one assumes that all contracts are debt?

Attempts to capture the "complexity" of the financial network have previously been made using various measures from graph theory, such as path length, degree, or concentration measures (Shin, 2010, Gai, Haldane and Kapadia, 2011, Arinaminpathy, Kapadia and May, 2012, Battiston et al., 2016). As these measures require ordinary graphs as their inputs, where edges cannot contain more information than weights, they only apply to debt networks (unless some kind of transformation is applied first; such a transformation would always lose information).

One way to operationalize complexity of financial networks is by means of the *clearing problem*: we are given a financial network where some of the banks have been exposed to a shock, i.e., a loss on their assets. Each bank now makes payments to its creditors based on its own *external* assets and its *interbank* assets, i.e., the payments it receives from other banks. This implies a constraint reminiscent of a flow identity: banks with sufficient (total) assets to pay their liabilities in full must do so; the other banks are in default and must pay out all their assets to creditors in proportion to the respective liability. Banks may further incur *default costs* and lose a percentage of their assets upon

¹⁰Elliott, Golub and Jackson (2014) studied *cross-holdings*, encoded as the percentage of one institution owned by another institution. Kusnetsov and Veraart (2019) studied a model of debt contracts with multiple maturities. These variations of the standard debt contract all behave in a very similar way for my purposes in this thesis. I use "debt contract" as an umbrella term for any model of financial contracts where the liability from one bank to the other is a fixed number.

defaulting.¹¹

The clearing problem serves as a model for how a financial crisis will turn out following the initial shock. Note that it is based entirely on simple rules of accounting rather than, e.g., the order of defaults and payments. This makes the clearing problem robust to errors in the details of the contracts, such as maturities. Due to its simple structure, the clearing problem is analytically tractable, making it a useful tool for theoretical research.¹² Central banks' stress tests, currently in the process of moving to a macroprudential (i.e., system-wide) regime, will likely include a variant of the clearing problem in the future as well.¹³

One interpretation of the clearing problem is that in a financial crisis, a clearing authority (e.g., a central bank) observes the whole network of contracts, seeks to solve the clearing problem, and prescribes to each bank how much it has to pay to every other bank. Such a scenario is not a mere theoretical device. Indeed, then world bank chief economist Joseph Stiglitz described his attempt at a resolution of the 1997 East Asia financial crisis as “an extraordinarily complex general equilibrium problem that had not been solved” akin to the clearing problem. Not being able to solve the clearing problem resulted in a “paralysis” and a costly delay in restructuring (Stiglitz, 2016, at 0:51).

It would take another four years after the East Asia crisis until Eisenberg and Noe (2001) provided the first formalization of the clearing problem in debt networks and proved that it always has a solution. While there may be several solutions, there is always one that maximizes the equity (i.e., the money left for shareholders after clearing) of each individual bank. The equity-maximal solution is the obvious choice for the clearing authority to implement because it is preferred by each bank to every other solution. It can further be computed in polynomial time. Rogers and Veraart (2013) extended this result to a situation where there may be default costs. Importantly, these results only apply to debt networks.

When we think about the suitability of debt networks as a model for derivatives networks, we should distinguish two types of derivatives: if the obligation to pay only depends on variables that are external to the financial system, then these variables can be assumed to be fixed for the purpose of clearing, which essentially gives rise to a debt

¹¹While the exact rules of clearing vary across the literature (Eisenberg and Noe, 2001, Cifuentes, Ferrucci and Shin, 2005, Rogers and Veraart, 2013, Acemoglu, Ozdaglar and Tahbaz-Salehi, 2015), they all share as a common feature that payments (or, in some cases, contract valuations) happen simultaneously and an input-output identity needs to hold at each bank. The essential properties of these models are very similar. In most parts of this thesis, I consider the very influential clearing model by Eisenberg and Noe (2001) and its extension to default costs by Rogers and Veraart (2013).

¹²Most of the theoretical pieces of work cited at the beginning of this sub-section are based on some variant of the clearing problem.

¹³The European Central Bank's recent STAMP€ framework, which was developed based on “top-down models used to support EU-wide stress-testing exercises” (Constâncio, 2017), includes network effects as one of its central elements. Specifically, a variant of the clearing problem very close to Eisenberg and Noe (2001) is solved 20,000 times in the context of a Monte Carlo simulation to obtain a probability distribution of contagion losses (Dees, Henry and Martin, 2017, Chapter 12).

network. If, however, the obligation to pay depends on a variable that is fundamentally endogenous to the financial system, then such a derivative cannot be modeled as a debt contract. It is here where we should look first to find new effects due to derivatives.

I therefore study financial networks that contain credit default swaps (CDSs) written *on other banks* in addition to debt. Recall from Section 1.1 that a CDS is a financial derivative where the obligation to pay depends on the default of a third party, the *reference entity*. Market participants use CDSs to insure themselves against a default of the reference entity or to place a speculative bet on this event. The situation between Lehman Brothers and AIG mentioned in Section 1.1 was one where a large amount of CDSs was written on other banks.¹⁴

The little prior work that has studied financial networks with CDSs on other banks has not employed the clearing problem, but has rather resorted to models where either some of the effects of CDSs are ignored or CDSs are evaluated in some order (Heise and Kühn, 2012, Leduc, Poledna and Thurner, 2017, Banerjee and Feinstein, 2019). While such models always produce a solution, they also *rely* on their assumptions or the respective evaluation order for their result and therefore do not provide a definition of “the unique and well-defined outcome” of a crisis. The question now arises if this was a coincidence.

Research Question 1 Under which conditions can financial networks with debt and credit default swaps be cleared?

Recall from above that this is always the case in debt networks. Thus, if the answer to this question is not “always,” then I have captured a way in which CDSs make a financial network more “complex.” The possibility of a situation where the financial network cannot be cleared should also be considered a new systemic risk, as we have learned from Stiglitz’ account of the East Asia crisis.

Recall that in debt networks, not only does an (equity-maximizing) solution to the clearing problem always exist, but it can also be computed in polynomial time. This is important: if regulators only knew that a solution *exists*, but could not find it quickly enough, then this would be almost equivalent to a situation where no solution exists in the first place. Clearing algorithms for debt networks do not extend to CDSs on other banks,¹⁵ and therefore the question regarding computational complexity of clearing

¹⁴The market for CDSs on financial firms alone currently has a size of about USD 900 billion. In the years following the 2008 crisis, this number was as high as USD 5 trillion. See Bank for International Settlements (2018, Section Single-name instruments, Subsection Financial firms) and the graph linked there.

¹⁵I provide a discussion on this in Chapter 3. Note that brute-force approaches are impractical due to the size of the financial systems considered. For example, the 2014 European stress test considered 123 banks (European Banking Authority, 2014). In the ECB’s stress testing methodology in Dees, Henry and Martin (2017, Chapter 12), the authors consider 144 banks, only few of which are trivial, i.e., sources or sinks in the network (see Chart 12.1 in that paper). If one were to include all actors in the CDS market

in this setting has remained open. Besides the practical relevance of efficient clearing algorithms, I argue that computational complexity provides a useful measure for the “complexity” of financial networks. Researchers have previously used this approach in the context of derivatives. Arora et al. (2011) and Zuckerman (2011) studied the hardness of detecting rigged *collateralized debt obligations (CDOs)*. Braverman and Pasricha (2014) showed that pricing a class of complex financial derivatives called *compound options* is PSPACE-hard. These pieces of work considered individual financial products that are “complex” by themselves. In contrast, a single CDS is a rather simple derivative and only develops “complexity” in a network context.

Hemenway and Khanna (2016) studied computational complexity in debt networks and showed that it is computationally hard¹⁶ to, given a cross-holdership network and a “budget” for negative shocks, determine the distribution of this budget to the banks that does the worst damage in terms of defaults. In contrast, the clearing problem is concerned with determining the impact of a particular *known* distribution of shocks to banks, which is likely an easier problem for any type of network. My second research question concerns the computational complexity of the clearing problem in financial networks with CDSs on banks. Depending on the answer to research question 1, this question needs to be asked in two parts.

Research Question 2 What is the computational complexity of clearing financial networks with credit default swaps? Specifically, (1) what is the computational complexity of deciding whether a solution to the clearing problem exists, and (2) when a solution is guaranteed to exist, what is the computational complexity of computing an (approximate) solution?

I study both research questions 1 and 2 in general networks and under restrictions imposed by the regulatory changes after the financial crisis. This provides an opportunity to evaluate these policies, to understand if they help reduce “complexity” of the financial network in the context of derivatives. The following are policies that were put into place specifically for the regulation of OTC derivatives and that seem particularly relevant from a network perspective:

Naked CDSs on European sovereign states were banned by the European Union in 2012.

This means that a CDS on a European sovereign can only be bought if a sufficiently high (debt) exposure towards that sovereign is present as well. In this case, the CDS functions as insurance, offsetting any losses in the debt contract. In a *naked*

in a stress test, one would end up with almost 1000 institutions (Peltonen, Scheicher and Vuillemy, 2014). Note further that in (Dees, Henry and Martin, 2017, Chapter 12), the clearing problem is solved in the inner loop of a Monte Carlo simulation, making running time concerns even more important.

¹⁶Their result is relative to the Balanced Complete Bipartite Subgraph problem, which is the subject of various hardness conjectures.

CDSs, in contrast, the holder only holds a CDS and thus benefits if (other things equal) the reference entity is in financial distress. A ban on *all* naked CDSs has been part of the public debate following the 2008 crisis (Soros, 2009, Reuters, 2009), but was never implemented.

Central counterparty clearing is mandated for a significant part of the OTC derivative market. In its most extreme form, this means that all contracts have to be made via a central node, the *central clearing counterparty (CCP)*. A bank A would not write a contract to a bank B directly, but rather bank A would write a contract to the CCP and the CCP would write a contract to bank B. One of the desired effects is that the CCP would absorb a shock on the banks, prevent it from spreading through the network, and thus prevent financial contagion.¹⁷

Portfolio Compression is multilateral netting, i.e., the elimination of cycles in the network. Netting is performed for cycles of the same type of derivative, e.g. CDSs with the same reference entity. Compression is used in markets for OTC derivatives where insufficient standardization prohibits the use of a CCP. This includes single-name credit default swaps, i.e., the financial derivatives with which research questions 1 and 2 are concerned.¹⁸

Out of the three, portfolio compression stands out as particularly “complex.” For example, a non-trivial choice needs to be made regarding which cycles should be compressed. If cycles overlap, it may be the case that not *all* cycles can be compressed and trade-offs need to be made. The involved banks further need to agree for compression to be performed, giving rise to a potential incentive problem. Financial networks with derivatives may therefore be more “complex” than debt networks by virtue of a complex process being applied exclusively to them.

While the impact of central counterparty clearing on systemic risk has been studied intensely (e.g., Duffie and Zhu, 2011, Loon and Zhong, 2014, Duffie, Scheicher and Vuilleumey, 2015), very little is known about the analogous question for portfolio compression (the only prior piece of work dealing specifically with this question being Veraart (2019)). Given this little prior knowledge, it appears sensible to study portfolio compression in the simpler debt-only model first. To further simplify the analysis, one should consider shocks that are arbitrary but fixed, rather than a random distribution of shocks.

¹⁷Both the regulatory framework EMIR (in Europe) and Dodd-Frank (in the US) mandate the use of a CCP for certain types of derivatives (interest rate swaps and index CDSs), but not yet for the kind of CDSs I study in research questions 1 and 2 (single-name CDSs).

¹⁸EMIR regulations include an “obligation to have procedures to analyse the possibility to conduct the exercise” of portfolio compression when counterparties have more than 500 contracts with each other that are not centrally cleared (European Securities and Markets Authority, 2017). Portfolio compression can also be applied in the context of a CCP; however, I do not explicitly study this use case in this thesis.

Research Question 3 In debt-only financial networks, what are the banks' incentives to engage in portfolio compression and what are the effects of portfolio compression on systemic risk?

1.3 Publications Contained in This Thesis

This thesis consists of four papers, which together address the three research questions: paper 1 addresses research question 1, paper 2 and 3 address research question 2, and paper 4 addresses research question 3. I restate the research questions and provide a list of papers that address the respective research question. For working papers, the dates listed here refer to the most recent publicly available version. The chapters of this thesis include small changes to the exposition relative to these publicly available versions.

Research Question 1 Under which conditions can financial networks with debt and credit default swaps be cleared?

Publications

- “Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks.” Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. *Management Science*. Published in Articles in Advance, June 2019.

See also: “Clearing Payments in Financial Networks with Credit Default Swaps.” Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Extended abstract in *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, Maastricht, The Netherlands, July 2016.

Research Question 2 What is the computational complexity of clearing financial networks with credit default swaps? Specifically, (1) what is the computational complexity of deciding whether a solution to the clearing problem exists, and (2) when a solution is guaranteed to exist, what is the computational complexity of computing an (approximate) solution?

Publications

- “The Computational Complexity of Clearing Financial Networks with Credit Default Swaps.” Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Working Paper, May 2019.

See also: “Finding Clearing Payments in Financial Networks with Credit Default Swaps is PPAD-complete.” Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. In *Proceedings of the 8th Innovations in Theoretical Computer Science (ITCS) Conference*, Berkeley, USA, January 2017.

See also: “Clearing Payments in Financial Networks with Credit Default Swaps.” Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Extended abstract in *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, Maastricht, The Netherlands, July 2016.

- “Monotonic and Non-Monotonic Solution Concepts for Generalized Circuits.” Steffen Schuldenzucker and Sven Seuken. Working Paper, July 2019.

Research Question 3 In debt-only financial networks, what are the banks’ incentives to engage in portfolio compression and what are the effects of portfolio compression on systemic risk?

Publications

- “Portfolio Compression in Financial Networks: Incentives and Systemic Risk.” Steffen Schuldenzucker and Sven Seuken. Working Paper, November 2019.

1.4 Summary of Contributions

In the following, I provide a brief summary of all four papers and explain how they answer the three research questions.

1.4.1 Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks

In the first paper of my thesis (Chapter 2), I answer my first research question in the negative for general networks: in financial networks with CDSs, it may indeed be the case that it is not well-defined which banks are in default in terms of the clearing problem, a situation I call *default ambiguity*. This can happen in two different ways. *Non-existence* refers to a situation where the clearing problem has no solution. *Non-maximality* means that there is a solution, but no solution maximizes the equity of each bank simultaneously: banks disagree on which solution they prefer. Recall that both situations are impossible in debt networks.

If the clearing authority was facing a situation of non-existence in a crisis, a “paralysis” like in the East Asia crisis may ensue because it would not be clear how to proceed. In a situation of non-maximality, the clearing authority would have to choose among the different solutions, which would imply favoring the equity (and thus shareholders’ profits) of one bank over that of another one. This in turn might lead to major lobbying activities, as banks would have an incentive to influence the clearing authority to select a solution that is favorable to them. If default ambiguity came up during a stress test, it would lead to an inconclusive outcome. Default Ambiguity thus constitutes a fundamental

barrier to effective stress testing in financial networks with CDSs. These reasons are why I consider default ambiguity a new systemic risk that is specific to financial networks with derivatives like CDSs.

The intuition for my non-existence result is that with CDSs, a bank A can hold a *short position* on another bank B , i.e., A is better off if B is worse off. In a dense network of debt and CDS contracts, a bank may easily find itself *indirectly* holding a short position on itself, i.e., bank A is better off if bank A is worse off, which intuitively leads to a contradiction. The non-maximality case is similar. In contrast, in a debt-only network, banks only hold *long* positions on each other (if one bank is worse off, then the other is also worse off), so that this phenomenon does not exist.

To understand which kinds of networks are exposed to default ambiguity, I develop a new analysis framework. More in detail, I define the *colored dependency graph*, where the nodes are banks and colored edges indicate long and short positions. By restricting the cycles in the colored dependency graph, I receive sufficient conditions under which existence and/or maximality are restored. Specifically, I show that default ambiguity hinges on the presence of *naked CDSs*. Recall that naked CDSs are CDSs that are held without also holding a corresponding debt contract (or holding an insufficient amount of debt) so that the holder of a naked CDS benefits from financial distress at the reference entity. My results imply that a ban on all naked CDSs would eliminate default ambiguity.

In contrast, I find that the policy of using a central clearing counterparty (CCP) does *not* eliminate default ambiguity. This may be surprising since it may look like a CCP transforms the financial network into a trivial star network. However, a CCP in fact only protects against *counterparty risk* (i.e., the dependence of a bank on its debtors) while *fundamental risk* (i.e., the dependence of CDS holders and writers on the respective reference entity) still passes directly between the banks, essentially “around” the CCP. This is enough to lead to non-existence of a solution.

The results in this paper provide a first answer regarding the “complexity” of financial networks with derivatives: yes, in the case of CDSs on other banks, derivatives make a financial network more complex. Default ambiguity is a new systemic risk that does not exist in debt networks. Specifically, naked CDSs are to blame.

1.4.2 The Computational Complexity of Financial Networks with CDSs

In the next paper (Chapter 3), I study the “complexity” of financial networks with CDSs through the lens of computational complexity. Recall that in debt networks, the clearing problem can always be solved in polynomial time. At the same time, we know from the previous paper that with CDSs, the clearing problem may not even have a solution. Recall from research question 2 that this immediately raises two questions regarding the computational aspects of the clearing problem with CDSs:

1. Given a financial network, can we efficiently determine whether a solution to the

clearing problem exists?

2. Given a financial network in which a solution is known to exist, can we efficiently compute it?

In this paper, I answer both questions in the negative. Towards the first question, I show that deciding if a solution exists is NP-hard. An appropriate relaxation to ε -approximate solutions is NP-complete, for a sufficiently small constant ε .¹⁹ Towards the second question, I restrict my attention to the special case where banks do not incur default costs and where it is known that a solution always exists.²⁰ Here, the clearing problem gives rise to the total search problem of, given a financial network, finding an ε -approximate solution. I show that this problem is PPAD-complete for a sufficiently small constant ε . Thus, no polynomial-time approximation scheme (PTAS) exists unless P=PPAD.

I then attempt to isolate an “origin” of this hardness. In a first step, I show that already obtaining the most basic information about the solutions of the clearing problem, namely which banks default, is hard. More in detail, it is already NP-hard to decide if some given bank will default in some ε -solution (an appropriate relaxation being NP-complete) and in the case without default costs, it is already PPAD-complete to find a set of banks that will default in some ε -solution. These results suggest that the newfound “complexity” of the clearing problem is not an artifact of the problem formulation, but is fundamental to financial networks with CDSs.

In a second step towards the “origin” of the hardness, I study restrictions on the contract space. My results echo the findings from the previous paper. Computational complexity arises from fundamental risk, not counterparty risk, so CCPs do not help. Banning naked CDSs on the other hand *does* help: in this case, we receive a fully polynomial-time approximation scheme (FPTAS).

Computational complexity was likely not the regulators’ first concern during and after the 2008 crisis. However, I have demonstrated in this paper that it is useful measure of the fundamental “complexity” of financial networks and that it constitutes a new systemic risk in financial networks with CDSs. My results on the “origin” of the complexity provide a tool to guide regulatory policy.

1.4.3 Monotonic and Non-Monotonic Solution Concepts for Generalized Circuits

The third paper of this thesis (Chapter 4) provides the technical foundation for some of the proofs in the previous paper. The most important insight for the PPAD-hardness

¹⁹The (exact) solutions to the clearing problem with CDSs can be all irrational, so that finding an *exact* solution is not a well-defined computational problem and the exact decision problem is likely not in NP. This is another difference to debt-only networks, where solutions are always rational and of polynomial length.

²⁰I prove this in Chapter 2 of this thesis.

results in Chapter 3 is that financial networks with CDSs can perform calculations. There are financial networks where the recovery rate of one bank (i.e., the percentage of its liabilities it can pay) is the sum, difference, etc., of the recovery rates of two other banks. More in detail, my PPAD-hardness proofs in Chapter 3 are reductions from *generalized circuits* to financial networks with CDSs. Originally introduced for the study of the Nash equilibrium approximation problem (Daskalakis, Goldberg and Papadimitriou, 2009, Chen, Deng and Teng, 2009, Rubinstein, 2018), generalized circuits have been used to show PPAD-hardness of many other equilibrium approximation problems since.²¹

When I study the “origin” of the computational complexity in Chapter 3 and I show that it is already hard to compute a set of banks that default in some ε -solution, I perform reduction from a new discrete “support finding” variant of the generalized circuit problem. This proof requires a particularly close correspondence between the generalized circuit and the financial network I reduce it to. Such a correspondence was, however, hindered by a conceptual flaw in the generalized circuit concept itself, namely that the solution concept is not *monotonic*. By this I mean that if $\varepsilon < \varepsilon'$, then an ε -approximate solution for a certain generalized circuit is not necessarily also an ε' -approximate solution. This very unintuitive property, which had not been discussed before, creates subtle technical issues, including in prior work, that require intricate additional arguments to circumvent.

To overcome this problem of non-monotonicity, in my third paper (Chapter 4), I introduce two new computationally equivalent variants of the generalized circuit problem that are monotonic, serve as a drop-in replacement in prior work, eliminate the aforementioned issues in a natural way, and enable my above hardness proof. I hope that my results will enable new studies of sub-classes of generalized circuits as well as simpler and more natural reductions from generalized circuits to other equilibrium search problems in the future.

1.4.4 Portfolio Compression in Financial Networks: Incentives and Systemic Risk

The market for CDSs and other OTC derivatives is subject to several regulatory policies, some of which I have discussed in the previous papers. In the final paper in this thesis (Chapter 5), I focus on one specific policy: mandatory *portfolio compression*, i.e., the practice of removing cycles in the financial network. Note that portfolio compression is only applied in OTC derivatives networks (see Section 1.2).

Compression originated in the private sector and was only later endorsed by regulators. It proceeds in three steps: first, participating institutions submit the trades they would like to compress to a financial service provider. Second, the service provider combines the information submitted by all participants to construct the network and it calculates

²¹See, for instance: Babichenko, Papadimitriou and Rubinstein (2016), Chen, Paparas and Yannakakis (2017), Othman, Papadimitriou and Rubinstein (2016)

an *unwind proposal*, i.e., a collection of contracts to be removed. Third and crucially, all involved banks need to agree to the unwind proposal before the compression is implemented.²²

One might assume that compression universally reduces systemic risk because it reduces interconnectedness. However, in a pre-study for this paper I have shown that this is not the case: there are networks where compression *increases* losses following a particular shock (Schuldenzucker, Seuken and Battiston, 2018). This immediately raises two questions:

1. Under which conditions is a particular compression socially beneficial or even beneficial for each individual bank?
2. Under which conditions do involved banks have an incentive to agree to compression?
Is there a misalignment with the previous question?

With this paper, I am among the first to conduct a principled theoretical study of these questions and of the impact of portfolio compression on systemic risk in general.²³ As explained in research question 3, I study a particularly simple model: only debt contracts and an arbitrary but fixed shock.

I find that compression can be socially and individually detrimental, and it can even hurt the banks that participate in it. Furthermore, incentives to agree to compression may be misaligned with social welfare. I show that this effect depends on the parameters of the financial system and on the compression in a complex and non-monotonic way. This reveals another degree of “complexity” in OTC derivatives networks that arises from the business practices and regulatory policies imposed in these markets.

Based on my findings in the previous paragraph, it is a complex strategic decision for banks whether or not to agree to compression. In practice, however, banks generally seem to agree to unwind proposals without further deliberation. An explanation for this might be local information. I show that the incentives for banks to agree to compression depend on the presence of *feedback paths*, i.e., paths of liabilities that are not compressed and that lead from an involved bank to another involved bank. If banks do not take the possibility of feedback paths into account (and make an additional normality assumption), they would always consider it in their best interest to agree to compression.

I then present sufficient conditions under which compression is beneficial for all banks in a Pareto sense. This is the case if the recovery rates of involved banks are relatively high or when their balance sheets are sufficiently homogeneous. These effects depend on the default costs: if interbank payments are subject to lower default costs,

²²I present a simplified description of the process at one of the largest compression providers, TriOptima (see TriOptima, “triReduce Overview”, <https://www.trioptima.com/resources/>). I also copy their terminology. The basic process is the same for all compression providers.

²³To the best of my knowledge, the only piece of prior work dealing explicitly with this question is Veraart (2019). All of the results from that paper can be viewed as special cases of the results I obtain in Chapter 5 of the present thesis.

more homogeneity seems to be required, for example. This finding may be taken as an argument for compression in the pre-crisis financial system, where inefficient resolution processes kept default costs high while a convergence of investment and risk management strategies led to homogeneous risk profiles across the financial system (Plosser, 2009, Haldane, 2009); of course, further research is needed before any definite conclusions can be drawn.

The results from this paper reveal that, contrary to conventional wisdom, the effects of portfolio compressions are by no means straightforward. This gives rise to “complexity” in financial networks with derivatives. Both the context of the compression and incentives matter. The analytical tools developed in this paper enable a better understanding of these effects.

1.5 Conclusion and Future Work

Financial derivatives are fundamental to the functioning of today’s financial system. However, the 2008 financial crisis has shown to us that banks’ use of derivatives also generates *systemic risk*, which endangers the financial system as a whole, through the creation of the *financial network*. Policymakers have attributed the crisis to “excessive complexity” of the financial network, while the term “complexity” has remained informal.

In this thesis, I have shown that financial networks with derivatives are fundamentally “more complex” compared to those without. I have captured this complexity formally using tools from finance and theoretical computer science. I have shown that this “complexity” implies new systemic risks that are specific to financial networks with derivatives. More in detail, credit default swaps (CDSs), if they occur in a network, create the new systemic risk of *default ambiguity*, where it may no longer be well-defined which banks default following a shock. CDSs also increase the computational complexity of network clearing, which is a direct barrier to stress-testing while taking all network effects into account. I have shown that this complexity can be attributed to the presence of *naked* CDSs. The regulatory policies for OTC derivatives may also be a source of complexity, as exemplified by portfolio compression: rather than universally reducing systemic risk, the systemic effects of portfolio compression depend on various properties of the financial system, such as default costs and homogeneity.

Future Work

I see two promising, but also challenging, research threads for future work. The first is to study financial networks with derivatives from a perspective ex-ante to a random shock. While in this thesis, I have always considered arbitrary, but fixed shocks, regulators and market participants are often interested in a *valuation* of contracts under a random distribution of future shocks to banks. Under such an extension of the model, we may

expect a wealth of new phenomena. Randomness might “smooth out” default ambiguity such that ex-ante valuations are always well-defined. Incentives for portfolio compression may change as well if banks act from an ex-ante perspective. Some prior work has approached the problem of consistent ex-ante valuations in networks (Barucca et al., 2016, Veraart, 2018, Bertschinger and Stobbe, 2018), but the problem is still open, especially for derivatives.

The second thread is to consider the process of strategic formation of financial networks with derivatives. Ultimately, we as a society want to incentivize banks to use derivatives in such a way as to *reduce* network-induced systemic risk. While prior work has studied financial network formation with debt (Leitner, 2005, Farboodi, 2014, Acemoglu, Ozdaglar and Tahbaz-Salehi, 2014) and even restricted cases of CDSs (Zawadowski, 2013, Babus and Hu, 2017, Leduc, Poledna and Thurner, 2017), a general, analytically tractable model has remained elusive. A study of strategic network formation with derivatives would first have to answer questions regarding banks’ incentives to enter into them. Note that a derivative is only incentivized for both parties if there are differences in beliefs or differences in risk preferences and if there is some uncertainty. Thus, this thread will likely also include some aspect of the first research thread I have discussed above.

Bibliography

- Acemoglu, Daron, Asuman E. Ozdaglar, and Alireza Tahbaz-Salehi.** 2014. "Systemic Risk in Endogenous Financial Networks." Working Paper, <https://ssrn.com/abstract=2553900>.
- Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi.** 2015. "Systemic Risk and Stability in Financial Networks." *American Economic Review*, 105(2): 564–608.
- Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi.** 2012. "The network origins of aggregate fluctuations." *Econometrica*, 80(5): 1977–2016.
- Allen, Franklin, and Douglas Gale.** 2000. "Financial contagion." *Journal of Political Economy*, 108(1): 1–33.
- Arinaminpathy, Nimalan, Sujit Kapadia, and Robert M. May.** 2012. "Size and complexity in model financial systems." *Proceedings of the National Academy of Sciences*, 109(45): 18338–18343.
- Arora, Sanjeev, Boaz Barak, Markus Brunnermeier, and Rong Ge.** 2011. "Computational complexity and information asymmetry in financial products." *Communications of the ACM*, 54(5): 101–107.
- Babichenko, Yakov, Christos Papadimitriou, and Aviad Rubinstein.** 2016. "Can Almost Everybody Be Almost Happy?" In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*. Cambridge, Massachusetts, USA:ACM.
- Babus, Ana, and Tai-Wei Hu.** 2017. "Endogenous intermediation in over-the-counter markets." *Journal of Financial Economics*, 125(1): 200–215.
- Banerjee, Tathagata, and Zachary Feinstein.** 2019. "Impact of contingent payments on systemic risk in financial networks." *Mathematics and Financial Economics*, Published online.
- Bank for International Settlements.** 2018. "Semiannual OTC derivatives statistics, credit default swaps, by sector of reference, H1 2018." <http://stats.bis.org/statx/srs/table/d10.4?p=20181>, Retrieved November 23, 2019.
- Barucca, Paolo, Marco Bardoscia, Fabio Caccioli, Marco D'Errico, Gabriele Visentin, Stefano Battiston, and Guido Caldarelli.** 2016. "Network Valuation in Financial Systems." Working Paper, <https://ssrn.com/abstract=2795583>.
- Battiston, Stefano, Guido Caldarelli, Robert M May, Tarik Roukny, and Joseph E Stiglitz.** 2016. "The price of complexity in financial networks." *Proceedings of the National Academy of Sciences*, 113(36): 10031–10036.
- Battiston, Stefano, Michelangelo Puliga, Rahul Kaushik, Paolo Tasca, and Guido Caldarelli.** 2012. "DebtRank: Too central to fail? Financial networks, the FED and systemic risk." *Scientific reports*, 2.

- Bertschinger, Nils, and Julian Stobbe.** 2018. “Systemic Greeks: Measuring risk in financial networks.” Working Paper, <https://arxiv.org/abs/1810.11849>.
- Braverman, Mark, and Kanika Pasricha.** 2014. “The computational hardness of pricing compound options.” In *Proceedings of the 5th conference on Innovations in Theoretical Computer Science*.
- Chen, Xi, Dimitris Paparas, and Mihalis Yannakakis.** 2017. “The Complexity of Non-Monotone Markets.” *Journal of the ACM (JACM)*, 64(3): 20:1–20:56.
- Chen, Xi, Xiaotie Deng, and Shang-Hua Teng.** 2009. “Settling the complexity of computing two-player Nash equilibria.” *Journal of the ACM (JACM)*, 56(3): 14.
- Cifuentes, Rodrigo, Gianluigi Ferrucci, and Hyun Song Shin.** 2005. “Liquidity Risk and Contagion.” *Journal of the European Economic Association*, 3(2–3): 556–566.
- Constâncio, Vítor.** 2017. “Macroprudential stress tests: A new analytical tool.” *VoxEU.org*. Retrieved November 23, 2019, <http://voxeu.org/article/macroprudential-stress-tests-new-analytical-tool>.
- Crawford, George, and Bidyut Sen.** 1996. *Derivatives for decision makers: Strategic management issues*. Vol. 1, John Wiley & Sons.
- Daskalakis, Constantinos, Paul W Goldberg, and Christos H Papadimitriou.** 2009. “The complexity of computing a Nash equilibrium.” *SIAM Journal on Computing*, 39(1): 195–259.
- Dees, Stéphane, Jérôme Henry, and Reiner Martin**, ed. 2017. *STAMP€: Stress-Test Analytics for Macroprudential Purposes in the euro area*. Frankfurt:European Central Bank.
- Demange, Gabrielle.** 2016. “Contagion in Financial Networks: A Threat Index.” *Management Science*, 64(2): 955–970.
- D’Errico, Marco, Stefano Battiston, Tuomas Peltonen, and Martin Scheicher.** 2018. “How does risk flow in the credit default swap market?” *Journal of Financial Stability*, 35: 53–74.
- Duffie, Darrell, and Haoxiang Zhu.** 2011. “Does a central clearing counterparty reduce counterparty risk?” *Review of Asset Pricing Studies*, 1(1): 74–95.
- Duffie, Darrell, Martin Scheicher, and Guillaume Vuillemey.** 2015. “Central clearing and collateral demand.” *Journal of Financial Economics*, 116(2): 237–256.
- Eisenberg, Larry, and Thomas H Noe.** 2001. “Systemic risk in financial systems.” *Management Science*, 47(2): 236–249.
- Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson.** 2014. “Financial Networks and Contagion.” *American Economic Review*, 104(10): 3115–53.
- Elsinger, Helmut, Alfred Lehar, and Martin Summer.** 2006. “Risk assessment for banking systems.” *Management Science*, 52(9): 1301–1314.
- European Banking Authority.** 2014. “EBA publishes 2014 EU-wide stress test results.” *Press Release*, <http://www.eba.europa.eu/-/eba-publishes-2014-eu-wide-stress-test-results>.
- European Securities and Markets Authority.** 2017. “OTC Derivatives and Clearing Obligation.” Retrieved November 23, 2019, <https://www.esma.europa.eu/regulation/post-trading/otc-derivatives-and-clearing-obligation>.

- Farboodi, Maryam.** 2014. “Intermediation and Voluntary Exposure to Counterparty Risk.” Working Paper, <https://ssrn.com/abstract=2535900>.
- Fender, Ingo, Allen Frankel, and Jacob Gyntelberg.** 2008. “Three market implications of the Lehman bankruptcy.” *BIS Quarterly Review*, 1(1): 6–7.
- Financial Crisis Inquiry Commission.** 2011. “The Financial Crisis Inquiry Report.” <http://fcic.law.stanford.edu/report>.
- Fitch Ratings.** 2007. “CDx Survey — Market Volumes Continue Growing while New Concerns Emerge.”
- Gai, Prasanna, Andrew Haldane, and Sujit Kapadia.** 2011. “Complexity, concentration and contagion.” *Journal of Monetary Economics*, 58(5): 453–470.
- Glasserman, Paul, and H. Peyton Young.** 2015. “How likely is contagion in financial networks?” *Journal of Banking & Finance*, 50(0): 383–399.
- Haldane, Andrew G.** 2009. “Rethinking the Financial Network.” Speech at the Financial Student Association, Amsterdam, April 28, 2009, <https://www.bankofengland.co.uk/speech/2009/rethinking-the-financial-network>.
- Heise, S., and R. Kühn.** 2012. “Derivatives and credit contagion in interconnected networks.” *The European Physical Journal B*, 85(4): 1–19.
- Hemenway, Brett, and Sanjeev Khanna.** 2016. “Sensitivity and computational complexity in financial networks.” *Algorithmic Finance*, 5(3–4).
- Hu, Daning, J Leon Zhao, Zhimin Hua, and Michael CS Wong.** 2012. “Network-based modeling and analysis of systemic risk in banking systems.” *MIS Quarterly*, 36(4): 1269–1291.
- Kusnetsov, Michael, and Luitgard Anna Maria Veraart.** 2019. “Interbank clearing in financial networks with multiple maturities.” *SIAM Journal on Financial Mathematics*, 10(1): 37–67.
- Leduc, Matt V, Sebastian Poledna, and Stefan Thurner.** 2017. “Systemic Risk Management in Financial Networks with Credit Default Swaps.” *Journal of Network Theory in Finance*, 3(3).
- Leitner, Yaron.** 2005. “Financial Networks: Contagion, Commitment, and Private Sector Bailouts.” *The Journal of Finance*, 60(6): 2925–2953.
- Loon, Yee Cheng, and Zhaodong Ken Zhong.** 2014. “The impact of central clearing on counterparty risk, liquidity, and trading: Evidence from the credit default swap market.” *Journal of Financial Economics*, 112(1): 91–115.
- Mehrling, Perry.** 2010. “Credit Default Swaps: The Key to Financial Reform.” In *Time for a Visible Hand.* , ed. Stephany Griffith-Jones, José Antonio Ocampo and Joseph E. Stiglitz. Oxford, UK:Oxford University Press.
- Mehrling, Perry, Zoltan Pozsar, James Sweeney, and Daniel H Neilson.** 2013. “Bagehot was a shadow banker: shadow banking, central banking, and the future of global finance.” Working Paper, <https://ssrn.com/abstract=2232016>.
- Othman, Abraham, Christos Papadimitriou, and Aviad Rubinstein.** 2016. “The complexity of fairness through equilibrium.” *ACM Transactions on Economics and Computation (TEAC)*, 4(4): 20.

- Peltonen, Tuomas A, Martin Scheicher, and Guillaume Vuillemey.** 2014. “The network structure of the CDS market and its determinants.” *Journal of Financial Stability*, 13: 118–133.
- Plosser, Charles I.** 2009. “Redesigning Financial System Regulation.” Speech at the New York University Conference “Restoring Financial Stability: How to Repair a Failed System”, March 6, 2009, <https://www.bankofengland.co.uk/speech/2009/rethinking-the-financial-network>.
- Reuters.** 2009. “Geithner testifies on financial regulation reform.” *Reuters*. News Item, Business News, 26 March, 2009, <https://uk.reuters.com/article/uk-financial-regulation-highlights-idUKTRE52P4UX20090326>.
- Rogers, LCG, and Luitgard AM Veraart.** 2013. “Failure and rescue in an interbank network.” *Management Science*, 59(4): 882–898.
- Rubinstein, Aviad.** 2018. “Inapproximability of Nash equilibrium.” *SIAM Journal on Computing*, 47(3): 917–959.
- Schuldenzucker, Steffen, and Sven Seuken.** 2019a. “Monotonicity and Non-Monotonicity of Solution Concepts in Generalized Circuits.” Working Paper, <https://arxiv.org/abs/1907.12854>.
- Schuldenzucker, Steffen, and Sven Seuken.** 2019b. “Portfolio Compression in Financial Networks: Incentives and Systemic Risk.” Working Paper, <https://ssrn.com/abstract=3483919>.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2018. “Portfolio Compression: Positive and Negative Effects on Systemic Risk.” Research Note, <https://ssrn.com/abstract=3135960>.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2019a. “The Computational Complexity of Financial Networks with Credit Default Swaps.” Working Paper, <https://arxiv.org/abs/1710.01578>.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2019b. “Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks.” *Management Science*, Published in Articles in Advance.
- Shin, Hyun Song.** 2010. “Financial Intermediation and the Post-Crisis Financial System.” *BIS Working Paper*, 304.
- Soros, George.** 2009. “One Way to Stop Bear Raids.” *The Wall Street Journal*. Op-ed, https://www.georgesoros.com/2009/03/23/one_way_to_stop_bear_raids/.
- Stiglitz, Joseph E.** 2016. “Financial Complexity and Climate Change.” Public Lecture at the University of Zurich on January 18, 2016, <https://www.youtube.com/watch?v=J36werCxpa>.
- Veraart, Luitgard AM.** 2018. “Distress and default contagion in financial networks.” Working Paper, <https://ssrn.com/abstract=3029702>.
- Veraart, Luitgard AM.** 2019. “When Does Portfolio Compression Reduce Systemic Risk?” Working Paper, <https://ssrn.com/abstract=3311176>.
- Yellen, Janet L.** 2013. “Interconnectedness and Systemic Risk: Lessons from the Financial Crisis and Policy Implications.” Speech at the American Economic Association / American

- Finance Association Joint Luncheon, San Diego, California, January 4, 2013, <https://www.federalreserve.gov/newsevents/speech/yellen20130104a.htm>.
- Zawadowski, Adam.** 2013. “Entangled financial systems.” *The Review of Financial Studies*, 26(5): 1291–1323.
- Zuckerman, David.** 2011. “Pseudorandom Financial Derivatives.” In *Proceedings of the 12th ACM Conference on Electronic Commerce*. New York, NY, USA:ACM.

2 Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks

This is ten percent luck,
Twenty percent skill,
Fifteen percent concentrated power of will,
Five percent pleasure,
Fifty percent pain,
And a hundred percent reason to remember the name.

Fort Minor, *Remember The Name*

The content of this chapter has previously appeared in:

“Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks.” Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. *Management Science*. Published in Articles in Advance, June 2019.

See also: “Clearing Payments in Financial Networks with Credit Default Swaps.” Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. Extended abstract in *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, Maastricht, The Netherlands, July 2016.

Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks

Steffen Schuldenzucker,^a Sven Seuken,^a Stefano Battiston^{b,c}

^a Department of Informatics, University of Zurich, 8050 Zurich, Switzerland; ^b Department of Banking and Finance, University of Zurich, 8050 Zurich, Switzerland; ^c Swiss Finance Institute, 8006 Zurich, Switzerland

Contact: schuldenzucker@ifi.uzh.ch,  <http://orcid.org/0000-0003-2161-9948> (StS); seuken@ifi.uzh.ch,  <http://orcid.org/0000-0001-8525-8120> (SvS); stefano.battiston@uzh.ch,  <http://orcid.org/0000-0002-0051-973X> (SB)

Received: September 12, 2017

Revised: September 6, 2018;
November 25, 2018

Accepted: December 21, 2018

Published Online in Articles in Advance:
June 26, 2019

<https://doi.org/10.1287/mnsc.2019.3304>

Copyright: © 2019 The Author(s)

Abstract. We study financial networks and reveal a new kind of systemic risk arising from what we call *default ambiguity*—that is, a situation where it is impossible to decide which banks are in default. Specifically, we study the *clearing problem*: given a network of banks interconnected by financial contracts, determine which banks are in default and what percentage of their liabilities they can pay. Prior work has shown that when banks can only enter into debt contracts with each other, this problem always has a unique maximal solution. We first prove that when banks can also enter into *credit default swaps* (CDSs), the clearing problem may have no solution or multiple conflicting solutions, thus leading to default ambiguity. We then derive sufficient conditions on the network structure to eliminate these issues. Finally, we discuss policy implications for the CDS market.

History: Accepted by Gustavo Manso, finance.

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution-NoDerivatives 4.0 International License. You are free to download this work and share with others commercially or noncommercially, but cannot change in any way, and you must attribute this work as "Management Science. Copyright © 2019 The Author(s). <https://doi.org/10.1287/mnsc.2019.3304>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by-nd/4.0/>."

Funding: All authors gratefully acknowledge financial support from the European Union's FP7 and Horizon 2020 research and innovation programme [Future and Emerging Technologies Grant Agreements 610704 (SIMPOL) and 640772 (DOLFINS)]. Additionally, S. Battiston acknowledges funding from the Swiss National Fund Professorship [Grant PP00P1-144689] and from the Institute of New Economic Thinking through the Task Force in Macroeconomic Efficiency and Stability.

Supplemental Material: The electronic companion is available at <https://doi.org/10.1287/mnsc.2019.3304>.

Keywords: financial networks • credit default swaps • systemic risk • clearing systems

1. Introduction

During the 1996 East Asia crisis, “[i]n Indonesia, . . . 75 percent of all businesses were put into distress, while in Thailand close to 50 percent of bank loans became nonperforming” (Stiglitz 2002, p. 112). All of these firms were interconnected, and as a result of the complexity of this network, regulators were facing a phenomenon we call *default ambiguity*. As then World Bank Chief Economist Joseph Stiglitz describes it,

Every firm owed money to every other firm. But . . . you couldn't tell whether they were bankrupt or not, because that depended on whether they got paid money that was owed to them by other firms who might or might not be in default, depending on whether the firms that owed them money went bankrupt. (Stiglitz 2016, at 0:51h)

In other words, *default ambiguity* is a situation where one cannot tell which banks are in default. Stiglitz (2016) points out that this led to a paralysis (“it took years to resolve it”), resulting in large welfare losses because banks’ resolution could not be carried out quickly.

It may be intuitive to expect that default ambiguity can arise when the financial authority only has imperfect information about banks’ contractual obligations. For instance, Haldane (2009) described a related effect on asset prices in the 2008 financial crisis. In this paper, we show that, remarkably, default ambiguity can also arise in a *perfect information setting*, where the whole financial network is known to the financial authority.

In the perfect information setting, default ambiguity can be studied in terms of the *clearing problem*: given a network of banks (or other financial institutions) interconnected by financial contracts, determine which banks are in default and for the defaulting banks what percentage of their liabilities they can still pay to their creditors (i.e., we are looking for the *recovery rate* of each bank). As in Eisenberg and Noe (2001), we assume that all payments are made simultaneously and in accordance with standard bankruptcy regulations. The banks’ assets may lose part of their value when banks default (i.e., the banks incur *default costs*).¹

An interpretation of the clearing problem is that in a financial crisis, a *clearing authority* (e.g., a central bank) observes the whole network of contracts, seeks to solve the clearing problem, and prescribes to each bank how much it has to pay to every other bank. The clearing problem is challenging because banks typically rely on payments they receive from other banks to meet their obligations, and banks can form an intricate web of contractual relations with each other. Default ambiguity arises when the clearing problem has no solution or when there are multiple conflicting solutions (i.e., none of which is simultaneously best for all banks).

Eisenberg and Noe (2001) and Rogers and Veraart (2013) showed that financial networks where banks can only enter into simple *debt contracts* (i.e., loans from one bank to another) have two very desirable properties from a clearing perspective: First, the clearing problem always has a solution (we call this property *existence*). Second, there is always a solution that maximizes the equity of each bank simultaneously (we call this property *maximality*).² Thus, although there may be multiple solutions, the maximal solution is the obvious choice for the clearing authority to implement (because it is simultaneously best for all banks).³

In this work, we study financial networks that contain debt contracts as well as *credit default swaps* (CDSs). A CDS is a financial derivative in which the writer insures the holder of the contract against the default of a third party, the *reference entity*. The holder may or may not have an exposure to the reference entity. Prior work has shown that the network structure of CDSs has a significant effect on systemic risk (Duffie and Zhu 2011, Loon and Zhong 2014). A large part of the CDS market is made up of CDSs where the reference entity is itself a financial institution.⁴ An analysis of CDS transaction data by D'Errico et al. (2018) has shown that the financial institutions (including reference entities) in the CDS market are tightly connected, implying the presence of circular relationships involving holders, writers, and reference entities.

We ask, under which conditions can financial systems still be cleared when they contain such CDSs in addition to debt? We take existence and maximality as desiderata for the design of a financial system. We then derive constraints on the network structure under which the financial system is guaranteed to satisfy these two desiderata, independent of banks' external assets (i.e., assets that do not depend on other banks).⁵ Similar to prior work on the clearing problem, our approach is agnostic to how networks have formed. Thus, our results apply to *any* network, including those that could arise in equilibrium from a decentralized process of network formation.

In this work, we are the first to present an analytically tractable model for the clearing problem with CDSs on financial institutions (Section 2).⁶ Our first major finding is that, in networks with debt and CDSs, default ambiguity can occur (Section 3). We first show that if there are default costs, then existence is not always satisfied.⁷ The intuition for this is that with CDSs, a bank A can hold a position on another bank B whereby A is better off if B is worse off, referred to as a *short position* hereafter. In a dense network of debt and CDS contracts, a bank may easily find itself *indirectly* holding a short position on itself (i.e., bank A is better off if bank A is worse off), which intuitively leads to a contradiction.⁸ By contrast, in a debt-only network, banks only hold *long positions* on each other (if one bank is worse off, then the other is also worse off) so that this phenomenon does not exist. If the clearing authority was facing a situation where no solution exists in a crisis, a "paralysis" such as in the East Asia crisis may ensue because it would not be clear how to proceed. One might wonder which changes to our model might restore existence in the general case. We provide a discussion on this at the end of Section 3.1.

Second, we show that even in situations where existence is satisfied, maximality may not be satisfied. This resolves an open question by Demange (2016), who conjectured that if one extended the Eisenberg/Noe model to CDSs, "multiple and noncomparable ratios might then clear the market" (p. 967). The intuition for our result is that CDSs can give rise to a situation in which two banks happen to hold a short position on each other. In this case, exactly one of the two banks can be well off while the other one is doing poorly, but it is not possible to maximize both equities at the same time. Again, because networks of debt obligations contain only long positions, this effect can only be observed in networks with CDSs, with or without default costs. In a situation where no maximal solution exists, the clearing authority would have to choose among the different solutions, which would imply favoring the equity (and thus shareholders' profits) of one bank over that of another one. This, in turn, might lead to major lobbying activities, as banks would have an incentive to influence the clearing authority to select a solution that is favorable to them.

Note that solving the clearing problem is not only relevant in a financial crisis. Regulators such as the European Central Bank regularly conduct *stress tests* to evaluate how likely certain banks are to default given adverse economic scenarios. As regulators progressively take on a macroprudential (i.e., systemwide) perspective, stress tests increasingly take network effects into account.⁹ In the future, it seems prudent to also include CDSs in network-based stress tests, given the important role they played in the 2008 financial crisis.¹⁰ Our work shows that the inclusion of CDSs

may lead to an inconclusive outcome of a stress test due to default ambiguity. Another real-world application that illustrates the importance of our findings is the recent provision to resolve a failing bank within one weekend (Single Resolution Board 2016). If default ambiguity arose in this application, this would hinder the quick resolution of the bank.

To eliminate these issues regarding default ambiguity, we next study what constraints on the network structure are sufficient to guarantee existence and maximality. To this end, we first introduce the *colored dependency graph*, a new analysis framework to capture the dependencies among banks, in particular among the three parties (holder, writer, and reference entity) involved in a CDS (Section 4). By restricting the cycles in this dependency graph, we then derive sufficient conditions under which existence and/or maximality are satisfied (Section 5). Furthermore, we provide an algorithm to compute a solution in this case. The conditions we derive provide *ex ante* guarantees; that is, they are robust to any possible future shock on the banks' external assets. Ex ante guarantees are important for practical applications because the mere possibility that the market could not be cleared in the future could undermine trust of market participants and bring about a liquidity crisis today. Furthermore, if a bank anticipated a future incentive to influence the clearing authority's choice of a solution, then the bank would have a motivation to already start lobbying today.

We last discuss potential policy implications. We show within our model that the policy of routing all contracts through a central counterparty does not guarantee existence. By contrast, when "naked" CDSs (i.e., CDSs that are held without also holding a corresponding debt contract) are not allowed, then existence and maximality are always fulfilled. Our results thus contribute to the debate on a possible regulation of the CDS market (Section 6).

Prior work on financial networks has primarily focused on *financial contagion* (i.e., how local shocks to market participants' portfolios spread through the network and cause systemic crises). Researchers have considered two questions in particular: First, what is the impact of network topology on contagion compared with other factors such as correlation between banks' asset portfolios (Allen and Gale 2000, Elsinger et al. 2006, Acemoglu et al. 2012, Glasserman and Young 2015)? And second, how can the likelihood of an individual bank to trigger contagion be measured (Hu et al. 2012, Acemoglu et al. 2015, Battiston et al. 2016, Demange 2016)? Bardoscia et al. (2017) have shown how specific closed chains in networks of credit contracts are a sufficient condition for instability. This prior work has shown that financial contagion can amplify the effect of a small shock

leading to a large loss. By contrast, default ambiguity describes a situation in which the effect of a shock on a financial network is not even mathematically well defined. This means that neither the interbank payments nor the systemwide losses can be determined. In this sense, the risk of a financial system to experience default ambiguity is more fundamental than the risk of financial contagion. Our dependency analysis framework constitutes a new tool to study this risk and inform regulatory policy.

2. Formal Model and Visual Representation

Our model is based on the model by Eisenberg and Noe (2001) and its extension to default costs by Rogers and Veraart (2013). Both of these prior models were restricted to debt contracts. We define an extension to credit default swaps. Following said prior work, we assume a static model where a financial system is given exogenously and all contracts are evaluated simultaneously. We adjust the notation where necessary.

2.1. The Model

Banks and External Assets. Let N denote a finite set of banks. Each bank $i \in N$ holds a certain amount of *external assets*, denoted by $e_i \geq 0$. Let $e = (e_i)_{i \in N}$ denote the vector of all external assets.

Contracts. There are two types of contracts: *debt contracts* and *CDSs*. Every contract gives rise to a conditional or unconditional obligation to pay a certain amount, called a *liability*, from its *writer* to its *holder*. Banks that cannot fulfill this obligation are said to be *in default*. The *recovery rate* r_i of a bank i is the share of its liabilities it is able to pay. Thus, $r_i = 1$ if i is not in default and $r_i < 1$ if i is in default. Let $r = (r_i)_{i \in N}$ denote the vector of all recovery rates.

A *debt contract* obliges the writer i to unconditionally pay a certain amount to the holder j . The amount is called the *notional* of the contract and is denoted by $c_{i,j}^0$. A *credit default swap* obliges the writer i to make a conditional payment to the holder j . The amount of this payment depends on the recovery rate of a third bank k , called the *reference entity*. Specifically, the payment amount of the CDS from i to j with reference entity k and notional $c_{i,j}^k$ is $c_{i,j}^k \cdot (1 - r_k)$. The contractual relationships between all banks are represented by a three-dimensional matrix $c = (c_{i,j}^k)_{i \in N, j \in N, k \in N \cup \{0\}}$. Zero entries indicate the absence of the respective contract.

Note that when banks enter contracts, there typically is an initial payment. For example, debt contracts arise because the holder lends an amount of money to the writer, and the holder of a CDS pays a premium to obtain the CDS. In our model, we assume that any such initial payments have been made at an

earlier time and are implicitly reflected in the external assets.

We make two *sanity assumptions* to rule out pathological cases. First, we require that no bank enters into a contract with itself or on itself (i.e., $c_{i,i}^{\emptyset} = c_{i,i}^j = c_{i,j}^j = c_{i,j}^i = 0$ for all $i, j \in N$). Second, as CDSs are defined as insurance on debt, we require that any bank that is a reference entity in a CDS must also be the writer of a debt contract (i.e., if $\sum_{k,l \in N} c_{k,l}^i > 0$, then $\sum_{j \in N} c_{i,j}^{\emptyset} > 0$ for all $i \in N$).

For any bank i , the *creditors of i* are those banks that are holders of contracts for which i is the writer (i.e., the banks to which i owes money). Conversely, the *debtors of i* are the writers of contracts of which i is the holder (i.e., the banks that owe money to i). Note that the two sets can overlap: for example, a bank could hold a CDS on one reference entity while writing a CDS on another reference entity, both with the same counterparty.

Default Costs. We model default costs following Rogers and Veraart (2013): there are two *default cost parameters* $\alpha, \beta \in [0, 1]$. Defaulting banks are only able to pay to their creditors a share of α of their external assets and a share of β of their incoming payments. Thus, $\alpha = \beta = 1$ means that there are no default costs, and $\alpha = \beta = 0$ means that assets held by defaulting banks are worthless. The values $1 - \alpha$ and $1 - \beta$ are the *default costs*.¹¹

Financial System. A *financial system* is a tuple (N, e, c, α, β) where N is a set of banks, e is a vector of external assets, c is a three-dimensional matrix of contracts, and α and β are default cost parameters.

Liabilities, Payments, and Assets. For two banks i, j and a vector of recovery rates r , the *liability of i to j at r* is the amount of money that i has to pay to j if recovery rates in the financial system are given by r , denoted by $l_{i,j}(r)$. It arises from the aggregate of any debt contract and all CDSs from i to j :

$$l_{i,j}(r) := c_{i,j}^{\emptyset} + \sum_{k \in N} (1 - r_k) \cdot c_{i,j}^k.$$

The *total liabilities of i at r* are the aggregate liabilities that i has toward other banks given the recovery rates r , denoted by $l_i(r)$:

$$l_i(r) := \sum_{j \in N} l_{i,j}(r).$$

The actual *payment* $p_{i,j}(r)$ from i to j at r can be lower than $l_{i,j}(r)$ if i is in default. By the *principle of proportionality* (discussed below), a bank that is in default makes payments for its contracts in proportion to the respective liability:

$$p_{i,j}(r) := r_i \cdot l_{i,j}(r).$$

The *total assets* $a_i(r)$ of a bank i at r consist of its external assets e_i and the incoming payments:

$$a_i(r) := e_i + \sum_{j \in N} p_{j,i}(r).$$

In case bank i is in default, its *assets after default costs* $a'_i(r)$ are the assets reduced according to the factors α and β . This is the amount that will be paid out to creditors:

$$a'_i(r) := \alpha e_i + \beta \sum_{j \in N} p_{j,i}(r).$$

Clearing Recovery Rate Vector. Following Eisenberg and Noe (2001), we call a recovery rate vector r *clearing* if it is in accordance with the following three principles of bankruptcy law:

1. *Absolute priority*: Banks with sufficient assets pay their liabilities in full. Thus, these banks have recovery rate 1.

2. *Limited liability*: Banks with insufficient assets to pay their liabilities are in default and pay all of their assets to creditors after default costs have been subtracted. Thus, these banks have recovery rate $a'_i(r)/l_i(r) < 1$.

3. *Proportionality*: In case of default, payments to creditors are made in proportion to the respective liability.

The principle of proportionality is automatically fulfilled in our model by the definition of the payments $p_{i,j}(r)$. The other two principles lead to the following definition.

Definition 1 (Clearing Recovery Rate Vector). Let $X = (N, e, c, \alpha, \beta)$ be a financial system. A *recovery rate vector* is a vector of values $r_i \in [0, 1]$ for each $i \in N$. We denote by $[0, 1]^N$ the space of all possible recovery rate vectors. Define the *update function*

$$F : [0, 1]^N \rightarrow [0, 1]^N,$$

$$F_i(r) := \begin{cases} 1 & \text{if } a_i(r) \geq l_i(r) \\ \frac{a'_i(r)}{l_i(r)}, & \text{if } a_i(r) < l_i(r). \end{cases} \quad (1)$$

A recovery rate vector r is called *clearing* for X if it is a fixed point of the update function (i.e., if $F_i(r) = r_i$ for all i). We also call a clearing recovery rate vector a *solution* to the clearing problem.

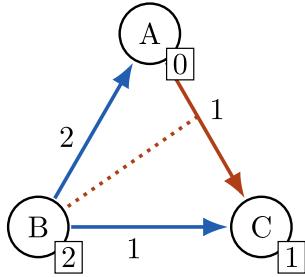
Equity. For any bank i , its *equity* $E_i(r)$ is the positive difference between assets and liabilities. This is the profit that the owners of bank i get to keep after clearing:

$$E_i(r) := \max(0, a_i(r) - l_i(r)).$$

2.2. Example and Visual Representation

Figure 1 shows a visual representation of an example financial system. There are three banks $N = \{A, B, C\}$,

Figure 1. (Color online) Example Financial System with $\alpha = \beta = 0.5$



drawn as circles, with external assets of $e_A = 0$, $e_B = 2$, and $e_C = 1$, drawn as rectangles on top of the banks. Debt contracts are drawn as blue arrows from the writer to the holder, and they are annotated with the notionals $c_{B,A}^0 = 2$ and $c_{B,C}^0 = 1$. CDSs are drawn as orange arrows with a dashed line connecting to the reference entity, and they are also annotated with the notionals: $c_{A,C}^B = 1$. Default cost parameters $\alpha = \beta = 0.5$ are given in addition to the picture. A solution for this example is $r_A = 1$, $r_B = 1/3$, and $r_C = 1$. The liabilities at this recovery rate vector are $l_{B,A}(r) = 2$, $l_{B,C}(r) = 1$, and $l_{A,C}(r) = 2/3$. Payments are $p_{B,A}(r) = 2/3$, $p_{B,C}(r) = 1/3$, and $p_{A,C}(r) = 2/3$, and equities are $E_A(r) = 0$, $E_B(r) = 0$, and $E_C(r) = 1$. This is the only solution.

2.3. Discussion of Our Formal Model

Note that our addition of CDSs to the Rogers and Veraart (2013) model substantially changes its mathematical properties. The liabilities $l_{ij}(r)$ now depend on the recovery rate vector r , and the assets $a_i(r)$ contain terms of the form $c_{j,i}^k \cdot r_j \cdot (1 - r_k)$. Thus, the update function $F_i(r)$ depends on r in a way that is both *nonlinear* and *nonmonotonic*: an increase in some recovery rate r_l could lead to a higher or lower value of $F_i(r)$ for another bank i . Because $F_i(r)$ is nonmonotonic, we cannot in general find a solution to the clearing problem by simply iterating the function F . The iteration sequence may cycle among different recovery rate vectors without even getting near a solution (see Appendix A for an example). For the same reason, Eisenberg and Noe's (2001) algorithm for computing clearing payments in debt-only systems cannot be applied to systems with CDSs.

Prior work has modeled financial networks almost exclusively as weighted binary graphs where edges reflect binary long relationships such as debt (Eisenberg and Noe 2001, Cifuentes et al. 2005, Rogers and Veraart 2013) and cross-ownership (Vitali et al. 2011, Elliott et al. 2014). Barucca et al. (2016) presented a unified framework for such models. However, CDSs give rise to *ternary* relationships because the holder is affected by the financial health of both the writer and the reference entity, and they imply both long and

short positions. Weighted-graph models cannot accurately represent these features, whereas our model captures them well.¹²

3. Existence and Maximality in General Financial Systems

In this section, we explore the possible shapes of the set of solutions for a financial system with debt and CDSs. We construct financial systems that have no solution or multiple conflicting solutions. Consequently, neither existence nor maximality is guaranteed in general.

At the heart of our constructions lies the following lemma, which may be of independent interest to some readers. The lemma demonstrates a gap in the space of possible solutions: the recovery rate of any bank is either 1 or below α or β .

Lemma 1. Let $X = (N, e, c, \alpha, \beta)$ be a financial system, r clearing for X , and let $i \in N$ be a bank. If $r_i < 1$, then the following hold:

- (1) If i has only external assets (i.e., $\sum_j p_{j,i}(r) = 0$), then $r_i \leq \alpha$. If, in addition, $\alpha > 0$, then $r_i < \alpha$.
- (2) If i has only interbank assets (i.e., $e_i = 0$), then $r_i \leq \beta$. If, in addition, $\beta > 0$, then $r_i < \beta$.
- (3) In any case, $r_i \leq \max(\alpha, \beta)$. If $\alpha > 0$ or $\beta > 0$, then $r_i < \max(\alpha, \beta)$.

Proof. We prove part (3). From the definition, it follows that $a'_i(r) \leq \max(\alpha, \beta) \cdot a_i(r)$. Because $r_i < 1$, we must have $a_i(r) < l_i(r)$ (in particular, $l_i(r) > 0$) and $r_i = F_i(r) = a'_i(r)/l_i(r) \leq \max(\alpha, \beta) \cdot a_i(r)/l_i(r) \leq \max(\alpha, \beta)$. If $\alpha > 0$ or $\beta > 0$, then the last inequality is strict.

The proofs of parts (1) and (2) are analogous. \square

3.1. Existence of a Solution

What is perhaps most surprising about financial networks with CDSs is that as soon as there are any default costs, the existence of a solution can no longer be guaranteed.

Theorem 1 (No Solution with Default Costs). For any pair (α, β) with $\alpha < 1$ or $\beta < 1$ there exists a financial system (N, e, c, α, β) that has no clearing recovery rate vector.

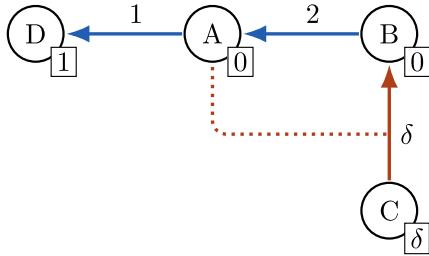
Proof. If $\beta < 1$, consider the system in Figure 2. Let $\delta = 3 \cdot 1/(1 - \beta)$. Assume, toward a contradiction, that there is a clearing recovery rate vector r .

• If $r_A = 1$, then $p_{C,B}(r) = l_{C,B}(r) = \delta(1 - r_A) = 0$; hence $a_B(r) = 0$ and $p_{B,A}(r) = 0$. This implies $a_A(r) = 0 < 1 = l_A(r)$, and thus $r_A = 0$ —a contradiction.

• If $r_A < 1$, then $r_A \leq \beta$ by Lemma 1. Thus, $p_{C,B}(r) = l_{C,B}(r) = \delta(1 - r_A) \geq \delta(1 - \beta) = 3$. Now, $a_B(r) = 3 \geq 2 = l_B(r)$, so $p_{B,A}(r) = l_{B,A}(r) = 2$. Hence $a_A(r) \geq l_A(r)$, and so $r_A = 1$ —a contradiction.

The proof for the case $\alpha < \beta = 1$ is provided in Appendix B. It uses a similar construction but where A has positive external assets. \square

Figure 2. (Color online) Financial System with No Solution for $\beta < 1$



The system in Figure 2 is paradoxical because A is implicitly holding a CDS (and is thus short) on itself: if A is in default, it receives a payment because of the CDS written on it, so it is not in default, and vice versa. Although A *actually* holding a CDS on itself would be absurd, having B in between makes the paradox much less obvious. Supervisory authorities could only notice that the two scenarios are, in fact, equivalent once they are aware of network effects and have detailed knowledge about the contract structure, including the ternary relationships introduced by CDSs.

In addition, although it is hard to imagine why a bank would ever buy a CDS on itself, Figure 2 could have formed in an entirely sensible way. For example, B could have borrowed money from A and later placed a speculative bet on A's default before both banks were hit by a shock that wiped out their external assets. With only knowledge of their own assets and liabilities, none of the banks would have noticed any problem.

Figure 2 is a particularly simple example to show nonexistence because of its small size and zero external assets for all relevant banks. Note that these features are not essential for nonexistence. We present a larger example where nonexistence arises in a much more indirect way in the electronic companion of this paper.

Remark 1. We know from Rogers and Veraart (2013) that no example such as in Theorem 1 can be constructed using only debt contracts. Note that it can also not be constructed using only CDSs because in a financial system consisting of *only* CDSs, the recovery rate vector $(1, \dots, 1)$ (nobody defaults) is always clearing: under this recovery rate vector, no liabilities arise, and thus, indeed, no bank defaults. Therefore, nonexistence can only arise in systems with debt *and* CDSs.

It turns out that the nonexistence of a solution hinges on the presence of default costs.

Theorem 2 (Existence of a Solution Without Default Costs). *Any financial system $(N, e, c, \alpha = 1, \beta = 1)$ has a clearing recovery rate vector.*

Proof. Because $\alpha = \beta = 1$, we can simplify the update function F from Definition 1 to $F_i(r) = \min(1, a_i(r)/l_i(r))$ on the set $L_i := \{r \mid l_i(r) > 0\}$. Note that L_i is an open set because l_i is continuous, and note that F_i is continuous

on L_i . We use this fact and apply a fixed-point theorem. Care must be taken because F_i is not, in general, continuous on $[0, 1]^N \setminus L_i$. Consider the set-valued function ρ defined by

$$\rho : [0, 1]^N \rightarrow \mathcal{P}([0, 1]^N), \text{ where } \mathcal{P}(S) \text{ denotes the power set of } S;$$

$$\rho(r) := \bigtimes_{i \in N} \rho_i(r), \text{ where } \rho_i(r) := \begin{cases} F_i(r) & \text{if } r \in L_i, \\ [0, 1] & \text{if } r \notin L_i. \end{cases}$$

If there is an r such that $r \in \rho(r)$, then r can be made clearing by setting the recovery rates of banks with zero liabilities to 1. This is because for all i , if $r \in L_i$, then $F_i(r) = r_i$ by choice of r , and if $r \notin L_i$, then $F_i(r) = 1$, and no other bank depends on i because of our sanity assumptions.

It remains to show that an r with $r \in \rho(r)$ exists. By the Kakutani (1941) fixed-point theorem, this is the case if (1) the domain of ρ is compact and convex, (2) the set $\rho(r)$ is convex for each r , and (3) the graph of ρ , $G_\rho := \{(r, s) \mid s \in \rho(r)\}$, is a closed set. Properties (1) and (2) are obvious.

To prove (3), it suffices to show that for each i , the graph of ρ_i , $G_{\rho_i} := \{(r, s_i) \mid s_i \in \rho_i(r)\}$, is closed. To this end, let $(r^k, s_i^k)_{k \in \mathbb{N}}$ be a sequence in $[0, 1]^N \times [0, 1]$ converging to some point (r, s_i) such that $s_i^k \in \rho_i(r^k)$ for each k . We need to show that $s_i \in \rho_i(r)$. If $r \notin L_i$, then trivially, $s_i \in \rho_i(r) = [0, 1]$. If $r \in L_i$, then $s_i \in \rho_i(r) = \{F_i(r)\}$ because F_i is continuous on the open set L_i . \square

In financial systems without default costs, money is never lost, just redistributed. Theorem 2 shows that these systems always have a solution. It does not apply once default costs are present because the update function F then has a discontinuity where the assets of a bank are equal to its liabilities (i.e., when a bank is just on the verge of defaulting). This discontinuity creates a gap in the space of possible recovery rates (see Lemma 1) and can give rise to nonexistence.

One might wonder what changes to our model might restore existence in the general case. Our model differs from Eisenberg and Noe (2001) only in that we allow for default costs and CDSs. Thus, if one seeks to represent these two features and aims to guarantee existence, the only option would be to use a clearing model different from simultaneous clearing. Perhaps the first alternative that comes to mind is *sequential clearing*, where the contracts are not evaluated at the same time but in some order. The result of this procedure, however, heavily depends on the order of evaluation, as the following example shows.

Example 1 (Sequential Clearing). We define a natural sequential clearing procedure: The debt contracts are evaluated in a predetermined order, and banks pay their liabilities based on their external assets and

payments received so far (i.e., their “cash” holdings). If a bank cannot pay a liability, then it enters bankruptcy. Default costs are subtracted from the bank’s cash holdings, and the recovery rate is computed based on the remaining cash. Then all CDSs written on the bank are triggered and are evaluated next. The process ends when all debt contracts have been evaluated.

Now assume that this procedure is applied to Figure 2 for $\beta = 0.5$, so that $\delta = 6$.

- If the debt contract from A to D is evaluated first, then A defaults with $r_A = 0$, B receives six in the CDS, and A receives two from B. The resulting equities are $E_A(r) = 2$ and $E_B(r) = 4$.
- If the debt contract from B to A is evaluated first, then B defaults with $r_B = 0$, A receives nothing and defaults with $r_A = 0$, and B receives six in the CDS. We have $E_A(r) = 0$ and $E_B(r) = 6$.

The order of evaluation in sequential clearing could be chosen at random or based on some objective criterion, such as the maturity of the contracts. In any case, it would introduce an element of arbitrariness and an opportunity for strategic manipulation. Although other sequential clearing procedures could be defined, it seems unlikely that this problem could be fully avoided.

Alternative clearing models that go beyond sequential clearing have been proposed in the literature; each has its own limitations. Csóka and Herings (2018) studied a *decentral clearing* procedure where payments are made incrementally in an arbitrary order. In a setting with CDSs, the result of this procedure depends on the order in which payments are made, similar to sequential clearing.¹³ Banerjee and Feinstein (2018) defined a *dynamic clearing* procedure where multiple rounds of simultaneous payments are performed while a CDS is triggered with a delay of one round after its reference entity has defaulted. A solution always exists and can be chosen in a natural way. However, CDSs in this model cannot represent a complete insurance on a debt exposure. This is because, if the writer of a debt contract fails, the holder still incurs a loss *in that round*, irrespective of any CDS they might hold. This loss may be enough to send the holder into permanent bankruptcy. Acemoglu et al. (2015) studied a simultaneous clearing model where default costs arise exclusively from the partial liquidation of illiquid projects. This assumption would ensure existence even with CDSs via continuity in a similar way to Theorem 2, but it also precludes modeling any kind of *discontinuous* default costs such as time delays or operational losses.

3.2. Multiplicity of Solutions

We now show that even when the clearing problem has a solution, there can be multiple ones, and the structure of the set of solutions may not be economically desirable. We discuss this structure in terms

of the banks’ aggregate preferences. Recall that we denote the equity of a bank i by $E_i(r)$. We assume that, when there are multiple solutions, banks prefer those that maximize their equity.

Definition 2 (Preferred and Maximal Solution). Fix a financial system X . A bank i is said to *weakly prefer* a solution r over another solution r' if $E_i(r) \geq E_i(r')$. A solution r is called *maximal* if it is weakly preferred to all other solutions by all banks.¹⁴

Our second desideratum, *maximality*, requires that a maximal solution exists. Otherwise, any solution the clearing authority could select would be opposed by at least one bank because this bank could achieve strictly higher equity in a different solution.¹⁵ Such a situation is illustrated in the following theorem.

Theorem 3 (No Maximal Solution). *For any α and β , there exists a financial system (N, e, c, α, β) that has a clearing recovery rate vector but no maximal one.*

Proof. We use the financial system in Figure 3 with $\delta = 1/(1 - \beta)$ if $\beta < 1$ and $\delta > 1$ arbitrary if $\beta = 1$. It is easy to verify that $r^0 := (0, 1, 1, 1)$ and $r^1 := (1, 0, 1, 1)$ (where entries are in alphabetical order) are clearing. In any potential other solution, we must have $r_C = r_D = 1$ and $0 < r_A, r_B < 1$.

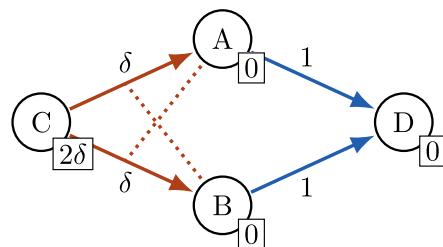
For $\beta < 1$, no other solution exists: if r was another one, then because $r_A < 1$, by Lemma 1, we have $r_A \leq \beta$, so $a_B(r) = \delta(1 - r_A) \geq \delta(1 - \beta) = 1$. Thus, $r_B = 1$ —a contradiction.

For $\beta = 1$, there is exactly one other solution $r^2 = (\zeta, \zeta, 1, 1)$, where $\zeta = (\delta^2 - \delta)/(\delta^2 - 1)$. This is because r is a solution with $r_A, r_B < 1$ iff $r_A = \delta(1 - r_B)$ and $r_B = \delta(1 - r_A)$. It is easy to verify that $r_B = r_A = \zeta$ is the unique solution of this linear equation system.

For any value of β , bank A has a positive equity of $\delta - 1$ in r^1 and equity 0 (because it is in default) in the other solution(s). Thus, A strictly prefers r^1 . Analogously, B strictly prefers r^0 . This implies that none of the solutions of this system is maximal.¹⁶ □

To see why the solution structure in the previous theorem is economically undesirable, consider the $\beta < 1$ case in the above proof and imagine a clearing authority faced with the problem of actually clearing

Figure 3. (Color online) Financial System with No Maximal Solution



the market: there are two solutions, one where A defaults and one where B defaults. Choosing among the solutions means giving preference to one of the banks. It is not clear how this decision should be made, and the clearing authority may even be legally prohibited from making such a trade-off. If a choice among nonmaximal solutions were legally allowed, then a bank may have a large incentive to lobby for the implementation of a solution that it prefers most. Note that, in contrast to nonexistence, nonmaximality can even occur in systems without default costs.¹⁷

If clearing were done sequentially in the scenario from Theorem 3, one of the two solutions, r^0 or r^1 , would be chosen based on which of the two debt contracts is evaluated first. In practice, such a scenario could lead to severe incentive problems. In today's financial practice, whether a CDS is triggered is decided by so-called *determinations committees*, which consist of the most active dealers in addition to nondealer members (International Swaps and Derivatives Association 2012). In Figure 3, A, B, and C would be members of the determinations committee. Taking the perspective of A, it would be rational to try to convince the other members that B's financial situation qualifies as a default. This triggers the CDS, A receives the payment and does not default, and B receives nothing. Thus, in hindsight, it appears as though A made a correct objective assessment about B. Of course, B would argue exactly the opposite of A. By contrast, when a maximal solution exists, it can be implemented without having to make any choices that could be manipulated.

Remark 2. Note that Rogers and Veraart (2013) have previously observed multiple solutions in debt-only networks stemming from default costs. However, the form of multiplicity they observed is much less problematic because in debt-only systems, there always exists a maximal solution.

4. Dependency Analysis Framework: The Colored Dependency Graph

In Section 3, we have shown that introducing CDSs into the well-established clearing model by Eisenberg and Noe (2001) has the effect that existence and maximality are no longer guaranteed. In this section, we develop an analysis framework, which we call the "colored dependency graph," to better understand how and when this effect arises. In Section 5, we then show how to use the colored dependency graph to derive sufficient conditions under which the two desiderata are satisfied.

4.1. Covered and Naked CDS Positions

At the level of an individual bank, we need to distinguish between two fundamentally different uses of

CDSs. For the purposes of illustration, consider a financial system with a single CDS where the CDS writer cannot default. If the holder of the CDS also holds at least an equal amount of debt written by the reference entity, then the CDS holder is long on the reference entity: a worse situation of the reference entity would, at most, be offset by the CDS payment, but it could never be beneficial for the holder. This use of a CDS is called *covered*. By contrast, if the holder holds no or not enough debt written by the reference entity, then it is short on the reference entity: a worse financial situation of the reference entity would benefit the holder. This use of a CDS is called *naked*. See the top row of Figure 4 for a depiction of a prototypical (a) debt contract, (b) naked CDS, and (c) covered CDS. For the formal definition in general financial systems, we must consider the notional of *all* CDSs that a bank holds on a reference entity to classify a CDS position as covered or naked.

Definition 3 (Covered and Naked CDS Position). Let $X = (N, e, c, \alpha, \beta)$ be a financial system. A bank j has a *covered CDS position* toward another bank k if

$$\sum_{i \in N} c_{i,j}^k \leq c_{k,j}^\emptyset.$$

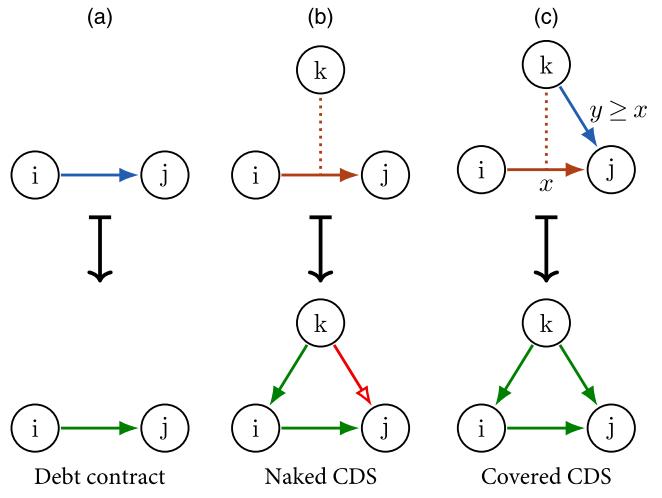
Otherwise, j has a *naked CDS position* toward k .

4.2. The Colored Dependency Graph

We can now define the colored dependency graph (or just the "dependency graph"), in which long and short positions among banks are represented by green edges (with filled arrow tips) and red edges (with hollow arrow tips), respectively.

Definition 4 (Colored Dependency Graph). Let $X = (N, e, c, \alpha, \beta)$ be a financial system. The *colored dependency*

Figure 4. (Color online) Prototypical Financial Systems (Top) and Their Colored Dependency Graphs (Bottom)



Note. Green edges with filled arrow tips indicate long positions, and red edges with hollow arrow tips indicate short positions.

graph $CD(X)$ is the graph with nodes N and edges of colors red and green constructed as follows:

- (1) For each $i, j \in N$, if $c_{i,j}^\emptyset > 0$ or $c_{i,j}^k > 0$ for any $k \in N$, then add a green edge $i \rightarrow j$.
 - (2) For each $i, k \in N$, if $c_{i,j}^k > 0$ for any $j \in N$, then add a green edge $k \rightarrow i$.
 - (3) For each $j, k \in N$, if j has a naked CDS position toward k , then add a red edge $k \rightarrow j$.

The definition of the colored dependency graph can be understood in terms of the three primitive contract patterns illustrated in Figure 4: debt contracts, naked CDSs, and covered CDSs. In each case, the holder of any contract is long on the writer because, in case the writer defaults, the lower the recovery rate of the writer, the lower the payment that the holder receives. This is expressed by rule (1) in Definition 4. In case of a debt contract, this is the only dependency that is induced, whereas a CDS gives rise to two additional dependencies. The writer of a CDS is always long on the reference entity because, the lower the recovery rate of the reference entity, the higher the liability for the writer. This is expressed by rule (2) in Definition 4. The position of the holder of a CDS toward the reference entity depends on whether it is a naked or a covered CDS: only the holder of a naked CDS is short on the reference entity, expressed by rule (3) in Definition 4. A covered CDS, on the other hand, only gives rise to a long position together with the debt contract.

The following proposition shows the usefulness of the framework in capturing the directional behavior of the update function F . We will repeatedly use it in Section 5 when deriving sufficient conditions. The proof is straightforward and thus omitted.

Proposition 1 (The Colored Dependency Graph and the Update Function). *For any two banks i and j , we let r_{-ij} denote a vector of recovery rates of all banks excluding i and j . Then the following holds:*

- (1) If there exists an r_{-ij} such that, holding r_{-ij} fixed, the function F_j is increasing¹⁸ in r_i , then there is a green edge from i to j in $\text{CD}(X)$.
 - (2) If there exists an r_{-ij} such that, holding r_{-ij} fixed, the function F_j is decreasing in r_i , then there is a red edge from i to j in $\text{CD}(X)$.
 - (3) If there is no edge from i to j of any color, then F_j is independent of r_i . The converse is not necessarily the case.

Remark 3 (Parallel Edges). Both a red and a green edge can be present in the dependency graph in the same direction between the same two banks. In this case, whether a long or a short effect is present depends on the recovery rates of the other banks. The two edges do not cancel out.

If a financial system contains only debt contracts, then the colored dependency graph only has green edges; specifically, it has a green edge $i \rightarrow j$ whenever

$c_{ij}^\emptyset > 0$. Notice that this graph coincides with the “financial structure graph” introduced by Eisenberg and Noe (2001). For systems with debt and CDSs, our colored dependency graph provides an elegant conversion from the ternary relations introduced by CDSs to binary relations, making them amenable to graph-theoretic analysis.¹⁹

Figure 5 depicts the colored dependency graphs of two financial systems that exhibit very different behavior: Figure 5, panel (a) corresponds to the example financial system from Figure 1, which has a unique solution. Figure 5, panel (b) corresponds to the financial system from Figure 2, which has no solution. We immediately see some similarities and differences: both graphs have a red edge; panel (a) has no *directed cycle*, whereas (b) has two of them, A–B–A and A–C–B–A; and the former cycle contains a red edge. All of these features will be of importance in the analysis in Section 5. Note that, although the cycles in Figure 5, panel (b) happen to be very short, this is not a necessary condition for the nonexistence of a solution. The electronic companion provides a more involved example with longer cycles.

5. Analysis of Restricted Network Structures

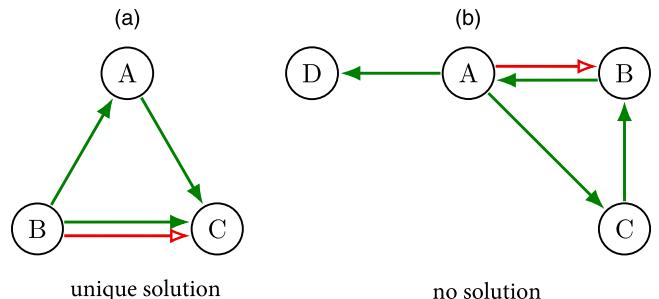
With our analysis framework in place, we now use it to describe sufficient conditions under which our desiderata are fulfilled. We show that one can guarantee our desiderata by restricting the ways in which the edges in the dependency graph may form cycles. We present three domain restrictions where we successively allow more cycles and receive successively fewer guarantees.

5.1. Acyclic Financial Systems

If there are no cycles in the colored dependency graph, then the clearing problem has a *unique* solution. As this solution is trivially maximal, both desiderata are fulfilled.

Theorem 4 (Existence and Uniqueness in Acyclic Financial Systems). *Let X be a financial system such that $\text{CD}(X)$ has no cycles. Then X has a unique clearing recovery rate vector.*

Figure 5. (Color online) The Colored Dependency Graphs of the Financial Systems from (a) Figure 1 and (b) Figure 2



Proof. Without loss of generality, assume that $N = \{1 \dots n\}$ and banks are sorted in *topological order*; that is, whenever there is an edge $i \rightarrow j$ in $CD(X)$, we have $i \leq j$. This is possible because $CD(X)$ has no cycles by assumption. To find a solution r , iterate over banks i in order. In each step, set $r_i := F_i(r_1, \dots, r_{i-1})$, where r_1, \dots, r_{i-1} have already been computed. This is well defined by Proposition 1. In the end, r is clearing by construction.

Toward uniqueness, if r and r' are both clearing, it follows by induction on i that $r_i = F_i(r_1, \dots, r_{i-1}) = F_i(r'_1, \dots, r'_{i-1}) = r'_i$ for all i , where the middle equality is by induction hypothesis. Note that F_1 is a constant function. \square

Theorem 4 shows formally that default ambiguity in financial systems with CDSs is due to cycles in the dependency graph. Note that we must consider *all* dependency edges here, including those originating at reference entities of CDSs. It is not sufficient to consider the graph of *liabilities*, where an edge exists from the writer of each contract to the holder, corresponding to only rule (1) in Definition 4. This graph would be acyclic for all our counterexamples in Section 3, although they clearly did not fulfill our desiderata. Thus, the more sophisticated colored dependency graph is necessary to capture the behavior of a financial system with CDSs.

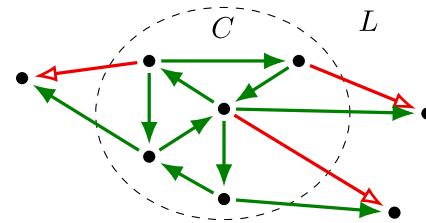
5.2. Green Core Systems

The previous theorem required a very strong assumption; in reality, financial systems *do* contain cycles in the dependency graph, but not all of them pose a problem. In fact, we know from Rogers and Veraart (2013) that debt-only financial systems, even if they contain cycles, always satisfy existence and maximality. At the same time, debt-only systems always have a completely green dependency graph; that is, banks are only long on each other. In this section, we show that *all* financial systems with a completely green dependency graph satisfy existence and maximality, thus generalizing Rogers and Veraart's result. We consider a slightly more general class of financial systems that we call *green core systems*.

Definition 5 (Green Core System). A financial system X is called a *green core system* if in $CD(X)$, banks with an incoming red edge (i.e., the holders of naked CDS positions) have no outgoing edges. We call the set of these banks the *leaf set* and the other banks the *core*.

An example of a green core system is shown in Figure 6. Banks in the leaf set have no liabilities (otherwise, they would have an outgoing green edge) and hence always have recovery rate 1. This does not render the leaf set obsolete: allowing a leaf set keeps the definition of green core systems general enough so that

Figure 6. (Color online) Colored Dependency Graph of a Green Core System with Core C and Leaf Set L



banks in the core can be writers of naked CDSs. This feature will also be essential in Section 5.3, where we consider even more general network structures that are composed of multiple green core systems that can be connected by red edges.

Green core systems always have a solution that is best for all banks *in the core*. We call such a solution *core maximal*. Our proof is constructive.

Theorem 5 (Existence and Core Maximality in Green Core Systems). *In any green core system, the following holds:*

- (1) *There exists a recovery rate vector that maximizes both the recovery rate and the equity of all banks in the core.*
- (2) *The iteration sequence (r^n) defined by $r^0 = (1, \dots, 1)$ and $r^{n+1} = F(r^n)$ converges to this recovery rate vector.*

Proof. The main technical challenge lies in proving the following lemma.

Lemma 2. *Consider a green core system with core C and leaf set L .*

- (1) *The update function F is monotonic and continuous from above, where the order relation is a pointwise comparison of recovery rate vectors.*
- (2) *If $i \in C$, then the equity E_i is monotonic, also with respect to a pointwise comparison.*

The proof of Lemma 2 is given in Appendix C. The lemma formalizes the fact that because all relevant dependency edges are green, a decrease in any bank's recovery rate can only affect the other banks in the core in a negative way. In addition, this happens in a continuous fashion.

From part (1) of the lemma, it follows via a standard technique from lattice theory (see Lemma 3 in Appendix C) that the sequence (r^n) converges to a solution that maximizes the recovery rate of each bank. By part (2) of the lemma, this solution also maximizes all equities in the core. \square

Theorem 5 shows that green core systems always satisfy existence and core maximality. Furthermore, the proof of the theorem tells us that green core systems are structurally very similar to debt-only systems. They share the following properties, which have previously been observed for debt-only systems by Rogers and Veraart (2013). First, the update function is monotonic and continuous from above. Second, the set of

solutions even forms a *complete lattice* (which follows from monotonicity of F via the Knaster–Tarski fixed-point theorem; see, e.g., Granas and Dugundji 2003). Third, a core-maximal solution can be found via the iteration sequence provided in part (2) of Theorem 5. A subtle difference to the debt-only case is that a core-maximal solution of a green core system can contain irrational numbers, whereas the maximal solution of a debt-only system is always rational.²⁰

A special case of Theorem 5 is a situation in which naked CDSs are not present.

Corollary 1 (Existence and Maximality Without Naked CDSs). *If no bank in a financial system has a naked CDS position toward another bank, then there exists a maximal clearing recovery rate vector.*

Proof. In this case, the colored dependency graph contains only green edges. The financial system is hence trivially a green core system where the core consists of all banks. \square

Corollary 1 has important implications for regulatory policy regarding naked CDSs, which we discuss in detail in Section 6.

5.3. Systems Without Red-Containing Cycles

We know from Theorem 5 and Corollary 1 that default ambiguity can be attributed to the presence of red edges in the dependency graph. Green core systems restrict these edges in an extreme way, only allowing them to leaf banks. But we know from Theorem 4 (acyclic systems) that red edges to nonleaf banks do not *always* pose a problem. In this section, we study in which situations they do. Our main result is that, regarding existence, only red edges that are part of a cycle of dependencies can pose a problem.

Theorem 6 (Existence Without Red-Containing Cycles). *Assume that in the colored dependency graph of a financial system, there is no cycle that contains a red edge. Then a clearing recovery rate vector exists.*

Our proof of the theorem is constructive by using an algorithm. Unfortunately, simply iterating the function F such as in the green core case does not work anymore in the more general no-red-containing-cycle case (we provide an example in Appendix D). Instead, our algorithm exploits the structure of the dependency graph. Recall from above that the solutions of a financial system with debt and CDSs may be irrational. Given this, it is impossible to design an algorithm that can compute an exact solution in finite time. Instead, we devise an *approximation algorithm* that computes an arbitrarily accurate *approximate* solution. We now first describe our approximate solution concept and our algorithm. We then prove correctness of the algorithm and Theorem 6.

Definition 6 (Approximately Clearing Recovery Rate Vector). Let X be a financial system, and let $\varepsilon \geq 0$. A recovery rate vector r is called ε -approximately clearing or an ε -solution for X if $\|F(r) - r\| \leq \varepsilon$, where $\|r\| := \max_i |r_i|$ is the maximum norm.

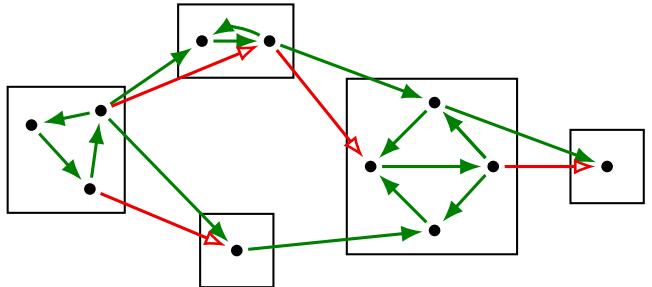
We now describe our *core iteration algorithm* to compute an ε -solution in a financial system X when no cycle in $CD(X)$ contains a red edge. Given are ε and X . We begin by partitioning the dependency graph into strongly connected components; or *cores*. Each of these corresponds to the core of a green core system. A core is a minimal set of banks such that all banks with which these banks are in cycles are also part of the core. By partitioning the graph in this way, the connections between different cores form an acyclic graph, so we can sort them in topological order (i.e., edges only go from earlier to later cores in the order, but never in the other direction).²¹ Figure 7 provides an example for such a dependency graph. We now iterate over cores. By assumption, all edges within a core are green, so we can use the iteration sequence from Section 5.2 to compute an ε -solution for each of them.

More in detail, let C_1, \dots, C_m be the cores in topological order. We store recovery rates in a vector r . Initially, r is the empty vector. In step $k \in \{1, \dots, m\}$, we define a function

$$F^k : [0, 1]^{C_k} \rightarrow [0, 1]^{C_k}, \\ F^k(s) := F_i(r \frown s),$$

where the symbol “ \frown ” denotes a concatenation of vectors. This corresponds to the update function F restricted to C_k with previously computed recovery rates of the previous cores C_1, \dots, C_{k-1} given by r . Function F^k is well defined because, by the topological ordering, each bank $i \in C_k$ depends only on the banks in C_1, \dots, C_k . We iterate the function F^k starting at $s = (1, \dots, 1)$ until $\|F^k(s) - s\| \leq \varepsilon$. We then add the recovery rates computed in s to r and continue with the next core. The algorithm stops when all cores have been visited.

Figure 7. (Color online) Dependency Graph Where No Cycle Contains a Red Edge



Notes. Cores are marked by black rectangles. The topological ordering of cores is from left to right; the two cores second from the left can be visited in any order.

Proposition 2 (Correctness of the Core Iteration Algorithm). *Assume that in the colored dependency graph of a financial system, no cycle contains a red edge. Then for any $\varepsilon > 0$, the core iteration algorithm computes an ε -approximately clearing recovery rate vector.*

Proof. To see that r is an ε -solution when the algorithm terminates, let i be a bank, let k be such that $i \in C_k$, and let s be r restricted to the indices in C_k . By the topological ordering, $F_i(r)$ depends only on the r_j with $j \in \bigcup_{l \leq k} C_l$. Hence, $F_i(r) = F_i^k(s)$, and therefore $|F_i(r) - r_i| = |F_i^k(s) - s_i| \leq \|F^k(s) - s\| \leq \varepsilon$ as required, where the last inequality holds by the algorithm's stopping criterion.

It remains to show that the algorithm terminates—that is, that the iteration sequence for F^k reaches the stopping criterion $\|F^k(s) - s\| \leq \varepsilon$ after finitely many steps for each k . First note that F^k is monotonic and continuous from above. This follows just like in Lemma 2, where we also need to account for the effects of earlier cores on the financial subsystem C_k : CDSs written by banks in C_k on banks in earlier cores give rise to additional fixed liabilities, and incoming payments from earlier cores to C_k give rise to additional assets. These manifest as constants that do not affect the argument in the proof. Now the iteration sequence converges to a maximal fixed point of F^k like in Theorem 5. In particular, we reach the stopping criterion after finitely many steps.²² □

Given the core iteration algorithm, it is now straightforward to prove existence in systems without red-containing cycles.

Proof of Theorem 6. We “run” the algorithm with $\varepsilon = 0$ to receive a constructive proof of existence. The stopping criterion $\|F^k(s) - s\| = 0$ is not attained after finitely many steps but in the limit of the iteration sequence. All other steps of the proof of Proposition 2 remain the same. □

Theorem 6 generalizes and unifies the existence statements of Theorems 4 and 5: individually, neither cycles nor red edges going to nonleaf nodes are a problem; only red-containing cycles can cause nonexistence. Thus, the no-red-containing-cycle condition is the most general (weakest) condition we have derived for existence (as it also covers acyclic and green core systems). Regarding maximality, the weakest condition we have derived is “acyclic or no naked CDSs” (Theorem 4 and Corollary 1). The absence of red-containing cycles does not guarantee maximality because cores with an incoming red edge may be made worse off when the recovery rates of earlier cores are maximized. It is an open question whether a condition exists that is weaker than acyclic or no naked CDSs and guarantees maximality. However, because our analysis has shown that cycles and red edges in the dependency graph are essential

factors for nonmaximality, it seems unlikely that a simple condition that fulfills this requirement can be found.

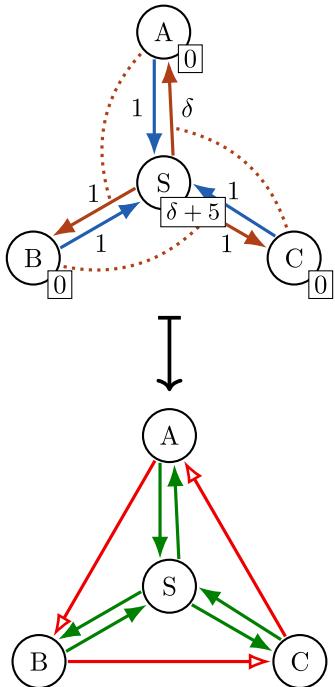
In this section, we have shown that our dependency analysis framework can be used to derive sufficient conditions for existence and maximality. However, they are not *necessary conditions*. Although it may be possible to derive stronger guarantees by taking even more information about the contract structure into account, we should not expect to obtain *equivalence conditions*: Our computational complexity results in Schuldenzucker et al. (2019) imply that any condition that is equivalent to existence or maximality would be NP-hard to check (informally, this would take exponential runtime) and would therefore be of limited use. By contrast, our framework has yielded sufficient conditions that are simple and easy to check. Thus, we argue that our colored dependency graph hits a “sweet spot” by capturing the most important interactions among contracts, enabling us to distinguish between long and short positions as well as between covered and naked CDSs.

6. Discussion: Policy Relevance

We evaluate two recent policies regarding their effectiveness for protecting against default ambiguity under the assumptions of our model: *central counterparty clearing* and *banning naked CDSs*.

The regulatory frameworks EMIR (in Europe) and Dodd–Frank (in the United States) mandate the use of a *central clearing counterparty* (CCP) for a large part of the over-the-counter derivatives market.²³ In its most extreme form, this means that all contracts are routed via a central node: a bank A would not write a contract to a bank B directly, but rather bank A would write a contract to a highly capitalized central entity S , and S would write a contract to bank B. One of the desired effects is that the CCP would absorb a shock on the banks, prevent it from spreading through the network, and thus prevent financial contagion. Although using a CCP simplifies the network of liabilities, surprisingly, it is not effective for protecting against default ambiguity in our model. Figure 8 provides an example: there are three banks that hold CDSs and write debt together with a CCP S .²⁴ Note that S has very high external assets such that it cannot default. This system does not have a solution (the proof is given in Appendix E). Indeed, when we look at the colored dependency graph (Figure 8, bottom panel), we see that there is still a red-containing cycle A–B–C–A. At a higher level, we see that although a CCP can help reduce *counterparty risk* (i.e., the risk to a bank that a debtor cannot pay its liability), the *flow of fundamental risk* (i.e., the risk that the reference entity in a CDS has a higher or lower recovery rate than expected; see D’Errico et al. 2018) still takes place

Figure 8. (Color online) Financial System with a CCP and No Solution and Its Colored Dependency Graph



Notes. Top Panel: Financial system with a CCP S where $\beta < 1$ is arbitrary and $\delta := 1/(1 - \beta + \beta^2 - \beta^3)$. There is no clearing recovery rate vector. Bottom Panel: The corresponding colored dependency graph (note that we omit offsetting positions that the CCP S has with other banks as they do not affect our results).

directly between the banks, essentially “around” the CCP. This is enough to lead to nonexistence of a solution. Overall, our example shows that requiring banks to trade all CDSs via a CCP, even if the CCP is very well capitalized, does not guarantee existence of a solution to the clearing problem. This result is formally proven in Appendix E.

Another policy that has seen adoption in Europe since the European sovereign debt crisis in 2011 is *banning naked CDSs*. A CDS on a European sovereign state can

only be bought if a corresponding (debt) exposure is present as well (European Commission 2011; also see European Securities and Markets Authority 2017b). Corollary 1 shows that if *all* naked CDSs are banned, not only those on sovereigns, then the clearing problem is guaranteed to have a maximal solution. Thus, under the assumptions of our model, this policy is effective against default ambiguity. Note that in this paper, we refrain from recommending the adoption of any particular policy. Instead, our findings illustrate how our framework can be used to help inform regulatory policy.

7. Conclusion

In this paper, we have shown that financial networks that contain debt contracts and CDSs are prone to a phenomenon we call *default ambiguity*—that is, a situation where it is impossible to decide which banks are in default. For many years, the total notional of CDSs written on financial institutions has exceeded USD1 trillion worldwide. Although the European Commission has previously acknowledged that CDSs can give rise to new kinds of systemic risk, they have not yet considered the risk of default ambiguity. Our new dependency analysis framework reveals that default ambiguity hinges on the presence of cycles in the colored dependency graph. Table 1 summarizes our findings. As we have shown, the more we relax the restrictions on the type of these cycles, the weaker the guarantees we obtain for our desiderata. To find a solution for the restricted network structures we have studied, one can use the *core iteration algorithm* we have provided.

Our results illustrate that, to understand the behavior of financial systems with CDSs, it is essential to consider the *ternary* relations they introduce, including the reference entities. If we had instead only considered the writer-holder relationships, all of our counterexamples in Section 3 would have looked like simple acyclic graphs, and we would only have captured one of the three dependencies arising from a CDS. Our insights may help bring about a paradigm

Table 1. Summary of Our Results

Contracts	Network structure	Existence	Maximality	Reference
Debt only	Any network structure	✓	✓	Eisenberg and Noe (2001), Rogers and Veraart (2013)
Debt + CDSs	External reference entities only	✓	✓	See “Debt only” above ^a
	Acyclic dependency graph	✓ (unique)	✓	This paper (Section 5.1)
	No naked CDSs	✓	✓	This paper (Section 5.2)
	Green core systems	✓	(✓) (in the core)	This paper (Section 5.2)
	No red-containing cycles	✓	✗	This paper (Section 5.3)
	Red-containing cycles	✓ (no default costs) ^b /✗ (with default costs)	✗	This paper (Section 3)

^aWhen reference entities are external to the financial system, then CDS liabilities can be considered constant for the purpose of clearing, and the results for financial systems without CDSs carry over to this case.

^bRemember that our proof of this result is nonconstructive. Indeed, we have shown in a separate piece of work (Schuldenzucker et al. 2019) that, in general financial systems without default costs, finding a solution is PPAD-hard (informally, any algorithm would need exponential runtime in the worst case). This implies that any *practical* algorithm would have to use heuristics.

shift in the literature on systemic risk in CDS markets where, so far, either the reference entities were aggregated or the interactions across different reference entities were not taken into account.

From a conceptual perspective, the reason why CDSs can give rise to default ambiguity is that the holder of a naked CDS may profit from financial distress of another market participant. Note that this phenomenon is not exclusive to CDSs. For example, the holder of a bond put option and the writer of a bond call option both benefit if the issuer of the underlying bond is in financial distress, and therefore the price of the bond declines. Stock options exhibit similar behavior. Thus, we expect that these markets would also be susceptible to default ambiguity. Our framework can be extended to these other derivative markets in a straightforward way (in particular, to options).

Our dependency analysis framework enables ex ante guarantees that hold irrespective of a shock to banks' external assets. If the external assets are known or can be bounded, future work may be able to derive stronger sufficient conditions for existence and maximality. To do this, one could extend the colored dependency graph with weights that represent the "strength" of the dependency. This will be a challenging task, however, because in contrast to standard weighted-graph models, this "weighted dependency graph" would have to represent highly nonlinear effects.

An important problem that is closely related to clearing is *network valuation* of contracts. Studying this problem requires a model with uncertainty about the future value of banks' external assets. Barucca et al. (2016) designed such a model for debt-only networks by extending the Eisenberg and Noe (2001) model. Researchers interested in network valuation with CDSs could similarly extend our new model. This would raise new questions regarding whether a consistent vector of CDS valuations exists and what is needed for market prices to reflect these true values.

Acknowledgments

The authors thank (in alphabetical order) Vitor Bosshard, Marc Chesney, Marco D'Errico, Helmut Elsinger, Timo Mennle, Thomas Noe, Joseph Stiglitz, and Martin Summer for helpful comments on this work. They are further thankful for the feedback they received from various participants at the 17th ACM Conference on Economics and Computation (EC'16), the 5th World Congress of the Game Theory Society (GAMES'16), and the Second Conference on Financial Networks and Sustainability (FINEXUS'18). They also thank the editor Gustavo Manso, the anonymous associate editor, and two anonymous reviewers for useful comments. Some of the ideas presented in this paper were also described in a one-page abstract that was published in the conference proceedings of EC'16 (Schuldenzucker et al. 2016).

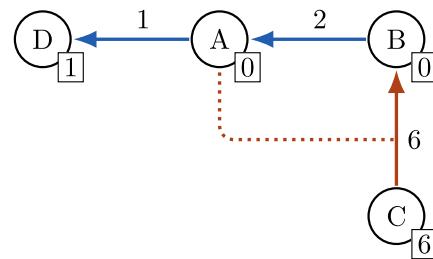
Appendix A. Example That Iterating the Update Function Does Not Generally Converge to a Solution

Consider Figure A.1.²⁵ The unique solution of this system is $r_A = 6/7$, $r_B = 3/7$, and $r_C = r_D = 1$. However, the iteration sequence defined by $r^0 = (1, 1, 1, 1)$ and $r^{n+1} = F(r^n)$ does not converge to this solution but rather exhibits cycling behavior: we have $r^1 = (1, 0, 1, 1)$, $r^2 = (0, 0, 1, 1)$, $r^3 = (0, 1, 1, 1)$, $r^4 = (1, 1, 1, 1) = r^0$, etc. One may think that the cycling behavior is due to an unfortunate choice of the starting point r^0 , but this is not the case: the iteration sequence does not converge for *any* starting point other than the solution itself. To see this, let $\Delta \neq 0$ and $r_B = 3/7 + \Delta$. It is easy to see from the definition of F that

$$\begin{aligned} F_B(F(r)) &= \min(1, \max(0, 3(1 - 2r_B))) \\ &= \min\left(1, \max\left(0, \frac{3}{7} - 6\Delta\right)\right). \end{aligned}$$

Thus, after two iterations, the distance to the solution has increased sixfold until the sequence again enters the infinite loop above.

Figure A.1. (Color online) Financial System Where Iterating F Does Not Converge to a Solution; Let $\alpha = \beta = 1$ (No Default Costs)



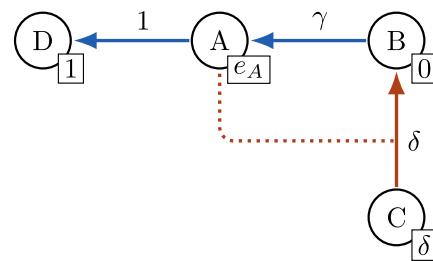
Appendix B. Omitted Proofs from Section 3

We describe a financial system that does not have a solution and where $\alpha < \beta = 1$, thus completing the proof of Theorem 1.

Proof of Theorem 1 ($\alpha < \beta = 1$). Consider Figure B.1, a variant of Figure 2, with values for e_A , γ , and δ chosen as follows: let $e_A \in (0, 1)$ arbitrary, set $\gamma = 1 - (1 + \alpha)/2 \cdot e_A$, and let $\delta \geq \gamma/(1 - \alpha e_A - \gamma)$. It is easy to see that (1) $e_A < 1$, (2) $e_A + \gamma > 1$, and (3) $\alpha e_A + \gamma < 1$. We have $\gamma > 0$ by definition and $\delta > 0$ by (3), so this is a well-defined financial system.

We perform a case distinction such as in the proof for $\beta < 1$. Assume toward a contradiction that r is clearing.

Figure B.1. (Color online) Financial System with No Solution for $\alpha < \beta = 1$



- If $r_A = 1$, then $p_{C,B}(r) = 0$, so $a_B(r) = 0$ and $p_{B,A}(r) = 0$. Thus, $a_A(r) = e_A < 1$, which implies that $r_A < 1$ —a contradiction.
- If $r_A < 1$, then A is in default, so $r_A = (\alpha e_A + p_{B,A}(r))/1 \leq \alpha e_A + \gamma$. Thus, $p_{C,B}(r) = \delta(1 - r_A) \geq \delta(1 - \alpha e_A - \gamma) \geq \gamma$, so B is not in default and $p_{B,A}(r) = \gamma$. Now $a_A(r) = e_A + \gamma > 1$ by (2), so A is not in default, and $r_A = 1$ —a contradiction. \square

Appendix C. Omitted Proofs from Section 5

We show that in a green core system, the update function is monotonic and continuous from above and the equities of core banks are monotonic, which is the main ingredient to the proof of Theorem 5.

Proof of Lemma 2. As the main step of the proof, we show that for all $i \in C$, the assets $a_i(r)$ and the assets after default costs $a'_i(r)$ are monotonically increasing in r .

- *The assets a_i and a'_i are monotonic:* It suffices to show that the total incoming payments of bank i , $\sum_j p_{j,i}(r)$, are monotonically increasing in r . To this end, let

$$q_{k,i}(r) := r_k c_{k,i}^\emptyset + (1 - r_k) \sum_j r_j c_{j,i}^k.$$

Observe that $\sum_j p_{j,i}(r) = \sum_k q_{k,i}(r)$. Each individual summand $q_{k,i}(r)$ is monotonically increasing in r by the green core property, which can be seen as follows. Let $r \leq r'$ pointwise. Then

$$\begin{aligned} q_{k,i}(r') - q_{k,i}(r) &= r'_k c_{k,i}^\emptyset - r_k c_{k,i}^\emptyset + (1 - r'_k) \sum_j r'_j c_{j,i}^k \\ &\quad - (1 - r_k) \sum_j r_j c_{j,i}^k \\ &\geq r'_k c_{k,i}^\emptyset - r_k c_{k,i}^\emptyset + (1 - r'_k) \sum_j r_j c_{j,i}^k \\ &\quad - (1 - r_k) \sum_j r_j c_{j,i}^k \\ &= (r'_k - r_k) \cdot (c_{k,i}^\emptyset - \sum_j r_j c_{j,i}^k) \\ &\geq (r'_k - r_k) \cdot (c_{k,i}^\emptyset - \sum_j c_{j,i}^k) \geq 0, \end{aligned}$$

where the last inequality holds because $r'_k - r_k \geq 0$ by assumption and $c_{k,i}^\emptyset - \sum_j c_{j,i}^k \geq 0$ because we are in a green core system, so i must have a covered CDS position toward k .

- *The equity E_i is monotonic for $i \in C$:* First note that the liabilities $l_i(r)$ are monotonically decreasing in r in any financial system, as can be seen directly from the definition. As $a_i(r)$ is monotonically increasing by the above argument, $E_i(r) = \max(0, a_i(r)) - l_i(r)$ is monotonically increasing.

- *The function F is monotonic and continuous from above:* Because F_i is constant 1 for $i \in L$, it suffices to show the statement for each F_i with $i \in C$. To this end, note that F_i is of form

$$F_i(r) = \begin{cases} f(r) & \text{if } h(r) \geq 0, \\ g(r) & \text{if } h(r) < 0, \end{cases}$$

where $f(r) := 1$, $g(r) := a'_i(r)/l_i(r)$, and $h(r) := a_i(r) - l_i(r)$ are all monotonic and continuous. It is easy to see that this implies that F_i is monotonic and continuous from above. \square

The following lemma has become a standard proof technique in financial network theory, for example in Rogers and Veraart (2013). It can be viewed as a special case of the Kleene

or Tarski–Kantorovitch fixed-point theorems (see Granas and Dugundji 2003). We restate and prove it here because there is no standard reference for it.

Lemma 3. Let N be any finite set, and let $F : [0, 1]^N \rightarrow [0, 1]^N$ be any function that is monotonic and continuous from above, where the order relation is given by pointwise ordering. Then F has a pointwise maximal fixed point, and the iteration sequence (r^n) defined by $r^0 = (1, \dots, 1)$ and $r^{n+1} = F(r^n)$ converges to this maximal fixed point.

Proof. We proceed in three steps.

(i) *The sequence (r^n) is descending and convergent:* We show by induction that (r^n) is a descending sequence; that is, $r^n \geq r^{n+1}$ pointwise. For $n = 0$, this is trivial because r^0 is the maximal element of $[0, 1]^N$. For $n > 0$, and assuming $r^{n-1} \geq r^n$, we have $r^n = F(r^{n-1}) \geq F(r^n) = r^{n+1}$ by monotonicity of F . Because (r^n) is also bounded from below by $(0, \dots, 0)$, it must be convergent. Call the limit of the sequence r .

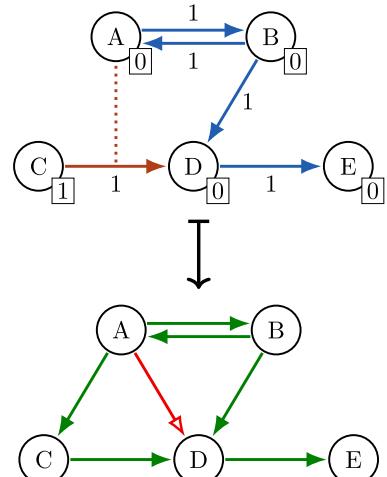
(ii) *The point r is greater or equal to any fixed point of F :* It suffices to show that any r^n is greater or equal to any fixed point r^* of F . We proceed by induction: for $n = 0$, the statement is obvious; for $n > 0$, and assuming $r^{n-1} \geq r^*$, we receive by monotonicity of F that $r^n = F(r^{n-1}) \geq F(r^*) = r^*$.

(iii) *The point r is a fixed point of F :* Because F is continuous from above and (r^n) is descending, we have $F(r) = F(\lim_n r^n) = \lim_n F(r^n) = \lim_n r^{n+1} = \lim_n r^n = r$. \square

Appendix D. Example That Iterating the Update Function Is Not Effective in the No-Red-Containing-Cycle Case

Consider Figure D.1 and let $\alpha = \beta = 0.5$. The cores of the dependency graph in topological order are $\{A, B\}$, $\{C\}$, $\{D\}$, and $\{E\}$. The unique solution of this system is given by $r_A = r_B = 0$ and $r_C = r_D = r_E = 1$. This is because C and E cannot default, and A and B must default with recovery rate 0 as they together have no assets but an outgoing liability; from this it follows that D has assets exactly equal to its liabilities.

Figure D.1. (Color online) Financial System Where Iterating the Update Function Does Not Converge



Notes. Top Panel: Financial system with no red-containing cycle where iterating F does not converge to a solution (let $\alpha = \beta = 0.5$). Bottom Panel: Its colored dependency graph.

Simply iterating the update function F does not converge in this system. To see this, let $r^0 = 1, \dots, 1$ and $r^{n+1} = F(r^n)$ for each n . We first consider the recovery rates of bank A and B as the iteration sequence proceeds. Note that B defaults in step 1, and then, following the default of B , A defaults in step 2. We thus have from the definition of F for r_A^n and r_B^n (recall that default costs are 0.5):

$$\begin{aligned} r_A^n &= \frac{1}{2}r_B^{n-1} \text{ for } n \geq 2 \text{ and } r_A^0 = r_A^1 = 1, \\ r_B^n &= \frac{1}{4}r_A^{n-1} \text{ for } n \geq 1 \text{ and } r_B^0 = 1. \end{aligned}$$

Solving these recursive equations yields $r_A^n = 2^{-3[n/2]}$ for all n and $r_B^n = 2^{-3[n/2]+1}$ for $n \geq 1$ and $r_B^0 = 1$. We observe that for $n \geq 2$, $r_A^n < r_B^n$ if n is even and $r_A^n > r_B^n$ if n is odd. This is because $\log(r_A^n) - \log(r_B^n) = -3[n/2] + 3[n/2] - 1 = 3[n/2] - [n/2] - 1$ and $[n/2] - [n/2]$ is 0 if n is even and 1 if n is odd.

Now consider bank D . The assets of D consist of a CDS on A and debt from B , so $a_D(r) = 1 - r_A + r_B$, and $l_D(r) = 1$. Thus, D is in default iff $r_A > r_B$. Over the course of the iteration, whenever n is even, we have $r_A^n < r_B^n$, so D is not in default and $r_D^{n+1} = F_D(r^n) = 1$. Whenever n is odd, we have $r_A^n > r_B^n$, so D is in default and $r_D^{n+1} = F_D(r^n) < \max(\alpha, \beta) = 0.5$. Hence, r_D^n changes by at least 0.5 from each iteration to the next. In particular, the sequence does not converge.

Our example may appear artificial because bank D is just on the verge of defaulting in the solution. We, indeed, expect that the iteration sequence converges if this is not the case for any bank. However, it is not clear how one would detect this property if the exact solution is not yet known.

Appendix E. Nonexistence with a CCP

We formalize and prove the fact that the presence of a highly capitalized central clearing counterparty (CCP) does not guarantee existence of a clearing recovery rate vector (see Section 6).

Proposition 3. *There exists a financial system (N, e, c, α, β) with a distinguished bank $S \in N$ (the CCP) such that the following holds:*

- (1) *The CCP is a counterparty to each contract: For any $k \in N \cup \{\emptyset\}$ and $i, j \in N$, if $c_{i,j}^k > 0$, then $S \in \{i, j\}$.*
- (2) *The CCP is running a balanced book: For any $k \in N \cup \{\emptyset\}$, $\sum_i c_{i,S}^k = \sum_j c_{S,j}^k$.*
- (3) *The CCP is so highly capitalized that it cannot default: $e_S \geq \sum_i c_{S,i}^\emptyset + \sum_{i,k} c_{S,i}^k$.*
- (4) *The financial system has no clearing recovery rate vector.*

Proof. Let $\beta < 1$. Consider a financial system with banks A , B , C , and S and contracts such as in Figure 8 together with an additional bank D (for offsetting positions). Choose e_D arbitrarily, and for any $k \in N \cup \{\emptyset\}$, let $c_{D,S}^k = \sum_{i \in N \setminus \{D\}} c_{S,i}^k$ and $c_{S,D}^k = \sum_{i \in N \setminus \{D\}} c_{i,S}^k$. The system is well defined because $1 - \beta + \beta^2 - \beta^3 = (1 - \beta) + \beta^2(1 - \beta) > 0$, so $\delta > 0$. It is easy to see that it fulfills conditions (1)–(3). To see that it also fulfills condition (4), assume toward a contradiction that $r \in [0, 1]^N$ is clearing. As S is highly capitalized, $r_S = 1$.

If $r_A = 1$, then $r_B = 0$; thus $r_C = 1$, and thus $r_A = 0$, as is easily seen from the contracts—a contradiction.

If $r_A < 1$, then $r_A \leq \beta$ by Lemma 1. Then $a_B(r) \geq 1 - \beta$, so $r_B \geq \beta(1 - \beta)$. Then $a_C(r) \leq 1 - \beta(1 - \beta)$, so $r_C \leq \beta(1 - \beta(1 - \beta))$.

Thus, $a_A(r) \geq \delta(1 - \beta(1 - \beta(1 - \beta))) = \delta(1 - \beta + \beta^2 - \beta^3) = 1 = l_A(r)$: a contradiction. \square

Endnotes

¹ As payments are made simultaneously and there is no timing, *default* and *technical insolvency* are equivalent conditions in our model. We use the term “default” throughout this paper.

² Note that the models in Eisenberg and Noe (2001) and Rogers and Veraart (2013) are based on the payments between banks instead of recovery rates. It is easy to see that the two points of view are equivalent. In debt-only financial systems, maximizing payments, recovery rates, and equities are equivalent objectives.

³ One could also include the interests of “society” (i.e., the real economy) in our analysis by introducing it as an additional node in the network. In Section 5.2 (Corollary 1), we derive sufficient conditions for maximality that guarantee that a solution is simultaneously best for all banks and society, assuming that banks’ defaults can only have a negative effect on society.

⁴ The total notional of these CDSs was USD1 trillion in the second half of 2017. See Bank for International Settlements (2018, section “Single-name instruments,” subsection “Financial firms”).

⁵ An orthogonal question is whether we can efficiently compute a solution to the clearing problem or determine algorithmically whether a solution exists. We have found in a separate stream of work (Schuldenzucker et al. 2019) where, in contrast to debt-only networks, both problems are computationally infeasible in general financial networks with CDSs. It is an open question to which extent algorithms exist that are guaranteed to be efficient in restricted cases.

⁶ See our discussion at the end of Section 2.3 for a comparison with prior approaches toward modeling these networks.

⁷ Following Rogers and Veraart (2013), we model default costs as multiplicative discounts $\alpha \in [0, 1]$ on the external assets and $\beta \in [0, 1]$ on the interbank assets that apply only in case of default. If $\alpha = \beta = 1$, no default costs are present, and we show using a fixed-point theorem that a solution always exists by continuity. By contrast, we show that if either $\alpha < 1$ or $\beta < 1$ (or both), default costs introduce a discontinuity, and existence is not guaranteed anymore.

⁸ A similar kind of contradiction was observed by Sundaresan and Wang (2015), who considered a setting with a single bank that issues *contingent capital*—that is, debt that is converted into equity as soon as the bank’s stock price falls below a threshold. For certain contract parameters, this led to nonexistence or multiplicity of an equilibrium stock price. Their situation bears some resemblance to a bank that has gone “short on itself” by buying a pathological CDS where it itself is the reference entity or “long on itself” by selling such a CDS. Sundaresan and Wang’s scenario thus concerns the balance sheet of an *individual* bank. By contrast, the ambiguity we illustrate is due to the *interactions* among different contracts in a network, although each individual bank may look innocuous.

⁹ The ECB’s recent STAMP\ euro framework, which was developed based on “top-down models used to support EU-wide stress-testing exercises” (Constâncio 2017), includes network effects as one of its central elements. Specifically, a variant of the clearing problem very close to Eisenberg and Noe (2001) is solved 20,000 times in the context of a Monte Carlo simulation to obtain a probability distribution of contagion losses (Dees et al. 2017, chapter 12).

¹⁰ For example, Fender et al. (2008) described how the default of Lehman Brothers, which was both a major counterparty and reference entity in CDSs, had significant repercussions in money markets. Further distress in these markets could only be averted by the government rescue of AIG, another major CDS trader.

¹¹ Default costs could result from legal and administrative costs, from a delay in payments, or from fire sales when defaulting banks need to

sell off their assets quickly. Details can be found in Rogers and Veraart (2013).

¹² Some prior work has employed graph-based models for CDS networks when simplifying assumptions reduce a CDS to a binary relationship. It was implicitly assumed that either the default of a reference entity is an event exogenous to the network (Duffie and Zhu 2011, Markose et al. 2012, Brunnermeier et al. 2013) or that CDSs do not carry counterparty risk (Puliga et al. 2014, Leduc et al. 2017). Heise and Kühn (2012) studied a model of CDS networks with ternary relationships and short positions. However, they only considered recovery rates 0 and 1 and a fixed number of update steps. We are the first to study simultaneous clearing in networks where CDSs are modeled as ternary relationships.

¹³ This is the case, for example, in any variant of Figure 2 where B has positive external assets. Here, B might first receive the full amount in the CDS and then pay A, or B might pay A first, which reduces the payment in the CDS.

¹⁴ Note that a maximal solution is not necessarily unique, but all maximal solutions lead to the same equities $E_i(r)$.

¹⁵ One could define a notion of *utility* for banks equal to their equity. Then a solution to the clearing problem would be maximal if and only if it Pareto dominates any other solution, and the situation in Theorem 3 would be one of multiple Pareto optima. Note, however, that the multiplicity we reveal here is different from a multiplicity of equilibria that is often observed in strategic games. Recovery rates are not actions chosen by banks. Rather, they are mathematically implied by the network of obligations and the rules of bankruptcy. Thus, instead of banks “choosing an equilibrium,” the clearing authority chooses a solution to be implemented, which implies a requirement to treat all market participants fairly (in particular, not to advantage one bank over another one). This is only possible if maximality is fulfilled.

¹⁶ The solution r^2 that exists only in the $\beta = 1$ case is strictly disfavored by A and B over all other solutions (but strictly preferred by C and D).

¹⁷ If there exists a solution, there also exists one that maximizes total equity (except in pathological cases). One might be tempted to simply select such a solution. However, without maximality, there would still be banks that prefer another solution. For *complete markets*, we know from the second welfare theorem that one can impose lump sum transfers to move from one (less desirable) Pareto-optimal outcome to any other (more desirable) one. By contrast, without maximality, all solutions of the clearing problem that are Pareto optimal pose the issue that the clearing authority would have to favor one bank over another.

¹⁸ It follows from the structure of F that in this situation, holding r_{-ij} fixed, F_j is increasing at *some* point r_i if and only if it is increasing at *all* points r_i .

¹⁹ Leduc et al. (2017) presented a mapping from a financial system with CDSs to a weighted graph where they distinguished between naked and covered CDSs in a similar way as we do. However, they made simplifying assumptions regarding the regulatory environment such that only a subset of the possible dependencies need to be considered. For example, their model does not represent default by CDS writers or short dependencies so that naked CDSs cannot be captured. This makes their model unsuitable to study general financial systems with naked and covered CDSs.

²⁰ See Schuldenzucker et al. (2019, appendix B) for a financial system with CDSs and a unique and irrational solution. Inspection shows that it is a green core system. The maximal solution of a debt-only system is rational because it can be computed exactly in finite time using Eisenberg and Noe’s (2001) *fictitious default algorithm*.

²¹ Both computing strongly connected components and sorting in topological order can be done easily using well-known algorithms (see, e.g., Korte and Vygen 2012). Note that the topological order may not be unique.

²² It is possible to show that the number of steps of the algorithm is bounded by a function that is polynomial with respect to the size of the financial system and $1/\epsilon$. This makes the core iteration algorithm a computationally efficient approximation scheme. See Schuldenzucker et al. (2019) for details.

²³ Both frameworks mandate use of a CCP for certain types of derivatives (interest rate swaps and index CDSs) but not for the kind of CDSs we discuss in this paper (single-name CDSs). See European Securities and Markets Authority (2017a, c) for EMIR and the documents linked at U.S. Commodity Futures Trading Commission (2017) for Dodd-Frank.

²⁴ For simplicity, we assume that all CDSs and debt contracts are cleared via the same CCP. This is not necessary for our result. We further assume that S has offsetting positions with other banks (i.e., it is running a balanced book). These positions do not affect our result and are therefore omitted from the figure.

²⁵ Figure A.1 corresponds to Figure 2 for $\beta = 0.5$; however, in Figure A.1, we set $\alpha = \beta = 1$. That is why, in contrast to Figure 2, this system has a solution.

References

- Acemoglu D, Ozdaglar A, Tahbaz-Salehi A (2015) Systemic risk and stability in financial networks. *Amer. Econom. Rev.* 105(2):564–608.
- Acemoglu D, Carvalho VM, Ozdaglar A, Tahbaz-Salehi A (2012) The network origins of aggregate fluctuations. *Econometrica* 80(5): 1977–2016.
- Allen F, Gale D (2000) Financial contagion. *J. Political Econom.* 108(1):1–33.
- Banerjee T, Feinstein Z (2018) Impact of contingent payments on systemic risk in financial networks. Working paper, Washington University in St. Louis, St. Louis.
- Bank for International Settlements (2018) Semiannual OTC derivatives statistics, credit default swaps, by sector of reference, H2 2017. Accessed March 26, 2019, <http://stats.bis.org/statx/srs/table/d10.4?p=20172>.
- Bardoscia M, Battiston S, Caccioli F, Caldarelli G (2017) Pathways towards instability in financial networks. *Nature Comm.* 8:Article 14416.
- Barucca P, Bardoscia M, Caccioli F, D’Errico M, Visentin G, Battiston S, Caldarelli G (2016) Network valuation in financial systems. Working paper, University of Zurich, Zurich.
- Battiston S, Caldarelli G, May RM, Roukny T, Stiglitz JE (2016) The price of complexity in financial networks. *Proc. Natl. Acad. Sci. USA* 113(36):10031–10036.
- Brunnermeier M, Clerc L, Scheicher M (2013) Assessing contagion risks in the CDS market. *Banque France Financial Stability Rev.* (17):123–134.
- Cifuentes R, Ferrucci G, Shin HS (2005) Liquidity risk and contagion. *J. Eur. Econom. Assoc.* 3(2–3):556–566.
- Constance V (2017) Macroprudential stress tests: A new analytical tool. Retrieved July 21, 2017, <http://voxeu.org/article/macroprudential-stress-tests-new-analytical-tool>.
- Csóka P, Herings PJ-J (2018) Decentralized clearing in financial networks. *Management Sci.* 64(10):4681–4699.
- Dees S, Henry J, Martin R, eds. (2017) STAMP\ euro: Stress-Test Analytics for Macroprudential Purposes in the Euro Area (European Central Bank, Frankfurt, Germany).
- Demange G (2016) Contagion in financial networks: A threat index. *Management Sci.* 64(2):955–970.
- D’Errico M, Battiston S, Peltonen T, Scheicher M (2018) How does risk flow in the credit default swap market? *J. Financial Stability* 35(April):53–74.
- Duffie D, Zhu H (2011) Does a central clearing counterparty reduce counterparty risk? *Rev. Asset Pricing Stud.* 1(1):74–95.

- Eisenberg L, Noe TH (2001) Systemic risk in financial systems. *Management Sci.* 47(2):236–249.
- Elliott M, Golub B, Jackson MO (2014) Financial networks and contagion. *Amer. Econom. Rev.* 104(10):3115–53.
- Elsinger H, Lehar A, Summer M (2006) Risk assessment for banking systems. *Management Sci.* 52(9):1301–1314.
- European Commission (2011) Regulation on short selling and credit default swaps—Frequently asked questions. Press release (October 19), http://europa.eu/rapid/press-release_MEMO-11-713_en.htm.
- European Securities and Markets Authority (2017a) OTC derivatives and clearing obligation. Retrieved July 21, 2017, <https://www.esma.europa.eu/regulation/post-trading/otc-derivatives-and-clearing-obligation>.
- European Securities and Markets Authority (2017b) Policy activities/short selling. Retrieved June 27, 2017, <https://www.esma.europa.eu/regulation/trading/short-selling>.
- European Securities and Markets Authority (2017c) Public register for the clearing obligation under EMIR. Retrieved July 21, 2017, https://www.esma.europa.eu/sites/default/files/library/public_register_for_the_clearing_obligation_under_emir.pdf.
- Fender I, Frankel A, Gyntelberg J (2008) Three market implications of the Lehman bankruptcy. *BIS Quart. Rev.* 1(1):6–7.
- Glasserman P, Young HP (2015) How likely is contagion in financial networks? *J. Banking Finance* 50(January):383–399.
- Granas A, Dugundji J (2003) *Fixed Point Theory* (Springer, New York).
- Haldane AG (2009) Rethinking the financial network. Speech, Financial Student Association, April 28, Amsterdam. Accessed March 26, 2019, <https://www.bankofengland.co.uk/speech/2009/rethinking-the-financial-network>.
- Heise S, Kühn R (2012) Derivatives and credit contagion in interconnected networks. *Eur. Phys. J. B* 85(4):1–19.
- Hu D, Zhao JL, Hua Z, Wong MC (2012) Network-based modeling and analysis of systemic risk in banking systems. *MIS Quart.* 36(4):1269–1291.
- International Swaps and Derivatives Association (2012) The ISDA Credit Derivatives Determinations Committees. Retrieved July 21, 2017, http://www2.isda.org/attachment/NDM1NA==AGM%202012_DC%20anniversary_appendix_043012.pdf.
- Kakutani S (1941) A generalization of Brouwer's fixed point theorem. *Duke Math. J.* 8(3):457–459.
- Korte B, Vygen J (2012) *Combinatorial Optimization: Theory and Algorithms*, 5th ed. (Springer, Berlin).
- Leduc MV, Poledna S, Thurner S (2017) Systemic risk management in financial networks with credit default swaps. *J. Network Theory Finance* 3(3):29–39.
- Loon YC, Zhong ZK (2014) The impact of central clearing on counterparty risk, liquidity, and trading: Evidence from the credit default swap market. *J. Financial Econom.* 112(1):91–115.
- Markose SM, Giansante S, Shaghaghi AR (2012) ‘Too interconnected to fail’ financial network of US CDS market: Topological fragility and systemic risk. *J. Econom. Behav. Organ.* 83(3):627–646.
- Puliga M, Caldarelli G, Battiston S (2014) Credit default swaps networks and systemic risk. *Sci. Rep.* 4:Article 6822.
- Rogers L, Veraart LA (2013) Failure and rescue in an interbank network. *Management Sci.* 59(4):882–898.
- Schuldenzucker S, Seuken S, Battiston S (2016) Clearing payments in financial networks with credit default swaps. (Extended abstract.) *Proc. 17th ACM Conf. Econom. Comput. (EC '16)* (ACM, New York), 759.
- Schuldenzucker S, Seuken S, Battiston S (2019) The computational complexity of financial networks with credit default swaps. Working paper, University of Zurich, Zurich.
- Single Resolution Board (2016) The single resolution mechanism: Introduction to resolution planning. Accessed July 21, 2017, <https://srbeuropa.eu/en/content/introduction-resolution-planning>.
- Stiglitz JE (2002) *Globalization and Its Discontents* (Penguin Books, New York).
- Stiglitz JE (2016) Climate change and financial complexity. Public lecture, January 18, University of Zurich, Zurich. Available at https://www.youtube.com/watch?v=J36werC_XpA.
- Sundaresan S, Wang Z (2015) On the design of contingent capital with a market trigger. *J. Finance* 70(2):881–920.
- U.S. Commodity Futures Trading Commission (2017) Clearing requirement. Accessed July 21, 2017, <http://www.cftc.gov/LawRegulation/DoddFrankAct/Rulemakings/ClearingRequirement/index.htm>.
- Vitali S, Glattfelder JB, Battiston S (2011) The network of global corporate control. *PLoS One* 6(10):e25995.

3 The Computational Complexity of Financial Networks with Credit Default Swaps

Oh, simple thing, where have you gone?
I'm getting old and I need something to rely on.
So tell me when you're gonna let me in.
I'm getting tired and I need somewhere to begin.

Keen, *Somewhere Only We Know*

The content of this chapter has previously appeared in:

“The Computational Complexity of Clearing Financial Networks with Credit Default Swaps.” Steffen Schulenzucker, Sven Seuken, and Stefano Battiston. Working Paper, May 2019.

See also: “Finding Clearing Payments in Financial Networks with Credit Default Swaps is PPAD-complete.” Steffen Schulenzucker, Sven Seuken, and Stefano Battiston. In *Proceedings of the 8th Innovations in Theoretical Computer Science (ITCS) Conference*, Berkeley, USA, January 2017.

See also: “Clearing Payments in Financial Networks with Credit Default Swaps.” Steffen Schulenzucker, Sven Seuken, and Stefano Battiston. Extended abstract in *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, Maastricht, The Netherlands, July 2016.

The Computational Complexity of Financial Networks with Credit Default Swaps*

Steffen Schuldenzucker
University of Zurich

Sven Seuken
University of Zurich

Stefano Battiston
University of Zurich

First version: October 5, 2017

This version: November 23, 2019

Abstract

The 2008 financial crisis has been attributed to “excessive complexity” of the financial system due to financial innovation. We employ computational complexity theory to make this notion precise. Specifically, we consider the problem of clearing a financial network after a shock. Prior work has shown that when banks can only enter into simple *debt contracts* with each other, then this problem can be solved in polynomial time. In contrast, if they can also enter into *credit default swaps (CDSs)*, i.e., financial derivative contracts that depend on the default of another bank, a solution may not even exist. In this work, we show that deciding if a solution exists is NP-complete if CDSs are allowed. This remains true if we relax the problem to ε -approximate solutions, for a constant ε . We further show that, under sufficient conditions where a solution is guaranteed to exist, the approximate search problem is PPAD-complete for constant ε . We then isolate the “origin” of the complexity. We show that already determining which banks default is hard. Further, we show that the complexity is not driven by the dependence of counterparties on each other, but rather hinges on the presence of so-called *naked* CDSs. If naked CDSs are not present, we receive a simple polynomial-time algorithm. Our results are of practical importance for regulators’ stress tests and regulatory policy.

1 Introduction

The year 2008 has provided a painful example of how moderate losses in a comparatively small financial market can spread and amplify in the financial system to create the worst economic crisis since the Great Depression. It has since become widely accepted that this was not just the result of financial institutions’ individual

*Steffen Schuldenzucker, Sven Seuken: Department of Informatics, University of Zurich, Switzerland, Email: {schuldenzucker,seuken}@ifi.uzh.ch. Stefano Battiston: Department of Banking and Finance, University of Zurich, and Swiss Finance Institute, Switzerland, Email: stefano.battiston@uzh.ch. Some of the ideas presented in this paper were previously described in a one-page abstract that was published in the conference proceedings of EC’16 (Schuldenzucker, Seuken and Battiston, 2016) and in a paper that was published in the conference proceedings of ITCS’17 (Schuldenzucker, Seuken and Battiston, 2017).

risk-taking, but a consequence of the overall architecture of the financial system at the time. Haldane (2009), then Executive Director of Financial Stability at the Bank of England, described the crisis as a manifestation of “the behaviour under stress of a complex, adaptive network,” in which “financial innovation [had] increased further network dimensionality, complexity, and uncertainty.” Yellen (2013), then Vice Chair of the Federal Reserve, described regulatory changes after the crisis that are explicitly targeted at “excessive systemic risk arising from the complexity and interconnectedness that characterize our financial system.”

The financial system has undergone rapid change since the 1990s. *Financial derivatives* (i.e., financial contracts where an obligation to pay depends on some other event or market variable) today allow to trade and reallocate individual components of risk. For example, an investment bank may bundle a loan in a foreign currency with a derivative that will pay the difference between the domestic and the foreign exchange rate. A domestic investor could then buy the bundle without having to worry about a devaluation of the foreign currency. Another trader would take on the other side of the derivative and thus the exchange rate risk. Of course, now both the bank and the investor depend on the trader meeting her obligation to pay. If the trader or the investor is another financial institution, the process can continue over any number of stages, with each party buying, rebundling, and reselling risk. A network of obligations arises: a graph where the nodes are financial institutions (“banks” for short) and the edges are financial contracts. We call this the *financial network*.

The above accounts by policymakers attribute the financial crisis to excessive “complexity” of the financial network. The question remains, though, how exactly we should understand the term “complexity” here. In particular, while a financial network could arise in a variety of ways,¹ most people share an intuition that financial networks with derivatives are “more complex” than ones without. In this paper, we will show that this informal notion materializes in the form of *computational* complexity of a concrete problem that regulators need to solve.

More in detail, we study the *clearing problem*. We are given a financial network consisting of banks and contracts between banks. Each contract defines an obligation to pay a certain amount of money under certain conditions. We assume that some of the banks experienced a shock on their assets, which may render them unable to meet their obligations towards other banks and force them into bankruptcy (or *default*). Defaults may trigger defaults of other banks downstream. For each bank,

¹Short-term loans between banks and *securization*, i.e., the pooling and re-selling of debt, are two other ways how a financial network can form. The products resulting from securization (most prominently *collateralized debt obligations (CDOs)*) are sometimes called “derivatives,” but in contrast to the kinds of derivatives we discuss in this paper, they are defined using a priority structure and do not depend on any market variable except for the debt they are based on. For the purpose of our discussion, we therefore consider these products a form of debt.

we are now looking for its *recovery rate*, i.e., the percentage of its liabilities the bank can pay to its creditors. These payments are made from its own (external) assets and the money it receives from other banks. Recovery rates must be in accordance with standard bankruptcy regulations, which imply a constraint reminiscent of a flow identity: defaulting banks must pay out all their assets to creditors and must do so in proportion to the respective obligation. Banks may further incur *default costs* and lose a percentage of their assets upon defaulting. The clearing problem is non-trivial because the contractual relationships can form cycles in the network.²

The clearing problem serves as a model for how a financial crisis will turn out following the initial shock. Once a solution to the clearing problem has been found, the effect of the initial shock can be judged by metrics like the number of defaulted banks or the total loss of money due to default costs. Researchers have used this approach to study the effect of network structure on systemic risk, such as financial diversification and integration (Elliott, Golub and Jackson, 2014), the interplay of shock size and network structure (Acemoglu, Ozdaglar and Tahbaz-Salehi, 2015), and to determine bounds on the extent of a crisis (Elsinger, Lehar and Summer, 2006; Glasserman and Young, 2015), among many others.³

The clearing problem has practical relevance in the context of *stress tests*, where regulators such as the European Central Bank (ECB) evaluate the stability of the financial system under an array of adverse economic scenarios. While today, the official stress tests still operate at a *microprudential* (individual bank) level, efforts are underway to transition to a *macroprudential* point of view, where the financial system is considered as a whole (Constâncio, 2017). These new stress tests need to take network effects into account. For example, the ECB's *STAMP€* framework (Dees, Henry and Martin, 2017) includes a model for the assessment of interbank contagion that is very close to the literature on clearing. The clearing problem is solved many thousands of times in the context of a Monte Carlo simulation. Therefore, these stress tests crucially depend on having access to efficient algorithms for the clearing problem.⁴

Eisenberg and Noe (2001) showed that, if banks can only enter into simple *debt*

²Perhaps contrary to intuition, cyclic structures are overwhelmingly common in real financial networks. See D'Errico et al. (2018), for example, for an empirical study of the network structure of credit default swap markets, which will be discussed below.

³While the exact rules of clearing vary across the literature, they all share as a common feature that payments (or, in some cases, contract valuations) happen simultaneously and some kind of input-output identity needs to hold at each bank. The essential properties of these models with respect to existence and computation of a solution are very similar. In this paper, unless specifically indicated, the term “clearing” refers to the model by Eisenberg and Noe (2001), discussed below, and its extensions.

⁴The size of the financial systems analyzed in stress tests must be expected to lie in the order of 100 banks, making brute-force approaches impractical. For example, Dees, Henry and Martin (2017, Chapter 12) consider 144 banks and the European Banking Authority (2014) stress tests covered 123 banks. Only few of these banks are trivial, i.e., sources or sinks in the network Dees, Henry and Martin (2017, Chart 12.1).

contracts with each other, then the clearing problem always has a solution and it can be computed in polynomial time. A debt contract is any contract where the obligation to pay is a fixed number. This could be just a loan from one bank to another, but it can also serve as a model for, e.g., a derivative when the obligation to pay only depends on variables that are external to the financial system and that can be assumed to be fixed for the purpose of clearing. Rogers and Veraart (2013) extended the result to default costs. These clearing models have seen widespread adoption in research and stress testing, such as in the above-mentioned *STAMP* framework. We argue, however, that it is necessary to consider extensions of the model, where the obligation to pay may not be fixed.

Specifically, we study financial networks that contain *credit default swaps (CDSs)* in addition to debt. A CDS is a financial derivative where the obligation to pay depends on the default of a third party, the *reference entity*. Market participants use CDSs to insure themselves against a default of the reference entity or to place a speculative bet on this event. CDSs have played a major role in the default of Lehman Brothers and the bailout of AIG during the 2008 crisis (Fender and Gyntelberg, 2008). Both firms were among the most important institutions in this market, both as counterparties and as reference entities (Fitch Ratings, 2007). It has hence become conventional wisdom that they were counterparties in significant amounts of CDSs where the respective other bank was the reference entity. Such CDSs on other banks cannot accurately be modeled as debt contracts because they depend on an event that is fundamentally endogenous to the financial system. At the same time, future stress tests cannot afford to neglect the dependencies implied by a Lehman–AIG type situation. That is why it is necessary to consider an extension of the existing debt-only to CDSs where reference entities can be other financial institutions.⁵

We consider such a model that extends the Rogers and Veraart (2013) clearing model to CDSs in a straightforward way. In our own recent work (Schuldenzucker, Seuken and Battiston, 2019), we have studied existence of a solution in this model. We have found that there are financial systems where the clearing problem has no solution. At the same time, the clearing algorithms for debt-only networks do not extend to CDSs even in cases where a solution is known to exist. This immediately raises two questions regarding the computational aspects of the clearing problem with CDSs:

1. Given a financial network, can we efficiently determine whether a solution to the clearing problem exists?
2. Given a financial network in which a solution is known to exist, can we efficiently

⁵The market for CDSs on financial firms alone currently has a size of about USD 900 billion. In the years following the 2008 crisis, this number was as high as USD 5 trillion. See Bank for International Settlements (2018, Section Single-name instruments, Subsection Financial firms) and the graph linked there.

compute it?

In this paper, we answer both questions in the negative. Towards the first question, we show that it is NP-complete to distinguish networks that have an (exact) solution from those that have no ε -approximate solution, for a natural approximate solution concept and sufficiently small constant ε . In particular, deciding existence of an exact solution or an ε -approximate solution is NP-hard (Section 3).

Towards the second question regarding the computation of a solution, we restrict our attention to the special case where banks do not incur default costs. Here, it is known that a solution always exists (Schuldenzucker, Seuken and Battiston, 2019), but the only known proof of this statement is non-constructive via a fixed-point theorem and so the question regarding computation has remained open so far. As exact solutions can be irrational, we need to consider an approximation problem. We show that the total search problem of finding an ε -approximate solution in a financial system without default costs is PPAD-complete if ε is a sufficiently small constant. Thus, no polynomial-time approximation scheme (PTAS) exists unless P=PPAD (Section 4).

At this point, we have shown that financial networks with CDSs are indeed “more complex” than those without. However, we should also be able to explain where this newfound complexity comes from. This is important to be able to inform future decisions on regulatory policy beyond an overly simple statement like “CDSs are problematic.”

In our quest for an “origin” of the computational complexity, we proceed in two steps. In a first step, we ask what aspect of a solution to the clearing problem is hard to compute. We show that hardness in the decision and search problems does not arise exclusively from the need to compute precise numerical values for the recovery rates. Instead, it is already NP-hard to decide if some given bank will default in some ε -solution (an appropriate distinction variant is NP-complete) and in the case without default costs, it is already PPAD-complete to find a set of banks that will default in some ε -solution (Section 5).

In a second step, we study restrictions on the network structure to discern what economic aspects of financial networks the computational complexity might originate from. It follows from our reductions that the problems are still hard in a model where *counterparty risk* (i.e., the dependence of a bank on its debtors) is neglected. Thus, we can say that the complexity originates from *fundamental risk* (i.e., the dependence of CDS counterparties on the reference entity). Finally, we obtain an upper bound on complexity. We show that hardness hinges on the presence of *naked CDSs*, i.e., CDSs

that are held without also holding a corresponding debt contract.⁶ If naked CDSs are not allowed, a solution always exists and we show that a simple iterative algorithm first presented in Schuldenzucker, Seuken and Battiston (2019) constitutes a fully polynomial-time approximation scheme (FPTAS; Section 6). These insights will allow us to frame a rather complete picture regarding the “origin of the complexity” and they have various implications for regulatory policy (Section 7).

Attempts to capture the “complexity” of the financial network have, of course, been made before. We differentiate between informal complexity due to i) the structure of interconnections and ii) the nature of the contracts themselves. Complexity due to the network structure has previously been approached using various measures from graph theory, such as the length of a path between ultimate borrowers and lenders (Shin, 2010), average degree (Gai, Haldane and Kapadia, 2011), network concentration (Arinaminpathy, Kapadia and May, 2012), network entropy (Battiston et al., 2016), or spectral measures (Bardoscia et al., 2017). As these measures require ordinary graphs as their inputs, where edges cannot contain more information than weights, they need to abstract over details of the contracts, such as the dependence of a financial derivative on its underlying market variable. Sensitivity results (Hemenway and Khanna, 2016; Liu and Staum, 2010; Feinstein et al., 2017) are another way to capture “complexity due to interconnectedness” and are also related to computational complexity. These results have so far only been obtained for networks of debt or cross-holdership. Basel III regulations measure a bank’s “interconnectedness” by the size of its intra-financial assets, while its “complexity” (of individual contracts) is measured by its amount of OTC⁷ derivatives, among others (Basel Committee on Banking Supervision, 2014).

The second kind of complexity, due to the nature of individual contracts, has begun to receive attention from theoretical computer science. Arora et al. (2011) and Zuckerman (2011) studied the cost of asymmetric information in financial derivatives markets with computationally bounded agents. Braverman and Pasricha (2014) showed that *compound options*⁸ are computationally hard to price correctly. These pieces of work study types of contracts that are “complex” even in isolation. In

⁶Naked CDSs are a common phenomenon in practice. While we are not aware of any empirical studies that quantify the share of CDSs that are naked, there seems to be a broad consensus that they form the majority of CDS positions. Kiff et al. (2009) noted that the (gross) notional of CDSs “continues to far exceed the stock of corporate bonds and loans on which most contracts are written.” Crotty (2009) quotes Eric Dinallo, then Superintendent of Insurance for New York State, saying that 80 percent of the CDSs outstanding are speculative (i.e., naked). Regulatory changes after 2009, like central clearing and portfolio compression (see Section 7), may have reduced the share of naked CDSs, but they cannot eliminate it below a significant level.

⁷Over-the-counter, i.e., traded directly with other banks rather than through an exchange. In this paper we only consider OTC derivatives.

⁸An option is a derivative that grants the holder the right to buy (*call option*) or sell (*put option*) an asset A at a specified time in the future for a previously agreed-upon price K . A compound option is an option where A is itself an option.

contrast, a single CDS is a very simple contract.⁹ Hence, we show in this paper that otherwise simple derivatives, if they occur as part of an otherwise simple network structure, create a financial system of high (computational) complexity.

The only other computational complexity result for financial networks we are aware of is by Hemenway and Khanna (2016), who studied the clearing model by Elliott, Golub and Jackson (2014). The authors showed that it is computationally hard to determine the distribution of a given total negative shock to the banks that does the worst damage in terms of value. In contrast, we prove in this work that in financial networks with CDSs, it is already hard to determine the impact of a *known* distribution of shocks to banks. To the best of our knowledge, we are the first to present a computational hardness result for the clearing problem.

Note that we do not claim that the “complexity” we reveal in this paper is necessarily what the regulatory authorities had in mind during and after the 2008 crisis. Neither do we claim that it was the most pressing issue or the cause of the crisis. There were many issues, such as counterparty risk, opacity, and lack of regulatory constraints for the OTC derivatives market (Financial Crisis Inquiry Commission, 2011). However, the “complexity” we reveal separates financial networks with derivatives from those without in a fundamental way. Today, new regulatory requirements have reduced counterparty risk and new reporting obligations have reduced opacity, while algorithmic tools to analyze the reported data are still under development. This is why we believe that the “complexity” we illustrate in this paper is even more relevant today.

Techniques Used

Since the clearing problem refers to an explicit network, it is natural for us to employ reduction from circuit problems to prove our hardness results.

To prove that deciding existence of a solution is NP-hard, we perform reduction from the CIRCUIT SATISFIABILITY problem. We encode Boolean circuits in a way reminiscent of electrical circuits. Boolean values are represented by recovery rates that are bounded away from $1/2$ by a constant and we define two *financial system gadgets*: one that allows a bank to have either a low or a high recovery rate, for the inputs, and one that implements a NAND operation, for the gates. We prevent accumulation of errors via a special *reset gadget* that maps low values to $0 \pm \varepsilon$ high values to $1 \pm \varepsilon$. Finally, we add a financial sub-network that has no solution *iff* the recovery rate of the “output bank” of the circuit is low.

We frame the decision problem as a *promise problem*: algorithms are only required to show any useful behavior on “clear-cut” instances where either an exact solution exists or not even an ε -approximate solution exists, for some small ε . For

⁹Valuation of a CDS is straightforward if distributions of recovery rates are known for the reference entity and counterparty. See Duffie (1999).

intermediate instances (*only* an approximate solution exists), any behavior including non-termination is acceptable. Promise problems are useful for problems where solutions may not be of polynomial length. For example, Schoenebeck and Vadhan (2012) used this approach in the context of certain classes of Nash equilibria. See Goldreich (2005) for a further discussion.

We show PPAD-hardness of the search problem via reduction from generalized circuits. Originally developed for the analysis of the complexity of finding a Nash equilibrium (Daskalakis, Goldberg and Papadimitriou, 2009; Chen, Deng and Teng, 2009; Rubinstein, 2018), generalized circuits have found application in the study of other total search problems. A generalized circuit consists of arithmetic gates and comparison gates,¹⁰ and it can have cycles. The associated search problem asks for a vector of values that is approximately consistent with each gate. Our reduction is straightforward: for each type of gate, we define a gadget that (approximately) performs the respective operation on the recovery rates. PPAD-hardness then follows from hardness of generalized circuits for constant ε (Rubinstein, 2018). Hardness for *constant* ε is the strongest kind of hardness result one can obtain here and precludes existence of a PTAS unless P=PPAD. If ε shrinks polynomially as the input grows, this only precludes an FPTAS. If it shrinks exponentially, it only precludes membership in P.

To show that already finding a set of defaulting banks is PPAD-hard, we define a new discrete variant of the generalized circuit problem, which may be of independent interest. In this problem, we only ask for one of three *states* for each gate: *high*, *medium*, or *low*. These states correspond to “decision” or “truncation points” in the definition of the gates. For example, the addition gate is in a *high* state if and only if its inputs sum to more than one and its output is therefore truncated at one. States also allow for ε errors. It is PPAD-complete to find a collection of states consistent with some ε -solution of the circuit because with states fixed, the constraints on the gates are linear and one can reconstruct an ε -solution via linear programming. We then show that in our above reduction from financial networks to generalized circuits, the set of defaulting banks already determines the states of the gates. We hope that our technique may be useful to prove hardness of discrete versions of other search problems in the future.

2 Preliminaries

We now describe the formal model for simultaneous clearing in financial networks with CDSs (Schuldenzucker, Seuken and Battiston, 2019). We present a new relaxation

¹⁰Traditionally, generalized circuits have also supported Boolean gates that operate on approximately Boolean values similar to above. Schuldenzucker and Seuken (2019) recently showed that the Boolean gates are in fact redundant and can therefore be omitted.

of the solution concept, which is necessary to be able to receive an (approximate) solution of finite, polynomial length. We then discuss how the addition of CDSs changes the mathematical properties of the model compared to debt-only financial networks.

2.1 Basic Notation

Throughout this paper, we say that ε is *sufficiently small* (in symbols: $\varepsilon \ll 1$) if it is below a certain positive threshold, where the exact value of the threshold is not relevant in the following. The threshold may depend on parameters that are arbitrary, but fixed, but it will *never* depend on the input to any computational problem and should therefore treated as a constant. We sometimes write $\varepsilon \ll \beta$ to indicate that the threshold is a monotonic (not necessarily linear) function of a term β . We write $\Theta(\varepsilon)$ for $c\varepsilon$ where $c > 0$ is a certain constant that is not relevant in the following and does not depend on any values or parameters. We define $[x] := \min(1, \max(0, x))$, the *truncation* of x to the interval $[0, 1]$. We write $x = y \pm \varepsilon$ for $|x - y| \leq \varepsilon$ if x and y are scalars and for $\|x - y\|_\infty \leq \varepsilon$ if they are vectors. We also use the notation “ $\pm \varepsilon$ ” in compound expression like $[x \pm \varepsilon]$ to indicate a range of values. This notation formally corresponds to interval arithmetic.

2.2 Financial Systems and Clearing Recovery Rates¹¹

Financial System. Let N denote a finite set of *banks*. Each bank $i \in N$ holds a certain amount of *external assets*, denoted by $e_i \geq 0$. Between any two banks i and j , $|N| + 1$ numbers capture the contracts from the contract *writer* i to the *holder* j . Let $c_{i,j}^\emptyset \geq 0$ be the total *notional amount* of debt that i owes to j and for $k \in N$ let $c_{i,j}^k \geq 0$ be the total notional amount of CDSs from i to j with *reference entity* k . If $c_{i,j}^k > 0$ for some $k \in N \cup \{\emptyset\}$, we call j a *creditor* of i and j a *debtor* of i .

We make two *sanity assumptions* to rule out pathological cases. First, no bank may enter into a contract with itself and no bank may enter into a CDS on itself (i.e., $c_{i,i}^\emptyset = c_{i,i}^j = c_{i,j}^i = c_{i,j}^i = 0$ for all $i, j \in N$). Second, as CDSs are defined as insurance on debt, we require that any bank that is a reference entity in a CDS must also be a writer of some debt contract (i.e., if $\sum_{k,l \in N} c_{k,l}^i > 0$, then $\sum_{j \in N} c_{i,j}^\emptyset > 0$, for all $i \in N$).¹²

We model default costs following Rogers and Veraart (2013): there are two *default cost parameters* $\alpha, \beta \in [0, 1]$. Defaulting banks are only able to pay to their creditors a share of α of their external assets and a share of β of their incoming payments.

¹¹Parts of this sub-section until and excluding the definition of approximate solutions have previously appeared in our prior work (Schuldenzucker, Seuken and Battiston, 2019, Section 2). Also see the aforementioned paper for a discussion of alternative models for financial networks with CDSs.

¹²For technical reasons, we will allow our *financial system gadgets* in Sections 3 and 4 to violate the second assumption. In this case, the violating banks will be “dummy banks” that hold and write no contracts and are ignored when considering solutions.

Thus, $\alpha = \beta = 1$ means that there are no default costs and $\alpha = \beta = 0$ means that assets held by defaulting banks are worthless. The values $1 - \alpha$ and $1 - \beta$ are the *default costs*.¹³

A *financial system* is a tuple (N, e, c, α, β) where N is a set of banks, e is a vector of external assets, c is a 3-dimensional matrix of contracts, and α and β are default cost parameters. Note that, even though α and β are part of the definition of a financial system, our results in this paper will be for restrictions of the respective problems to arbitrary but fixed values of α and β . The parameters can therefore be considered constant.

Note that when banks enter into a contract, there is typically an initial payment. In addition, the values of banks' assets are usually assumed to be subject to random fluctuation. We do not model these complications explicitly but rather, we assume that they are implicitly reflected in the external assets.

Assets and Liabilities. We are ultimately looking for a vector of *recovery rates* $r_i \in [0, 1]$. For any two banks i and j , the contracts from i to j give rise to a *liability* from i to j . This is the amount of money that i has to pay to j . A debt contract gives rise to an unconditional liability equal to its notional, while the liability in a CDS with reference entity k depends on the recovery rate of k and is proportional to $1 - r_k$. The total liability from i to j at r is therefore:

$$l_{i,j}(r) := c_{i,j}^\emptyset + \sum_{k \in N} (1 - r_k) \cdot c_{i,j}^k$$

The *total liabilities* of i at r are the aggregate liabilities that i has toward all other banks, denoted by

$$l_i(r) := \sum_{j \in N} l_{i,j}(r).$$

The actual *payment* $p_{i,j}(r)$ from i to j at r can be lower than $l_{i,j}(r)$ if i is in default. A bank that is in default makes payments in its contracts in proportion to the respective liability;

$$p_{i,j}(r) := r_i \cdot l_{i,j}(r).$$

The *total assets* $a_i(r)$ of a bank i at r consist of its external assets e_i and the incoming payments;

$$a_i(r) := e_i + \sum_{j \in N} p_{j,i}(r).$$

In case bank i is in default, its *assets after default costs* $a'_i(r)$ are the assets reduced

¹³Default costs could arise from a variety of sources, such as legal costs, operational losses, or from a discount in prices when assets need to be sold off quickly during bankruptcy. See Rogers and Veraart (2013) for details.

according to the factors α and β . This is the amount that will be paid out to creditors;

$$a'_i(r) := \alpha e_i + \beta \sum_{j \in N} p_{j,i}(r).$$

Remark 1. To see that CDSs indeed act as insurance on default, let $i, j, k \in N$ and assume that bank j holds both debt from k and a CDS on k , both with the same notional: $c_{i,j}^k = c_{k,j}^\emptyset =: \delta > 0$. Then the assets of j contain the term $r_k \cdot \delta + r_i \cdot (1 - r_k) \cdot \delta$. As long as $r_i = 1$, this term is equal to δ independently of r_k ; thus, i is insured against k 's default.

Clearing Recovery Rate Vector. Following Eisenberg and Noe (2001), we call a recovery rate vector $r \in [0, 1]^N$ *clearing* if it satisfies the essential principles of bankruptcy law:

1. Banks with sufficient assets to pay their liabilities in full must do so.
2. Banks with insufficient assets to pay their liabilities in full are in default and must pay out all their assets to creditors after default costs have been subtracted.

This leads to the following formal definition:

Definition 1 (Clearing Recovery Rate Vector). Let $X = (N, e, c, \alpha, \beta)$ be a financial system. Define the *update function*

$$F : [0, 1]^N \rightarrow [0, 1]^N$$

$$F_i(r) := \begin{cases} 1 & \text{if } a_i(r) \geq l_i(r) \\ \frac{a'_i(r)}{l_i(r)} & \text{if } a_i(r) < l_i(r). \end{cases}$$

A recovery rate vector $r \in [0, 1]^N$ is called *clearing* for X if it is a fixed point of the update function, i.e., if $F_i(r) = r_i$ for all i . We also call a clearing recovery rate vector a *solution* to the clearing problem.

Approximate Solutions. There exist financial systems with CDSs where all solutions contain irrational numbers (see Appendix B for an example). To receive finite (polynomial-length) objects for our computational considerations in the present paper, we relax Definition 1 to an approximate solution as follows.

Definition 2 (Approximately Clearing Recovery Rate Vector). Let $X = (N, e, c, \alpha, \beta)$ be a financial system and let $\varepsilon \geq 0$. A recovery rate vector r is called ε -approximately *clearing* (or an ε -solution) for X if for each $i \in N$ at least one of the following two

conditions is satisfied:

$$\begin{aligned} r_i &= 1 \pm \varepsilon \text{ and } a_i(r) \geq (1 - \varepsilon)l_i(r) \\ r_i &= \frac{a'_i(r)}{l_i(r)} \pm \varepsilon \text{ and } a_i(r) < (1 + \varepsilon)l_i(r) \end{aligned}$$

Note how both the “case selection part” and the “output part” of the definition of the function F in Definition 1 are relaxed. Banks with much higher assets than liabilities are unambiguously not in default ($r_i = 1 \pm \varepsilon$) and those with much lower assets than liabilities are in default ($r_i = a'_i(r)/l_i(r) \pm \varepsilon$). But when assets approximately equal liabilities, either of the two states are possible. This is an appropriate model for the real world, where default is not a knife-edge decision, but has some tolerance.¹⁴ Therefore, the above definition likely reflects what (say) a regulator running a stress test will be interested in. Recall that the precision ε is defined in the space of *recovery rates*. That is why ε is used as an additive error in the output recovery rate, but a multiplicative one when comparing assets and liabilities.

We take the following elementary properties as evidence that our approximate solution concept is natural. The proof is straightforward and hence omitted.

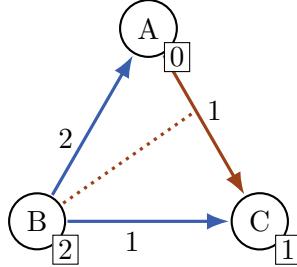
Proposition 1. *Let $X = (N, e, c, \alpha, \beta)$ be a financial system.*

1. *An exact solution is the same as a 0-solution. If $\varepsilon' > \varepsilon \geq 0$, then any ε -solution is also an ε' -solution.*
2. *If $F(r) = r \pm \varepsilon$, then r is an ε -solution.*
3. *If $\alpha = \beta = 1$ (i.e., there are no default costs), then $F(r) = r \pm \varepsilon$ if and only if r is an ε -solution.*
4. *If r is an ε -solution, then we have:*

$$\begin{aligned} a_i(r) &\geq (1 + \varepsilon)l_i(r) \Rightarrow r_i = 1 \pm \varepsilon \\ a_i(r) &< (1 - \varepsilon)l_i(r) \Rightarrow r_i = \frac{a'_i(r)}{l_i(r)} \pm \varepsilon \end{aligned}$$

5. *If r is an ε -solution and $l_i(r) > 0$, then $r_i \leq \frac{a_i(r)}{l_i(r)} + \varepsilon$.*
6. *If r is an ε -solution and $r_i < 1 - \varepsilon$, then $r_i \leq \max(\alpha, \beta) + \varepsilon$. If in addition $e_i = 0$, then $r_i \leq \beta + \varepsilon$.*

Figure 1 Example financial system. Let $\alpha = \beta = 0.5$



2.3 Example and Visual Representation¹⁵

Figure 1 shows a visual representation of an example financial system. There are three banks $N = \{A, B, C\}$, drawn as circles, with external assets of $e_A = 0$, $e_B = 2$, and $e_C = 1$, drawn as rectangles on top of the banks. Debt contracts are drawn as blue arrows from the writer to the holder and they are annotated with the notional values $c_{B,A}^\emptyset = 2$ and $c_{B,C}^\emptyset = 1$. CDSs are drawn as orange arrows, where a dashed line connects to the reference entity, and are also annotated with the notional value: $c_{A,C}^B = 1$. Default cost parameters $\alpha = \beta = 0.5$ are given in addition to the picture.

A clearing recovery rate vector for this example is given by $r_A = 1$, $r_B = \frac{1}{3}$, and $r_C = 1$. The liabilities arising from this recovery rate vector are $l_{B,A}(r) = 2$, $l_{B,C}(r) = 1$, and $l_{A,C}(r) = \frac{2}{3}$. Payments are $p_{B,A}(r) = \frac{2}{3}$, $p_{B,C} = \frac{1}{3}$, and $p_{A,C}(r) = \frac{2}{3}$. This is the only solution for this system.

Let now $\varepsilon = 0.1$. The ε -approximate solutions are exactly the recovery rate vector r that can be chosen according to the following process. Note that $a_B(r) < (1 - \varepsilon)l_B(r) \forall r$ and thus choose $r_B \in \left[\frac{1}{3} - 0.1, \frac{1}{3} + 0.1\right] \approx [0.23, 0.43]$ arbitrary. Let $r_C \in [0.9, 1]$ arbitrary. To choose r_A , perform case distinction over the chosen value of r_B :

- If $r_B \geq \frac{1+\varepsilon}{3+\varepsilon} \approx 0.35$, then $a_A(r) \geq (1 + \varepsilon)l_A(r)$ and thus choose $r_A \in [1 - \varepsilon, 1] = [0.9, 1]$ arbitrary.
- If $r_B < \frac{1-\varepsilon}{3-\varepsilon} \approx 0.31$, then $a_A(r) < (1 - \varepsilon)l_A(r)$ and thus choose $r_A = \frac{a'_A(r)}{l_A(r)} \pm \varepsilon = \frac{0.5r_B}{2(1-r_B)} \pm 0.1 = \frac{1}{4} \cdot \frac{r_B}{1-r_B} \pm 0.1$ arbitrary. Note that in this case, r_A can take on exactly the values in $\left[\frac{0.5(1/3-0.1)}{1/3-0.1} - 0.1, \frac{0.5(1/3+0.1)}{1/3+0.1} + 0.1\right] = [0.4, 0.6]$.
- If neither of the two above conditions applies, choose r_A according to either of the two above rules.

Note that some of the ε -approximate solutions are close to the exact solution while

¹⁴In the real world, edge cases are decided by *determinations committees*, which are panels consisting of representatives of financial institutions (specifically, the most active dealers and non-dealer members). See International Swaps and Derivatives Association (2012).

¹⁵This example has been adopted and extended from our prior work (Schuldenzucker, Seuken and Battiston, 2019).

others are not. We will discuss this phenomenon in Section 3.

2.4 Effects of Allowing CDSs

If there are no CDSs (i.e., $c_{i,j}^k = 0 \forall k \in N$), then a solution to the clearing problem always exists (Eisenberg and Noe, 2001; Rogers and Veraart, 2013). In this case, the assets a_i of each bank i are linear and monotonic functions of r and the liabilities are constant. This implies that the update function F is piecewise linear and monotonic, where the linear segments are given by regions where the set of defaulting banks does not change. Eisenberg and Noe (2001) exploited this for their *fictitious default algorithm*: we keep track of a candidate set of defaulting banks, beginning with the empty set. In each step, we solve a linear equation system to compute clearing recovery rates assuming this set of banks defaults. We then update the set of defaulting banks. By monotonicity, this set can only grow over time and we terminate at a solution after at most $|N|$ steps. In particular, a rational solution of polynomial length always exists. If there are default costs, a discontinuity appears at the boundary of the default regions, but, as Rogers and Veraart (2013) have shown, the algorithm still works.

The fictitious default algorithm does not extend to CDSs. This is because the assets of a bank i can then contain terms of form $c_{j,i}^k \cdot r_j \cdot (1 - r_k)$ and can thus be *non-linear* and *non-monotonic*, and so can the update function F . Since F is no longer linear (and also not convex or concave), the individual steps of the fictitious default algorithm cannot easily be performed in polynomial time. Further, by non-monotonicity, the set of defaulting banks need not grow monotonically over time and thus the algorithm would not necessarily terminate.¹⁶

If there are default costs and CDSs are allowed, a solution to the clearing problem need not even exist (Schuldenzucker, Seuken and Battiston, 2019). The intuition for this is that if a bank holds a CDS on itself, then this can lead to a situation where a default of this bank implies non-default and vice versa. In a network, an equivalent situation can arise indirectly. The discontinuity due to default costs implies that a “middle ground” is not attainable.

In two important special cases, existence of a solution is guaranteed. First, if there are no default costs ($\alpha = \beta = 1$), continuity is restored¹⁷ and one can show using a fixed-point theorem that a solution always exists. It has not been studied how computationally hard it is to actually *find* (an approximation of) this solution. The second special case is when *naked CDSs* (see Section 3) are not allowed and where *monotonicity* is restored. In this case, a solution can be described as the

¹⁶Cycling behavior is easy to construct, see for instance Schuldenzucker, Seuken and Battiston (2019, Appendix A).

¹⁷The function F can still contain a discontinuity at the boundaries of the sets $\{r \mid l_i(r) = 0\}$. This is easy to circumvent, though.

limit of an iteration sequence. No easy-to-check condition is known that is *equivalent* to the existence of a solution to the clearing problem. The question regarding the computational complexity of this decision problem has thus remained open.

3 The Complexity of Deciding Existence of a Solution

In this section, we prove our first main result: It is hard to distinguish between financial systems that have an exact solution and those that have no ε -solution. To make this problem well-defined, we make the following additional technical assumption.

Definition 3 (Non-degenerate Financial System). A financial system $X = (N, e, c, \alpha, \beta)$ is called *non-degenerate* if every bank that writes any contracts also writes a debt contract. That is, for all $i \in N$, if $\sum_{j,k \in N} c_{i,j}^k > 0$, then $\sum_{j \in N} c_{i,j}^\emptyset > 0$.

Note that non-degeneracy is a very weak requirement in the real world. All we demand is that every bank has *some* constant liabilities, for example to its customers. The following lemma shows that non-degenerate financial systems are sufficiently “smooth” to ensure that ε -solutions are not too long.

Lemma 1. *Let $\varepsilon > 0$ and let X be a non-degenerate financial system. If X has an exact solution, then X has an ε -solution of size polynomial in the sizes of ε and X .*

The proof can be found in Appendix A and works by rounding. Note that, since the update function F is not usually continuous, this does not necessarily lead to an approximate fixed point. Instead, we make use of the fact that we relaxed both sides of the definition of F in our definition of an ε -solution. Note further that the lemma does not imply that *all* ε -solutions are close to an exact solution (if one exists) or have polynomial length. This is common for approximate solution concepts.¹⁸

We can now state our first main result in the language of promise problems.

Theorem 1. *For any fixed α and β such that $\alpha < 1$ or $\beta < 1$ and for ε sufficiently small depending on α and β , the promise problem $\varepsilon\text{-HASCLEARING}^{\alpha,\beta}$, defined as follows, is NP-complete: given a non-degenerate financial system $X = (N, e, c, \alpha, \beta)$,*

- *if X has an exact solution, return YES.*
- *if X has no ε -solution, return NO.*

¹⁸Using a theorem by Anderson (1986), it follows from the syntactic structure of the definition that for any X and δ there is an ε such that any ε -solution for X is δ -close to an r that is *almost* an exact solution: $r_i = F_i(r)$ unless $a_i(r) = l_i(r)$, in which case we can have $r_i = 1$ or $r_i = a'_i(r)/l_i(r) = a'_i(r)/a_i(r)$. However, ε depends on X in this case, while we consider a constant ε in this paper. It is easy to construct examples where for constant ε , exact and approximate solutions are far apart. Search problems for *strong* approximate fixed points, which are defined as being close to an exact fixed point, have previously been studied by Etessami and Yannakakis (2010). Computational complexity is markedly higher than for the regular approximate variants. For example, it is an open question if finding a strong approximate Nash equilibrium is even in NP.

Note that if $\alpha = \beta = 1$, then we know from prior work that an exact solution always exists and so $\varepsilon\text{-HASCLEARING}^{\alpha,\beta}$ is trivial (see Section 2.4). For general α and β , membership in NP follows immediately via Lemma 1:

Proof of Theorem 1, membership. Given X and ε , let, via Lemma 1, $L = \text{poly}(\text{size}(X), \text{size}(\varepsilon))$ be such that if an exact solution exists, then so does an ε -solution of size $\leq L$. Check if an ε -solution of size $\leq L$ exists; return YES if so and NO otherwise. This is in NP via complete enumeration. If now there is an exact solution, by Lemma 1 there is also an ε -solution of length $\leq L$ and we correctly return YES. If there is no ε -solution, we correctly return NO. \square

Remark 2 (A Non-Promise Variant of the Decision Problem). If none of the two conditions in Theorem 1 is satisfied, any behavior including non-termination is allowed. One might argue that it would be more natural to instead consider the problem “Given X , return YES if an ε -solution exists and NO otherwise.” Theorem 1 implies that this problem is NP-hard for $0 \leq \varepsilon \ll 1$, but it does not imply that it is a member of NP. This is because Lemma 1 does not imply that there is *always* an ε -solution of polynomial length if an ε -solution exists, but only when an exact solution exists.¹⁹ Note further that $\varepsilon\text{-HASCLEARING}^{\alpha,\beta}$ becomes (weakly) easier as we increase ε , while this is not clear for the non-promise variant. Framing the problem as a promise problem avoids these issues.

The remainder of this section is dedicated to showing NP-hardness of the $\varepsilon\text{-HASCLEARING}^{\alpha,\beta}$ problem. Our reduction is from CIRCUIT SATISFIABILITY. We represent Boolean values by recovery rates that are contained in the set $[0, 1/4] \cup [3/4, 1]$, with the low part of this set representing FALSE and the high part representing TRUE. We then encode the inputs and the gates using *financial system gadgets* and we force the output to TRUE by adding another special financial sub-system. We prevent error accumulation using a special *reset gadget* reminiscent of the *brittle comparison* gadget in Daskalakis, Goldberg and Papadimitriou (2009).

3.1 Financial System Gadgets

We first introduce two technical tools that we will use for the proofs in this section as well as in Section 4. We define a *financial system gadget* as a small financial system where the recovery rate of an *output bank* depends on a collection of *input banks* in a certain way. We can “apply” a gadget into another financial system by identifying its input banks with some other banks, thereby extending the existing financial system with new banks. We will use this to successively build up financial systems.

¹⁹The proof of the lemma implies that an ε -solution of length polynomial in the length of $\varepsilon - \varepsilon_0$ exists if an ε_0 -solution exists, for some $0 \leq \varepsilon_0 < \varepsilon$. To derive membership in NP for the non-promise variant, we would need to bound $\inf \{\varepsilon_0 \mid \text{an } \varepsilon_0\text{-solution exists}\}$ from above relative to the length of X . We leave it to future work to explore to which extent this might be possible.

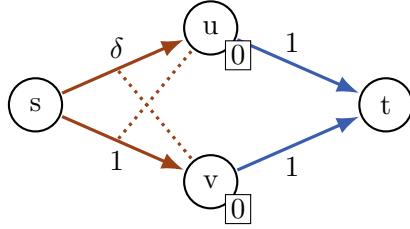
Definition 4 (Financial System Extension). An *extension* of a financial system $X = (N, e, c, \alpha, \beta)$ is a financial system $X' = (N', e', c', \alpha, \beta)$ such that $N \subseteq N'$ and the assets and liabilities of each bank $i \in N$ are the same in X and X' . That is, we have i) $e'_i = e_i \forall i \in N$ and ii) if $c_{i,j}^k > 0$ and one of i, j is in N , then *both* $i, j \in N$, $k \in N \cup \{\emptyset\}$, and $c_{i,j}^k = c_{i,j}^k$. We call X' an *extension* of X *on* $N_0 \subseteq N$ if this holds whenever one of i, j is in N_0 . We call $r \in [0, 1]^N$ an ε -solution *on* N_0 if the condition from Definition 2 holds for all $i \in N_0$.

Definition 5 (Financial System Gadget). A *financial system gadget* is a financial system $G = (N, e, c, \alpha, \beta)$ with a set of distinguished *input banks* $A := \{a_1, \dots, a_m\} \subseteq N$ that have no assets or liabilities (i.e., $e_{a_i} = c_{a_i,j}^k = c_{j,a_i}^k = 0$ for $i = 1, \dots, m$, $j \in N$, and $k \in N \cup \{\emptyset\}$) such that the following property holds: for any $r_A \in [0, 1]^A$ there exists an $r_{N \setminus A} \in [0, 1]^{N \setminus A}$ such that $r_A \cup r_{N \setminus A}$ is an exact solution on $N \setminus A$. We say that G *implements* a property $P : [0, 1]^N \rightarrow \{\text{TRUE}, \text{FALSE}\}$ if for any sufficiently small ε and any ε -solution r of G on $N \setminus A$, $P(r)$ holds. If X is a financial system and a'_1, \dots, a'_m are banks in X that each write some debt contract, then the *application* of G to X and a'_1, \dots, a'_m is a new financial system X' obtained as the union of X and G where we identify a_i and a'_i for $i = 1, \dots, m$. Note that X' is an extension of X and an extension of G on $N \setminus A$.

Remark 3 (Applying Several Gadgets). The properties implemented by gadgets are preserved under further extensions of the resulting financial system, in particular by other gadgets. To see this, first note that if X' is an extension of X on N_0 and r' is an ε -solution for X' , then the restriction $r'|_N$ is an ε -solution for X on N_0 . Assume now that a gadget $G = (N, e, c, \alpha, \beta)$ with inputs A implements a property P and X'' is an extension of G on $N \setminus A$. X'' could result, for example, from applying G to some financial system and then applying arbitrary other gadgets on top of it. Let $\varepsilon \ll 1$ and let r be an ε -solution for X'' . Then, by the extension, $r|_N$ is an ε -solution for G on $N \setminus A$ and thus, $P(r|_N)$ holds.

Our financial system gadgets will have between zero and two input banks, which we will call a and b for convenience. They will also have an *output bank*, which we will call v , and they will implement properties that make the recovery rate of the output bank equal to a certain function of the recovery rates of the input banks, up to errors. Each gadget will contain a *source bank* s that holds no contracts and a *sink bank* t that writes no contracts. The other banks hold CDSs from s and write a debt contract of notional 1 to t , so that $l_i(r) = 1$ for each of these banks i . The connection to the inputs is established via CDS references. We set $c_{s,t}^\emptyset = 1$ to ensure non-degeneracy and we further always set $e_s \geq 2 \sum_{i \in N, k \in N \cup \{\emptyset\}} c_{s,i}^k$. This implies that s cannot default and thus $r_s, r_t = 1 \pm \varepsilon$ in any ε -solution. For the sake of conciseness, we will leave out these contracts and external assets in our descriptions of the gadgets. We mark the input banks via dashed circles in our figures.

Figure 2 Zero-One Gadget. Variant of Schuldenzucker, Seuken and Battiston (2019, Figure 3).



3.2 Reducing Boolean Circuits to Financial Systems

We begin with a financial system gadget where the output bank can have recovery rates approximately 0 or 1. We will use this to encode variables.

Lemma 2 (Zero-One Gadget). *For all $\alpha, \beta \in [0, 1]$ there is a financial system gadget with no input banks such that the following hold:*

- There exists an exact solution where $r_v = 0$ and there exists an exact solution where $r_v = 1$.
- The gadget implements the following property: $r_v = 0 \pm \Theta(\varepsilon)$ or $r_v = 1 \pm \Theta(\varepsilon)$.

Proof. Consider the financial system in Figure 2. We distinguish the cases $\beta < 1$ and $\beta = 1$.

If $\beta < 1$, let $\delta = 2\frac{1}{1-\beta}$. It is easy to see that $(r_u, r_v) \in \{(0, 1), (1, 0)\}$ are exact solutions. We show that in any ε -solution r we have $r_v = 1 \pm \varepsilon$ or $r_v = 0 \pm 2\varepsilon$. To see this, let $r_v < 1 - \varepsilon$. Then $r_v \leq \beta + \varepsilon$ by Proposition 1 and thus $a_u(r) = \delta r_s(1 - r_v) \geq \delta(1 - \varepsilon)(1 - \beta - \varepsilon) = 2(1 - \varepsilon)(1 - \frac{\varepsilon}{1-\beta}) \geq 1 + \varepsilon = (1 + \varepsilon)l_u(r)$ where the last inequality holds for $\varepsilon \ll 1 - \beta$. Thus, $r_u = 1 \pm \varepsilon$ and thus $a_v(r) \leq \varepsilon$, so $r_v \leq \frac{a_v(r)}{l_v(r)} + \varepsilon = 2\varepsilon$.

If $\beta = 1$, let $\delta = 2$. It is again easy to see that $(r_u, r_v) \in \{(0, 1), (1, 0)\}$ are exact solutions. Let r be an ε -solution. For $i = u, v$ we have $a_i(r) = a'_i(r) \forall r$. This follows from the definition of a'_i because $\beta = 1$ and $e_i = 0$. Like in Proposition 1 part 3, this implies $r_i = F_i(r) \pm \varepsilon = [a_i(r)] \pm \varepsilon$. That is:

$$\begin{aligned} r_v &= [(1 \pm \varepsilon)(1 - r_u)] \pm \varepsilon = 1 - r_u \pm 2\varepsilon \\ r_u &= [2(1 \pm \varepsilon)(1 - r_v)] \pm \varepsilon = [2(1 - r_v)] \pm 3\varepsilon \end{aligned}$$

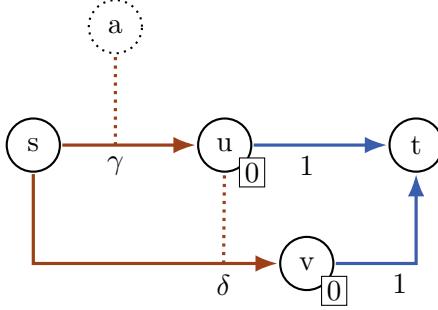
Taken together, these imply:

$$r_v = 1 - [2(1 - r_v)] \pm 5\varepsilon = [1 - 2(1 - r_v)] \pm 5\varepsilon = [2r_v - 1] \pm 5\varepsilon \quad (*)$$

We now perform a case distinction on r_v .

- If $r_v \geq 1/2$, then $[2r_v - 1] = 2r_v - 1$ and thus by $(*)$, $r_v = 1 \pm 5\varepsilon$.

Figure 3 Cutoff Gadget



- If $r_v < 1/2$, then $[2r_v - 1] = 0$ and thus by (*), $r_v = 0 \pm 5\varepsilon$. \square

We next work towards a NAND gadget to encode Boolean gates. We begin by introducing a versatile tool that can be used to map values significantly above or below certain thresholds to approximately 0 or 1. We will use this *cutoff gadget* in this section and Section 4.

Lemma 3 (Cutoff Gadget). *Let $0 < K < L < 1$. There exists a financial system gadget with one input bank that implements the following property for all $\alpha, \beta \in [0, 1]$:*

$$\begin{aligned} r_a \leq K - \Theta(\varepsilon) &\Rightarrow r_v = 0 \pm \Theta\left(\frac{\varepsilon}{L-K}\right) \\ r_a \geq L + \Theta(\varepsilon) &\Rightarrow r_v = 1 \pm \Theta\left(\frac{\varepsilon}{L-K}\right). \end{aligned}$$

Proof. Consider the gadget in Figure 3 where we set:²⁰

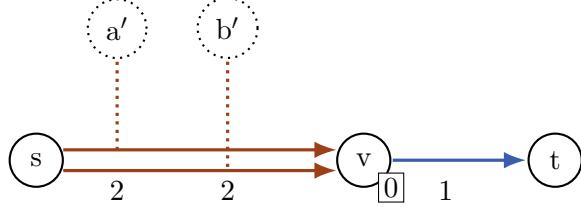
$$\begin{aligned} \gamma &:= \frac{1}{1-K} \\ \delta &:= \frac{1-K}{L-K} \end{aligned}$$

Consider an ε -solution. Let first $r_a \leq K - 3\varepsilon$. Then $a_u(r) = \gamma r_s(1 - r_a) \geq \gamma(1 - \varepsilon)(1 - K + 3\varepsilon) = (1 - \varepsilon)(1 + \frac{3\varepsilon}{1-K}) \geq (1 - \varepsilon)(1 + 3\varepsilon) = 1 + 2\varepsilon - 3\varepsilon^2 \geq 1 + \varepsilon$ if $\varepsilon \ll 1$. Thus, $r_u = 1 \pm \varepsilon$. Now $a_v(r) = \delta r_s(1 - r_u) \leq \delta\varepsilon \leq \frac{2\varepsilon}{L-K}$. And $r_v \leq \frac{a_v(r)}{l_v(r)} + \varepsilon = a_v(r) + \varepsilon \leq a_v(r) + \frac{3\varepsilon}{L-K}$.

Let now $r_a \geq L + 4\varepsilon$. Then $a_u(r) \leq \gamma(1 - L - 4\varepsilon) = \frac{1-L}{1-K} - \frac{4\varepsilon}{1-K}$ and thus $r_u \leq \frac{1-L}{1-K} - \frac{4\varepsilon}{1-K} + \varepsilon$. Now $a_v(r) \geq \delta(1 - \varepsilon)\left(1 - \frac{1-L}{1-K} + \frac{4\varepsilon}{1-K} - \varepsilon\right) = (1 - \varepsilon)\left(1 + \frac{(3+K)\varepsilon}{L-K}\right) \geq (1 - \varepsilon)(1 + 3\varepsilon) = 1 + 2\varepsilon - 3\varepsilon^2 \geq 1 + \varepsilon$ for $\varepsilon \ll 1$. Thus, $r_v = 1 \pm \varepsilon$. \square

²⁰Note that the gadget violates our sanity assumptions because the input bank is a CDS reference entity, but not writers of any debt contracts. In this case, this does not cause any problems. Note in particular that, as soon as the gadget is applied to some other system, the sanity assumption will hold.

Figure 4 NAND Gadget



The cutoff gadget has two kinds of errors. First, there is always a region $r_a \in (K - \Theta(\varepsilon), L + \Theta(\varepsilon)) \supset (K, L)$ where the output value is unspecified and depends on error terms more than on the input. This is reminiscent of the *brittle comparison* gate in generalized circuits (see Daskalakis, Goldberg and Papadimitriou (2009) and Section 4 in the present paper). Second, there is an error in the output that becomes larger the less “brittleness” we are willing to tolerate. This trade-off in errors is not fundamental and we will use in Section 4 an extension of our construction where both errors are small. In this section, we use a large brittleness to receive a gadget that will prevent error accumulation in the Boolean gadgets we will introduce afterwards.

Lemma 4 (Reset Gadget). *There exists a financial system gadget with one input bank that implements the following property for all $\alpha, \beta \in [0, 1]$:*

$$\begin{aligned} r_a \leq 1/4 &\Rightarrow r_v = 0 \pm \Theta(\varepsilon) \\ r_a \geq 3/4 &\Rightarrow r_v = 1 \pm \Theta(\varepsilon). \end{aligned}$$

Proof. Use a cutoff gadget with $K = 2/5$ and $L = 3/5$. For $\varepsilon \ll 1$ we have $1/4 \leq K - \Theta(\varepsilon)$ and $3/4 \geq L + \Theta(\varepsilon)$ and $L - K = 1/5$ is a constant. \square

To represent Boolean gates, we introduce a gadget that mirrors the Boolean operation $x \text{ NAND } y = \neg(x \wedge y)$.

Lemma 5 (NAND Gadget). *There exists a financial system gadget with two input banks that implements the following property for all $\alpha, \beta \in [0, 1]$:*

$$\begin{aligned} r_a \leq 1/4 \text{ or } r_b \leq 1/4 &\Rightarrow r_v = 1 \pm \Theta(\varepsilon) \\ r_a \geq 3/4 \text{ and } r_b \geq 3/4 &\Rightarrow r_v = 0 \pm \Theta(\varepsilon) \end{aligned}$$

Proof. Apply the reset gadget each to a and b and call the output banks a' and b' , respectively. Then apply the gadget in Figure 4.

Assume first that $r_a \leq 1/4$ or $r_b \leq 1/4$. Then $r_{a'}, r_{b'} = 0 \pm \Theta(\varepsilon)$ and thus $a_v(r) \geq 2(1-\varepsilon)(1-\Theta(\varepsilon)) \geq 1+\varepsilon$ if $\varepsilon \ll 1$. Thus, $r_v = 1 \pm \varepsilon$. Assume next that $r_a \geq 3/4$ and $r_b \geq 3/4$. Then $r_{a'}, r_{b'} = 1 \pm \Theta(\varepsilon)$ and thus $r_v \leq a_v(r) + \varepsilon \leq 4\Theta(\varepsilon) + \varepsilon = \Theta(\varepsilon)$. \square

Using the previous lemma, we easily construct gadgets for all Boolean functions. Note in particular that we can chain NAND gadgets without having to worry about error accumulation because $r_v = 1 \pm \Theta(\varepsilon) \Rightarrow r_v \geq 3/4$ if $\varepsilon \ll 1$, and likewise for 0. A chain of NAND gadgets will thus produce the appropriate output if the inputs to the chain are in $[0, 1/4] \cup [3/4, 1]$. This is why we use a reset gadget. We can combine a collection of zero-one and NAND gadgets to represent a Boolean circuit.

Proposition 2 (Financial Boolean Circuit). *Let C be a Boolean circuit with m inputs. For $\chi \in \{0, 1\}^m$ write $C(\chi) \in \{0, 1\}$ for the value of the output of C given values χ at the inputs. For any $\alpha, \beta \in [0, 1]$ and ε sufficiently small there exists a financial system $X = (N, e, c, \alpha, \beta)$ with $m + 1$ distinguished banks $V := \{a_1, \dots, a_m, v\}$ such that the following hold:*

1. *For any assignment $\chi \in \{0, 1\}^m$ there exists an exact solution r such that $r_{a_i} = \chi_i$ for $i = 1, \dots, m$.*
2. *If r is an ε -solution, then $r_i = 0 \pm \Theta(\varepsilon)$ or $r_i = 1 \pm \Theta(\varepsilon)$ for all $i \in V$.*
3. *If r is an ε -solution and $i \in V$, let $\chi_i = 0$ if $r_i \leq 1/4$ and $\chi_i = 1$ if $r_i \geq 3/4$. Then $\chi_v = C(\chi_{a_1}, \dots, \chi_{a_m})$.*

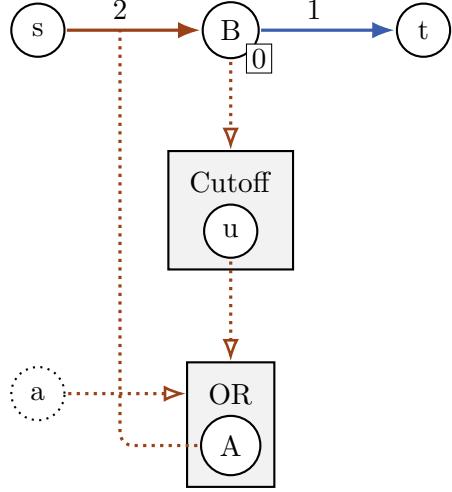
Proof. Assume WLOG that C consists only of NAND gates. We will identify the nodes of C with certain banks in the to-be-constructed financial system. We begin our construction with an empty financial system and build it up iteratively. First apply the zero-one gadget (Lemma 2) m times and identify the output banks of these gadgets with the input nodes of C . Now iterate over the gates of C in topological order. For each NAND gate connecting two inputs to an output, by the topological order, the inputs are already nodes in the financial system and the output is not. Apply the NAND gadget (Lemma 5) to the inputs and identify the output bank with the output node of the gate.

Property 1 is satisfied by the zero-one gadget, as is Property 2 for $i = a_1, \dots, a_m$. Property 2 for $i = v$ and property 3 follow by induction on the number of gates by the NAND gadgets. \square

3.3 Reducing Satisfiability to Existence of a Solution

The final step in our construction is a way to “destroy” certain unwanted ε -solutions. We use this to remove exactly those solutions that correspond to falsifying assignments of the Boolean circuit in our previous construction, leaving only those corresponding to satisfying assignments, if any. Note that this cannot be done using a financial system gadget because it does not preserve all existing solutions. The approach is otherwise exactly the same, though.

Figure 5 High-level structure of the financial system in Lemma 6, case $\beta < 1$. Gray boxes indicate gadgets with their output banks. A dashed line with a hollow arrow tip connects a bank to a gadget of which it is an input bank. The parameters of the cutoff gadget are chosen such that $\beta < K < \frac{\beta+1}{2} < L < 1$ evenly spaced.



Lemma 6 (Removing Solutions). *Let α, β be such that $\alpha < 1$ or $\beta < 1$. There exists a financial system $G = (N, e, c, \alpha, \beta)$ with a distinguished input bank $a \in N$ with no assets or liabilities such that the following hold:*

1. *For any $r_v \geq 3/4$, there exists an $r_{N \setminus \{a\}} \in [0, 1]^{N \setminus \{a\}}$ such that $r_a \cup r_{N \setminus \{a\}}$ is an exact solution on $N \setminus \{a\}$.*
2. *For $\varepsilon \ll 1$, there is no ε -solution r on $N \setminus \{a\}$ where $r_a \leq 1/4$.*

Proof. We distinguish the cases $\beta < 1$ and $\alpha < \beta = 1$.

If $\beta < 1$, perform the construction outlined in Figure 5: assume we have a source and a sink bank as usual. Add a new bank B and let $e_B = 0$ and $c_{B,t}^\emptyset = 1$. Apply a cutoff gadget to B with $K = \frac{3\beta+1}{4}$ and $L = \frac{\beta+3}{4}$. Note that $\beta < K < \frac{\beta+1}{2} < L < 1$ evenly spaced and the output error of the cutoff gadget is $\Theta(\frac{\varepsilon}{L-K}) = \Theta(\frac{\varepsilon}{1-\beta})$. Call the output bank of the cutoff gadget u , apply an OR gadget to a and u and call the output A . Finally, add CDS $c_{s,B}^A = 2$.

Towards property 1, if $r_a \geq 3/4$, then by the OR gadget (and this by the NAND gadget), we can set $r_A = 1$ independently of r_u . We can then extend r to the other banks via $r_B = 0$ and by setting the recovery rates for the intermediate nodes of the gadgets accordingly.

Towards property 2, if $r_a \leq 1/4$, assume towards a contradiction that r is an ε -solution on $N \setminus \{a\}$. We perform case distinction on r_B .

- If $r_B \geq 1 - \varepsilon$, then in particular $r_B \geq L + \Theta(\frac{\varepsilon}{1-\beta})$ if $\varepsilon \ll 1 - \beta$.²¹ Thus, by the

²¹The threshold for ε is $\Theta((1 - \beta)^2)$. The fact that this is nonlinear in $1 - \beta$ is not a problem. Recall that we did *not* assume that $\varepsilon = \Theta(1 - \beta)$ if $\varepsilon \ll 1 - \beta$.

cutoff gadget and the OR gadget, $r_A = 1 \pm \Theta(\varepsilon)$. But then $r_B \leq a_B(r) + \varepsilon \leq 2\Theta(\varepsilon) + \varepsilon < 1 - \varepsilon$ for $\varepsilon \ll 1$. Contradiction.

- If $r_B < 1 - \varepsilon$, then $r_B \leq \beta + \varepsilon \leq K - \Theta(\frac{\varepsilon}{1-\beta})$, where the first inequality is by Proposition 1 and the second inequality holds for $\varepsilon \ll 1 - \beta$. Therefore, $r_u = 0 \pm \Theta(\varepsilon)$ and since $r_v \leq 1/4$ we have $r_a = 0 \pm \Theta(\varepsilon)$. This implies $a_B(r) \geq 2(1 - \varepsilon)(1 - \Theta(\varepsilon)) \geq 1 + \varepsilon$ for $\varepsilon \ll 1$. This implies $r_B \geq 1 - \varepsilon$. Contradiction.

If $\alpha < \beta = 1$, let $e_B = c_{s,B}^A = 4/5$ and $K = \frac{3\alpha+1}{4}$ and $L = \frac{\alpha+3}{4}$, and keep everything else the same as above. Note that $\alpha < K < \frac{\alpha+1}{2} < L < 1$ evenly spaced and the output error of the cutoff gadget is $\Theta(\frac{\varepsilon}{1-\alpha})$. It is clear that property 1 follows just like above. We show that property 2 follows in a similar way to above. Assume that $r_a \leq 1/4$.

- If $r_B \geq 1 - \varepsilon$, then this implies $r_B \geq L + \Theta(\varepsilon)$ for $\varepsilon \ll 1 - \alpha$, so $r_A = 1 \pm \Theta(\varepsilon)$ and thus $r_B \leq a_B(r) + \varepsilon \leq 4/5 + 4/5 \cdot \Theta(\varepsilon) + \varepsilon < 1 - \varepsilon$ for $\varepsilon \ll 1$. Contradiction.
- If $r_B < 1 - \varepsilon$, we must have $a_B(r) < 1 + \varepsilon$ and $r_B = a'_B(r) \pm \varepsilon \leq a_B(r) - (1 - \alpha)e_B + \varepsilon < 1 - (1 - \alpha)e_B + 2\varepsilon$. The middle inequality is by definition of a'_B in case $\beta = 1$. We further receive

$$\begin{aligned} 1 - (1 - \alpha)e_B + 2\varepsilon &= 1 - \frac{4}{5}(1 - \alpha) + 2\varepsilon \\ &= \frac{3\alpha+1}{4} - \frac{1}{20}(1 - \alpha) - 2\varepsilon \\ &= K - \frac{1}{20}(1 - \alpha) - 2\varepsilon \leq K - \Theta\left(\frac{\varepsilon}{1 - \alpha}\right) \end{aligned}$$

where the last line holds for $\varepsilon \ll 1 - \alpha$. This implies $r_A = 0 \pm \Theta(\varepsilon)$ and $a_B(r) \geq 4/5 + 4/5(1 - \varepsilon)(1 - \Theta(\varepsilon)) \geq 1 + \varepsilon$ for $\varepsilon \ll 1$ and this implies $r_B \geq 1 - \varepsilon$. Contradiction. \square

Theorem 1 now follows by application of Lemma 6 to our financial Boolean circuit.

Proof of Theorem 1, hardness. Reduction from CIRCUIT SATISFIABILITY. Given a Boolean circuit C , apply Proposition 2 to construct the financial system X with output bank v corresponding to C . Apply the system from Lemma 6 to X and v (where by ‘‘application’’ we mean the same like for gadgets) to construct an extended system X' . Let $\varepsilon \ll 1$ and solve ε -HASCLEARING $^{\alpha,\beta}$ for X' .

If C has a satisfiable assignment χ , this yields an exact solution for χ where $r_v = 1 \geq 3/4$, so we can extend this to an exact solution of X' . If C has no satisfying assignment, then any solution to X satisfies $r_v = 0 \pm \Theta(\varepsilon)$ and therefore, no ε -solution for X' exists. As these are the only two cases, any algorithm for ε -HASCLEARING $^{\alpha,\beta}$ must terminate on this instance and return YES if C is satisfiable and NO otherwise. \square

4 The Complexity of the Search Problem without Default Costs

We now focus on financial systems without default costs, i.e., where $\alpha = \beta = 1$. In these systems, we know that a solution always exists:

Theorem (Schuldenzucker, Seuken and Battiston, 2019, Theorem 2). *Any financial system $(N, e, c, \alpha = 1, \beta = 1)$ has an exact solution.*

The proof of the above theorem is by Kakutani's fixed point theorem and thus not constructive. In this section, we study the associated search problem. Since it may still be the case that all solutions are irrational (see Appendix B), we study the associated approximation problem of computing an ε -solution.

Theorem 2. *For $\varepsilon \ll 1$, the total search problem ε -FINDCLEARING, defined as follows, is PPAD-complete: Given a non-degenerate financial system $X = (N, e, c, \alpha = 1, \beta = 1)$, compute an ε -solution.*

The theorem immediately implies that no polynomial-time approximation scheme (PTAS) exists, unless P=PPAD.

ε -FINDCLEARING is a well-defined total search problem because, by Lemma 1 and existence of a solution, there is always an ε -solution of polynomial length. Recall from Proposition 1 that for $\alpha = \beta = 1$, an ε -solution is the same as an ε -approximate fixed point of the update function, i.e., an $r \in [0, 1]^N$ such that $F(r) = r \pm \varepsilon$. Note further that, since $\alpha = \beta = 1$ and we assume non-degeneracy, F simplifies to:

$$F_i(r) = \begin{cases} 1 & \text{if } c_{i,j}^k = 0 \forall j \in N, k \in N \cup \{\emptyset\} \\ \left[\frac{a_i(r)}{l_i(r)} \right] & \text{otherwise.} \end{cases}$$

Recall that $[x] := \min(1, \max(0, x))$ is the truncation of x to $[0, 1]$. Membership in PPAD easily follows.

Proof of Theorem 2, membership. The proof is similar to the proof of Lemma 1. By the above considerations, F is polynomially continuous. And finding a (Brouwer) fixed point of a polynomially continuous function is in PPAD (Papadimitriou, 1994). \square

The remainder of this section is dedicated to showing PPAD-hardness via a reduction from generalized circuits.

4.1 Generalized Circuits

A generalized circuit (Chen, Deng and Teng, 2009) consists of nodes interconnected by arithmetic or comparison gates. In contrast to regular arithmetic or Boolean circuits, generalized circuits may contain cycles, which turns finding a consistent assignment

Figure 6 Conditions that should hold at a gate g for an ε -solution x of a generalized circuit. Assume that the inputs of g are called a and b (if any) and the output is called v . The gates C_ζ , $C_{\times\zeta}$, and $C_{>\zeta}$ take an additional parameter $\zeta \in [0, 1]$.

$$\begin{aligned}
g = C_\zeta &\Rightarrow x[v] = \zeta \pm \varepsilon \\
g = C_+ &\Rightarrow x[v] = [x[a] + x[b]] \pm \varepsilon \\
g = C_- &\Rightarrow x[v] = [x[a] - x[b]] \pm \varepsilon \\
g = C_{\times\zeta} &\Rightarrow x[v] = \zeta \cdot x[a] \pm \varepsilon \\
g = C_{>\zeta} &\Rightarrow x[a] < \zeta - \varepsilon \Rightarrow x[v] = 0 \pm \varepsilon \\
&\quad x[a] > \zeta + \varepsilon \Rightarrow x[v] = 1 \pm \varepsilon
\end{aligned}$$

of node values into a non-trivial fixed point problem. Rubinstein (2018) introduced a variant of generalized circuits that is well-suited for our purposes. To make our reduction to financial systems as simple as possible, we consider a reduced set of gates, which does not change the computational complexity of the problem (see Appendix C for a detailed comparison).

Definition 6 (Generalized Circuit and Approximate Solution). A *generalized circuit* is a collection of *nodes* and *gates*, where each node is labeled *input* of any number of gates (including zero) and *output* of at most one gate. Inputs to the same gate are distinguishable from each other. Each gate has one of the following types: C_ζ (constant, no inputs), $C_{\times\zeta}$ (scaling, one input), C_+ or C_- (addition and subtraction, two inputs), or $C_{>\zeta}$ (comparison to a constant, one input). For the gate types C_ζ , $C_{\times\zeta}$, and $C_{>\zeta}$, a numeric parameter $\zeta \in [0, 1]$ is specified in addition to the input and output nodes of the gate. The *length* of a generalized circuit is the number of bits needed to describe the circuit, including the nodes, the mapping from nodes to inputs and outputs of gates, and numeric parameters ζ involved.

For $\varepsilon \geq 0$, an ε -solution of a generalized circuit is a mapping x that assigns to each node v a value $x[v] \in [0, 1]$ such that the constraints in Figure 6 hold at each gate of type g with inputs a and b (if any) and output v .

We know from prior work that finding an ε -solution is hard:

Theorem (Essentially Rubinstein (2018)). *For a (constant) sufficiently small ε , the total search problem ε -GCIRCUIT, defined as follows, is PPAD-complete: Given a generalized circuit, find an ε -solution.*

4.2 Reducing Generalized Circuits to Financial Systems

We now show how to encode a generalized circuit into a financial system via financial system gadgets corresponding to the five gate types. Any ε -solution to the financial system will give rise to a $\Theta(\varepsilon)$ -solution of the generalized circuit. Compared to

Figure 7 Constant Gadget

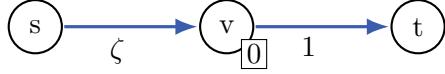
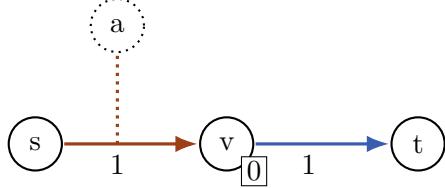


Figure 8 Inverter Gadget



Section 3, our gadgets need to be more precise because we do not only have to map between the appropriate parts of the set $[0, 1/4] \cup [3/4, 1]$ to each other, but we have to encode arithmetic operations on continuous inputs in $[0, 1]$. The fact that we do not have default costs in this section will help us achieve this higher precision. In our gadgets, all banks except for the source and sink banks will have liabilities constant 1. We therefore have $r_i = F_i(r) \pm \varepsilon = [a_i(r)] \pm \varepsilon$ in any ε -solution.

Our simplest gadget establishes a constant recovery rate at the output bank:

Lemma 7 (Constant Gadget). *Let $\zeta \in [0, 1]$. If $\alpha = \beta = 1$, there is a financial system gadget with no input banks that implements the property $r_v = \zeta \pm \Theta(\varepsilon)$.*

Proof. Consider Figure 7. Clearly, $a_v(r) = \zeta(1 \pm \varepsilon) = \zeta \pm \varepsilon$ and $r_v = [a_v(r)] \pm \varepsilon = \zeta \pm 2\varepsilon$. \square

An important building block for the following constructions is a gadget that “inverts” the recovery rate of a bank.

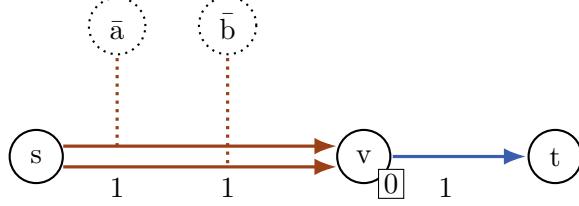
Lemma 8 (Inverter Gadget). *If $\alpha = \beta = 1$, there is a financial system gadget with one input bank that implements the property $r_v = 1 - r_a \pm \Theta(\varepsilon)$.*

Proof. Consider Figure 8. Clearly, $a_v(r) = (1 \pm \varepsilon)(1 - r_a) = 1 - r_a \pm \varepsilon$ and $r_v = [a_v(r)] \pm \varepsilon = [1 - r_a] \pm 2\varepsilon = 1 - r_a \pm 2\varepsilon$. \square

Note that we could *not* have used the inverter gadget as a Boolean NOT gadget in Section 3 because i) it relies on the assumption $\alpha = \beta = 1$ and ii) it accumulates errors, i.e., $2n$ inverters in a row yield $r_a \pm \Theta(n\varepsilon)$, not $r_a \pm \varepsilon$. We proceed with the addition and subtraction gadgets, which are slight variants of each other.

Lemma 9 (Sum Gadget). *If $\alpha = \beta = 1$, there is a financial system gadget with two input banks that implements the property $r_v = [r_a + r_b] \pm \Theta(\varepsilon)$.*

Figure 9 Sum Gadget. \bar{a} and \bar{b} are the outputs of inverters applied to a and b , respectively.



Proof. Apply inverter gadgets (Lemma 8) to both a and b and call the output banks \bar{a} and \bar{b} , respectively. Now apply the gadget in Figure 9. Then:

$$r_v = [(1 \pm \varepsilon)(1 - r_{\bar{a}}) + (1 \pm \varepsilon)(1 - r_{\bar{b}})] \pm \varepsilon = [r_a + r_b] \pm 3\varepsilon \quad \square$$

Note the similarity of Figure 9 with Figure 4, which was used in the construction of the NAND gadget in Section 3. Indeed, a similar operation is performed given that $a \text{ NAND } b = \neg a \vee \neg b$ and addition is somewhat similar to Boolean “ \vee ”. Note however that these two constructions need to deal with different challenges: the NAND gadget needs to work with default costs and the sum gadget needs to provide a correct sum across all input values, not just approximately Boolean values.

Lemma 10 (Difference Gadget). *There is a financial system gadget without default costs with two input banks that implements the property $r_v = [r_a - r_b] \pm \Theta(\varepsilon)$.*

Proof. Apply an inverter gadget (Lemma 8) to a and call the output bank \bar{a} . Apply the gadget in Figure 9 to \bar{a} and $\bar{b} := b$ and call the output bank u . From the proof of the previous lemma we know that $r_u = [1 - r_a + r_b] \pm \Theta(\varepsilon)$. Now apply an inverter to u and call the output bank v . It follows that

$$r_v = 1 - [1 - r_a + r_b] \pm \Theta(\varepsilon) = [r_a - r_b] \pm \Theta(\varepsilon),$$

where the last equality follows by case distinction. \square

Our last two gadgets are scaling and comparison. Scaling is easily achieved by noting that a chain of two inverters approximately copies the input. We then adjust the notional in one of those gadgets. This happens to be a degenerate variant of the cutoff gadget (Lemma 3).

Lemma 11 (Scaling Gadget). *Let $\zeta \in [0, 1]$. If $\alpha = \beta = 1$, there exists a financial system gadget that implements the property $r_v = \zeta r_a \pm \Theta(\varepsilon)$.*

Proof. We use the financial system in Figure 3 from Section 3 with $\gamma = 1$ and $\delta = \zeta$. We have $r_u = 1 - r_a \pm \Theta(\varepsilon)$ like in the inverter gadget and thus $a_v(r) = \zeta(1 \pm \varepsilon)(1 -$

$(1 - r_a \pm \Theta(\varepsilon)) = \zeta(1 \pm \varepsilon)(r_a \pm \Theta(\varepsilon)) = \zeta r_a \pm \Theta(\varepsilon)$. And $r_v = [a_v(r)] \pm \varepsilon = \zeta r_a \pm \Theta(\varepsilon)$. Here we use $\zeta, r_a \leq 1$ to bound the error. \square

For the comparison gate, the cutoff gadget (Lemma 3) almost does what we want. However, the comparison gate makes demands to low brittleness (at the order of ε) and low output error (*also* at the order of ε) that the cutoff gadget is not able to achieve. Fortunately, we can use our previously introduced reset gadget (which itself happens to be another incarnation of the cutoff gadget) to fix the output error.

Lemma 12 (Comparison Gadget). *Let $\zeta \in [0, 1]$. If $\alpha = \beta = 1$, there exists a financial system gadget with one input bank that implements the following property:*

$$\begin{aligned} r_a \leq \zeta - \Theta(\varepsilon) &\Rightarrow r_v = 0 \pm \Theta(\varepsilon) \\ r_a \geq \zeta + \Theta(\varepsilon) &\Rightarrow r_v = 1 \pm \Theta(\varepsilon) \end{aligned}$$

Proof. Let $C > 0$ be a constant such that C is an upper bound on the implicit constant factors in the expressions $\Theta(\frac{\varepsilon}{L-K})$ on the right-hand sides of Lemma 3 (i.e., the output error; we can choose $C = 3$). Let $c = 2C$. Assume WLOG that $c\varepsilon < \zeta < 1 - c\varepsilon$. If this is not the case, we can simply give v recovery rate constant 0 or 1. Apply now a cutoff gadget (Lemma 3) with $K = \zeta - c\varepsilon$ and $L = \zeta + c\varepsilon$ and call the output gate u . Then apply a reset gadget (Lemma 4) to u and call the output v .

By the cutoff gadget, if $r_a \leq K - \Theta(\varepsilon) = \zeta - \Theta(\varepsilon)$, then $r_u = 0 \pm C \frac{\varepsilon}{L-K} = 0 \pm C \frac{\varepsilon}{2c\varepsilon} = 0 \pm 1/4$. Thus, by the reset gadget, $r_v = 0 \pm \Theta(\varepsilon)$. Likewise for $r_a \geq \zeta + \Theta(\varepsilon)$. \square

With all gadgets in place, we can connect our gadgets to represent a generalized circuit and prove PPAD-hardness.

Proof of Theorem 2, hardness. Reduction from $\Theta(\varepsilon)$ -GCIRCUIT. Let C be a generalized circuit. We construct a financial system. For each node v of C , add a bank with no assets or liabilities and identify that bank with v . For each gate g of C , execute the corresponding gadget from this section with the appropriate input banks and identify the output bank of the gadget with g . Finally, if a node v is the output of gate g , take a copy of the scaling gadget with $\zeta = 1$ (this is the same as two inverters connected) and identify the input of the gadget with g and the output of the gadget with v .²²

²²In the language of gadgets, this operation can be interpreted as applying the constructed financial system, say X_0 , and the $\times 1$ -scaling gadget “to each other.” Note that in X_0 , v has no assets and liabilities and can thus be considered an input bank. The result of the “application” will be an extension of X_0 at all banks but v and of the scaling gadget at all banks but its input. By Remark 3., this implies that all gadgets still behave as expected. In particular, we have $r_v = r_g \pm \Theta(\varepsilon)$ in any ε -solution.

Now, by the gadgets, if $\varepsilon \ll 1$ and r is an ε -solution for the financial system, setting $x[v] = r_v$ for all nodes yields a $\Theta(\varepsilon)$ -solution of the generalized circuit. And finding this is hard for $\varepsilon \ll 1$. \square

5 Origin of the Complexity: Determining Defaults

Given the results in the previous two sections, one may wonder if we can pin down the “origin” of the computational complexity. What is it really that CDSs *do* to a financial system that makes the decision and search problems so much harder to solve? In this section and the next we explore this question.

To understand where the computational complexity comes from, we look for ways to circumvent it. One way to do this might be to ask for less information than the recovery rates themselves. If it is easy to find some bounds on the recovery rates, for example, this could already be very useful. The minimum level of detail we will likely be interested in is which banks default. To this end, we define a *default set* as a collection of banks that (approximately) default in some (approximate) solution. A default set thus provides a kind of “coarse representation” of a solution to the clearing problem.

Definition 7. Let $X = (N, e, c, \alpha, \beta)$ be a financial system and let $\varepsilon \geq 0$. A set $D \subseteq N$ is called an ε -*default set* for X if there is an ε -solution r for X such that for all $i \in N$:

$$\begin{aligned} i \notin D &\Rightarrow a_i(r) \geq (1 - \varepsilon)l_i(r) \\ i \in D &\Rightarrow a_i(r) < (1 + \varepsilon)l_i(r). \end{aligned}$$

In this case, we call r and D ε -*compatible*.

We have relaxed the notion of being in default in the same way as in the definition of an ε -solution. Again, this ensures that the problem will not be hard for the wrong reasons, namely because of “knife-edge” defaults, where a small error in the assets or liabilities could otherwise determine whether a bank defaults and thus lead to a large error in the recovery rate. Banks at the edge of default can instead be considered either in default or not in default. A side effect of this freedom is that, say, $i \notin D$ does *not* generally imply $r_i = 1 \pm \varepsilon$. It could also be that $a_i(r) \in [(1 - \varepsilon)l_i(r), (1 + \varepsilon)l_i(r))$ and $r_i = \frac{a'_i(r)}{l'_i(r)}$. That is, the decision about default in D and r need not coincide.

In the remainder of this section, we show that computational complexity does not arise exclusively from the need to compute precise numeric values for the recovery rates. Rather, the computational problems discussed in Sections 3 and 4 are already hard at the level of default sets.

5.1 Deciding the Default of a Given Bank

In the decision variant of our problem, we ask if there is a solution where a specified bank defaults. This is a basic question a regulator might ask in a stress test: for example, given a certain shock scenario, will this make it necessary to save AIG again? If there are multiple solutions and AIG only defaults in some of them, our answer should still be “Yes,” i.e., we should consider the worst case.

We can consider this decision problem for any value of the default cost parameters. If there are default costs, a naive framing of the problem would contain the question if there is any solution in the first place. Since this is not what we are interested in at this point, we exclude it by a promise. We then consider a distinction variant similar to ε -HASCLEARING $^{\alpha,\beta}$ (Theorem 1). The proof is rather straightforward using our financial Boolean circuits from Section 3. The key technical step is to notice that assets are either very small or very large compared to liabilities, so that defaults are never ambiguous.

Theorem 3. *For any fixed $\alpha, \beta \in [0, 1]$ and any $\varepsilon \ll 1$, the following promise problem is NP-complete: Given a non-degenerate financial system $X = (N, e, c, \alpha, \beta)$, and a bank $i \in N$...*

- *if X has an exact solution and there is an exact default set D such that $i \in D$, return YES.*
- *if X has an exact solution and there is no ε -default set D such that $i \in D$, return No.*

Proof. Membership: By the proof of Lemma 1, if r is an exact solution and D exactly compatible, we receive a polynomial-length ε -solution with which D is ε -compatible with via rounding. Hence, it is enough to check all ε -solutions of a certain polynomial maximum length. Note that if r is an ε -solution, then $i \in D$ for some ε -compatible default set D iff $a_i(r) < (1 + \varepsilon)l_i(r)$.

Hardness: Reduction from CIRCUIT FALSIFIABILITY. Given a Boolean circuit C , consider the financial Boolean circuit system X from Proposition 2 and the output node $v =: i$. Recall that v is the output of a NAND gadget (Lemma 5). If C is falsifiable, let r be the exact solution corresponding to a falsifying assignment and D its exactly compatible default set. By the proof of Lemma 5, we then have $a_v(r) = 0 < 1 = l_i(r)$ and thus $v \in D$. If C is not falsifiable, let $\varepsilon \ll 1$, let r be any ε -solution and D ε -compatible. By Proposition 2, the inputs to v 's NAND gadget are approximately Boolean and not both TRUE and again by the proof of the lemma, we have $a_v(r) \geq 1 + \varepsilon$, so $v \notin D$. \square

The theorem immediately implies that the following problem is NP-hard: Given a non-degenerate financial system with the promise that it has an exact solution and a bank i , decide if $i \in D$ for some ε -default set D . The problem may not be in NP

because it is not guaranteed that every ε -solution has polynomial length. See the discussion after Theorem 1. Note in particular that it is not clear how to check if a given $D \subseteq N$ is an ε -default set.

It is easy to see that the distinction problem is still NP-complete if we replace “ $\in D$ ” by “ $\notin D$ ” (proof by negation of the circuit) and it is coNP-complete if we replace “there exists D ” by “for all D ” (proof by reduction from CIRCUIT CONTRADICTION).

5.2 Finding Default Sets Without Default Costs

We return to the search problem from Section 4, where we are given a financial system without default costs and we are looking for an ε -solution. Given an ε -solution, it is trivial to compute a default set compatible with it, while the converse is not so clear. Thus, finding ε -default sets may a priori be easier than finding the ε -solution itself. We will show that this is not the case.

Theorem 4. *For $\varepsilon \ll 1$, the following problem is PPAD-complete: Given a non-degenerate financial system $X = (N, e, c, \alpha = 1, \beta = 1)$, compute an ε -default set.*

For some search problems, statements like the above follow trivially from hardness of the respective continuous variant. For example, in a two-player normal-form game, the supports of the strategies of the two players (i.e., the strategies that are played with nonzero probability) could be taken as a “coarse representation” of a Nash equilibrium, similar to default sets in our case. Finding the supports of a Nash equilibrium is trivially PPAD-complete because a Nash equilibrium can be reconstructed from its supports via linear programming, and finding Nash equilibria in two-player games is hard (Chen, Deng and Teng, 2009).²³ Given an ε -default set however, there does not seem to be an easy way to reconstruct a corresponding ε -solution, given that already the assets a_i can contain terms like $r_j(1 - r_k)$, which are non-linear, non-convex/concave, and non-monotonic. This remains true for the particular construction we perform in Section 4. Here, the assets of the relevant banks contain terms of form $r_s(1 - r_a)$ where s is the source bank and $r_s \in [1 - \varepsilon, 1]$. It may be tempting to just assume $r_s = 1$ to make the problem linear. However, given a default set of an ε -solution where (say) $r_s = 1 - \varepsilon$, assuming $r_s = 1$ implicitly introduces an error of ε into r_s . As the notional in the system can be at the order of $1/\varepsilon$,²⁴ this may in turn imply a large (constant in ε) change in the recovery rate of some other bank. No ε -solution with such a recovery rate may exist.

Rather than reconstructing an ε -solution to the financial system from a default set, we introduce a new discrete variant of the ε -GCIRCUIT problem that is still PPAD-complete. We then show that any ε -default set gives rise to a solution for the discrete $\Theta(\varepsilon)$ -GCIRCUIT problem. Our approach may be of independent interest for

²³ ε must be polynomial, not constant, in the size of the game in this particular example.

²⁴The comparison gadget (Lemma 12) introduces notionals of order $1/\varepsilon$.

Figure 10 Constraints that need to hold at each gate g when x is an ε -solution for a generalized circuit compatible with an assignment d . Let a and b be the inputs (if any) and v the output of g .

$$\begin{aligned}
g = C_+ \Rightarrow d_g \in \{L, M\} &\Rightarrow x[a] + x[b] \leq 1 + \varepsilon \wedge x[v] = x[a] + x[b] \pm \varepsilon \\
d_g = H &\Rightarrow x[a] + x[b] \geq 1 - \varepsilon \wedge x[v] = 1 \pm \varepsilon \\
g = C_- \Rightarrow d_g = L &\Rightarrow x[a] - x[b] \leq \varepsilon \wedge x[v] = 0 \pm \varepsilon \\
d_g \in \{M, H\} &\Rightarrow x[a] - x[b] \geq -\varepsilon \wedge x[v] = x[a] - x[b] \pm \varepsilon \\
g = C_{>\zeta} \Rightarrow d_g = L &\Rightarrow x[a] \leq \zeta + \varepsilon \wedge x[v] = 0 \pm \varepsilon \\
d_g = M &\Rightarrow x[a] = \zeta \pm \varepsilon \\
d_g = H &\Rightarrow x[a] \geq \zeta - \varepsilon \wedge x[v] = 1 \pm \varepsilon
\end{aligned}$$

the study of other “support finding” problems where hardness does not follow from hardness of the continuous variant.

5.3 The discrete ε -GCircuit problem

The main idea for our discrete variant of ε -GCIRCUIT is that the constraints for a generalized circuit (Figure 6) are piecewise linear with a finite number of “cutoff” or “decision points.” For example, the C_+ gate corresponds to either a sum or a constant depending on whether the sum of its inputs lies in $[0, 1]$ or $(1, \infty)$. Once we know on which side of these “decision points,” defined in an appropriate way, each constraint lies, we can reconstruct a solution by solving a linear feasibility problem. Thus, already obtaining this information is hard. Our approach is similar in spirit to Vazirani and Yannakakis (2011), where the authors split equilibrium computation in Fisher markets into two steps: first, a PPAD-complete problem is solved to determine the combinatorial structure of the equilibrium. Second, the exact numbers are computed in polynomial time.

In our definition of the discrete ε -GCIRCUIT problem, we relax the constraints in a way that mirrors the definition of an ε -default set, allowing ambiguity when we are close to the respective “decision point.”²⁵ Note that C_ζ and $C_{\times\zeta}$ gates have linear constraints and therefore do not need any “decision” specified, which leaves us with three gate types.

Definition 8. Let C be a generalized circuit and $\varepsilon \geq 0$. A *discrete ε -solution* for C assigns to each gate g of C a value $d_g \in \{H, M, L\}$ such that there exists an ε -solution x for C such that the constraints in Figure 10 hold for each gate g with inputs a and b and output v . In this case, we call x and d ε -compatible.

²⁵This relaxation is crucial for our following reduction from ε -default sets. This is not just because ε -default sets have this kind of relaxation, but due to ε errors in the solution itself. That is, even if we required exact compatibility in the definition of an ε -default set, we would have to make this relaxation here.

Note that our definition is *monotonic* in ε : if d is a discrete ε -solution and $\varepsilon' > \varepsilon$, then d is also a discrete ε' -solution. This is what we typically expect of an approximate solution concept and we take it as an indication that our concept of a discrete ε -solution is natural. We will exploit monotonicity in our proofs below.

Remark 4. If x is an ε -solution for C , we can define a discrete ε -solution d_g as follows.

- If $g = C_+$, let $d_g = H$ if $x[a] + x[b] \geq 1$ and $d_g = M$ otherwise.
- If $g = C_-$, let $d_g = L$ if $x[a] - x[b] \leq 0$ and $d_g = M$ otherwise.
- If $g = C_{>\zeta}$, let

$$d_g = \begin{cases} L & \text{if } x[a] < \zeta - \varepsilon \\ M & \text{if } x[a] = \zeta \pm \varepsilon \\ H & \text{if } x[a] > \zeta + \varepsilon. \end{cases}$$

The fact that x is an ε -solution will now ensure that the outputs match, too, and thus d is a discrete ε -solution compatible with x .

Theorem 5. *For a (constant) $\varepsilon \ll 1$, the following total search problem, which we call the discrete ε -GCIRCUIT problem, is PPAD-complete: Given a generalized circuit, find a discrete ε -solution.*

Proof. Membership is obvious by reduction to the (continuous) ε -GCIRCUIT problem. For hardness, we perform reduction *from* the continuous problem. Let C be a generalized circuit and let d be a discrete ε -solution for C . Consider the linear feasibility problem (LFP) with a variable $x[v]$ for each node v of C and the following constraints:

1. For each v , add the constraint $0 \leq x[v] \leq 1$.
2. For each gate g of type C_+ , C_- , or $C_{>\zeta}$, add the constraint from Figure 6 corresponding to the value of d_g .
3. For each gate g of type C_ζ or $C_{\times\zeta}$, add the constraint from Figure 6 corresponding to the type of g .

Note that all constraints are linear. It is easy to verify that a vector x is feasible for the LFP iff it is an ε -solution ε -compatible with d . By assumption, such an x exists and we can find it in polynomial time via the LFP. \square

The proof of the theorem implies that a generalized circuit always has an exact solution of polynomial length. This was previously noted by Etessami and Yannakakis (2010) in their study of their FIXP complexity class. To see it from our proof, let d be the default set of an exact solution (which exists by Kakutani's fixed-point theorem) and solve the LFP for $\varepsilon = 0$.

5.4 ε -Default Sets Define a Discrete $\Theta(\varepsilon)$ -Solution

We now re-examine the sum, difference, and comparison gadget from Section 4. We show that the default states of some banks can be used to define a discrete $\Theta(\varepsilon)$ -solution for the respective gate type that will be $\Theta(\varepsilon)$ -compatible with any ε -solution of the gadget.

Remark 5 (Bounding Recovery Rates via Default Sets). In contrast to the case with default costs, if $\alpha = \beta = 1$, we can bound the recovery rates based on the default set. More in detail, if r is an ε -solution ε -compatible with D , then:

$$\begin{aligned} i \notin D &\Rightarrow r_i = 1 \pm 2\varepsilon \\ i \in D &\Rightarrow r_i = \frac{a_i(r)}{l_i(r)} \pm 2\varepsilon \end{aligned}$$

This follows by case distinction using the equivalence between an ε -solution and an ε -approximate fixed point of F .

If G is a gadget with input banks A and $P : 2^N \times [0, 1]^N \rightarrow \{\text{TRUE}, \text{FALSE}\}$ is a property, we say that G implements P on default sets if for $\varepsilon \ll 1$, any ε -solution r on $N \setminus A$, and any $D \subseteq N$ ε -compatible with r , $P(D, r)$ holds.

For the sum gadget, we can simply consider the default state of the output bank.

Lemma 13 (Default Set of the Sum Gadget). *If $\alpha = \beta = 1$, then the sum gadget (Lemma 9) implements the following property on default sets:*

$$\begin{aligned} v \notin D &\Rightarrow r_a + r_b \geq 1 - \Theta(\varepsilon) \wedge r_v = 1 \pm \Theta(\varepsilon) \\ v \in D &\Rightarrow r_a + r_b < 1 + \Theta(\varepsilon) \wedge r_v = r_a + r_b \pm \Theta(\varepsilon) \end{aligned}$$

Proof. Recall that $a_v(r) = r_a + r_b \pm \Theta(\varepsilon)$ and $l_v(r) = 1$. The statement now follows from Remark 5 and the fact that $r_v = [r_a + r_b] \pm \Theta(\varepsilon)$. \square

For the difference gadget, the default state of the output bank is not informative. To see this, recall that this bank is the output of an inverter gadget. In the inverter gadget, though, the output bank is never unambiguously not in default because its assets are at most 1 (see Figure 8). Thus, we could have $v \in D$ independently of the input or output recovery rates. To extract information from the default set, we need to consider the *input* bank to the inverter gadget instead.

Lemma 14 (Default Set of the Difference Gadget). *Let $\alpha = \beta = 1$ and consider the difference gadget (Lemma 10). Let u be the intermediate bank in that gadget. Then the gadget implements the following property on default sets:*

$$\begin{aligned} u \notin D &\Rightarrow r_a - r_b \leq \Theta(\varepsilon) \wedge r_v = 0 \pm \Theta(\varepsilon) \\ u \in D &\Rightarrow r_a - r_b > -\Theta(\varepsilon) \wedge r_v = r_a - r_b \pm \Theta(\varepsilon) \end{aligned}$$

Proof. Recall that we have $a_u(r) = 1 - r_a + r_b \pm \Theta(\varepsilon)$ and $r_v = 1 - r_u \pm \Theta(\varepsilon)$. The statement now follows just like in Lemma 13. \square

For the comparison gadget, we proceed in a similar way and consider the default states in the first cutoff gadget (before the reset gadget). We need to consider the default states of both of the banks in this gadget to determine in which of the three possible states the gadget is. Compared to the other two gadgets, we need to consider the details of the construction in much greater detail.

Lemma 15 (Default Set of the Comparison Gadget). *Let $\alpha = \beta = 1$ and consider the comparison gadget (Lemma 12). Let u_1 and v_1 be the intermediate banks that correspond to the first cutoff gadget. Then the comparison gadget implements the following property on default sets:*

$$\begin{aligned} u_1 \notin D \wedge v_1 \in D &\Rightarrow r_a \leq \zeta + \Theta(\varepsilon) \wedge r_v = 0 \pm \Theta(\varepsilon) \\ u_1 \in D \wedge v_1 \notin D &\Rightarrow r_a \geq \zeta - \Theta(\varepsilon) \wedge r_v = 1 \pm \Theta(\varepsilon) \\ u_1 \in D \wedge v_1 \in D &\Rightarrow r_a = \zeta \pm \Theta(\varepsilon) \\ u_1 \notin D \wedge v_1 \notin D &\text{ is impossible.} \end{aligned}$$

Proof. Recall from the proof of Lemma 12 that the first cutoff gadget has parameters $K = \zeta - c\varepsilon$ and $L = \zeta + c\varepsilon$ where $c > 4$ is a sufficiently large constant. Recall from the definition of the cutoff gadget (Lemma 3) that this implies for the notionals in Figure 3 that

$$\begin{aligned} \gamma &= \frac{1}{1 - K} = \frac{1}{1 - \zeta + c\varepsilon} \\ \delta &= \frac{1 - K}{L - K} = \frac{1 - \zeta + c\varepsilon}{2c\varepsilon} = \frac{1 - \zeta}{2c\varepsilon} + \frac{1}{2}. \end{aligned}$$

Assume first that $u_1 \notin D$. By definition of an ε -default set and $l_{u_1}(r) = 1$, we have $1 - \varepsilon \leq a_{u_1}(r) = \gamma(1 - r_a)$. Rearranging yields $r_a \leq \zeta - c\varepsilon + \varepsilon(1 - \zeta + c\varepsilon) \leq \zeta - (c - 2)\varepsilon$ if $\varepsilon \ll c$. We further must have $r_{u_1} \geq 1 - 2\varepsilon$, so $r_{v_1} \leq a_{v_1}(r) + \varepsilon \leq \delta \cdot 2\varepsilon + \varepsilon = (1 - \zeta)/c + 2\varepsilon \leq 1/4$ for $\varepsilon \ll 1$. This implies $v_1 \in D$ and, as v_1 is input to a reset gadget with output v , $r_v = 0 \pm \Theta(\varepsilon)$.

Assume next that $u_1 \in D$. Then by Remark 5, $r_{u_1} = a_{u_1}(r) \pm 2\varepsilon = \gamma(1 - r_a) \pm (\gamma + 2)\varepsilon$.

If now $u_1 \in D \wedge v_1 \notin D$, then

$$\begin{aligned} 1 - \varepsilon \leq a_{v_1}(r) &\leq \delta(1 - r_{u_1}) \\ &\leq \delta(1 - (\gamma(1 - r_a) - (\gamma + 2)\varepsilon)) \\ &= \delta - \delta\gamma + \delta\gamma r_a + (\delta\gamma + 2\delta)\varepsilon. \end{aligned}$$

Rearranging yields: $r_a \geq \frac{1}{\delta\gamma} - \frac{1}{\gamma} + 1 - \left(1 + \frac{2}{\gamma} + \frac{1}{\delta\gamma}\right)\varepsilon \geq L - 4\varepsilon = \zeta + (c - 4)\varepsilon$, where the middle inequality follows using the identities $\frac{1}{\delta\gamma} = L - K$ and $\frac{1}{\gamma} = 1 - K$. Of course, $r_{v_1} = 1 \pm 2\varepsilon$ because $v_1 \notin D$ and thus $r_v = 1 \pm \varepsilon$.

If $u_1 \in D \wedge v_1 \in D$, then

$$\begin{aligned} 1 + \varepsilon &\geq a_{v_1}(r) \geq \delta(1 - \varepsilon)(1 - r_{u_1}) \\ &\geq \delta(1 - (\gamma(1 - r_a) + (\gamma + 2)\varepsilon) - \delta\varepsilon \\ &= \delta - \delta\gamma + \delta\gamma r_a - (\delta\gamma + 3\delta)\varepsilon. \end{aligned}$$

Rearranging like above yields: $r_a \leq L + 5\varepsilon = \zeta + (c + 5)\varepsilon$. This bounds r_a from above. To bound r_a from below, notice that, since $u_1 \in D$, we have $1 + \varepsilon > a_{u_1}(r) \geq \gamma(1 - \varepsilon)(1 - r_a)$. This implies $r_a \geq 1 - \frac{1}{\gamma} \cdot \frac{1+\varepsilon}{1-\varepsilon} \geq 1 - \frac{1}{\gamma}(1 + 3\varepsilon) \geq K - 3\varepsilon \geq \zeta - (c - 2)\varepsilon$, where the second inequality holds for $\varepsilon \ll 1$. \square

Note that the implicit constants in the $\Theta(\varepsilon)$ expressions in the above lemma are not the same. That is why the different cases in the previous lemma overlap.

Note that the proof of the previous lemma actually gives us slightly stronger bounds on r_a than stated. For example, in case $u_1 \notin D \wedge v_1 \in D$, we receive from the proof that $r_a \leq \zeta - (c - 2)\varepsilon = \zeta - \Theta(\varepsilon)$, not just $r_a \geq \zeta + \Theta(\varepsilon)$. Beyond uniformity with the definition of a discrete ε -solution, the weaker version of the conditions has the benefit of monotonicity: the conditions continue to hold if we increase ε . We will exploit this in the following proof.

Using the above three lemmas, we can define a discrete ε -solution to a generalized circuit given an ε -default set. This proves our theorem:

Proof of Theorem 4. Let C be a generalized circuit and let X be the financial system without default costs corresponding to C like in the proof of Theorem 2. Let $\varepsilon \ll 1$ and let D be an ε -default set of X . We show that for some $\varepsilon' = \Theta(\varepsilon)$, D induces a discrete ε' -solution d of C . This proves the theorem because finding the latter is hard for $\varepsilon' \ll 1$. We define d following the preceding lemmas. For each gate g of C , ...

- If $g = C_+$, consider the corresponding sum gadget and let $d_g = H$ if $v \in D$ and $d_g = M$ if $v \notin D$.
- If $g = C_-$, consider the corresponding difference gadget and let $d_g = L$ if $u \notin D$ and $d_g = M$ if $u \in D$.
- If $g = C_>$, consider the corresponding comparison gadget and let

$$d_g = \begin{cases} L & \text{if } u_1 \notin D \wedge v_1 \in D \\ H & \text{if } u_1 \in D \wedge v_1 \notin D \\ M & \text{if } u_1 \in D \wedge v_1 \in D. \end{cases}$$

Let r be an ε -solution of X ε -compatible with D and let $\varepsilon' = \Theta(\varepsilon)$ be the maximum of all the incarnations of $\Theta(\varepsilon)$ in Lemma 13–15, in the proof of Theorem 2, and ε itself. By the proof of Theorem 2, r induces an ε' -solution of C by restriction to the output banks of the gadgets corresponding to gates. Since the above lemmas still hold if one replaces every instance of $\Theta(\varepsilon)$ by ε' (due to monotonicity), they imply that this induced ε' -solution of C is ε' -compatible with d . Thus, d is a discrete ε' -solution of C . \square

In this section, we have shown that already finding an ε -default set of a financial system with CDSs is hard. En-route, we have developed a general methodology to show that “coarse” or “discrete” versions of PPAD-hard search problems are hard. We believe that our methodology can be applied to other problems to receive this type of result when the reduction is sufficiently faithful to the structure of a generalized circuit. For example, Daskalakis, Goldberg and Papadimitriou (2009) introduced gadgets that encode a generalized circuit in a binary, degree-3 graphical game. It is not hard to show that the supports of certain players in these gadgets inform a discrete ε -solution of the generalized circuit. Thus, already finding the supports of an ε -Nash equilibrium in such a game is hard.²⁶ Unlike for two-player, n -action games, this result is not trivial because graphical games can contain nonlinear interactions and two-player games are not an immediate special case of graphical games. However, in this particular instance, the result can be shown more directly. This is because the game gadgets can easily be modified to ensure that players’ utilities are linear combinations of other players’ strategies (Daskalakis, Goldberg and Papadimitriou, 2009, Section 6.1) and then an equilibrium can be reconstructed from the supports using linear programming. Future work may well encounter other domains where, like in financial networks with CDSs, no such modification is possible and where our methodology can be of use.

6 Origin of the Complexity: Structural Restrictions on the Contract Space

We continue our exploration of the “origin” of the computational complexity in financial networks with CDSs. In this section, we study under which restrictions on the network structure the distinction and search problems are still hard. This is important, following our original program of study, to understand how the informal “complexity” due to CDSs arises and to inform potential regulatory policies that reduce it.

²⁶Daskalakis, Goldberg and Papadimitriou (2009) showed hardness when ε decreases with the size of the game exponentially. Rubinstein (2018) extended their result to a constant ε using the same gadgets. From this, we receive hardness of finding supports for constant ε as well.

6.1 Counterparty Risk and Fundamental Risk

Inspection reveals that our gadgets, and thus all hard instances constructed in Sections 3 and 4, share three properties that make them particularly simple financial systems:

1. *Acyclic Liabilities*: The *liability graph*, where each writer of a contract is connected to the respective holder, is acyclic. In fact, this graph is a disjoint union of chains of form $s \rightarrow i \rightarrow t$, where s and t are the source and sink banks and i is some other bank.
2. *No Intermediation*: No bank both holds and writes a CDS on the same reference entity. The liability graphs for individual reference entities are therefore disjoint unions of (in- or out-) star graphs.
3. *No Counterparty Risk*: For each contract, either the holder or the writer is a highly capitalized bank, i.e., its external assets are significantly (by factor $2 \geq 1 + \varepsilon$, for any relevant ε) higher than its maximum liabilities and thus, they cannot default. Further, only highly capitalized banks are writers of CDSs.

Properties 1 and 2 are in stark contrast to much of the prior work on financial networks, which has often *only* considered the liability graph, where either reference entities were ignored altogether or they were treated as mere edge labels, but were not identified as nodes in the network (see our literature review in Section 1). No such approach would be able to capture the computational complexity we illustrate in this paper because the liability graph of our hard instances is always trivial.

Property 3 helps us discern the “origin of the complexity” from an economic point of view. The holder of a CDS is dependent on two banks: the reference entity (this is called *fundamental risk*) and the writer of the contract (this is called *counterparty risk*; see D’Errico et al. (2018)). By property 3 however, unless the holder of the CDS is highly capitalized, the recovery rate of the writer is fixed to 1 (up to ε errors) so that counterparty risk is only the risk of ε errors. Thus, counterparty risk does not significantly affect recovery rates.²⁷ The statement also holds for debt contracts. From this, it follows that computational complexity remains high if we neglect counterparty risk and it must therefore be driven by fundamental risk:

Proposition 3. *All our complexity results (Theorems 1–4) still hold in a variant of the clearing model where the assets of a bank i are defined as*

$$a_i(r) := e_i + \sum_{j \in N} c_{j,i}^\emptyset + \sum_{j,k \in N} c_{j,i}^k (1 - r_k).$$

²⁷In Section 5, we have argued that errors at source banks can have a large impact because notional assets are large in the order of $1/\varepsilon$. We have shown, however, that these errors do not have an impact on computational complexity because they do not significantly affect the outputs of our gadgets.

The above modified model corresponds to a world where a governmental agency like a central bank guarantees the payment in each and every contract while banks are still in default if they cannot pay their obligations.²⁸ The proof of the above proposition is by revisiting our gadget proofs and is omitted. The proof becomes slightly easier in the model without counterparty risk because we do not have to deal with ε errors at the source bank any more.

Overall, we have now seen that the computational complexity of the problems related to clearing with CDSs is not driven by counterparty risk, but by fundamental risk in CDSs. Since CDS writers are highly capitalized in our construction (property 3), it is not driven by fundamental risk on the liability side of banks' assets sheets and must thus be driven by fundamental risk on the asset side of their balance sheets. Mathematically, it does not arise from non-linearity and must therefore arise from non-monotonicity (see Section 2.4) This is because, in the above model without counterparty risk and with all liabilities of possibly-defaulting banks equal to 1, the update function F is piecewise linear and weakly decreasing in the point-wise ordering.

To eventually receive a polynomial-time algorithm and thus bound the computational complexity from above, since computational complexity is driven by non-monotonicity, it seems promising to study structural restrictions under which monotonicity is restored. This is what we do in the following.

6.2 Naked CDSs

Non-monotonicity of the update function emerges because a bank that holds a CDS and no other contracts profits from an ill-being of the reference entity. Economically, we say that it is *short* on the reference entity. This effect is only present when CDSs are held by banks in a *naked* fashion, i.e., without holding a corresponding debt contract from the reference entity. The opposite is called a *covered* CDS. In general networks, we need to consider all potential CDS writers to define what a covered and a naked CDS position are. We thus arrive at the following technical definition from our prior work.

Definition 9 (Covered and Naked CDS Position; Schuldenzucker, Seuken and Battiston (2019)). Let $X = (N, e, c, \alpha, \beta)$ be a financial system. A bank j has a *covered CDS position* towards another bank k if

$$\sum_{i \in N} c_{i,j}^k \leq c_{k,j}^\emptyset.$$

²⁸Leduc, Poledna and Thurner (2017) studied a model where payments are guaranteed for CDSs, but not for debt contracts. Our computational problems are hard under this assumption as well by the same argument as for Proposition 3. Leduc, Poledna and Thurner restricted their attention to *covered CDSs*, defined below. This is why the computational problems that emerged in the context of their paper were not hard.

Otherwise, j has a *naked CDS position* towards k . X has no naked CDSs if no bank has a naked CDS position towards another bank.

If j has a covered CDS position towards k and the recovery rate of k decreases, then j may receive a higher payment in the CDSs it holds on k (this depends on the recovery rates of the CDS writers), but it also receives a lower payment in the debt contract from k and the latter effect weakly dominates. Hence, j can never profit from the ill-being of k . A covered CDS thus functions as an insurance against default, while a naked CDS is often considered speculation on default.²⁹

For an example for a financial system without naked CDSs, see Appendix B. This also shows that it can still be the case that all solutions are irrational even without naked CDSs, so we still need to consider an approximation problem still.

In a financial system without naked CDSs, the update function is point-wise monotonically increasing. This implies that a solution always exists and a simple iteration sequence converges to a solution (Schuldenzucker, Seuken and Battiston, 2019). It is easy to see that it does so in polynomial time:

Theorem 6. *For any financial system $X = (N, e, c, \alpha, \beta)$ without naked CDSs and for any $\varepsilon > 0$, the iteration sequence (r^n) defined by $r^0 = (1, \dots, 1)$ and $r^{n+1} = F(r^n)$ reaches an ε -approximate fixed point of the update function F after $|N| \cdot 1/\varepsilon$ steps. In particular, this defines is a fully polynomial-time approximation scheme (FPTAS) for the total search problem of finding an ε -solution in a financial system with no naked CDSs.*

Proof. In each step where r^n is not an ε -approximate fixed point of F , some component r_i^n must decrease by at least ε in the next step. This follows from the definition of an ε -approximate fixed point and monotonicity of F . Since r is bounded below by $(0, \dots, 0)$, there can be at most $|N| \cdot 1/\varepsilon$ such steps. This defines an FPTAS because evaluating F and testing for an ε -approximate fixed point can obviously be done in polynomial time and any ε -approximate fixed point of F is an ε -solution (Proposition 1). \square

The above result extends to a slightly larger class of networks. In Schuldenzucker, Seuken and Battiston (2019), we have defined a structure called the *colored dependency graph* of a financial network. The nodes of this graph are the banks and an edge $i \rightarrow j$ exists whenever $F_j(r)$ depends on r_i (some of the edges may be false positives). Naked CDS positions are colored red and all other edges are colored green. We have shown in our prior work that, if no cycle in this graph contains a red edge, then

²⁹A covered CDS always acts as insurance, but a naked CDS need not be speculative per se. For example, a bank may hold a naked CDS on an entity that has very strong ties with one of its debtors, so that the CDS holder would still never profit from a default of the reference entity. It might also act as a mere intermediary. Detecting and appropriately handling these “indirectly covered” CDSs is a promising topic for future work, but beyond the scope of this paper.

a solution is still guaranteed to exist and we receive an approximation algorithm, essentially by iterating F on each strongly connected component in topological order. Theorem 6 implies that this algorithm is an FPTAS for the no-red-containing-cycle case.

7 Conclusion

In this paper, we have studied the clearing problem in financial networks that consist of debt and credit default swap (CDS) contracts. While in the debt-only case, the clearing problem can be solved in polynomial time, we have shown in this paper that the situation is markedly different if CDSs are allowed. Deciding if an (approximate) solution exists is NP-complete and finding an approximate solution when existence is guaranteed is PPAD-complete. In fact, already determining if a specific bank defaults or finding a consistent set of defaulting banks are hard problems. Hardness is preserved under various structural restrictions, but the case where no *naked* CDSs are present allows an FPTAS.

We can now answer our original question: Are financial networks with debt and CDSs “more complex” than those with only debt? Operationalizing informal “complexity” as computational complexity of the clearing problem, we can conclude: Yes, they are more complex, and in a precisely defined way so: understanding the interactions between banks in financial systems with CDSs is at least as challenging as understanding the structure of Boolean and generalized circuits. The complexity prevents us from even knowing which banks default following a shock. Complexity does not arise due to counterparty risk, but due to fundamental risk on the asset side of banks’ balance sheets. If anything like a structural “origin” of the complexity can be called out, it should be naked CDSs positions that occur as part of a cycle of dependencies.

These insights are relevant for regulatory policy. The post-2008 regulatory reforms related to the CDS market predominantly target counterparty risk. For example, *margin requirements* mandate counterparties to keep a “buffer account” from which fluctuations in the contract value are offset. Mandatory use of *central counterparties* (CCPs) re-routes all contracts via a highly capitalized central node. *Portfolio compression* eliminates cycles of liabilities for each individual reference entity.³⁰ All of these policies aim to reduce counterparty risk, but they do not affect fundamental risk. CCPs and portfolio compression modify the network structure, but they leave all reference entity–holder relationships of non-intermediaries as they are. Our results from Section 6.1 imply that this does not eliminate the kind of complexity we reveal in this paper.

³⁰See Financial Stability Board (2017) for details on the different market reforms. Benos, Wetherilt and Zikes (2013) provide an accessible introduction.

Another policy that will likely not affect the hardness of our problems are regulatory *capital constraints*. In our model, this would mean to require a minimum level $\gamma \in [0, 1)$ of external assets relative to maximum liabilities. Banks then have possible recovery rates in $[\alpha\gamma, 1]$ rather than $[0, 1]$. We believe that it will be straightforward to modify our constructions to re-map the latter to the former interval. This is why capital constraints do likely not eliminate the complexity we describe.

What *would* eliminate the complexity, by our results from Section 6.2, is banning all naked CDSs. This idea has been part of the public debate following the 2008 crisis (see, for instance, Soros (2009) and Reuters (2009)). During the European sovereign debt crisis in 2011, such a ban was in fact implemented for the subset of CDSs written on sovereign states. The ban is in effect until this day (European Commission, 2011; European Securities and Markets Authority, 2017). The policy implications we describe here echo earlier results regarding *existence* of a solution (Schuldenzucker, Seuken and Battiston, 2019).

Since the structure of our hard instances is so simple, our results are robust to changes to the details of the model. For example, our model abstracts over special provisions in bankruptcy code that essentially give derivatives priority over other contract types (debt in our model) in case of bankruptcy.³¹ As our constructions are not affected by counterparty risk and, in fact, relevant banks only ever write a single contract, priority is not relevant and our results persist. Our results *do* crucially depend on the assumption that all contracts are cleared at the same time. That is why they likely do not transfer to any variant of the *dynamic clearing* model in Banerjee, Bernstein and Feinstein (2018) or to a multi-maturity model (Kusnetsov and Veraart, 2019) when debt and CDSs mature at different points in time.

Future work should study which empirical properties of financial networks may make the clearing problem with CDSs feasible. For example, if the number of reference entities is small compared to the number of banks, we might be able to exploit the fact that with recovery rates of reference entities fixed, the update function is linear and monotonic. A similar approach may be feasible when the share of naked CDSs is positive, but small. All of these properties are incompatible with the constructions in our hardness proofs, which leaves hope that efficient algorithms might be available.

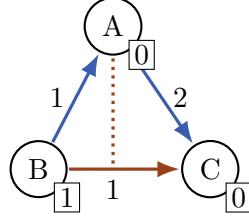
A Proofs from Section 3

Proof of Lemma 1. Assume WLOG that every bank writes a debt contract. If this is not true for some bank, no other bank depends on its recovery rate by non-degeneracy and our sanity assumptions. We can thus simply set its recovery rate to 1.

Note that the functions $\frac{a_i}{l_i}$ and $\frac{a'_i}{l_i}$ are *polynomially continuous* in X , i.e., continuous with a Lipschitz constant that is $O(2^{\text{poly}(\text{size}(X))})$. This is because a_i , a'_i , and l_i are

³¹For details see, for example, Bolton and Oehmke (2015).

Figure 11 Financial System where the unique solution is irrational. Let $\alpha = \beta = 1$ (no default costs).



polynomially continuous and l_i is bounded above $\sum_j c_{i,j}^\emptyset > 0$ because every bank writes a debt contract. In particular, $\frac{a_i(r)}{l_i(r)}$ and $\frac{a'_i(r)}{l_i(r)}$ is well-defined for all r . Let M be the maximum of the Lipschitz constants of these functions and 1.

Let r be an exact solution and let r' be r rounded to a multiple of $\delta := \varepsilon/(M+1)$, so that $r' = r \pm \delta$. By polynomial continuity, r' has a size as required. To see that r' is an ε -solution, we perform a case distinction for each i :

- If $r_i = 1$, then r' satisfies the first case in Definition 2. We have $r'_i = r_i \pm \delta = 1 \pm \varepsilon$. Further, since $r_i = 1$ we have $\frac{a_i(r)}{l_i(r)} \geq 1$, by choice of M and r' , $\frac{a_i(r')}{l_i(r')} = \frac{a_i(r)}{l_i(r)} \pm M\delta \geq 1 - \varepsilon$, and thus $a_i(r') \geq (1 - \varepsilon)l_i(r')$.
- If $r_i < 1$, then r' satisfies the second case. We have $r'_i = r_i \pm \delta = \frac{a'_i(r)}{l_i(r)} \pm \delta = \frac{a'_i(r')}{l_i(r')} \pm (M+1)\delta = \frac{a'_i(r')}{l_i(r')} \pm \varepsilon$. Since $r_i < 1$, we have $\frac{a_i(r)}{l_i(r)} < 1$ and thus $\frac{a_i(r')}{l_i(r')} = \frac{a_i(r)}{l_i(r)} \pm \varepsilon < 1 + \varepsilon$, that is, $a_i(r') < (1 + \varepsilon)l_i(r')$. \square

B Example That Financial Systems with CDSs May Have Only Irrational Exact Solutions

Figure 11 shows a financial system the unique exact solution of which is irrational. To see this, note that by the contract structure, $a_i(r) \leq l_i(r) \forall r, i = A, B$ and therefore r is clearing iff

$$r_A = \frac{r_B}{2}, \quad r_B = \frac{1}{2 - r_A},$$

and $r_C = 1$. One easily verifies that the unique solution in $[0, 1]^2$ to this system of equations is given by

$$r_A = 1 - \frac{1}{\sqrt{2}}, \quad r_B = 2 - \sqrt{2}.$$

C Comparison of our Generalized Circuit Definition to Rubinstein (2018)

Rubinstein's generalized circuits contain additional gates compared to ours. First, there is a C_+ gate that simply copies its input and can of course be replaced by a

$C_{\times 1}$ gate. Second, there are additional *Boolean* gates that operate on approximately Boolean values.³² While we could represent Boolean operations in a financial system using the gadgets from Section 3, Schuldenzucker and Seuken (2019) have shown in prior work that the Boolean operations are in fact redundant and can be represented using the comparison and arithmetic gates. To simplify our analysis, we omit these gates.

The third difference to Rubinstein (2018) is that Rubinstein assumed a binary comparison gate with *two* inputs where $x[v] = 0 \pm \varepsilon$ if $x[a] < x[b] - \varepsilon$ and $x[v] = 1 \pm \varepsilon$ if $x[a] > x[b] + \varepsilon$. One can emulate a binary comparison gate using our unary variant such that ε increases only by a constant factor. To see this, construct a sub-circuit corresponding to the expression

$$\left(\frac{1}{2} + (a - b) \right) - (b - a)$$

and call the output node u . Note that the order of operations matters due to truncation at 0 and 1. Then add a $C_{>1/2}$ gate with input u and output v . It follows immediately from the gates that if x is an ε -solution, then $x[u] = \tilde{u} \pm 5\varepsilon$ where

$$\tilde{u} := \left[\left[\frac{1}{2} + [x[a] - x[b]] \right] - [x[b] - x[a]] \right] = \left[\frac{1}{2} + x[a] - x[b] \right].$$

Note that $\tilde{u} - 1/2 = \min(1/2, \max(-1/2, x[a] - x[b]))$. From this, it follows that for $\varepsilon \ll 1$ ($\varepsilon < 1/10$ to be precise), v satisfies the definition of the binary comparison gadget for $\varepsilon' := 5\varepsilon$.

Acknowledgments

We would like thank (in alphabetical order) Vitor Bosshard, Yu Cheng, Constantinos Daskalakis, Timo Mennle, Noam Nisan, and Joseph Stiglitz for helpful comments on this work. Furthermore, we are grateful for the feedback we received from the anonymous reviewers and from various participants at EC 2016 and ITCS 2017.

All authors gratefully acknowledge financial support from the European Union's FP7 and Horizon 2020 research and innovation programme under Future and Emerging Technologies grant agreements No 610704 (SIMPOL) and No 640772 (DOLFINS). Additionally, Stefano Battiston acknowledges funding from the Swiss National Fund Professorship grant No PP00P1-144689 and from the Institute of New Economic Thinking through the Task Force in Macroeconomic Efficiency and Stability.

³²The definition of approximately Boolean values was weaker than what we did in Section 3, though. See Schuldenzucker and Seuken (2019) for a discussion.

References

- Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi.** 2015. “Systemic Risk and Stability in Financial Networks.” *American Economic Review*, 105(2): 564–608.
- Anderson, Robert M.** 1986. “‘Almost’ implies ‘near’” *Transactions of the American Mathematical Society*, 296(1): 229–237.
- Arinaminpathy, Nimalan, Sujit Kapadia, and Robert M. May.** 2012. “Size and complexity in model financial systems.” *Proceedings of the National Academy of Sciences*, 109(45): 18338–18343.
- Arora, Sanjeev, Boaz Barak, Markus Brunnermeier, and Rong Ge.** 2011. “Computational complexity and information asymmetry in financial products.” *Communications of the ACM*, 54(5): 101–107.
- Banerjee, Tathagata, Alex Bernstein, and Zachary Feinstein.** 2018. “Dynamic Clearing and Contagion in Financial Networks.” Working Paper, <https://arxiv.org/abs/1801.02091>.
- Bank for International Settlements.** 2018. “Semiannual OTC derivatives statistics, credit default swaps, by sector of reference, H1 2018.” <http://stats.bis.org/statx/srs/table/d10.4?p=20181>, Retrieved November 23, 2019.
- Bardoscia, Marco, Stefano Battiston, Fabio Caccioli, and Guido Caldarelli.** 2017. “Pathways towards instability in financial networks.” *Nature Communications*, 8: 14416.
- Basel Committee on Banking Supervision.** 2014. “The G-SIB assessment methodology — score calculation.” Retrieved November 23, 2019, <https://www.bis.org/bcbs/publ/d296.htm>.
- Battiston, Stefano, Guido Caldarelli, Robert M May, Tarik Roukny, and Joseph E Stiglitz.** 2016. “The price of complexity in financial networks.” *Proceedings of the National Academy of Sciences*, 113(36): 10031–10036.
- Benos, Evangelos, Anne Wetherilt, and Filip Zikes.** 2013. “The structure and dynamics of the UK credit default swap market.” *Bank of England Financial Stability Paper*, 25.
- Bolton, Patrick, and Martin Oehmke.** 2015. “Should derivatives be privileged in bankruptcy?” *The Journal of Finance*, 70(6): 2353–2394.
- Braverman, Mark, and Kanika Pasricha.** 2014. “The computational hardness of pricing compound options.” In *Proceedings of the 5th conference on Innovations in Theoretical Computer Science*.
- Chen, Xi, Xiaotie Deng, and Shang-Hua Teng.** 2009. “Settling the complexity of computing two-player Nash equilibria.” *Journal of the ACM (JACM)*, 56(3): 14.
- Constrâncio, Vítor.** 2017. “Macroprudential stress tests: A new analytical tool.” *VoxEU.org*. Retrieved November 23, 2019, <http://voxeu.org/article/macroprudential-stress-tests-new-analytical-tool>.
- Crotty, James.** 2009. “Structural causes of the global financial crisis: a critical assessment of the ‘new financial architecture’.” *Cambridge Journal of Economics*, 33(4): 563–580.
- Daskalakis, Constantinos, Paul W Goldberg, and Christos H Papadimitriou.** 2009. “The complexity of computing a Nash equilibrium.” *SIAM Journal on Computing*, 39(1): 195–259.
- Dees, Stéphane, Jérôme Henry, and Reiner Martin,** ed. 2017. *STAMP€: Stress-Test Analytics for Macroprudential Purposes in the euro area*. Frankfurt:European Central Bank.

- D'Errico, Marco, Stefano Battiston, Tuomas Peltonen, and Martin Scheicher.** 2018. "How does risk flow in the credit default swap market?" *Journal of Financial Stability*, 35: 53–74.
- Duffie, Darrell.** 1999. "Credit Swap Valuation." *Financial Analysts Journal*, 55(1): 73–87.
- Eisenberg, Larry, and Thomas H Noe.** 2001. "Systemic risk in financial systems." *Management Science*, 47(2): 236–249.
- Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson.** 2014. "Financial Networks and Contagion." *American Economic Review*, 104(10): 3115–53.
- Elsinger, Helmut, Alfred Lehar, and Martin Summer.** 2006. "Risk assessment for banking systems." *Management Science*, 52(9): 1301–1314.
- Etessami, Kousha, and Mihalis Yannakakis.** 2010. "On the complexity of Nash equilibria and other fixed points." *SIAM Journal on Computing*, 39(6): 2531–2597.
- European Banking Authority.** 2014. "EBA publishes 2014 EU-wide stress test results." *Press Release*, <http://www.eba.europa.eu/-/eba-publishes-2014-eu-wide-stress-test-results>.
- European Commission.** 2011. "Regulation on Short Selling and Credit Default Swaps — Frequently asked questions." *Press Release*, http://europa.eu/rapid/press-release_MEMO-11-713_en.htm, Retrieved November 23, 2019.
- European Securities and Markets Authority.** 2017. "Policy Activities / Short Selling." <https://www.esma.europa.eu/regulation/trading/short-selling>, Retrieved November 23, 2019.
- Feinstein, Zachary, Weijie Pang, Birgit Rudloff, Eric Schaanning, Stephan Sturm, and Mackenzie Wildman.** 2017. "Sensitivity of the Eisenberg-Noe clearing vector to individual interbank liabilities." Working Paper, <https://ssrn.com/abstract=3033476>.
- Fender, Ingo, and Jacob Gyntelberg.** 2008. "Overview: global financial crisis spurs unprecedented policy actions." *BIS Quarterly Review*, 1(1).
- Financial Crisis Inquiry Commission.** 2011. "The Financial Crisis Inquiry Report." <http://fcic.law.stanford.edu/report>.
- Financial Stability Board.** 2017. "OTC Derivatives Markets Reform — Twelfth Progress Report on Implementation." <http://www.fsb.org/2017/06/otc-derivatives-market-reforms-twelfth-progress-report-on-implementation/>.
- Fitch Ratings.** 2007. "CDx Survey — Market Volumes Continue Growing while New Concerns Emerge."
- Gai, Prasanna, Andrew Haldane, and Sujit Kapadia.** 2011. "Complexity, concentration and contagion." *Journal of Monetary Economics*, 58(5): 453–470.
- Glasserman, Paul, and H. Peyton Young.** 2015. "How likely is contagion in financial networks?" *Journal of Banking & Finance*, 50(0): 383–399.
- Goldreich, Oded.** 2005. "On promise problems (in memory of Shimon Even (1935–2004))." In *Electronic Colloquium on Computational Complexity*.
- Haldane, Andrew G.** 2009. "Rethinking the Financial Network." Speech at the Financial Student Association, Amsterdam, April 28, 2009, <https://www.bankofengland.co.uk/speech/2009/rethinking-the-financial-network>.
- Hemenway, Brett, and Sanjeev Khanna.** 2016. "Sensitivity and computational complexity in financial networks." *Algorithmic Finance*, 5(3–4).
- International Swaps and Derivatives Association.** 2012. "The ISDA Credit Deriva-

- tives Determinations Committees.” http://www2.isda.org/attachment/NDM1NA==/AGM%202012_DC%20anniversary_appendix_043012.pdf, Retrieved November 23, 2019.
- Kiff, John, Jennifer A. Elliott, Elias G. Kazarian, Jodi G. Scarlata, and Carolynne Spackman.** 2009. “Credit Derivatives: Systemic Risks and Policy Options.” *IMF Working Paper*, 9(254).
- Kusnetsov, Michael, and Luitgard Anna Maria Veraart.** 2019. “Interbank clearing in financial networks with multiple maturities.” *SIAM Journal on Financial Mathematics*, 10(1): 37–67.
- Leduc, Matt V, Sebastian Poledna, and Stefan Thurner.** 2017. “Systemic Risk Management in Financial Networks with Credit Default Swaps.” *Journal of Network Theory in Finance*, 3(3).
- Liu, Ming, and Jeremy Staum.** 2010. “Sensitivity analysis of the Eisenberg–Noe model of contagion.” *Operations Research Letters*, 38(5): 489–491.
- Papadimitriou, Christos H.** 1994. “On the complexity of the parity argument and other inefficient proofs of existence.” *Journal of Computer and System Sciences*, 48(3): 498–532.
- Reuters.** 2009. “Geithner testifies on financial regulation reform.” *Reuters*. News Item, Business News, 26 March, 2009, <https://uk.reuters.com/article/uk-financial-regulation-highlights-idUKTRE52P4UX20090326>.
- Rogers, LCG, and Luitgard AM Veraart.** 2013. “Failure and rescue in an interbank network.” *Management Science*, 59(4): 882–898.
- Rubinstein, Aviad.** 2018. “Inapproximability of Nash equilibrium.” *SIAM Journal on Computing*, 47(3): 917–959.
- Schoenebeck, Grant R., and Salil Vadhan.** 2012. “The Computational Complexity of Nash Equilibria in Concisely Represented Games.” *ACM Transactions on Computation Theory*, 4(2): 4:1–4:50.
- Schuldenzucker, Steffen, and Sven Seuken.** 2019. “Monotonicity and Non-Monotonicity of Solution Concepts in Generalized Circuits.” Working Paper, <https://arxiv.org/abs/1907.12854>.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2016. “Clearing Payments in Financial Networks with Credit Default Swaps [Extended Abstract].” In *Proceedings of the 17th ACM Conference on Economics and Computation*. Maastricht, The Netherlands.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2017. “Finding Clearing Payments in Financial Networks with Credit Default Swaps is PPAD-complete.” In *Proceedings of the 8th Innovations in Theoretical Computer Science (ITCS) Conference*. Berkeley, USA.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2019. “Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks.” *Management Science*, Published in Articles in Advance.
- Shin, Hyun Song.** 2010. “Financial Intermediation and the Post-Crisis Financial System.” *BIS Working Paper*, 304.
- Soros, George.** 2009. “One Way to Stop Bear Raids.” *The Wall Street Journal*. Op-ed, https://www.georgesoros.com/2009/03/23/one_way_to_stop_bear_raids/.
- Vazirani, Vijay V., and Mihalis Yannakakis.** 2011. “Market Equilibrium Under Separable, Piecewise-linear, Concave Utilities.” *Journal of the ACM (JACM)*, 58(3): 10:1–10:25.
- Yellen, Janet L.** 2013. “Interconnectedness and Systemic Risk: Lessons from the Financial

Crisis and Policy Implications.” Speech at the American Economic Association / American Finance Association Joint Luncheon, San Diego, California, January 4, 2013, <https://www.federalreserve.gov/newsevents/speech/yellen20130104a.htm>.

Zuckerman, David. 2011. “Pseudorandom Financial Derivatives.” In *Proceedings of the 12th ACM Conference on Electronic Commerce*. New York, NY, USA:ACM.

4 Monotonic and Non-Monotonic Solution Concepts for Generalized Circuits

I'm the operator with my pocket calculator.
I'm the operator with my pocket calculator.
I am adding and subtracting.
I'm controlling and composing.
I'm the operator with my pocket calculator.

Kraftwerk, *Pocket Calculator*

The content of this chapter has previously appeared in:

“Monotonic and Non-Monotonic Solution Concepts for Generalized Circuits.”
Steffen Schuldenzucker and Sven Seuken. Working Paper, July 2019.

Monotonic and Non-Monotonic Solution Concepts for Generalized Circuits

Steffen Schulenzucker
University of Zurich
schulenzucker@ifi.uzh.ch

Sven Seuken
University of Zurich
seuken@ifi.uzh.ch

First version: July 12, 2019
This version: November 23, 2019

Abstract

Generalized circuits are an important tool in the study of the computational complexity of equilibrium approximation problems. However, in this paper, we reveal that they have a conceptual flaw, namely that the solution concept is not *monotonic*. By this we mean that if $\varepsilon < \varepsilon'$, then an ε -approximate solution for a certain generalized circuit is not necessarily also an ε' -approximate solution. The reason for this non-monotonicity is the way Boolean operations are modeled. We illustrate that non-monotonicity creates subtle technical issues in prior work that require intricate additional arguments to circumvent. To eliminate this problem, we show that the Boolean gates are a redundant feature: one can simulate stronger, monotonic versions of the Boolean gates using the other gate types. Arguing at the level of these stronger Boolean gates eliminates all of the aforementioned issues in a natural way. We hope that our results will enable new studies of sub-classes of generalized circuits and enable simpler and more natural reductions from generalized circuits to other equilibrium search problems.

1 Introduction

Generalized circuits (Chen, Deng and Teng, 2009) have become a vital tool in the study of the computational complexity of equilibrium approximation problems. Reductions from generalized circuits have been used to show PPAD-completeness of a wide range of such problems, including the approximate search problems for: Nash equilibrium of a normal-form game (Daskalakis, Goldberg and Papadimitriou, 2009; Chen, Deng and Teng, 2009; Daskalakis, 2013; Babichenko, Papadimitriou and Rubinstein, 2016; Rubinstein, 2018), Arrow-Debreu market equilibrium (Chen, Paparas and Yannakakis, 2017), competitive equilibrium from equal incomes (Othman, Papadimitriou and Rubinstein, 2016), and payment equilibrium in a financial network (Schulenzucker, Seuken and Battiston, 2017, 2019).

A generalized circuit consists of *nodes* that are connected by *gates*. Nodes take values between 0 and 1 and each gate defines a constraint on the values of the nodes connected to it. Generalized circuits differ from regular algebraic circuits in

three important aspects. First, in addition to arithmetic gates (constants, addition, subtraction, and scaling by a constant), there are also a *comparison gate* and *Boolean gates* that implement the standard Boolean operations (AND, OR, NOT). Second, generalized circuits may contain *cycles*. Third, the constraints on the nodes are *approximate* depending on a precision parameter ε . This enables generalized circuits to express a large class of approximate-fixed-point problems. An ε -*solution* to a generalized circuit is an assignment of values in $[0, 1]$ to the nodes consistent with the constraints induced by the gates for precision ε . While an ε -solution always exists, the search problem ε -GCIRCUIT of finding such an ε -solution is PPAD-complete for a sufficiently small constant ε (Rubinstein, 2018).

In this paper, we reveal a conceptual flaw in the definition of the generalized circuit concept, namely that the solution concept is not *monotonic*. By this we mean that if $\varepsilon < \varepsilon'$, then an ε -solution to a given generalized circuit is not necessarily also an ε' -solution to the same circuit. The issue lies with the Boolean gates NOT, AND, and OR and the way how these gates operate on *approximately Boolean values* (Section 3).

Not having monotonicity violates our intuition for an approximate solution concept. For example, the simple idea that finding an ε -approximate solution gets (weakly) harder as ε gets smaller implicitly relies on the assumption of monotonicity.

To overcome this problem of non-monotonicity, we introduce a new variant of the generalized circuits problem that has stronger constraints for the Boolean gates that satisfy monotonicity. We call this variant ε -GCIRCUIT^{SB} (“SB” for “stronger Boolean”).¹ ε -GCIRCUIT^{SB} serves as a monotonic drop-in replacement for ε -GCIRCUIT in hardness proofs about generalized circuits themselves. In a second step, we show that Boolean gates (our stronger version or the original weaker version) are in fact a redundant feature: we can represent each of the Boolean gates using only the other (arithmetic and comparison) gate types. Our result implies that two new monotonic search problems are PPAD-complete: ε -GCIRCUIT^{SB} and the restriction ε -GCIRCUIT^{NB} of ε -GCIRCUIT where no Boolean gates are allowed (“NB” for “no Boolean”; see Section 4).²

We then illustrate that the lack of monotonicity in ε -GCIRCUIT has led to several subtle technical issues in prior work (to be precise, in Chen, Deng and Teng (2009)

¹The distinguishing feature of ε -GCIRCUIT^{SB}, which makes the solution concept monotonic, is shared by a variant of generalized circuits considered earlier by Othman, Papadimitriou and Rubinstein (2016). However, this variant is heavily customized to their application (fair allocation), while we aim to stay as faithful to the standard definition of generalized circuits as possible. The authors also did not discuss the relevance of monotonicity.

²The problem ε -GCIRCUIT^{NB} and its PPAD-completeness have been discussed before. See, for instance: Constantinos Daskalakis, “Algorithmic Lower Bounds: Fun with Hardness Proofs. Lecture 23. PPAD Reductions,” MIT Course, Fall 2014, available online: <https://www.youtube.com/watch?v=Ih0cPR745fM> (the problem is called ε -ARITHMCIRCUITSAT here). What was not discussed previously is the property of monotonicity, how it affects natural arguments and prior work, and that different variants of the generalized circuit problem differ in terms of monotonicity.

and Rubinstein (2018)) that, to the best of our knowledge, have been overlooked so far. While these issues are of a mere technical nature and can be circumvented using more careful argumentation, we demonstrate that ε -GCIRCUIT^{SB} can be used as a drop-in replacement for ε -GCIRCUIT in these pieces of work and provides a clean and conceptually simple way to eliminate these issues (Section 5).

We argue that, due to the desirability of monotonicity, the ε -GCIRCUIT problem in its current form is difficult to work with and future studies of generalized circuits should either consider the ε -GCIRCUIT^{SB} problem or the ε -GCIRCUIT^{NB} problem (i.e., leave out Boolean gates altogether). Monotonic solution concepts match our expectations and are thus easier to reason about. The fact that the Boolean gates are optional will simplify reductions from generalized circuits to other problems. We hope that this will enable new complexity results for equilibrium search problems in the future.

2 Preliminaries: Generalized Circuits

We follow the definition of a generalized circuit in Rubinstein (2018). A *generalized circuit* is a collection of *nodes* and *gates*, where each node is labeled as an *input* of any number of gates (including zero) and as an *output* of at most one gate. Inputs to the same gate are distinguishable from each other. Each gate has one of the types C_ζ , $C_{\times\zeta}$, $C_=_$, C_+ , C_- , $C_<$, C_V , C_\wedge , or C_\neg . For the gate types C_ζ and $C_{\times\zeta}$, a numeric parameter ζ is specified in addition to their input and output nodes. The *length* of a generalized circuit is the number of bits needed to describe the circuit, including the nodes, the mapping from nodes to inputs and outputs of gates, and numeric parameters ζ involved.

For any $\varepsilon > 0$, an ε -*approximate solution* (or ε -*solution* for short) of a generalized circuit is a mapping x that assigns to each node v a value $x[v] \in [0, 1]$ such that at each gate, the constraints in Table 1 hold. We write $[x] := \min(1, \max(0, x))$ and we write $y = x \pm \varepsilon$ to mean that $|x - y| \leq \varepsilon$. ε -GCIRCUIT is the search problem of finding an ε -solution of a given generalized circuit. It is easy to show that an ε -solution always exists (using Kakutani’s fixed-point theorem), has polynomial length, and that ε -GCIRCUIT is in PPAD. This is true even if ε decreases exponentially with the input size.

The gates can be grouped into three categories: the *arithmetic gates* C_ζ , $C_{\times\zeta}$, $C_=_$, C_+ , and C_- , the *comparison gate* $C_<$, and the *Boolean gates* C_V , C_\wedge , and C_\neg . Note from Table 1 how the comparison gate is *brittle*: its output value is unconstrained in $[0, 1]$ if $x[a_1]$ and $x[a_2]$ are ε -close to each other. This is crucial to guarantee existence of an ε -solution³ and it is also necessary to enable reductions from generalized circuits

³A brittle comparison gadget, and thus all gadgets, can be represented by a continuous function with Lipschitz constant $O(1/\varepsilon)$. By Brouwer’s fixed point theorem and rounding, an ε -approximate

Table 1 Conditions required to hold at a gate of the respective type with input nodes a_i and output node v in an ε -solution for a generalized circuit. For each gate, one output node and between 0 and 2 input nodes (depending on the gate type) are specified. For the gates C_ζ and $C_{\times\zeta}$, an additional numeric parameter $\zeta \in [0, 1]$ is specified.

Gate Type	Short	Constraint
Constant	C_ζ	$x[v] = \zeta \pm \varepsilon$
Scaling	$C_{\times\zeta}$	$x[v] = [\zeta \cdot x[a_1]] \pm \varepsilon$
Copy	$C_=$	$x[v] = x[a_1] \pm \varepsilon$
Addition	C_+	$x[v] = [x[a_1] + x[a_2]] \pm \varepsilon$
Subtraction	C_-	$x[v] = [x[a_1] - x[a_2]] \pm \varepsilon$
Comparison	$C_<$	$x[a_1] < x[a_2] - \varepsilon \Rightarrow x[v] \leq \varepsilon$ $x[a_1] > x[a_2] + \varepsilon \Rightarrow x[v] \geq 1 - \varepsilon$
OR	C_V	$x[a_1] \leq \varepsilon \quad \text{and} \quad x[a_2] \leq \varepsilon \quad \Rightarrow x[v] \leq \varepsilon$ $x[a_1] \geq 1 - \varepsilon \quad \text{or} \quad x[a_2] \geq 1 - \varepsilon \Rightarrow x[v] \geq 1 - \varepsilon$
AND	C_\wedge	$x[a_1] \leq \varepsilon \quad \text{or} \quad x[a_2] \leq \varepsilon \quad \Rightarrow x[v] \leq \varepsilon$ $x[a_1] \geq 1 - \varepsilon \quad \text{and} \quad x[a_2] \geq 1 - \varepsilon \Rightarrow x[v] \geq 1 - \varepsilon$
NOT	C_\neg	$x[a_1] \leq \varepsilon \quad \Rightarrow x[v] \geq 1 - \varepsilon$ $x[a_1] \geq 1 - \varepsilon \Rightarrow x[v] \leq \varepsilon$

to other approximate solution concepts like approximate Nash equilibrium, where an exact comparison gadget may not be attainable (see Daskalakis, Goldberg and Papadimitriou (2009) for a discussion). The Boolean gates are defined in a similar way, operating on *approximately Boolean values*. That is, we consider any value within $[0, \varepsilon]$ Boolean FALSE and any value within $[1 - \varepsilon, 1]$ Boolean TRUE. The Boolean gates are then only required to return an approximately Boolean value at their output if their inputs are also approximately Boolean. If the inputs are not approximately Boolean, i.e., if they lie in the interval $(\varepsilon, 1 - \varepsilon)$, any output value is allowed. For example, the C_\neg gate can map an input 0.5 to any number in $[0, 1]$. This provides a minimal specification of “approximate Boolean gates” and is important for reductions because the problem one wants to reduce to may only be able to express Boolean functions up to such errors (e.g., approximate fixed point problems, see Section 3 below). Note further how the arithmetic gates accumulate errors (a chain of, say n $C_=$ gates has a total error of $n\varepsilon$) while the Boolean gates do not. This is exploited, for example, in Rubinstein (2018).

ε -GCIRCUIT is known to be PPAD-complete for a sufficiently small constant ε (Rubinstein, 2018). Thus, no polynomial-time approximation scheme (PTAS) exists unless P=PPAD. This is the strongest hardness result for ε -GCIRCUIT known to date.

solution of length $O(1/\text{Length}(\varepsilon)^2)$ always exists.

3 ε -GCIRCUIT Does Not Satisfy Monotonicity

We now formally define monotonicity and we show that ε -GCIRCUIT does not in general satisfy it. Let X be a set and let $P_\varepsilon : X \rightarrow \{\text{TRUE}, \text{FALSE}\}$ for $0 < \varepsilon < 1$ be a *family of properties* of elements of X . We call the family P *monotonic* if for all $\varepsilon < \varepsilon'$ and all $x \in X$, $P_\varepsilon(x)$ implies $P_{\varepsilon'}(x)$.

Essentially anything we would call an “approximate solution concept” is monotonic. For example, if G is a game, X is the set of mixed strategy profiles of G , and $P_\varepsilon(x) = \text{TRUE}$ iff x is an ε -approximate Nash equilibrium of G , then the family P is monotonic. Likewise, well-supported approximate Nash equilibria (Daskalakis, Goldberg and Papadimitriou, 2009) and relative approximate Nash equilibria are monotonic. Another important family of monotonic properties are *approximate fixed points*. Let $n \geq 1$, $X = [0, 1]^n$, and let $F : X \rightarrow X$ be a function. Let $P_\varepsilon(x) = \text{TRUE}$ iff $F_i(x) = x_i \pm \varepsilon$ for all i . x is then called an ε -*approximate fixed point* of F . P is obviously monotonic.⁴

We now show that “ ε -solution to a certain generalized circuit” is *not* in general monotonic.

Proposition 1. *There exists a generalized circuit C such that the family of properties $P_\varepsilon(x) := “x \text{ is an } \varepsilon\text{-solution for } C”$ is not monotonic.*

Proof. Let C consist of two nodes a and v connected by a single C_\neg gate. Let $0 < \varepsilon < 1/4$ and let $x[a] = 1.5\varepsilon$ and $x[v] = 0.5$. Since $x[a] \notin [0; \varepsilon] \cup [1 - \varepsilon; 1]$, the C_\neg gate does not constrain the value of the output and thus x is an ε -solution. Let $\varepsilon' = 2\varepsilon$. Now $x[a] \leq \varepsilon'$, so the C_\neg gate requires that $x[v] \geq 1 - \varepsilon'$. But this is not the case. Thus, x is not an ε' -solution, which violates monotonicity. \square

Remark 1. In the specific, stylized example in the above proof, there are of course many ε -solutions that are also ε' -solutions for $\varepsilon' > \varepsilon$, like $(x[a] = 0, x[v] = 1)$. One might argue that one should only consider these “normal” solutions and that the ε -solution we discuss is pathological. If the gate is part of a larger circuit, however, it is not clear anymore how one would transition to a “normal” solution while still satisfying the constraints at all gates. We discuss this in Appendix B.

At a conceptual level, the reason why “ ε -solution to a generalized circuit” is not monotonic is because for the Boolean gates, the respective constraint is a collection of implications where conditions like $x \leq \varepsilon$ and $x \geq 1 - \varepsilon$ occur on both sides of each implication (see Table 1). Both sides get weaker as ε is increased and thus the effect on the overall constraint is ambiguous. Indeed, it is not hard to construct analogous counterexamples to the proof of Proposition 1 for the C_V and C_\wedge gates. Note that,

⁴Approximate fixed points occur in many search problems, where F is then defined in some way based on the input. A related concept are *strong approximate fixed points* (Etessami and Yannakakis, 2010), which need to be close to an exact fixed point.

in contrast to the Boolean gates, the comparison gate is not affected by this problem. This is because here, the left-hand side of the implication becomes *stronger* as ε is increased, so the whole implication becomes unambiguously weaker.

The fact that “ ε -solution to a generalized circuit” is not monotonic violates our intuition for approximation problems. For example, we would typically assume that the ε -GCIRCUIT search problem becomes (weakly) harder when we decrease ε . This is of course based on the assumption that a solution to ε -GCIRCUIT will also be a solution to ε' -GCIRCUIT if $\varepsilon < \varepsilon'$ (i.e., monotonicity). However, since monotonicity is not guaranteed, we cannot immediately exclude the possibility that the problem becomes *easier* again when we decrease ε far enough. This might be because for a low ε , many inputs to Boolean gates can be chosen to be not approximately Boolean and so the outputs of these gates can be arbitrary, giving us additional degrees of freedom to satisfy other constraints. We show in Appendix B that the problem does not actually become easier in a computational complexity sense, but this requires careful argumentation.

4 Restoring Monotonicity

We have just seen that the lack of monotonicity creates subtle pitfalls in otherwise trivial arguments. In fact, we will demonstrate in Section 5 that non-monotonicity can lead to many more issues, including in prior work. To overcome this problem, we now present a way to restore monotonicity. We will show in Section 5 that our approach eliminates the above-mentioned issues at a conceptual level.

Non-monotonicity arises due to the Boolean gates. Of course, we cannot simply remove the Boolean gates from consideration because the hardness proofs for ε -GCIRCUIT rely on having access to Boolean gates. Instead, to restore monotonicity, we define a new variant of the problem that has stronger constraints for the Boolean gates that satisfy monotonicity. We call this variant ε -GCIRCUIT^{SB} (“SB” for “stronger Boolean”). ε -GCIRCUIT^{SB} is very useful for hardness proofs about generalized circuits themselves (see Section 5). However, the fact that we have strengthened the Boolean gates creates two new problems: first, reductions from ε -GCIRCUIT to other problems do not automatically provide reductions for the stronger Boolean gates. Second, it is not clear at this point that ε -GCIRCUIT^{SB} is in PPAD. To resolve *these* problems, we show that Boolean gates (our stronger version or the original weaker version) are a redundant feature: we can represent each of the Boolean gates using only the other (arithmetic and comparison) gate types.⁵ This immediately implies that the restriction of ε -GCIRCUIT where no Boolean gates are allowed, and

⁵Of course, many more gates beyond the Boolean gates are redundant or could be replaced by simplified versions of the respective gate. For example, $C_<$ can be replaced by $C_{\times 1}$. However, only the Boolean gates are relevant for monotonicity. See Othman, Papadimitriou and Rubinstein (2016) and Schuldenzucker, Seuken and Battiston (2019) for reduced sets of gates.

Table 2 Conditions required to hold at a gate g with input nodes a_i and output node v in a strong ε -solution for a generalized circuit. The constraints differ from Table 1 only with regards to the Boolean gates (highlighted in gray).

Gate Type	Short	Constraint
Constant	C_ζ	$x[v] = \zeta \pm \varepsilon$
Scaling	$C_{\times\zeta}$	$x[v] = [\zeta \cdot x[a_1]] \pm \varepsilon$
Copy	C_*	$x[v] = x[a_1] \pm \varepsilon$
Addition	C_+	$x[v] = [x[a_1] + x[a_2]] \pm \varepsilon$
Subtraction	C_-	$x[v] = [x[a_1] - x[a_2]] \pm \varepsilon$
Comparison	$C_<$	$x[a_1] < x[a_2] - \varepsilon \Rightarrow x[v] \leq \varepsilon$ $x[a_1] > x[a_2] + \varepsilon \Rightarrow x[v] \geq 1 - \varepsilon$
OR	C_V	$x[a_1] < 1/2 - \varepsilon \text{ and } x[a_2] < 1/2 - \varepsilon \Rightarrow x[v] \leq \varepsilon$ $x[a_1] > 1/2 + \varepsilon \text{ or } x[a_2] > 1/2 + \varepsilon \Rightarrow x[v] \geq 1 - \varepsilon$
AND	C_\wedge	$x[a_1] < 1/2 - \varepsilon \text{ or } x[a_2] < 1/2 - \varepsilon \Rightarrow x[v] \leq \varepsilon$ $x[a_1] > 1/2 + \varepsilon \text{ and } x[a_2] > 1/2 + \varepsilon \Rightarrow x[v] \geq 1 - \varepsilon$
NOT	C_{\neg}	$x[a_1] < 1/2 - \varepsilon \Rightarrow x[v] \geq 1 - \varepsilon$ $x[a_1] > 1/2 + \varepsilon \Rightarrow x[v] \leq \varepsilon$

which we call ε -GCIRCUIT^{NB} (“NB” for “no Boolean”), is already PPAD-complete. Note that ε -GCIRCUIT^{NB} is also monotonic.

4.1 The ε -GCIRCUIT^{SB} Problem

Recall that non-monotonicity of the ε -solution concept arises because expressions of the form $x[a_1] \geq 1 - \varepsilon$ (which occur on the left-hand side of the constraints for Boolean gates) become *weaker* as ε increases. To resolve this, we replace these conditions so that they become *stronger* as ε increases, so that the whole implication becomes weaker. More in detail, we replace expressions of the form $x[a_1] \geq 1 - \varepsilon$ by expressions of the form $x[a_1] > 1/2 + \varepsilon$ on the left-hand side of the implication for Boolean gates.⁶ This yields the constraints in Table 2. We call an assignment x that satisfies these constraints a *strong ε -solution* and we call the corresponding search problem ε -GCIRCUIT^{SB}.⁷ This restores monotonicity.

Proposition 2. *Let C be a generalized circuit.*

⁶Note that we also replace weak by strict inequalities in the process. We do this to simplify the following arguments in this paper and to receive the *continuity* property discussed in Remark 2 below. It is not crucial for our construction, though.

⁷A similar variant of the generalized circuits problem was first studied by Othman, Papadimitriou and Rubinstein (2016). The authors introduced an additional parameter $\beta = \Theta(\varepsilon)$ and then specified the Boolean gates like in Table 2 if we replace ε by β on the left-hand sides of the Boolean gates. Their variant differs from the variant we describe here in other details of the problem. For example, scaling is only allowed by a factor $1/2$ and the definition of the C_- gate is not analogous to the two other Boolean gates. In this paper, we aim for the smallest deviation from Rubinstein’s (2018) variant that eliminates the aforementioned problems.

1. For any $\varepsilon < 1/4$, any strong ε -solution of C is also an ε -solution of C
2. The family of properties $P_\varepsilon(x) := "x \text{ is a strong } \varepsilon\text{-solution of } C"$ is monotonic.

Proof. 1: We can consider each gate separately and we only need to consider the Boolean gates, since the constraints for the other gates are the same between ε -GCIRCUIT and ε -GCIRCUIT^{SB}. For the Boolean gates, note that for all $z \in [0, 1]$ and $\varepsilon < 1/4$ we have $z \leq \varepsilon \Rightarrow z < 1/2 - \varepsilon$ and $z \geq 1 - \varepsilon \Rightarrow z > 1/2 + \varepsilon$. Thus, whenever we require $x[v] = 0 \pm \varepsilon$ or $x[v] = 1 \pm \varepsilon$ in an ε -solution, we make the same requirement in a strong ε -solution. Therefore, every strong ε -solution satisfies the constraints for an ε -solution.

2: Again, we only need to consider the constraints corresponding to Boolean gates since we have already seen that the others satisfy monotonicity. For the Boolean gates, consider the C_{\neg} gate, let x be an ε -solution and let $\varepsilon < \varepsilon'$. We distinguish the two cases in the constraint for the C_{\neg} gate for a strong ε' -solution.

- If $x[a_1] < 1/2 - \varepsilon'$, then $x[a_1] < 1/2 - \varepsilon$ and thus, since x is an ε -solution, $x[v] = 1 \pm \varepsilon = 1 \pm \varepsilon'$ as required.
- If $x[a_1] > 1/2 + \varepsilon'$, then likewise $x[a_1] > 1/2 + \varepsilon$ and thus $x[v] = 0 \pm \varepsilon = 0 \pm \varepsilon'$.

The proofs for the other two Boolean gates are analogous. \square

Due to monotonicity, the ε parameter of ε -GCIRCUIT^{SB} now behaves as we would intuitively expect. For example, the ε -GCIRCUIT^{SB} problem trivially becomes (weakly) harder as ε decreases.

Remark 2 (Continuity of the solution concept). ε -GCIRCUIT^{SB} is distinguished from ε -GCIRCUIT by another intuitive property that we call *continuity* of the solution concept.⁸ By *continuity* we mean the following. Fix a generalized circuit and let $x^n \rightarrow x$ and $\varepsilon^n \rightarrow \varepsilon$ be two convergent sequences such that x^n is a strong ε^n -solution for all n ; then x is a strong ε -solution. In particular, if $\varepsilon^n \rightarrow 0$, then x is a strong exact solution. Continuity holds for ε -GCIRCUIT^{SB} because only *strict* inequalities appear on the left-hand sides of the constraints for the Boolean gates.⁹ In ε -GCIRCUIT, these inequalities are weak and ε -GCIRCUIT does not satisfy continuity.¹⁰ Note that continuity of the solution concept is orthogonal to monotonicity and does not affect any of the other results, and in particular it does not affect hardness of the problem.

⁸We thank Xi Chen for bringing this property to our attention.

⁹More in detail, note that continuity for ε -GCIRCUIT^{SB} is equivalent to closedness of the set $S := \{(x, \varepsilon) \mid x \text{ is a strong } \varepsilon\text{-solution for } C\}$ for any generalized circuit C . This in turn holds because $S = \bigcap_{g \text{ gate}} \{(x, \varepsilon) \mid x \text{ satisfies the constraint in Table 2 for } g \text{ with tolerance } \varepsilon\}$ and all of these sets are closed. The first constraint of the C_{\neg} gate, for example, is equivalent to $x[a_1] \geq 1/2 - \varepsilon \vee x[v] \geq 1 - \varepsilon$, which obviously leads to a closed set. Also recall that we write $y = x \pm \varepsilon$ for the *weak* inequalities $x - \varepsilon \leq y \leq x + \varepsilon$. This is crucial for continuity at the arithmetic gates and differs from, e.g., Rubinstein (2018).

¹⁰To see that ε -GCIRCUIT is not continuous, consider a single C_{\neg} gate, let $\varepsilon^n = 1/n$, $x^n[a_1] = 2/n$, and $x^n[v] = 0$. WLOG assume $n \geq 4$. Note that $x^n[a_1] \not\leq \varepsilon^n$ and $x^n[a_1] \not\geq 1 - \varepsilon^n$, so x^n is an ε^n -solution for all n . However, $x[a_1] = 0 \leq 0 = \varepsilon$, but $x[v] = 0 \not\geq 1 = 1 - \varepsilon$, so x is not an ε -solution.

Note further that continuity does *not* imply any statement regarding the “speed of convergence.” That is, we do not receive an upper bound on $\|x - x^n\|$ dependent on $|\varepsilon - \varepsilon^n|$. In light of the hardness results in Etessami and Yannakakis (2010) regarding *strong fixed points*, such a result seems unlikely to be obtainable.

ε -GCIRCUIT^{SB} is monotonic and offers access to Boolean gates, which makes it useful to study the hardness of generalized circuit problems themselves. However, since we made the Boolean gates stronger, it might be the case that ε -GCIRCUIT^{SB} is a strictly harder problem than ε -GCIRCUIT. It is an immediate consequence of the discussion in the following section that this is not the case. It will turn out that (even our stronger) Boolean gates do not actually add any expressiveness on top of the other gates.

4.2 Redundancy of Boolean Gates

We can construct the Boolean gates in the definition of ε -GCIRCUIT^{SB} from the arithmetic and comparison gates. Thus, the Boolean gates are redundant as a feature and we receive a reduction from $\Theta(\varepsilon)$ -GCIRCUIT^{NB} to ε -GCIRCUIT^{SB}. Recall that ε -GCIRCUIT^{NB} is the restriction of ε -GCIRCUIT where no Boolean gates are allowed. In the following, we write “ $\varepsilon \ll 1$ ” (read: “ ε sufficiently small”) to mean that a statement holds for all ε below a certain positive threshold. Unless indicated, the threshold is a constant that does not depend on the context of the statement.

Lemma 1. *For any generalized circuit C we can construct in polynomial time a circuit C' such that i) the nodes of C' are a superset of the nodes of C , ii) C' does not contain any Boolean gates, and iii) for any $\varepsilon \ll 1$, any $\varepsilon/2$ -solution for C' induces a strong ε -solution for C via restriction to the nodes of C .*

Proof. We need to model the Boolean gates. The C_\wedge gate is redundant because it can be expressed using C_\neg and C_\vee (recall that Boolean gates do not accumulate errors!). Assume $\varepsilon < 1/3$, let $\varepsilon' = \varepsilon/2$, and consider an ε' -solution.

We model C_\neg as the expression $a_1 < 1/2$ using a $C_>$ and a C_ζ gate. Call the output of the C_ζ gate z . We have $x[z] = 1/2 \pm \varepsilon'$. If $x[a_1] < 1/2 - \varepsilon$, then $x[a_1] < x[z] - \varepsilon'$ and thus the output of the comparison gate is $1 \pm \varepsilon' = 1 \pm \varepsilon$. Likewise for $x[a_1] > 1/2 + \varepsilon$.

We model C_\vee as $(a_1 > 1/2) + (a_2 > 1/2) > 1/2$. If one of $x[a_1]$ or $x[a_2]$ is $> 1/2 + \varepsilon$, then like above, the respective inner $C_>$ gate will return $1 \pm \varepsilon'$ and thus the output of the C_+ gate is $1 \pm 2\varepsilon' > 1/2 + \varepsilon'$ and the outer $C_>$ gate returns $1 \pm \varepsilon' = 1 \pm \varepsilon$. If $x[a_1], x[a_2] < 1/2 - \varepsilon$, then both inner $C_>$ gates return $0 \pm \varepsilon'$, C_+ returns a value $\leq 2\varepsilon' < 1/2 - \varepsilon'$, and the final $C_>$ gate returns $0 \pm \varepsilon' = 0 \pm \varepsilon$. \square

Note that Daskalakis, Goldberg and Papadimitriou (2009) previously suggested that one could simulate Boolean gates using arithmetic gates, expressing $x \vee y$ as $[x+y]$ and $\neg x$ as $1-x$. However, they also noted that this would lead to various complications

regarding accuracy (specifically, we note that such gates would accumulate errors when several of them are put in a row). This is why, as a matter of convenience, Daskalakis, Goldberg and Papadimitriou (2009) use dedicated Boolean gadgets that do not accumulate errors. In Lemma 1, we show how the Boolean gates can be represented by arithmetic and comparison gates without error accumulation.¹¹

The lemma immediately implies:

Corollary 1. *The problems $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$, $\varepsilon\text{-GCIRCUIT}$, and $\varepsilon\text{-GCIRCUIT}^{\text{NB}}$ are all PPAD-complete for $\varepsilon \ll 1$:*

Proof. We have:

$$\varepsilon/2\text{-GCIRCUIT}^{\text{NB}} \geq_{\text{P}} \varepsilon\text{-GCIRCUIT}^{\text{SB}} \geq_{\text{P}} \varepsilon\text{-GCIRCUIT} \geq_{\text{P}} \varepsilon\text{-GCIRCUIT}^{\text{NB}}$$

where “ \geq_{P} ” stands for polynomial-time reducibility. The first relation is by Lemma 1 and the others are trivial.

It is well-known that $\varepsilon\text{-GCIRCUIT}$ is in PPAD for all $\varepsilon > 0$. Thus, all of the problems are in PPAD for all $\varepsilon > 0$. For PPAD-hardness for $\varepsilon \ll 1$, it is enough to show that $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$ or $\varepsilon\text{-GCIRCUIT}$ are PPAD-hard for $\varepsilon \ll 1$. This follows via Rubinstein’s (2018) hardness proof for $\varepsilon\text{-GCIRCUIT}$. We defer a discussion of this proof to Section 5, where we show that the proof is not affected by an implicit monotonicity assumption and further applies to $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$ without modification. \square

Corollary 1 is useful because it implies that, when performing reductions from generalized circuits, there is no need to provide a reduction for the Boolean gates. In particular, via Lemma 1, all reductions from $\varepsilon\text{-GCIRCUIT}$ to other problems in prior work also provide a reduction from $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$.

5 Eliminating Issues With Non-Monotonicity in Prior Work

To the best of our knowledge, monotonicity has not been discussed in any piece of prior work on generalized circuits. This raises the question whether or not it has been carefully considered in the past. As explained in the previous section, mere reductions *from* $\varepsilon\text{-GCIRCUIT}$ to other problems automatically provide reductions from $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$ and will therefore not be affected. We thus take a close look at those pieces of work where hardness of the $\varepsilon\text{-GCIRCUIT}$ problem itself and its variants is established. Specifically, we discuss the three foundational papers on

¹¹It should be noted that the Boolean game gadgets in Daskalakis, Goldberg and Papadimitriou (2009) satisfy a monotonic definition of the Boolean gates that is of intermediate strength between $\varepsilon\text{-GCIRCUIT}$ and $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$. For the OR game gadget, for example, we have that $x[a_1] + x[a_2] > 1/2 + \varepsilon \Rightarrow x[v] = 1$ and $x[a_1] + x[a_2] < 1/2 - \varepsilon \Rightarrow x[v] = 0$. This is not quite enough for a strong ε -solution, but it is a monotonic property by itself. We discuss another such intermediate definition of Boolean gates in Appendix C.

generalized circuits: Daskalakis, Goldberg and Papadimitriou (2009), Chen, Deng and Teng (2009), and Rubinstein (2018). We show that the first of these papers is not affected by non-monotonicity while in contrast, non-monotonicity does create subtle technical issues in the latter two. We then show how replacing ε -GCIRCUIT by ε -GCIRCUIT^{SB} eliminates these issues. ε -GCIRCUIT^{SB} serves as a drop-in replacement for ε -GCIRCUIT, allowing us to keep all unaffected arguments the same.

We would like to stress that the purpose of our discussion is not to diminish the contributions of these seminal papers. Rather, we find it instructive to demonstrate what problems non-monotonicity can create by using the proofs in the three seminal papers as examples, rather than inventing examples ourselves. Note that the issues in prior work that we point out are of a mere technical nature and could be circumvented by careful argumentation. We present a way how this could be done without relying on ε -GCIRCUIT^{SB} in Appendix B. However, as we will see, ε -GCIRCUIT^{SB} provides a particularly clean and conceptually simple solution to these problems. We will now go through the proof steps in the three papers one by one. We will label the issues #1–#5, to refer back to them in the appendix.

5.1 Daskalakis, Goldberg and Papadimitriou (2009)

Daskalakis, Goldberg and Papadimitriou (2009) were the first to prove PPAD-hardness for the problem of finding an approximate Nash equilibrium, for an exponentially small ε . The proof is by reduction from a variant of the Brouwer fixed-point problem using a collection of *game gadgets*. These game gadgets correspond to a variant of ε -GCIRCUIT where the Boolean gates are defined via *exact* rather than approximately Boolean values. For example, the output of the NOT gate is 1 if the input is 0, 0 if the input is 1, and unrestricted otherwise. The comparison gate also yields an exact Boolean value rather than an approximately Boolean one. In contrast to the (nowadays more standard) definition of generalized circuits we have presented in Section 2, their variant of ε -GCIRCUIT satisfies monotonicity. Thus, this paper is not affected.

5.2 Chen, Deng and Teng (2009)

Chen, Deng and Teng (2009) proved PPAD-hardness of finding an approximate Nash equilibrium in a two-player game for polynomially small ε . En-route to this result, the authors provide the first explicit definition of the generalized circuit concept. In this early variant, values of nodes are truncated to $[0, 1/K]$ rather than $[0, 1]$, where K is the number of nodes of the circuit. Note that ε has to decrease at least linearly in K , otherwise the error term ε would eventually become larger than the range of the solution and the problem would become trivial. We call this variant of the problem ε -GCIRCUIT_C to distinguish it from Rubinstein’s, nowadays more

standard, variant. It is easy to see that ε -GCIRCUIT is computationally equivalent to ε/K -GCIRCUIT_C via scaling. Like Chen, Deng and Teng, we write POLY^c-GCIRCUIT_C for K^{-c} -GCIRCUIT_C.

The hardness proof in Chen, Deng and Teng (2009) proceeds in three steps. (1) The authors establish hardness of a variant of the BROUWER approximate fixed-point problem. (2) They reduce this problem to POLY³-GCIRCUIT_C. (3) They reduce POLY³-GCIRCUIT_C to the problem of finding an n^{-12} -approximate Nash equilibrium in a two-player game, where n is the number of actions. The last step is carried out using a collection of two-player game gadgets. Two additional reductions establish that the exponents do not actually matter for the complexity of the problems.

A part that demands some scrutiny is step 2, where Brouwer is reduced to POLY³-GCIRCUIT_C. Fortunately, detailed examination of the proof shows that no implicit monotonicity assumption is made. This is because a single ε (namely *exactly* $\varepsilon = K^{-3}$) is considered over the whole course of the proof. The same is true for the description of the game gadgets (step 3).

A place that *does* suffer from an implicit monotonicity assumption is the “padding theorem” (Chen, Deng and Teng, 2009, Theorem 5.7), where the authors show that the hardness of the POLY^c-GCIRCUIT_C problem does not increase if we increase c , as long as $c \geq 3$. The proof is by reduction from POLY^{2b+1}-GCIRCUIT_C to POLY³-GCIRCUIT_C, for any integer $b > 1$. However, since we do not have monotonicity, this only implies the statement for *odd integer* values of c . POLY⁴-GCIRCUIT_C, for example, might still be a harder problem. Further, and again due to the lack of monotonicity, we only receive a statement for the ε -GCIRCUIT_C problem where ε is exactly of form $\varepsilon = n^{-c}$ for some c . We do not receive any statement for arbitrary polynomials like $2n^{-3} + n^{-2}$. We call this issue #1.

To resolve this issue, we can define an ε -GCIRCUIT^{SB} analog to Chen, Deng and Teng’s (2009) version of generalized circuits. To do this, we replace in Table 2 truncation to $[0, 1]$ by truncation to $[0, 1/K]$ and for the Boolean gates, we replace the constant $1/2$ by $1/(2K)$. We then consider the problem K^{-c} -GCIRCUIT'_C where $c \geq 1$ is a constant. Note that, like before, ε has to decrease in K at least linearly. When applied to this variant of the generalized circuit concept, the proof in the paper yields:

Proposition 3 (Chen, Deng and Teng (2009), Theorem 5.7 for ε -GCIRCUIT^{SB}). *For any $c \geq 3$, K^{-c} -GCIRCUIT'_C \leq_P K^{-3} -GCIRCUIT'_C.*

Proof. Since we now have monotonicity, it is enough to show the statement for every c of form $c = 2b + 1$ where $b > 0$ is an integer. To this end, let C be a generalized circuit with $K \geq 2$ nodes and let $\varepsilon = K^{-c}$. The proof of Theorem 5.7 in Chen, Deng and Teng (2009) constructs a new circuit with $K' := K^b = 1/K \cdot K^{1-b}$ nodes such that for $\varepsilon' := \varepsilon K^{1-b}$, any strong ε' -solution for the new circuit gives rise to a strong

ε -solution for the original circuit via scaling by K^{1-b} . And $\varepsilon' = K'^{-3}$. \square

5.3 Rubinstein (2018)

Rubinstein (2018) proved PPAD-hardness for the problem of finding an ε -approximate Nash equilibrium for a sufficiently small *constant* ε .¹² The proof proceeds in four steps. (1) The author establishes hardness of a new class of instances of the BROUWER problem with constant ε . (2) He reduces this problem to ε -GCIRCUIT for a certain constant ε . (3) The author shows, using an additional black-box reduction, that ε -GCIRCUIT is still hard for some ε when we limit the fan-out¹³ of each gate to 2. (4) The author employs the game gadgets from Daskalakis, Goldberg and Papadimitriou (2009) to reduce *this* problem to the problem of finding an approximate Nash equilibrium in a degree-3 graphical game. Based on the considerations at the start of this sub-section, we should now take a closer look at steps 2-4.

The main hardness proof for ε -GCIRCUIT (step 2) is a reduction from the problem of finding an $\varepsilon^{1/4}$ -approximate fixed point of a certain function to ε -GCIRCUIT, for any sufficiently small ε , where the constructed ε -GCIRCUIT instance depends on ε . Detailed examination of the proof shows that none of the arguments implicitly assume monotonicity. As the construction can be performed for arbitrarily small ε , this indeed shows hardness of ε -GCIRCUIT for any sufficiently small ε (and not just for *one specific* ε , which is not a priori clear when monotonicity is not given).

The lack of monotonicity *does* lead to several problems in step 3, a black-box reduction from any given generalized circuit to a circuit with fan-out 2 (Rubinstein, 2018, page 941). In this reduction, the outputs of each comparison or Boolean gate with a fan-out greater than 2 are distributed using binary trees of double negation gates. The outputs of arithmetic gates, in contrast, are first transformed into a unary bit representation, then the resulting Boolean values are distributed using the aforementioned trees of double negation gates, and finally each copy is converted back into its numeric form. This *distribution subroutine* has maximum fan-out 2 and guarantees that each of its outputs is equal to its input with an error of $\pm\varepsilon$ in any ε^2 -solution. One thus has to reduce the allowed error to $\varepsilon' \in \Theta(\varepsilon^2)$.¹⁴

There are three problems with this reduction, all of which arise from non-monotonicity at Boolean gates and all of which can lead to a situation where some

¹²To clarify the relationship between the three papers: the proof in Rubinstein (2018) is for the sub-class of polymatrix degree-3 graphical games, but does not extend to two-player games. It is therefore an unambiguous improvement upon the main result in Daskalakis, Goldberg and Papadimitriou (2009), who considered the same class of games and *exponentially small* ε , but not upon Chen, Deng and Teng (2009), who considered two-player games. In two-player games, the problem is likely not PPAD-hard for constant ε (Rubinstein, 2018).

¹³The *fan-out* of a gate g is the number of gates g' such that the output node of g is an input node of g' .

¹⁴Note that we have interchanged ε' and ε compared to Rubinstein's (2018) formulation of the theorem, notation-wise, to ensure consistency of notation within the present paper.

ε' -solution to the reduced circuit is not an ε -solution to the original circuit. We provide detailed examples for this in Appendix A. For the first issue, consider a Boolean gate that already has fan-out ≤ 2 . Since no changes are made to such gates and we do not have monotonicity, the ε' -solution to the reduced circuit may fail to be an ε -solution for the original circuit (call this issue #2). Next, there may be arithmetic gates with fan-out > 2 whose outputs feed into Boolean gates. Here, the distribution subroutine introduces an additional error, which may turn values from approximately-Boolean to not-approximately-Boolean and may thus not correctly copy these values (issue #3). Finally, there may be Boolean gates with fan-out > 2 that feed into arithmetic gates. For these gates, the fact that we use double negation gates to distribute the outputs creates additional degrees of freedom in the reduced circuit when the inputs to Boolean gates are not approximately Boolean and thus their outputs are arbitrary (issue #4).

When we replace ε -GCIRCUIT by ε -GCIRCUIT^{SB}, issue #2 is immediately resolved because ε -GCIRCUIT^{SB} has monotonicity. To eliminate issues #3 and #4, the fact that we use ε -GCIRCUIT^{SB} allows us to make a modification to Rubinstein's original proof to obtain the following lemma:

Proposition 4 (Rubinstein's (2018) fan-out 2 reduction for ε -GCIRCUIT^{SB}). *For any $\varepsilon \ll 1$, there is an $\varepsilon' \in \Theta(\varepsilon^2)$ such that there is a polynomial-time reduction from ε -GCIRCUIT^{SB} to the restriction of ε' -GCIRCUIT^{SB} to maximum fan-out 2.*

Proof. Let $\bar{\varepsilon} = \varepsilon/3$. Assume that $\varepsilon \ll 1$ in a way to be made precise below. We perform the construction in Rubinstein (2018, Theorem 1.6) with respect to $\bar{\varepsilon}$ where we make the following modification: instead of differentiating between the outputs of Boolean/comparison vs. arithmetic gates, we *always* apply the distribution subroutine to the output of any gate with fan-out > 2 . Recall that this subroutine has one input and any number of outputs and ensures that for a certain $\varepsilon' \in \Theta(\bar{\varepsilon}^2) = \Theta(\varepsilon^2)$, in an ε' -solution, each output equals the input up to an error of $\pm\bar{\varepsilon}$. It is easy to see that the distribution subroutine itself is not affected by any of the issues related to the fan-out 2 reduction. Assume that $\varepsilon' \leq \bar{\varepsilon}$.

Let C be the original circuit, C' the reduced circuit, and x' a strong ε' -solution to C' . We show that the restriction of x' to nodes in C is a strong ε -solution for C . Let g be a gate with inputs a_1 and a_2 (if any). Assume WLOG that distribution is applied to each of the inputs to g and let a'_i be the output of the respective distribution subroutine that is the new input to g in C' . Let v be the output of g in C and C' . We have $x'[a_i] = x'[a'_i] \pm \bar{\varepsilon}$ by the distribution subroutine. We perform case distinction over the type of g .

- If g is an arithmetic gate, then for a sufficiently small $\varepsilon' \in \Theta(\varepsilon^2)$ it follows from Lipschitz continuity that x is an $\bar{\varepsilon}$ -solution, and thus a strong ε -solution at g , just like in Rubinstein (2018).

- If g is a comparison gate, assume WLOG that $x[a_1] < x[a_2] - \varepsilon$, i.e., $x[a_1] < x[a_2] - 3\bar{\varepsilon}$. By the distribution subroutine, $x'[a'_1] < x'[a'_2] - \bar{\varepsilon} \leq x'[a'_2] - \varepsilon'$ and thus, since x' is a strong ε' -solution, $x'[v] \geq 1 - \varepsilon' \geq 1 - \varepsilon$ as required.
- If g is a Boolean gate, consider any input a_i to g . If $x'[a_i] < 1/2 - \varepsilon$, then $x'[a'_i] < 1/2 - \varepsilon + \bar{\varepsilon} \leq 1/2 - \varepsilon'$. Thus, if any input to g is approximately Boolean FALSE w.r.t. ε in C , then it is approximately Boolean FALSE w.r.t. ε' in C' , and likewise for TRUE. Thus, if we require, based on the constraints, that $x'[v] \leq \varepsilon$ in a strong ε -solution of C , we require $x'[v] \leq \varepsilon'$ in a strong ε' -solution of C' . And the latter implies the former. Likewise for $x'[v] \geq 1 - \varepsilon$. \square

Observe how the above proof eliminates issues #3 and #4 compared to Rubinstein's original proof. Issue #3 is eliminated in the last step (Boolean gates) and this crucially depends on the fact that we consider ε -GCIRCUIT^{SB} instead of ε -GCIRCUIT: by monotonicity, we can choose ε' sufficiently small to compensate for the additional error in the distribution subroutine. Issue #4 is eliminated because we use the distribution subroutine, which does not create additional degrees of freedom, for all gates.

The proposition together with Corollary 1 immediately yields:

Corollary 2. *For each of the problems ε -GCIRCUIT^{SB}, ε -GCIRCUIT, and ε -GCIRCUIT^{NB}, the restriction to maximum fan-out 2 is PPAD-complete for $\varepsilon \ll 1$.*

Proof. For ε -GCIRCUIT^{SB}, this follows from hardness of ε -GCIRCUIT^{SB} and Proposition 4. For ε -GCIRCUIT^{NB}, we observe that the reduction in Lemma 1 preserves the fan-out 2 property. For ε -GCIRCUIT, it now follows trivially. \square

Note that the lack of monotonicity introduces another subtle technical issue in Rubinstein (2018), specifically in the reduction from ε -GCIRCUIT to the problem of finding an approximate Nash equilibrium (step 4 in the outline of the proof above). Rubinstein uses the same game gadgets as Daskalakis, Goldberg and Papadimitriou (2009). However, for the Boolean game gadgets in that paper, we only know at this point that they work with exact Boolean values 0 and 1 in both the input and output, not necessarily approximately Boolean ones (see Section 5.1 above). And the former does not imply the latter because we do not have monotonicity. A priori, these gadgets might *rely* on receiving only values *exactly* 0 or 1 as their inputs (call this issue #5). Fortunately, this issue is eliminated immediately using Lemma 1: since the Boolean gates are redundant, it is not necessary to provide a reduction for them in the first place.¹⁵

¹⁵The way how we eliminated issue #5 may not feel satisfying to some readers because the graphical games that the Boolean gates are ultimately reduced to (via Lemma 1 and the game gadgets for the other gates) will be quite complicated. For those cases where a more direct representation is desired, we present a third PPAD-complete and monotonic variant of the ε -GCIRCUIT problem, called ε -GCIRCUIT ^{β} , that allows for this. The parameter β needs to be appropriately specified. See

Remark 3 (Exact Boolean Gates). For Rubinstein (2018), there is another solution to the problems with non-monotonicity: rather than using $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$, adopt the definition of generalized circuits from Daskalakis, Goldberg and Papadimitriou (2009), where only exact Boolean values are mapped to each other (see Section 5.1), and show hardness of this variant. For the case of graphical games, the game gadgets from Daskalakis, Goldberg and Papadimitriou (2009) provide a reduction from this variant with exact Boolean values. In many other applications, however, such a reduction is not possible. For example, any approximate fixed point problem inherently has ε errors in every dimension, so we cannot ever expect to receive values *exactly* equal to 0 or 1. The two-player game gadgets in Chen, Deng and Teng (2009) and the market gadgets in Othman, Papadimitriou and Rubinstein (2016) also have this inherent limitation. Note that exact Boolean gates could be represented using a variant of Lemma 1 only once we have access to a comparison gate that produces an *exact* Boolean output value, and such a gate does not seem to be attainable for the previously-mentioned applications. Thus, using exact Boolean values would greatly diminish the applicability of the generalized circuits framework.

6 Conclusion

Generalized circuits are a vital tool for reasoning about the computational complexity of equilibrium approximation problems. In this paper, we have revealed a conceptual issue in the generalized circuits framework, namely that it lacks *monotonicity* of its approximate solution concept. We have shown that this creates subtle technical issues, including in prior work. To overcome these issues, we have shown that the Boolean gate types in these circuits are redundant features and that stronger Boolean gates can be defined based on the other (arithmetic and comparison) gates. We have shown that the resulting (equivalent) $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$ problem satisfies monotonicity, serves as a drop-in replacement in prior work, and then eliminates the mentioned issues at a conceptual level. We have established monotonicity as a fundamental desideratum for any approximate solution concept.

Our results have implications for two potential future lines of research. First, future studies of generalized circuits (for example, hardness proofs for sub-classes of circuits) can consider either the $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$ problem or the $\varepsilon\text{-GCIRCUIT}^{\text{NB}}$ problem, i.e., ignore Boolean gates altogether. Both of these variants satisfy monotonicity, which makes for a much more natural way of reasoning and avoids the kinds of technical pitfalls we have discussed. This may lead to new insights about computational complexity in generalized circuits. One such area of research are “support finding”

Appendix C. Note that our proof in Appendix C implies that the Boolean game gadgets in *do* actually satisfy the constraints for $\varepsilon\text{-GCIRCUIT}$, even though this is not shown in the paper. $\varepsilon\text{-GCIRCUIT}$ is, however, not monotonic and the gadgets do not satisfy the stronger Boolean constraints in $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$.

problems, where we do not ask for numeric values, but only for a coarse discrete description of a solution. In our own recent work (Schuldenzucker, Seuken and Battiston, 2019, Section 5), we have studied one such PPAD-complete problem to prove hardness in the context of financial networks.

Another example where $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$ could be useful is a conjecture by Babichenko, Papadimitriou and Rubinstein (2016) that the following problem, termed $(\varepsilon, \delta)\text{-GCIRCUIT}$, is already PPAD-complete for $\varepsilon, \delta \ll 1$: given a generalized circuit, find an assignment where the constraints for an ε -solution hold at least at a fraction of $1 - \delta$ of the gates. This would settle various open questions regarding the Nash equilibrium search problem. Given the benefits of monotonicity illustrated in this paper and towards a proof of the conjecture, it might be useful to instead consider the $(\varepsilon, \delta)\text{-GCIRCUIT}^{\text{SB}}$ problem. Note that $(\varepsilon, \delta)\text{-GCIRCUIT}^{\text{SB}}$ is monotonic in both parameters and $(\varepsilon, \delta)\text{-GCIRCUIT}^{\text{SB}} \leq_{\text{P}} (\varepsilon/2, \Theta(\delta))\text{-GCIRCUIT}^{\text{NB}}$ by Lemma 1.

The second strand of future research concerns reductions from generalized circuits to other problems to show PPAD-hardness of these problems. The redundancy of Boolean gates implies that no reduction needs to be provided for them, which will hopefully simplify these kinds of proofs in the future. Since the reduction now happens between two monotonic problems, their connection may further become more natural and allow for a more detailed study of common features.

Appendix

A Examples for Issues #2–#4 in the Fan-Out 2 Reduction in Rubinstein (2018)

We present examples for issue #2–#4.

Issue #2 For issue #2, an example is given by our very first counterexample to monotonicity in Section 3.

Issue #3 For issue #3, assume that v is the output of some arithmetic gate, let $g = C_v$ with input v , and assume that v is input to at least two other gates so that its value needs to be distributed. Call this original circuit C and let C' be the circuit where a distribution subroutine is inserted after v . Let v' be an output of the distribution subroutine and the new input to g in C' . Let w be the output of g in C and C' . Assume that there exists an ε' -solution x' to C' such that $x'[v] \geq 1 - \varepsilon'$, $x'[v'] \in (\varepsilon', 1 - \varepsilon')$, and $x'[w] = 0.5$. The distribution subroutine does not prevent this, no matter what ε and ε' , and it is easy to construct C such that this actually happens. In C , we have for the input of g that $x'[v] \geq 1 - \varepsilon' \geq 1 - \varepsilon$, but for the output $x[w] = 0.5$. So x' is not an ε -solution for C . Note that this counterexample

does not depend on the fact that $\varepsilon \neq \varepsilon'$.

Issue #4 For issue #4, consider a generalized circuit C corresponding to the following definitions (where “=” assigns an output node to a gate):

$$\begin{aligned} z &= 0.3 \\ b &= (a < z) \\ c &= 1/2 \cdot b \\ d &= 1/3 \cdot b \\ e &= 1/4 \cdot b \end{aligned}$$

Note that node a is left unconstrained. We imagine that nodes $a\text{--}e$ are part of a larger circuit. Node b has fan-out $3 > 2$, so Rubinstein’s fan-out 2 reduction would attach a tree of double negation subroutines. The double negation subroutine is simply a chain of two negation gates connected by a new node. This turns an approximate TRUE into an approximate TRUE and an approximate FALSE into an approximate FALSE, but can return *any* value if its input is not approximately Boolean. The fan-out 2 reduction would now replace the definitions of nodes $c\text{--}e$ by the following to create a new reduced circuit C' :

$$\begin{aligned} b_1 &= \neg\neg b \\ b_2 &= \neg\neg b \\ b_{1,1} &= \neg\neg b_1 \\ b_{1,2} &= \neg\neg b_1 \\ b_{2,1} &= \neg\neg b_2 \\ c &= 1/2 \cdot b_{1,1} \\ d &= 1/3 \cdot b_{1,2} \\ e &= 1/4 \cdot b_{2,2} \end{aligned}$$

We now present a solution solution x' to C' that does not give rise to a solution to C . We will show that c and d can take on a combination of values in C' that is

not possible in C . Let $\varepsilon = 0.01$. Define x' as follows:

$$\begin{aligned}
x'[a] &:= x'[z] := 0.3 \\
x'[b] &:= 0.5 \\
x'[b_1] &:= 0.8 \\
x'[b_2] &:= 0.2 \\
x'[b_{1,1}] &:= x'[b_1] = 0.8 \\
x'[b_{1,2}] &:= x'[b_1] = 0.8 \\
x'[b_{2,1}] &:= x'[b_2] = 0.2 \\
x'[c] &:= 1/2 \cdot x'[b_{1,1}] = 1/2 \cdot 0.8 = 0.4 \\
x'[d] &:= 1/3 \cdot x'[b_{1,2}] = 1/3 \cdot 0.8 = 0.2\bar{6} \\
x'[e] &:= 1/4 \cdot x'[b_{2,1}] = 1/4 \cdot 0.2 = 0.05
\end{aligned}$$

For the interior nodes of the double negation subroutines, if the input node to the subroutine is v , set the interior node to value $1 - x'[v]$. This is always possible.

x' is an ε -solution for C' . Note that, by choice of $x'[a]$, any value is allowed for $x'[b]$. We chose a value that is not approximately Boolean w.r.t. ε . That is why the following double negation subroutines can each output an arbitrary value at $x'[b_1]$ and $x'[b_2]$. The key to our counterexample is that these values need not be the same. The other gates then copy and transform the values normally.

x' does not become an ε -solution for C if we restrict it to nodes in C . That is because, if x is any ε -solution to C , then $x[c] - x[e] = 1/2 \cdot x[b] \pm \varepsilon - 1/4 \cdot x[b] \pm \varepsilon = 1/4 \cdot x[b] \pm 2\varepsilon$. However, we have $x'[c] - x'[e] = 0.35 > 0.145 = 1/4 \cdot x'[b] + 2\varepsilon$.

Note further that i) the above value of $x'[c] - x'[e]$ would not be allowed in C' for *any* value of $x'[b]$ and ii) we cannot guarantee the ε -solution property by increasing ε by any constant factor.

B Minimal Modifications to Circumvent the Issues in Prior Work

We present a minimal set of modifications to Rubinstein (2018) and Chen, Deng and Teng (2009) that allow us to keep the current definition of the ε -GCIRCUIT problem and that eliminate the problems discussed above. Our modifications are based on careful examination of the details of the involved proofs.

To show that issue 1–5 in Section 5 are not critical for the results of the respective papers, we exploit a common feature of the generalized circuit constructions from prior work, namely that Boolean gates do not occur at arbitrary positions. Their inputs always come from gates that are meant to yield approximately Boolean values,

namely other Boolean gates and the comparison gate. Further, the output of each gate will be interpreted either as a Boolean value (by Boolean gates) or as a non-Boolean value (by other gates), but not both at the same time. Such circuits formally *still* do not satisfy monotonicity, but we can perform an additional normalization step after which they “essentially” do.

Lemma 2 (Boolean-regular circuit). *If g and g' are gates in a circuit such that the output of g is an input to g' , then g is called a predecessor of g' and g' is called a successor of g . We call a generalized circuit Boolean-regular if the following two conditions hold:*

1. *Any predecessor of any Boolean gate is either a Boolean gate itself or a comparison gate.*
2. *For any gate, if one of its successors is a Boolean gate, then all of its successors are Boolean gates.*

If C is Boolean-regular, then for any ε and any ε -solution x for C , we can compute in polynomial time an assignment x' such that for any $\varepsilon' \geq \varepsilon$, x' is an ε' -solution for C . We call an x' resulting from this procedure normalized.

Proof. Given x , define x' as follows. If v is not an input to any Boolean gate, then $x'[v] = x[v]$. If v is an input to a Boolean gate, then

$$x'[v] = \begin{cases} 0 & \text{if } x[v] \leq \varepsilon \\ 1/2 & \text{if } x[v] \in (\varepsilon, 1 - \varepsilon) \\ 1 & \text{if } x[v] \geq 1 - \varepsilon. \end{cases}$$

Let now $\varepsilon' \geq \varepsilon$. We show that x' is an ε' -solution. Let g be any gate with inputs a_1 and a_2 (if any) and output v . We distinguish three cases.

- If g is an arithmetic gate, then neither its output (by condition 1) nor any of its inputs (by condition 2) are input to any Boolean gate. Thus, $x' = x$ at these nodes. Since the constraints of arithmetic gates are monotonic in ε , the constraint at g is still satisfied for ε' .
- If g is a Boolean gate, then its constraints only distinguish the intervals $[0, \varepsilon']$, $(\varepsilon', 1 - \varepsilon')$, and $[1 - \varepsilon', 1]$. For each input a_i of g , by definition of x' it does not depend on ε' to which of these three intervals $x'[a_i]$ belongs. Therefore, we require $x'[v] = 0 \pm \varepsilon'$ in an ε' -solution iff we require $x[v] = 0 \pm \varepsilon$ in an ε -solution. And the latter implies the former, both if v is the input to another Boolean gate and if not. Likewise for $x'[v] = 1 \pm \varepsilon'$.
- If g is a comparison gate, by condition 2 its inputs are not also input to any Boolean gate and thus $x'[a_i] = x[a_i]$ for $i = 1, 2$. Now we apply the same

argument as for the outputs of Boolean gates to see that the constraint is still satisfied. \square

Detailed examination of the proofs in the aforementioned two pieces of prior work shows that almost all construction steps lead to a Boolean-regular circuit. The only exception we are aware of is the EXTRACTBITS subroutine in Chen, Deng and Teng (2009), where the output of a $C_<$ gate is fed into both Boolean gates (which simulate a given Boolean circuit) and a $C_{\times\zeta}$ gate. Here, Boolean regularity can be easily restored by a minor modification to the construction.¹⁶

Towards issue #1 in Chen, Deng and Teng (2009), we can now WLOG consider the restriction of the ε -GCIRCUIT_C problem where only Boolean-regular circuits are allowed as input and only normalized ε -solutions are allowed as output. Since this problem has monotonicity by definition of a normalized solution and the reduction in the proof of Theorem 5.7 in Chen, Deng and Teng (2009) preserves Boolean-regularity, issue #1 is eliminated.

Towards issue #2 and #3 in Rubinstein (2018), we notice that the fan-out 2 reduction preserves Boolean-regularity.¹⁷ The restriction of ε -GCIRCUIT to Boolean-regular circuits and normalized solutions then resolves issues #2 (because it has monotonicity) and #3 (because the described situation does not occur by Boolean-regularity).

To see that issue #4 does not invalidate hardness of ε -GCIRCUIT restricted to fan-out 2, we again perform detailed examination of the arguments that are used in the main hardness proof. Issue #4 arises because the values at outputs of Boolean gates with non-Boolean input are allowed to be arbitrary and different in the reduced instance while they are arbitrary, but must be equal in the original instance (see our example in Appendix A). However, such a property is never exploited in the proof of hardness of the ε -GCIRCUIT problem. Instead, whenever the output of a Boolean gate can be arbitrary, it is accounted for as an independent ± 1 error. Thus, if we apply the fan-out 2 reduction to the hard ε -GCIRCUIT instance, a solution to the reduced circuit is not necessarily a solution to the original circuit, but it *is* still a solution to the original hard BROUWER instance. And thus, the restriction to fan-out 2 is still hard.

Finally, to eliminate issue #5, one can study the Boolean game gadgets from Daskalakis, Goldberg and Papadimitriou (2009) to see that they in fact *do* satisfy the constraints for approximately Boolean values even though this is not stated explicitly in the paper. The proof is like in Proposition 6 in Appendix C, where we show it for ε -GCIRCUIT ^{β} .

¹⁶One way to restore Boolean-regularity is to insert a double negation in front of the $C_{\times\zeta}$ gate. This will, of course, create additional degrees of freedom like in issue #4 (see Section 5). These are not a problem in this case for the same reason why issue #4, discussed below, is not critical.

¹⁷Here we assume WLOG that the trees of double negations are constructed in such a way that all outputs of the tree are all at the same level.

C A More Direct Representation of Boolean Gates

The way how we eliminated issue #5 may not feel very satisfying. When we perform reduction from $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$ to other problems via Lemma 1, the representation of the Boolean gates will be rather indirect. Each Boolean gate is first represented by comparison gates, arithmetic gates, and using De Morgan’s laws. Then *these* gates are represented as (say) game gadgets. In some situations, a more direct representation of Boolean gates may be desirable. This could be useful, for example, if one seeks to further modify the generalized circuit concept in a way incompatible with Lemma 1.

In this section, we present a way how such a direct representation of monotonic Boolean gates can be achieved. For our discussion, we focus on the reduction from generalized circuits to graphical games via the game gadgets in Daskalakis, Goldberg and Papadimitriou (2009). These are the same gadgets used in Rubinstein (2018). We will show that these game gadgets do not provide a reduction from $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$. To overcome this, we will modify our solution concept again, which will lead to a new family of PPAD-complete search problems $\varepsilon\text{-GCIRCUIT}^\beta$, where $\beta \in (0, 1/2)$ is a parameter. We then show that the game gadgets provide a reduction from $\varepsilon\text{-GCIRCUIT}^\beta$ if β is not too small. A drawback of this variant is that the β parameter needs to be appropriately chosen for the individual application at hand.

Daskalakis, Goldberg and Papadimitriou (2009) and Rubinstein (2018) study binary graphical games in ε -*approximately well supported Nash equilibrium* ($\varepsilon\text{-WSNE}$ for short). This means that players only have two actions, called 0 and 1, and if both strategies are played with positive probability in equilibrium, then the expected utilities from both pure actions must be ε -close to each other.¹⁸ A mixed-strategy equilibrium of a binary game can be encoded by assigning to each player i the probability $p[i]$ with which player i plays action 1. *Game gadgets* are sub-games that in equilibrium enforce certain relationships, corresponding to the gates of a generalized circuit, on the $p[i]$ values of certain players.

The negation game gadget \mathcal{G}_\neg (Daskalakis, Goldberg and Papadimitriou, 2009, Lemma 5.5) satisfies the constraints for a strong ε -solution, but the other two, \mathcal{G}_\wedge and \mathcal{G}_\vee , do not. We consider \mathcal{G}_\wedge in the following. The proof for \mathcal{G}_\vee is analogous. Let a and b be two *input players* and let v be an *output player*. The utility function of player v in \mathcal{G}_\wedge is defined as follows:

$$u_v = \begin{cases} 1/2 & \text{if } v \text{ plays 0} \\ 1 & \text{if } v \text{ plays 1} \wedge a \text{ plays 1} \wedge b \text{ plays 1} \\ 0 & \text{if } v \text{ plays 1} \wedge (a \text{ plays 0} \vee b \text{ plays 0}) \end{cases}$$

¹⁸Daskalakis, Goldberg and Papadimitriou (2009) prove that $\varepsilon\text{-WSNE}$ and regular ε -approximate Nash equilibrium (where no deviation to any other mixed strategy can improve expected utility by more than ε) are equivalent if one scales ε appropriately.

If player v plays a pure strategy and the other players play mixed strategies according to p , the expected utility of v is

$$\mathbb{E}[u_v] = \begin{cases} 1/2 & \text{if } v \text{ plays 0} \\ p[a]p[b] & \text{if } v \text{ plays 1.} \end{cases}$$

This does not provide a reduction from $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$, no matter how much we reduce ε in the transition from generalized circuits to games:

Proposition 5. *There exists an $\varepsilon > 0$ such that there is no $\varepsilon' > 0$ such that, whenever \mathcal{G}_\wedge occurs as part of a larger game and p is an ε' -WSNE, $x := p$ satisfies the constraints for C_\wedge for a strong ε -solution.*

Proof. Consider an ε for which such an ε' does exist. Let $p[a] = p[b] = 1/2 + 2\varepsilon$. Then Table 2 prescribes that $p[v] \geq 1 - \varepsilon$. To guarantee *any* statement of form “ $p[v] \geq \dots$ ” in an ε' -WSNE, we require

$$\left(\frac{1}{2} + 2\varepsilon\right)^2 = p[a]p[b] = \mathbb{E}[u_v](1, p_{-v}) > \mathbb{E}[u_v](0, p_{-v}) = 1/2 + \varepsilon'.$$

By simple algebra, this implies that

$$\varepsilon > \frac{1}{24} + \frac{1}{6}\varepsilon' > \frac{1}{24}.$$

Therefore, for $\varepsilon \leq \frac{1}{24}$, we can always choose $p[v] = 0$ even though the constraints for a strong ε -solution prescribe $p[v] \geq 1 - \varepsilon$. Thus, $x := p$ is not a strong ε -solution. \square

The previous proposition shows that the game gadgets in Daskalakis, Goldberg and Papadimitriou (2009) do not imply sufficiently strong constraints to imply a direct representation of the Boolean gates in $\varepsilon\text{-GCIRCUIT}^{\text{SB}}$. However, we can make the solution concept itself slightly weaker to accommodate these gadgets while preserving monotonicity and hardness.

To do this, let $\beta < 1/2$ and $\varepsilon < \beta, 1/2 - \beta$. Given a generalized circuit, we call an assignment x an ε^β -solution¹⁹ if it satisfies the constraints in Table 2 where we replace ε by β in the preconditions of all Boolean gates. That is, for the Boolean gates we have the constraints in Table 3. We call the corresponding search problem $\varepsilon\text{-GCIRCUIT}^\beta$.²⁰

¹⁹We chose this notation to avoid confusion with the (unrelated) concept of an (ε, δ) -solution in Babichenko, Papadimitriou and Rubinstein (2016), where a $1 - \delta$ fraction of constraints needs to be satisfied up to precision ε .

²⁰Our definition of $\varepsilon\text{-GCIRCUIT}^\beta$ is inspired by Othman, Papadimitriou and Rubinstein (2016), where we however do *not* consider $\beta = \Theta(\varepsilon)$, but $\varepsilon \ll \beta$. Note further that we do not use β in the precondition of the comparison gate. This would make for an even weaker problem, but a *too* weak one: Rubinstein's (2018) hardness proof performs comparison with multiples of $\sqrt{\varepsilon}$ and the “brittleness” of the comparison gate needs to be significantly smaller than that.

Table 3 Constraints for the Boolean gates in an ε^β -solution. All other constraints are the same as in a (strong) ε -solution.

Gate Type	Short	Constraint
OR	C_\vee	$x[a_1] < 1/2 - \beta$ and $x[a_2] < 1/2 - \beta \Rightarrow x[v] \leq \varepsilon$ $x[a_1] > 1/2 + \beta$ or $x[a_2] > 1/2 + \beta \Rightarrow x[v] \geq 1 - \varepsilon$
AND	C_\wedge	$x[a_1] < 1/2 - \beta$ or $x[a_2] < 1/2 - \beta \Rightarrow x[v] \leq \varepsilon$ $x[a_1] > 1/2 + \beta$ and $x[a_2] > 1/2 + \beta \Rightarrow x[v] \geq 1 - \varepsilon$
NOT	C_{\neg}	$x[a_1] < 1/2 - \beta \Rightarrow x[v] \geq 1 - \varepsilon$ $x[a_1] > 1/2 + \beta \Rightarrow x[v] \leq \varepsilon$

The two parameters ε and β take on different roles. Typically, β will be fixed to an arbitrary, not necessarily small, constant, like $1/4$. Then ε is chosen arbitrarily small. It is easy to see that the solution concept is monotonic in both parameters and that, for any fixed β and sufficiently small ε depending on β , any strong ε -solution is an ε^β -solution and any ε^β -solution is an ε -solution. This immediately implies that ε -GCIRCUIT $^\beta$ is PPAD-complete for any $\beta < 1/2$ and $\varepsilon \ll 1$ depending on β .

The Boolean game gadgets satisfy the constraints for an ε^β -solution, and thus define a reduction from ε -GCIRCUIT $^\beta$, as long as β is not too small.

Proposition 6. *Let $1/4 < \beta < 1/2$ and let $\varepsilon \leq \beta - 1/4$. Let $o \in \{\vee, \wedge, \neg\}$ and consider the binary graphical game \mathcal{G}_o from Daskalakis, Goldberg and Papadimitriou (2009, Lemma 5.5) with input players a and b (if any) and output player v . Then any ε -WSNE p satisfies the constraint corresponding to the gate C_o and $x := p$ for an ε^β -solution.*

Proof. We show the statement for $o = \wedge$. The other operations are similar. Assume that $p[a] > 1/2 + \beta$ and $p[b] > 1/2 + \beta$. Then $\mathbb{E}[u_v](1, x_{-v}) = p[a]p[b] > (1/2 + \beta)^2 > 1/2 + \varepsilon = \mathbb{E}[u_v](0, x_{-v}) + \varepsilon$, where the middle inequality is by choice of β and ε . Since we are in an ε -WSNE, this implies $p[v] = 1$ and in particular $p[v] \geq 1 - \varepsilon$.

Vice versa, assume that $p[a] < 1/2 - \beta$ or $p[b] < 1/2 - \beta$. Then $\mathbb{E}[u_v](1, x_{-v}) = p[a]p[b] < 1/2 - \beta < 1/2 - \varepsilon = \mathbb{E}[u_v](0, x_{-v}) - \varepsilon$ and thus $x[v] = 0 \leq \varepsilon$ by the ε -WSNE. \square

By more careful analysis of the error terms in the previous proof, one can show that the gadget still works for all $\beta > (\sqrt{2} - 1)/4 \approx 0.10$ and $\varepsilon \ll 1$ (the threshold for ε depending on β), but not for smaller β .

References

- Babichenko, Yakov, Christos Papadimitriou, and Aviad Rubinstein. 2016. “Can Almost Everybody Be Almost Happy?” In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*. Cambridge, Massachusetts, USA:ACM.

- Chen, Xi, Dimitris Paparas, and Mihalis Yannakakis.** 2017. “The Complexity of Non-Monotone Markets.” *Journal of the ACM (JACM)*, 64(3): 20:1–20:56.
- Chen, Xi, Xiaotie Deng, and Shang-Hua Teng.** 2009. “Settling the complexity of computing two-player Nash equilibria.” *Journal of the ACM (JACM)*, 56(3): 14.
- Daskalakis, Constantinos.** 2013. “On the Complexity of Approximating a Nash Equilibrium.” *ACM Transactions on Algorithms (TALG)*. Special Issue for SODA 2011, Invited.
- Daskalakis, Constantinos, Paul W Goldberg, and Christos H Papadimitriou.** 2009. “The complexity of computing a Nash equilibrium.” *SIAM Journal on Computing*, 39(1): 195–259.
- Etessami, Kousha, and Mihalis Yannakakis.** 2010. “On the complexity of Nash equilibria and other fixed points.” *SIAM Journal on Computing*, 39(6): 2531–2597.
- Othman, Abraham, Christos Papadimitriou, and Aviad Rubinstein.** 2016. “The complexity of fairness through equilibrium.” *ACM Transactions on Economics and Computation (TEAC)*, 4(4): 20.
- Rubinstein, Aviad.** 2018. “Inapproximability of Nash equilibrium.” *SIAM Journal on Computing*, 47(3): 917–959.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2017. “Finding Clearing Payments in Financial Networks with Credit Default Swaps is PPAD-complete.” In *Proceedings of the 8th Innovations in Theoretical Computer Science (ITCS) Conference*. Berkeley, USA.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2019. “The Computational Complexity of Financial Networks with Credit Default Swaps.” Working Paper, <https://arxiv.org/abs/1710.01578>.

5 Portfolio Compression in Financial Networks: Incentives and Systemic Risk

I was just guessing at numbers and figures,
Pulling your puzzles apart.
Questions of science, science and progress,
Do not speak as loud as my heart.

Tell me you love me,
Come back and haunt me,
Oh and I rush to the start.
Running in circles, chasing our tails,
Coming back as we are.

Coldplay, *The Scientist*

The content of this chapter has previously appeared in:

“Portfolio Compression in Financial Networks: Incentives and Systemic Risk.”
Steffen Schuldenzucker and Sven Seuken. Working Paper, November 2019.

Portfolio Compression in Financial Networks: Incentives and Systemic Risk

Steffen Schuldenzucker

University of Zurich

schuldenzucker@ifi.uzh.ch

Sven Seuken

University of Zurich

seuken@ifi.uzh.ch

First version: November 9, 2019

This version: November 23, 2019

Abstract

We study *portfolio compression*, a post-trade mechanism that eliminates cycles in a financial network. We study the incentives for banks to engage in compression and its systemic effects in terms of all banks' equities. We show that, contrary to conventional wisdom, compression may be socially and individually detrimental and incentives may be misaligned with the social good. We show that these effects depends on the parameters of the financial system and the compression in a complex and non-monotonic way. We then present sufficient conditions under which compression is incentivized for participating banks or a Pareto improvement for all banks. More in detail, compression is universally beneficial when interbank payments are subject to high default costs, when recovery rates are high, or when the participating banks' balance sheets are sufficiently homogeneous. Furthermore, we show that banks only have an incentive to reject a compression if there are *feedback paths* of links that are not compressed. Our results contribute to a better understanding of the implications of recent regulatory policy.

1 Introduction

The 2008 financial crisis is widely regarded as the result of a shock on a complex, opaque, unregulated network of dependencies among financial institutions.¹ In normal times, connections to other institutions provide financial actors with means to secure funding and hedge risks. Since potential losses are distributed to many other institutions, *diversification* stabilizes the system as a whole. In the 2008 crisis, however, it quickly became clear that this network could just as well *amplify* and *spread* losses between markets and institutions. Like a disease, financial distress traveled through the network and “infected” institutions — an idea known as *financial contagion*.

¹For this paragraph, see the well-known speeches by then Executive Director of Financial Stability at the Bank of England Andrew Haldane (2009) and then Vice Chair of the Federal Reserve Janet Yellen (2013).

Financial regulators realized that, in the process of creating the financial network, financial institutions (just “banks” from now on) had not just taken on individual risk, but also created *systemic risk*, which endangered the financial system as a whole. After the crisis, different regulatory policies were put into place to reduce the “excessive systemic risk arising from the complexity and interconnectedness that characterize our financial system” (Yellen, 2013). Many of these policies were aimed directly at reducing interconnectedness. After all, reducing interconnectedness should also limit the “channels” through which financial contagion could spread.² For some financial products, this was achieved using *central clearing counterparties (CCPs)*: central nodes that act as middlemen to all trades and serve as buffers in case one of the participants in the trade is unable to pay its obligations.

For other products, *portfolio compression* was made mandatory. The idea is simple: banks’ daily business activities have the side effect of creating *cycles* of obligations in the network, where (say) a bank A has an obligation to a bank B, which has an obligation to a third bank C, which has an obligation back to bank A — all for the same financial product with the same parameters. Removing these cycles reduces interconnectedness without affecting any bank’s net position with respect to any product. Portfolio compression is the process of doing just that. This paper studies the effects of portfolio compression on systemic risk and the incentives for banks to engage in it.

Portfolio compression (just *compression* from now on) originated in the private sector and has only later been endorsed by regulators.³ Compression has several immediate benefits for the participating banks: it reduces their exposure to other banks, it reduces their operational costs to keep track of their claims and obligations, and, perhaps most importantly, it reduces the sizes of their balance sheets and thus the amount of capital they need to hold available.⁴

In practice, several financial service providers offer compression services.⁵ We explain the process at the example of TriOptima’s service triReduce.⁶ First, participating institutions submit the trades they would like to compress. Second, the

²Reducing interconnectedness has another important benefit, namely improving transparency. — A less dense network is easier to analyze for regulatory bodies when the next crisis strikes. In fact, many of the reforms after the crisis were geared towards only making the network of interconnections *visible* to regulatory bodies. See Financial Stability Board (2017) for an overview of the new regulatory measures for the OTC derivatives market, which is the main financial market relevant for the present piece of work.

³EMIR regulations include an “obligation to have procedures to analyse the possibility to conduct the exercise” of portfolio compression when counterparties have more than 500 contracts with each other (European Securities and Markets Authority, 2017).

⁴Basel III regulations require banks to hold liquid assets proportional to their *risk-weighted assets* (Basel Committee on Banking Supervision, 2011). This creates an incentive to, other things equal, keep assets and liabilities as low as possible.

⁵The most important ones are: TriOptima, LMRKTS, Markit, Catalyst, and SwapClear (D’Errico and Roukny, 2019).

⁶See TriOptima, “triReduce Overview”, <https://www.trioptima.com/resources/>

service provider combines the information submitted by all participants to construct the network and calculates an *unwind proposal*, i.e., a collection of suggested contract modifications and terminations that together constitute the compression. The unwind proposal is sent to the banks. Crucially, the unwind proposal itself carries no legal force. Instead, all involved banks need to agree to the proposal in a third step. Once this has happened, the compression is implemented. In practice, banks generally seem to agree to unwind proposals.⁷

Compression is used in markets for OTC⁸ derivatives where insufficient standardization prohibits the use of a CCP, specifically: interest rate swaps, currency swaps, and credit default swaps. According to TriOptima, their service has removed over USD 1.5 quadrillion in notional since its inception in 2003. Aldasoro and Ehlers (2018) cite compression as one of the main reasons for the drop in size of the credit default swap market after the 2008 financial crisis.

At first sight, one might think that compression is always (weakly) beneficial for all banks: compression reduces the liabilities of involved banks at least as much as their assets (at fair value) and thus benefits involved banks, and it has no effect on the balance sheets of the other banks. However, in this paper, we show that, in a networked system, this simple local view of compression is incorrect, as it does not capture systemic effects.

We study compression in a static setting ex-post to a shock, considering the following time line of events. Initially, all banks were solvent. Then compression was performed. At some point, a shock hit the banks, which caused some of the banks to default on their obligations and triggered a contagion process. We compare the final outcome of this process (modeled by the Rogers and Veraart (2013) clearing model) to the outcome had compression not been performed. In our analysis, we focus on two questions: first, was compression *economically efficient* in retrospect in terms of banks' equities? We consider both a social welfare as well as a Pareto perspective. Second, had the involved banks known the shock upfront, should they still have agreed to compression, i.e., was compression *incentivized* in retrospect?

By *incentivized* we mean that all banks involved in the compression would agree to it in step three of the aforementioned compression process. We assume that banks derive a utility from any given (compressed or uncompressed) network equal to their

⁷Further complications, which we not model in this paper, apply. The main reason for this is that compression is also performed among products whose characteristics (such as maturity, coupons, start and end dates) do not match exactly, but only approximately. This increases the amount of compression that can be done, but it may necessitate compensation payments. These payments are determined using a special market mechanism called a *compression auction* (Duffie, 2018). In this paper, we abstract away from these complications by assuming that all contracts in the network have the same characteristics. We further do not consider any exotic variants of compression, where one may be allowed to both add and remove liabilities. That is, we only consider *conservative compression* (D'Errico and Roukny, 2019) in this paper.

⁸*Over-the-counter*, i.e., traded directly between banks, rather than through an exchange. In this work, we only consider OTC derivatives markets.

equity and decide based on that. Note that in this paper, our goal is *not* to perform an equilibrium analysis. We do not consider compressions actions that banks perform. Rather, we assume that compressions are suggested by a central financial service provider that is not a bank.

For a given network, shocks, and compression, it is easy to answer the above questions regarding economic efficiency and incentives algorithmically: simply evaluate the clearing model on the uncompressed and the compressed network and compare the two outcomes. However, we want to answer these questions at a more general level: what determines whether or not compression will be efficient or incentivized? When are incentives misaligned with the social good? How do these effects depend on the structure of the financial system?

Portfolio compression is systemically important, as evidenced by regulators' attention to it. However, while there is an extensive literature on netting in the context of CCPs (for example, Duffie and Zhu, 2011; Duffie, Scheicher and Vuilleumey, 2015; Amini, Filipović and Minca, 2015; Cui et al., 2018), there has only been a surprisingly small amount of work on compression (without a CCP). Where compression was studied, authors mostly focused on algorithmic questions to achieve the optimal compressed amount subject to risk tolerances (O'Kane, 2017; D'Errico and Roukny, 2019). The impact of compression was therefore measured in terms of eliminated notional, rather than the consequences of a shock. Many researchers have studied the effect of different aspects of *network structure* in general on systemic risk (e.g., Elliott, Golub and Jackson, 2014; Glasserman and Young, 2015; Acemoglu, Ozdaglar and Tahbaz-Salehi, 2015; Demange, 2016). However, these results can at best provide a very broad estimate of what the specific *changes* to the network structure due to compression may imply. We discuss this in Appendix A. Feinstein et al. (2017) studied the sensitivity of the outcome of a crisis to changes in *relative* interbank liabilities, where the total liabilities of each bank remain the same. Compression, however, explicitly reduces the liabilities of each bank involved, which is why their theory does not apply.

In a recent research note, we were, to the best of our knowledge, the first to present an example where portfolio compression could be detrimental in terms of social welfare (Schuldenzucker, Seuken and Battiston, 2018). The only other piece of work we are aware of that deals with the effects of compression on systemic risk specifically is Veraart (2019). The author studied compression in much the same way as we do and obtained sufficient conditions under which compression is beneficial for all banks. The author also showed that compression can have negative effects both on the banks involved in it and on other banks and this can depend on assets external to the financial network. These results can be viewed as special cases of some of the results obtained in the present paper.

In the present paper we show that, in contrast to conventional wisdom, compression

is not in general socially beneficial in terms of banks’ equities. There are cases where compression is socially detrimental or even hurts every bank in the system in a Pareto sense. We further show that whether or not compression is socially beneficial depends on the parameters of the financial system and on the compression in a complex and non-monotonic way. The same applies to banks’ incentives to agree to compression. Their incentives may further be misaligned with the social good (Section 3).

We then derive sufficient conditions under which these complications do not apply, i.e., where compression is always a Pareto improvement with respect to banks’ equities. These conditions are *local* in the sense that they only depend on properties of individual banks and banks’ assets in the uncompressed system, but do not require any global network computation that depends on the particular compression in question. We find that compression is beneficial when banks are well capitalized or when all banks are involved in the compression to a certain degree. These conditions scale with the *default costs* in the system: our results become stronger when defaulting banks lose a higher share of their incoming payments due to frictions such as legal costs or early settlement fees (Section 4).

From these “local” conditions, we then turn to the network structure itself. We find that whether or not the involved banks have an incentive to agree to compression crucially depends on the presence of *feedback paths*, i.e., paths of liabilities that are not involved in the compression and that lead from an involved bank to another involved bank. If feedback paths do not exist and a normality condition is met, compression is always incentivized. This provides a possible explanation why banks virtually always agree to compression in practice while our results in Section 3 imply that this is a complex strategic decision: banks may not be taking the possibility of feedback paths into account (Section 5).

Finally, we turn to a high-level property of the financial network structure, namely *homogeneity* of the involved banks. We show that if a collection of asset and liability measures and the compressed amount of liabilities are equal across all involved banks, then compression is beneficial for all (involved and non-involved) banks. A study of an example network suggests a quantified version: the lower the default costs in the system, the more homogeneity is required to make compression beneficial for all banks (Section 6).

With the present paper, we are among the first to conduct a principled theoretical study of the effects of portfolio compression on systemic risk. We consider this paper a first step at the beginning of a larger program of research. Given that compression has already found its way into regulatory policy, we believe that a thorough understanding of compression is urgently needed. We discuss possible next steps in this direction in Section 7.

2 Preliminaries

We assume that we are given a financial network ex-post to a shock. We then perform *network clearing* in the compressed financial network and in the network had compression not been done and we compare the two outcomes by their efficiency and incentives. We now formally define the three elements of this approach: our model of (post-shock) financial systems and clearing payments, our formalization of portfolio compression, and our measure of utilities, efficiency, and incentives in a financial system.

2.1 Basic Notation

Throughout this paper, we employ the following notation. Matrices and vectors always range over arbitrary finite sets of indices. That is, we do not implicitly assume that vectors range over sets of form $\{1, \dots, n\}$. Ordering of matrices and vectors is always point-wise. That is, if N and M are sets and $p, q \in \mathbb{R}^{N \times M}$ are matrices, we write $p \leq q$ iff $p_{ij} \leq q_{ij} \forall i \in N, j \in M$. We write $[0, q]$ for the set of p such that $0 \leq p \leq q$, i.e., $0 \leq p_{ij} \leq q_{ij} \forall i, j$. Vectors are treated analogously. If $A \subseteq N \times M$, we write p_A for the restriction of p to pairs of indices in A , i.e., $p_A = (p_{ij})_{(i,j) \in A}$. Vectors are treated analogously. We use “ \cup ” to denote concatenation of vectors and matrices, respectively, with disjoint index sets. Thus, if $A, B \subseteq N \times M$, $A \cap B = \emptyset$, $p \subseteq \mathbb{R}^A$, $q \subseteq \mathbb{R}^B$, we write $p \cup q$ for the matrix with indices $A \cup B$ where $(p \cup q)_{ij} = p_{ij}$ if $(i, j) \in A$ and $(p \cup q)_{ij} = q_{ij}$ if $(i, j) \in B$.

2.2 Financial Systems and Clearing Payments

We use the clearing model from Rogers and Veraart (2013). In this sub-section, we provide a brief description of this model.

Financial System. Let N be a set of *banks*. We assume that each bank $i \in N$ holds an amount of *external assets* $e_i \geq 0$. For any two banks $i, j \in N$, let $l_{ij} \geq 0$ denote the amount that the *writer* i owes to the *holder* j of the contract (i.e., the liability). We also call this number the *notional* of the contract, where a notional of zero indicates the absence of a contract. If $l_{ij} > 0$, we also call j a *creditor* of i and i a *debtor* of j . Assume that $l_{ii} = 0 \forall i$, i.e., no bank has a contract with itself. We do not exclude the possibility that two banks are creditors of each other, i.e., $l_{ij}, l_{ji} > 0$. Note that the collection of all notionals $l = (l_{ij})_{i,j \in N}$ can be viewed as the adjacency matrix of a weighted graph with nodes N and $e = (e_i)_{i \in N}$ can be viewed as a vector of node weights.

Following Rogers and Veraart (2013), we assume that we are further given two *default cost parameters* $\alpha, \beta \in [0, 1]$. If a bank is in default, it is only able to pay to its creditors a share of α of its external assets and a share of β of its incoming

payments from other banks. That is, a share of $1 - \alpha$ and $1 - \beta$, respectively, is lost.⁹ If we set $\alpha = \beta = 1$ (no default costs), we receive the clearing model in Eisenberg and Noe (2001).

A tuple $X = (N, e, l, \alpha, \beta)$ is called a *financial system*.

Assets and Liabilities. Given such a financial system X , the *total liabilities* of a bank i are

$$l_i := \sum_{j \in N} l_{ij}$$

and the *relative liability* of i to j is

$$\pi_{ij} := \begin{cases} \frac{l_{ij}}{l_i} & \text{if } l_i > 0 \\ 0 & \text{if } l_i = 0. \end{cases}$$

Note that $\sum_j \pi_{ij} = 1$ unless $l_i = 0$.

The matrix of liabilities defines how much money every bank *is supposed to* pay to its creditors. The amount of money that said bank will actually *be able to* pay will be lower in case it defaults. We capture these amounts in a matrix of *payments* $p_{ij} \in [0, l_{ij}]$. We are ultimately looking for a payment matrix that is *clearing* in a sense that will be defined shortly.

Given p , the *total assets* $a_i(p)$ of a bank i at p consist of its external assets and the incoming payments from other banks:

$$a_i(p) := e_i + \sum_{j \in N} p_{ji}$$

If a bank's assets are insufficient to cover its liabilities, it is called *in default*. A defaulting bank i has its assets reduced according to the factors α and β . That is, its *assets after default costs* are:

$$a'_i(p) := \alpha e_i + \beta \sum_{j \in N} p_{ji}$$

Clearing Payment Matrix. We call p *clearing* if it follows the following fundamental principles of bankruptcy law:

1. Banks that are not in default pay their liabilities in full.

⁹Default costs may originate from various sources, corresponding to different combinations of the α and β parameters. For example, we can model a setting where the external assets are illiquid, which leads to a price discount when they need to be sold quickly in case of default, by choosing $\alpha < 1$. Smaller α values model lower liquidity of the external assets. If the external assets describe perfectly liquid “cash” holdings, we set $\alpha = 1$. Similarly, interbank payments may be lossless ($\beta = 1$) or they may be subject to time delays, legal costs, or early settlement costs ($\beta < 1$). It is easy to extend the model to allow for per-bank α and β values or several asset classes with different default cost levels. Our results easily generalize to these extensions with minor adjustments.

2. Banks that are in default pay out all their assets to creditors, after default costs have been subtracted.
3. By the *principle of proportionality*, the assets of defaulting banks are split up among creditors in proportion to the respective liability.

We thus call p *clearing* if it is the fixed point of the following function:

$$\begin{aligned} \Psi : [0, l] &\rightarrow [0, l] \\ \Psi_{ij}(p) &:= \begin{cases} l_{ij} & \text{if } a_i(p) \geq l_i \\ \pi_{ij}a'_i(p) & \text{if } a_i(p) < l_i \end{cases} \end{aligned} \quad (1)$$

In other words, p is clearing iff $\Psi(p) = p$. The following theorem shows that a clearing matrix of payments always exists and is essentially unique (unique up to a point-wise decrease in payments).¹⁰

Theorem (Rogers and Veraart (2013, Theorem 3.1)). *For any financial system $X = (N, e, l, \alpha, \beta)$ there is a matrix p of payments such that i) $p \in [0, l]$, ii) $p = \Psi(p)$, and iii) if p' is another matrix with these properties, then $p' \leq p$ point-wise.*

We call p like in the theorem the *maximal clearing payment matrix*. While there could be several clearing p , the point-wise maximal p is a canonical choice. Note in particular that maximizing p also maximizes the assets of each individual bank. Therefore, in this paper, we only consider the maximal clearing matrix of payments.

Remark 1 (Algorithms for Computing Clearing Payments). Two well-known algorithms can be used to compute the maximal clearing payment matrix. The simplest one is to consider the iteration sequence defined by $p^0 := l$ and $p^{n+1} := \Psi(p^n)$. Since Ψ is monotonic with respect to the point-wise ordering, this sequence converges from above to the maximal clearing payment matrix.¹¹ However, as soon as there is a cycle of defaulting banks, the sequence does not converge in finite time. An improved algorithm, called the *fictitious default algorithm* by Eisenberg and Noe (2001) and the *greatest clearing vector algorithm* by Rogers and Veraart (2013), skips over linear

¹⁰Note that this existence result is specific to financial networks where the liabilities between banks are fixed numbers. This model is appropriate when banks only enter into *debt contracts*, i.e., loans from one bank to another, or into financial derivatives that only depend on variables external to the financial system. In the latter case, liabilities can be assumed to be fixed or the purpose of clearing. In contrast, clearing payments need not exist (and need not be essentially unique) when banks can also enter into *credit default swaps*, i.e., default insurance, on other banks in the network (Schuldenzucker, Seuken and Battiston, 2019).

¹¹This is a consequence of the following theorem. If L is a complete lattice with maximal element \top and $F : L \rightarrow L$ is monotonic and continuous from above, then the iteration sequence defined by $x^0 := \top$ and $x^{n+1} = F(x^n)$ is decreasing and converges to a fixed point x^+ of F that dominates every other fixed point of F . This theorem is a special case of the Kleene and Tarski-Kantorovitch fixed-point theorems (see Granas and Dugundji (2003)) and has become a standard tool in the theory of financial networks, for example Rogers and Veraart (2013); Barucca et al. (2016). We provide a simple proof in Schuldenzucker, Seuken and Battiston (2019, Lemma 3). The statement for clearing payment matrices follows via $L := [0, l]$ with the point-wise ordering and $F := \Psi$.

stretches of the iteration sequence by solving a linear equation system and terminates with an exact solution after polynomially many steps.

Remark 2 (Individual payments, total payments, and recovery rates). Most pieces of prior work do not consider the matrix of individual payments p_{ij} as the fundamental object of clearing, but the vector of total payments p_i or the vector of *recovery rates*

$$r_i := \begin{cases} 1 & \text{if } a_i(p) \geq l_i \\ \frac{a'_i(p)}{l_i} & \text{if } a_i(p) < l_i. \end{cases}$$

It is easy to see that the definition (1) of Ψ implies that $p_{ij} = \pi_{ij} p_i = l_{ij} r_i$, so that these three objects all define each other and being “clearing” can be defined in terms of either of them. In this paper, we find that operating on the individual payments is most convenient to establish relationships between the uncompressed and the compressed network.

2.3 Portfolio Compression

Portfolio compression is the process of netting liabilities between any number of banks while preserving each bank i ’s *net position* $\sum_j l_{ji} - \sum_j l_{ij}$. We only consider *conservative compression* (D’Errico and Roukny, 2019) in this paper, i.e., point-wise reductions in liabilities. Again following D’Errico and Roukny (2019), we thus define a *compression* as any way how a given financial system can be compressed subject to these two constraints, i.e., as follows:

Definition 1 (Compression). Let $X = (N, e, l, \alpha, \beta)$ be a financial system. A *compression for X* is a circulation in the weighted graph associated to l in the sense of network flow theory. That is, a compression is a matrix $c \in [0, l]$ such that $c_i := \sum_j c_{ij} = \sum_j c_{ji}$. If c is a compression for X , the financial system X *compressed by c* is $X^c := (N, e, l - c, \alpha, \beta)$. We also use a superscript \cdot^c for the liabilities, assets, etc., evaluated in X^c . In particular, we write p^c for the maximal clearing payments in X^c . Let $N(c) = \{i \in N \mid c_i > 0\}$ be the set of banks involved in the compression.

Remark 3 (Cycles and Compressions). Veraart (2019) considered compression by cycles (i.e., closed directed paths in the liability graph) rather than circulations. A cycle can be viewed as a special case of a compression: if C is a cycle and $0 \leq \mu \leq \min_{(i,j) \in C} l_{ij}$, then c defined by

$$c_{ij} = \begin{cases} \mu & \text{if } (i, j) \in C \\ 0 & \text{otherwise} \end{cases}$$

is obviously a compression for X . In this case, $c_i = \mu$ if $i \in C$ and $c_i = 0$ otherwise. By slight abuse of notation, we write $c = (C, \mu)$ for this kind of compression. By the *flow decomposition theorem* (see, e.g., Korte and Vygen (2012, Chapter 8)), for any

compression c there exists a finite (not generally unique) set \mathcal{C} of pairs (C, μ) like above such that:

$$c_{ij} = \sum_{\substack{(C, \mu) \in \mathcal{C} \\ \text{s.t. } (i, j) \in C}} \mu$$

Our most important example in the following considerations will be compressions of form $(C, \min_{(i,j) \in C} l_{ij})$, where C is a cycle. However, we will show at the end of Section 4 that, to correctly capture the systemic effects of more complex compressions, it is *not* sufficient to apply flow decomposition and then consider individual cycles. Thus, our definition of compressions as circulations offers additional generality.

A compression of maximal *value* $\sum_i c_i$, where the largest possible amount of notional is eliminated, can be computed efficiently via linear programming (D'Errico and Roukny, 2019) or via combinatorial algorithms (e.g., Korte and Vygen (2012, Chapter 8)). Like for most flow problems, the naive greedy algorithm, where cycles are successively compressed until none are left, does *not* in general lead to a compression of maximal value.

2.4 Utilities, Efficiency, Welfare

Given a matrix of payments p , define the *balance* $B_i(p)$ of bank i as the difference between assets and liabilities:

$$B_i(p) := a_i(p) - l_i$$

Note that $B_i(p) < 0$ if and only if i defaults under p . Define the *equity* $E_i(p)$ as the balance if this is non-negative (i.e., if i is not in default):

$$E_i(p) := \max(0, B_i(p)) = \max(0, a_i(p) - l_i)$$

If there is no risk of confusion, we leave out the p argument and simply write a_i , B_i , and E_i for the respective values under the maximal clearing payment matrix. We likewise write a_i^c , B_i^c , and E_i^c for these values in the compressed financial system X^c under the maximal payment matrix p^c .

The equity is the profit that the owners (i.e., the shareholders) of a bank get to keep after clearing. Just like how a stock price can never be negative, the equity value of a bank can never be negative, too. This is a consequence of *limited liability* of equity holders. By the principle of *absolute priority* of debt, equity is zero in case of default.¹²

¹²There does not seem to be a standard for the definition of the term “equity”. While, for example, Barucca et al. (2016) and Veraart (2019) use the term “equity” to refer to what we call the balance, Eisenberg and Noe (2001) use the term in the same manner as we do. In this paper, it is important for us to differentiate between the two concepts of equity and balance.

We assume that banks derive a utility from a financial system that is equal to their equity under the maximal clearing matrix of payments. Recall that a compression is implemented if and only if each of the involved banks agrees to it, i.e., if it weakly increases the equity of each involved bank. We can now define Pareto efficiency, social welfare, and incentives in the usual way.

Definition 2. The *social welfare* in a financial system X is the total equity

$$E_{\Sigma} := \sum_{i \in N} E_i.$$

If c is a compression for X , then c is called a (weak) *Pareto improvement* if c weakly increases the equity of each bank in the system, i.e., if $E_i^c \geq E_i \forall i \in N$. The Pareto improvement is *strict* if the inequality is strict for at least one i . We call c (weakly) *incentivized* if it weakly increases the equity of each bank *involved in the compression* (but it may reduce the equity of other banks). That is, we call c incentivized if $E_i^c \geq E_i \forall i \in N(c)$. We call c *strictly incentivized* if the inequality is strict for at least one i . Obviously, every Pareto improvement is incentivized and weakly increases social welfare.

2.5 Discussion of Our Formal Model

If $\alpha = \beta = 1$ (no default costs), then we know from Eisenberg and Noe (2001) that $E_{\Sigma} = e_{\Sigma} = E_{\Sigma}^c$, where $e_{\Sigma} := \sum_i e_i$. Thus, in this case, compression cannot be welfare-improving and there can in particular be no strict Pareto improvements. However, strictly *incentivized* compressions can exist and they can reallocate equity from banks outside the compression to those inside. Prior work that studied systemic risk in the model without default costs (Glasserman and Young, 2015; Demange, 2016) considered the aggregate *payments* (equivalently, the total *balance*) as a welfare measure. This was necessary because total equity is a trivial measure if there are no default costs. For the present paper, where we *do* assume default costs, it is easy to see that $E_{\Sigma} < e_{\Sigma}$ unless no bank defaults or $\alpha = \beta = 1$, so that total equity is not a trivial measure. We consider total equity a measure of social welfare that is better motivated and better in line with banks' incentives than total payments or total balance.¹³

Recall that we do *not* assume that banks can actually make choices about compression *after* a shock has struck and clearing has already taken place. Banks that have already defaulted will be subject to severe limits on their capacity to contract.

¹³For example, it is not clear what immediate value society would derive from payments being high *per se*. In contrast, equity has a clear interpretation as the amount of money that "leaves" the financial system after all obligations have been settled. Note that the interests of society or the real economy can be represented explicitly in this framework by adding them as synthetic nodes to the network. These nodes will likely be leaves in the network and thus cannot be involved in any compression.

Instead, we take an ex-post perspective. A Pareto improvement is thus a compression where everyone would agree in retrospect that they should have compressed. An incentivized compression is one where the involved banks would agree on this.

Recall further that we do not perform an equilibrium analysis in this paper. A compression is not an action that banks perform as part of some game. Rather, we assume that a certain compression is provided exogenously by a central financial service provider (that is not a bank) and we are interested in efficiency and incentives for this single given compression. Our results in the present paper may form the basis for a future study of “stable” networks where no further incentivized compressions exist. It may further serve as a first step towards a study of the mechanism design problem where banks strategically report the liabilities to compress and the financial service provider chooses the compression according to a known set of rules.

To the best of our knowledge, the following are the only known general results on the effect of compression on efficiency in prior work:

Theorem (Essentially Veraart (2019)). *Let c be a compression for $X = (N, e, l, \alpha, \beta)$.*

1. *If $\alpha = \beta = 0$, then c is a Pareto improvement.*
2. *The following two conditions are equivalent: i) none of the banks in $N(c)$ default in X , ii) none of the banks in $N(c)$ default in X^c . If this condition holds, then $E = E^c$ point-wise, i.e., all banks are indifferent regarding compression.*

Proof. Veraart (2019) proved these statements when c is a cycle. It is easy to see that her arguments generalize to arbitrary compressions. \square

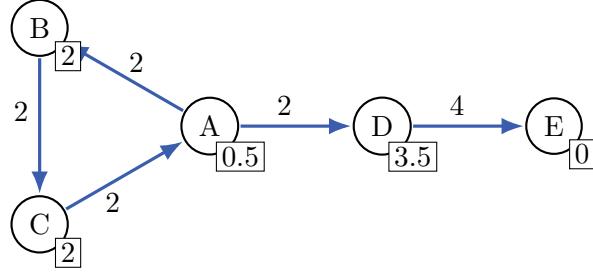
Beyond these results, we note that compression affects various measures of *network structure* and prior work has studied the effect of such measures on systemic risk quite extensively. By connecting these two effects, we can receive a first indication how compression may affect systemic risk. The respective predictions vary across different models. We provide a discussion in Appendix A.

An immediate implication of part 2 of this theorem is that compression cannot set *all* banks involved in it strictly better off: any incentivized compression leaves some of the involved banks indifferent. Likewise, no compression can set all involved banks strictly *worse* off.

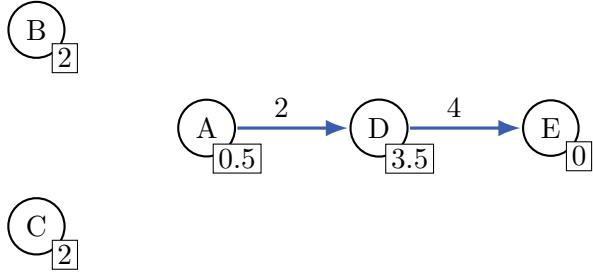
Corollary 1. *If c is a compression for X , then there exist $i, j \in N(c)$ such that $E_i \leq E_i^c$ and $E_j \geq E_j^c$.*

Proof. If no i like in the statement exists, then $E_i > E_i^c \forall i \in N(c)$. In particular, $E_i > 0 \forall i \in N(c)$ and thus none of the banks in $N(c)$ default in X . By the above theorem, part 2, this implies $E = E^c$. Contradiction. Likewise, if $E_j^c > E_j \forall j \in N(c)$, then none of those banks default in X^c and again, $E = E^c$. \square

Figure 1 Financial System where compressing a cycle decreases social welfare. Variant of Schuldenzucker, Seuken and Battiston (2018, Figure 1). Let $\alpha = \beta = 0.5$.



(a) Uncompressed Network



(b) Compressed Network

3 Detrimental Effects of Compression

In this section, we illustrate using a series of examples that compression can have beneficial as well as detrimental effects on efficiency, both in terms of social welfare and in a Pareto sense, and that this effect depends on the parameters of the financial system in a complex and non-monotonic way. Incentives to compress may further be misaligned with the social good.

3.1 Compression May Reduce Social Welfare

The basis for our examples in this section is a variant of a financial system first introduced in our prior work (Schuldenzucker, Seuken and Battiston, 2018); see Figure 1a. There are five banks A–E. Liabilities are depicted as blue arrows with notional next to them. External assets are depicted inside boxes on top of the banks. There is a single cycle, A–B–C. We are interested in the effect of compressing that cycle, i.e., we consider the compression $c = (A-B-C, 2)$. The result of the compression is depicted in Figure 1b. Let $\alpha = \beta = 0.5$. The way we would typically think about this example is that all banks were initially solvent and may or may not have

performed compression; then A was hit by a shock.¹⁴

Compression reduces social welfare. To see this, observe that in both the compressed and uncompressed case, banks B and C have external assets so high that they cannot default and bank A will always default. The key to our construction is that the recovery rate of A , and equivalently the payment p_{AD} , depends on whether or not compression is performed. Since there is no cycle of defaulting banks, clearing payments can be computed easily in topological order. In the uncompressed case, we have $a'_A = \alpha \cdot 0.5 + \beta \cdot 2 = 1.25$ and thus $p_{AD} = p_{AB} = 1/2 \cdot a'_A = 0.625$. In effect, $a_D = 4.125 > 4 = l_D$ and D does not default. In the compressed case, we have $a'^c_A = \alpha \cdot 0.5 = 0.25$ and thus $p_{AD}^c = a'^c_A = 0.25$ and $a^c_D = 3.75 < 4 = l_D^c$. Thus, D defaults. As the external assets of D are relatively large, the default costs due to α lead to a significant drop in social welfare compared to the uncompressed case. This dominates social welfare. More in detail, the equities in both cases are:

$$\begin{array}{ll} E_A = 0 & E_A^c = 0 \\ E_B = 0.625 & E_B^c = 2 \\ E_C = 2 & E_C^c = 2 \\ E_D = 0.125 & E_D^c = 0 \\ E_E = 4 & E_E^c = 1.75 \\ \boldsymbol{E_\Sigma = 6.75} & \boldsymbol{E_\Sigma^c = 5.75} \end{array}$$

The total share of value lost due to default costs is $\frac{E_\Sigma - e_\Sigma}{e_\Sigma} = \frac{1.25}{8} \approx 0.15$ in the uncompressed and $\frac{E_\Sigma^c - e_\Sigma}{e_\Sigma} \approx 0.28$ in the compressed case. This spread can be made arbitrarily close to $1 - \alpha$ by increasing e_C and l_{DE} by the same amount.¹⁵

3.2 Dependence on External Assets and Default Costs

One may wonder how our result depends on the choice of parameters, specifically how the difference in total equity $\Delta E_\Sigma := E_\Sigma^c - E_\Sigma$ depends on the default cost parameters and the level of the external assets e_A of A . We can determine two boundary cases analytically: if $\alpha = \beta = 0$, then Veraart (2019) has shown that compression constitutes a (weak) Pareto improvement so $\Delta E_\Sigma \geq 0$. In our case, $\Delta E_\Sigma = E_B^c = 2$ in case $\alpha = \beta = 0$. If $\alpha = \beta = 1$, we know from Eisenberg and Noe (2001) that $\Delta E_\Sigma = e_\Sigma - e_\Sigma = 0$. The existence of our counterexample (where $\Delta E_\Sigma < 1$) now implies that ΔE_Σ must be non-monotonic in the default cost parameters at least for $e_A = 0.5$.

¹⁴If $e_A \geq 2$ and one of the other banks is exposed to shock instead, then compression is always a Pareto improvement: for B and C , it is clear that it is better to isolate these banks from the others as long as A can pay. For D and E , compression makes no difference.

¹⁵It can in fact be made arbitrarily close to 1 by replacing the bank E by a cycle of intermediation.

Figure 2 Difference in total equity ΔE_{Σ} between the compressed and uncompressed network under variation in e_A and $\alpha = \beta$ in Figure 1. The blue plane is zero. Resolution: 0.01 steps in both parameters.

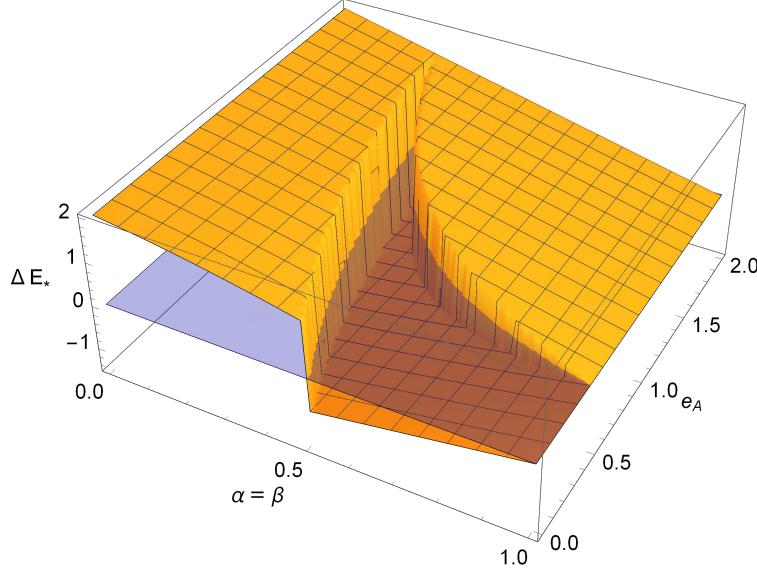


Figure 2 depicts the graph of ΔE_{Σ} for all relevant values of $\alpha = \beta$ and e_A ¹⁶. The blue transparent plane is zero. It is easy to see analytically that D does not default in X iff $\alpha e_A + \beta \cdot 2 = \alpha(e_A + 2) \geq 1$ and D defaults in X^c iff $\alpha e_A < 0.5$. The intersection of these two regions is the blue “canyon,” where compression is socially detrimental. Note how this set has a complex three-dimensional shape that depends on both parameters. Note in particular that if we fix a value of $e_A > 1.0$, the set $\{\alpha \mid \text{compression is socially detrimental for } e_A \text{ and } \alpha = \beta\}$ is non-contiguous. Analogously, for $\alpha = \beta \approx 0.6$, this is the case as we vary the e_A values.

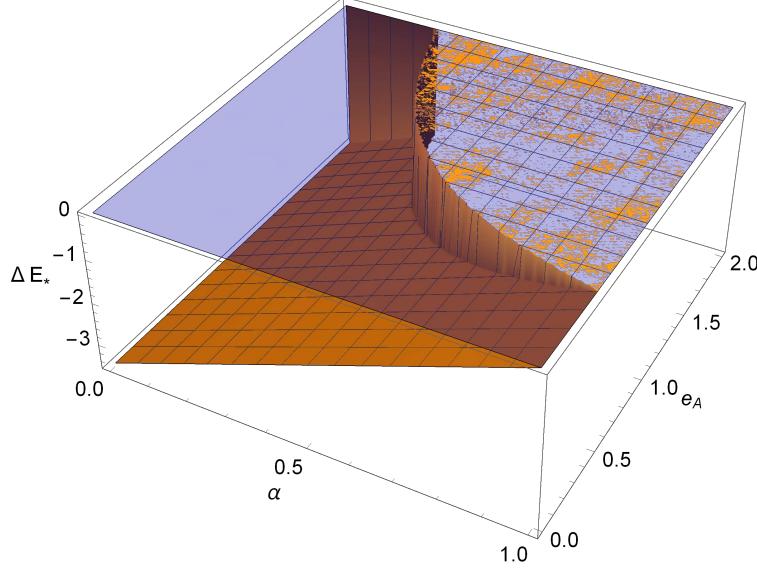
3.3 Individual Values for Default Cost Parameters

In the discussion so far, we have assumed that $\alpha = \beta$. One may wonder to which extent this is necessary for our result.

For the α parameter, it is easy to construct an equivalent example that does not depend on the value of α at all. To do this, add a *source bank* s and replace the external assets of all banks by an equivalent liability from s . Then give s sufficiently high external assets so that it cannot default. There is now no bank that holds any external assets *and* can possibly default, so the value of α is irrelevant. The resulting financial system will behave like our original example for $\alpha = \beta$. Thus, for any value of α there exists a financial system and a compression that decreases social welfare. Note that this technique is general and can transform *any* financial system where $\alpha = \beta$ into an equivalent one where α is arbitrary.

¹⁶Note that only values $e_A \in [0, 2)$ are relevant since no bank defaults if $e_A \geq 2$ and thus, $\Delta E_{\Sigma} = 0$ trivially

Figure 3 Difference in total equity between the compressed and uncompressed network under variations in e_A and $\alpha \leq \beta = 1$ in Figure 1. Resolution: 0.01 in both parameters.



The remaining two boundary cases to explore are thus $\beta = 0$ and $\beta = 1$. For $\beta = 0$, we will show in Section 4 below that this already implies that compression is a (weak) Pareto improvement, and is in particular (weakly) socially beneficial. For $\beta = 1$, compression is (weakly) socially detrimental in our example for any value of α and e_A . This is because then $E_\Sigma = e_\Sigma - (1 - \alpha) \sum_{i \in D} e_i$, where D is the set of defaulting banks, (Rogers and Veraart, 2013, Lemma 4.8) and, by the above discussion, only the default of D depends on any payments from others. It follows from the analytical expressions for the default of D in Section 3.2 that if D defaults in X , then also in X^c . Thus, $\Delta E_\Sigma \leq 0$ and $\Delta E_\Sigma = 0$ iff $\alpha = 1$ or $\alpha e_A \geq 0.5$. Figure 3 shows the corresponding plot.

3.4 Detrimental Effects for Involved Banks and Pareto Worsening

Note that the banks A, B, and C that are part of the compressed cycle in our running example benefit from compression (some weakly, some strongly). Thus, this compression is incentivized. We will show in Section 5 that this is always the case for this kind of network topology. For the parameter values discussed above, compression also reduces social welfare, so that incentives are misaligned with the social good.

This property is not universal. If we identify banks C and E, we receive the financial system in Figure 4. In this financial system, bank C suffers from compression. To see this note that C cannot default, so the outcome of clearing is the same as in Section 3.1, except for that C's equity is increased by what was E's equity before. We thus have $E_C^c = 3.75 < 6 = E_C$, i.e., compression would harm C and C would therefore not be willing to agree to it. Thus, C is not incentivized. Notice how C's

Figure 4 Financial system where compression hurts the cycle bank C.

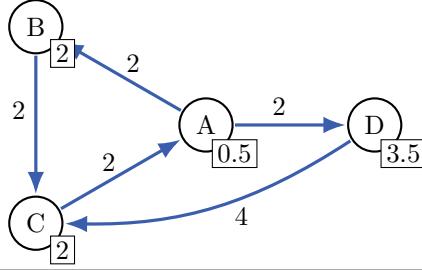
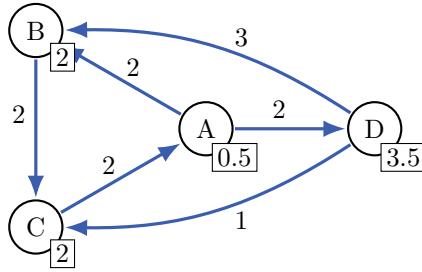


Figure 5 Financial system where compression hurts all banks (B, C, and D strictly).



veto to compression is at the expense of B, who would benefit from compression. Keeping the cycle does not cause an immediate cost for C because B does not default and thus allows C to support A “for free”. This eventually benefits C. We will study such phenomena in greater detail in Section 6 when we discuss homogeneity.

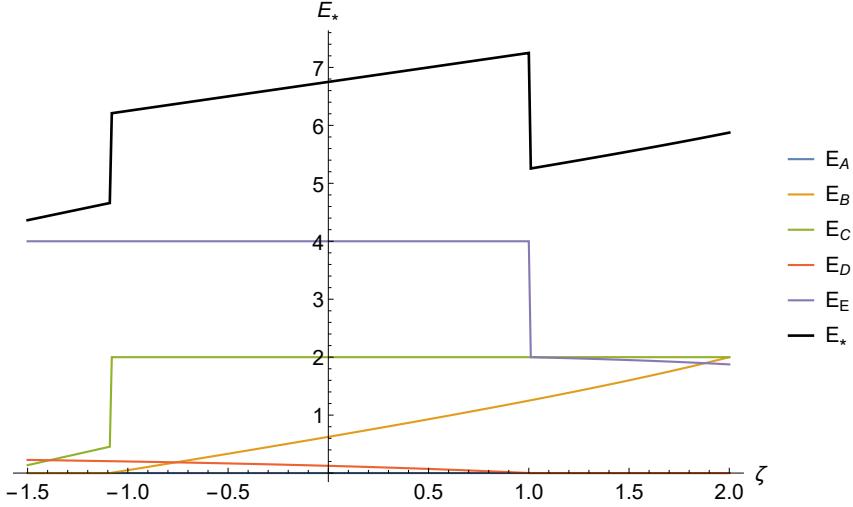
We can take the idea from the previous example one step further to receive a financial system where the equity of *every* bank decreases (strictly for some banks), i.e., compression is a Pareto worsening. To do this, consider Figure 1 again and notice that C is actually indifferent between compression and non-compression, B gains 1.375 and E loses 2.25 from it. By eliminating E and splitting up the liability from D to E into two separate liabilities to B and C, we can simultaneously split up the loss and make compression strictly harmful for both B and C. See Figure 5. We now have:

$$\begin{array}{ll}
 E_A = 0 & E_A^c = 0 \\
 E_B = 3.625 & E_B^c = 3.3125 \\
 E_C = 3 & E_C^c = 2.4375 \\
 E_D = 0.125 & E_D^c = 0 \\
 \mathbf{E}_\Sigma = \mathbf{6.75} & \mathbf{E}_\Sigma^c = \mathbf{5.75}
 \end{array}$$

Of course, the effect can be made arbitrarily much stronger by increasing e_D , l_{DB} , and l_{DC} like before.

Note that, by Corollary 1, it is not possible that *all* banks strictly suffer from compression.

Figure 6 Individual equities and total equity dependent on the compressed amount ζ in Figure 1. Resolution: 0.01



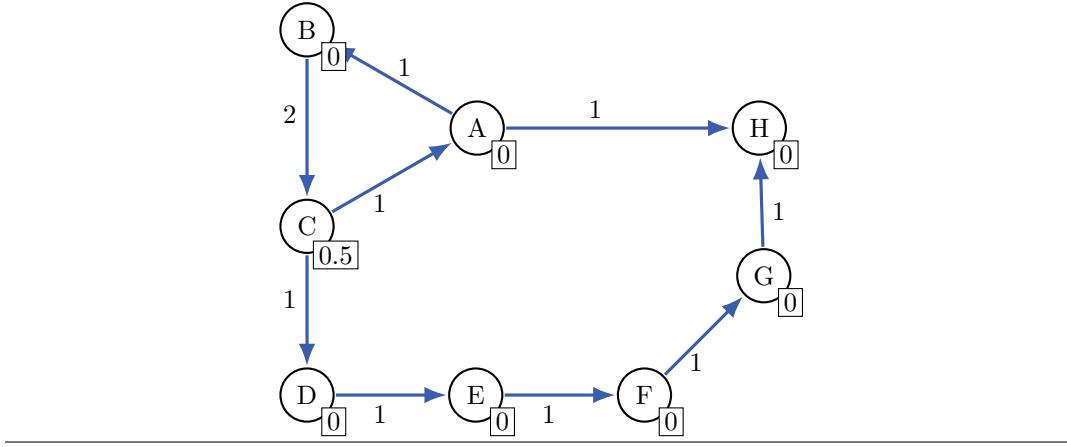
3.5 Partial Compression

One can imagine a situation where only part of a feasible compression is executed. This might be, for example, because in reality, opportunities for compression come up over time and are executed in an on-line fashion. One might assume that partial compression has a monotonic effect on social welfare: if compression is beneficial, performing part of the compression should be less beneficial, but still beneficial and vice versa. However, this is not the case even when we consider a single cycle. Figure 6 shows the total equity and individual banks' equities under the compression $c = (A-B-C, \zeta)$ in our original example from Figure 1a. $\zeta = 0$ corresponds to no compression and $\zeta = 2$ corresponds to full compression of the cycle. We also allow negative values of ζ , which we interpret as *increasing* all liabilities in the cycle by $-\zeta$. Observe that, as we proceed from 0 to 2 in the graph, the total equity first increases, then experiences a discontinuous drop as D defaults (and thus E_E drops sharply) and then the total equity increases again, but not to its previous level. Observe further that each of the individual equities is monotonic, only their sum E_Σ is not. Specifically, the banks A,B, and C involved in the compression monotonically profit from more compression while the other banks D and E are monotonically hurt by it. This is a consequence of the network topology in this example, which we will discuss in greater detail in Section 5.

3.6 Pareto Effects Without Additional Defaults

The effects presented in the previous subsections are driven by the fact that the default of bank D depends on whether or not compression is done. However, if $\beta < 1$, then compression can also Pareto-decrease equities while the set of defaulting banks

Figure 7 Financial system where compression hurts all banks (B, C, and D strictly).



stays the same, by changing the paths that money takes. To see this, let $\alpha = \beta = 0.5$ and consider Figure 7. We see immediately that C defaults and that this implies that all banks except H default. If we compress the only cycle $(A-B-C, 1)$, the same banks default. Thus, H is the only bank with positive equity and $E_H = p_{AH} + p_{FH}$ and $E_H^c = p_{AH}^c + p_{FH}^c$. To compute the payments in the uncompressed case, we first resolve the cycle. By $p = \Psi(p)$ and since all banks default, we have:

$$\begin{aligned} p_{AB} &= 1/2 \cdot 1/2 \cdot p_{CA} = 1/4 \cdot p_{CA} \\ p_{BC} &= 1/2 \cdot p_{AB} \\ p_{CA} &= 1/2 \cdot 1/2 \cdot (0.5 + p_{BC}) \end{aligned}$$

Solving this equation system yields $p_{AH} = p_{AB} = 1/31$, $p_{BC} = 1/62$, and $p_{CD} = p_{CA} = 4/31$. At each step of the path $D-E-F-G-H$, a factor of $1 - \beta = 0.5$ is lost due to default costs. We thus have $p_{FH} = (1/2)^4 \cdot 4/31 = 1/8 \cdot 1/31$. Overall, $E_H = (1 + 1/8) \cdot 1/31$. In the compressed network, all the money must take the less efficient path $C-D-E-F$, so $p_{AH}^c = 0$ and $E_H^c = p_{FH}^c = (1/2)^5 \cdot 1/2 = 1/64 < (1 + 1/8) \cdot 1/31 = E_H$. Thus, compression leads to a Pareto worsening without any new defaults. Of course, one can make the effect arbitrarily much stronger by increasing the length of the longer path.

If we give A positive external assets rather than C , setting $e_C = 0$ and $e_A = 0.5$, we can observe the opposite effect: compression leads to a Pareto *improvement* because it prevents money from flowing via the inefficient cycle and the longer path, and instead routes it directly to H .

4 Local Sufficient Conditions for a Pareto Improvement

In the previous section, we have seen that compression need not constitute a Pareto improvement, may in fact be a Pareto worsening, and that incentives for the involved

banks may be misaligned with the social good. In this section, we present conditions under which this cannot happen, i.e., where compression is always a Pareto improvement. Our conditions will be *local*, i.e., they will not require any network calculations beyond, in some cases, a lower bound for banks' assets under the maximal clearing values.

We begin by defining two formal tools that will be useful in the rest of our analytical examination.

Definition 3. Let c be a compression for $X = (N, e, l, \alpha, \beta)$, the *change in relative liability* due to c is $\Delta\pi_{ij} := \pi_{ij}^c - \pi_{ij} \in [-1, 1]$. The *adjusted compressed payments* are the matrix $p' := p^c + c$.

Remark 4. The following facts are easy to verify. The proof is omitted.

- If $c_i = 0$, then $\Delta\pi_{ij} = 0 \forall j$. That is, only relative liabilities of nodes involved in the compression change.
- $\Delta\pi_{ij} > 0$ iff $c_{ij} < \pi_{ij}c_i$. Likewise, $\Delta\pi_{ij} \leq 0$ iff $c_{ij} \geq \pi_{ij}c_i$. That is, the “inside” nodes of a compression are exactly those where the compressed amount is at least in proportion to the relative liabilities.
- $\sum_j \Delta\pi_{ij} = 0$, except for when $c_i = l_i > 0$, in which case $\Delta\pi_{ij} = -\pi_{ij} \forall j$ and thus $\sum_j \Delta\pi_{ij} = -1$.
- If $c = (C, \mu)$ is a cycle and $\mu < l_i$, then $\Delta\pi_{ij} = -(1 - \pi_{ij})\frac{\mu}{l_i - \mu} < 0$ if $(i, j) \in C$ and $\Delta\pi_{ij} = \pi_{ij}\frac{\mu}{l_i - \mu} > 0$ if $(i, j) \notin C$. If $\mu = l_i$, then $\Delta\pi_{ij} = -\pi_{ij} \forall j$
- $p' \in [0, l]$. $p' = l$ iff no bank defaults in X^c .
- $B_i(p') = B_i^c(p^c) = B_i^c$. In particular, $E_i(p') = E_i^c$ and if $p' \geq p$, then c is a Pareto improvement.

The change in relative liabilities allows us to distinguish between banks “inside” and “outside” the compression in a mathematically meaningful way. Intuitively, if $\Delta\pi_{ij} \leq 0$, we consider the edge (i, j) “rather inside” c and if $\Delta\pi_{ij} > 0$ we consider (i, j) “rather outside” c . Note that simply testing whether $c_{ij} > 0$ is not a good way to do this for general compressions. For example, for any financial system and compression, we could simply increase both l_{ij} and c_{ij} by a sufficiently small positive number. Then we would suddenly have $c_{ij} > 0$ for all (i, j) and $N(c) = N$ even though the mathematical properties of the compression have not changed.

The last property in Remark 4 means that we can “transfer” the clearing payments from the compressed financial system X^c to the original system X by considering the adjusted compressed payments instead. The payment matrix p' is not usually clearing because i) proportionality is not generally satisfied and ii) the c part of p' is not subjected to default costs due to β . Instead, we can interpret p' as payments in X that have a priority structure: payments in c are privileged before other payments

and also avoid default costs. The following lemma provides a technical sufficient condition under which this compression is beneficial along all dimensions. The lemma will be fundamental for all other, more conceptual sufficient conditions presented in this section.

Lemma 1. *Let c be a compression for a financial system $X = (N, e, l, \alpha, \beta)$. Let p be the maximal clearing payment matrix in X and let $p \vee c$ be the point-wise maximum of p and c , i.e., $(p \vee c)_{ij} = \max(p_{ij}, c_{ij})$. Assume that for all $i, j \in N$ where $\Delta\pi_{ij} > 0$ and i defaults under $p \vee c$ in X , we have:*

$$\Delta\pi_{ij}a'_i(p \vee c) + c_{ij} \geq \beta\pi_{ij}^c c_i$$

Then $p' \geq p$. In particular, compression by c weakly increases the balance of each bank and is a Pareto improvement.

The proof of the lemma can be found in Appendix B. Recall that $\Delta\pi_{ij} = 0$ if $i \notin N(c)$. Thus, to verify the precondition of the theorem, we only need to look at banks i that are involved in the compression. The precondition of the lemma is not completely local because it depends on the maximal clearing payments p in X . However, it *only* depends on p : once p is known, $a'_i(p \vee c) = \alpha e_i + \beta \sum_j \max(p_{ji}, c_{ji})$ can be easily computed. Often, we can bound p from below, which may already be enough to apply Lemma 1. We demonstrate some of these applications in the following.

We call c *immediately beneficial at (i, j) for X* if the preconditions of the lemma for the pair (i, j) are satisfied. Note that this alone does not imply that $p_{ij}^c \geq p_{ij}$ or that bank j benefits from compression. That is because compression may be detrimental to other banks, which may negatively affect i and j .¹⁷ — The precondition needs to hold for *all* pairs (i, j) for the lemma to have any implications. We will now describe a few simpler and more interpretable conditions that imply immediate benefits for some pair. By definition, if for each pair (i, j) where $\Delta\pi_{ij} > 0$ and i defaults (under p or $p \vee c$), one of these conditions hold, then c is a Pareto improvement.

An immediate implication of the theorem is that compression is always a Pareto improvement if $\beta = 0$. This generalizes Veraart (2019, Theorem 3.7), where it was shown that compression is always a Pareto improvement if $\alpha = \beta = 0$. We thus learn that the condition $\alpha = 0$ was not actually necessary.

Corollary 2. *For any financial system $X = (N, e, l, \alpha, \beta = 0)$, any compression is immediately beneficial for all pairs. In particular, $p' \geq p$ and c is a Pareto improvement.*

Proof. If $\beta = 0$, the right-hand side in the precondition of Lemma 1 is 0 and the left-hand side is always non-negative. Thus, the precondition holds trivially. \square

¹⁷We explore the idea of these *feedback effects* further in Section 6.

Compression is immediately beneficial when the recovery rate of the contract writer is not too low.

Theorem 1. *Let $X = (N, e, l, \alpha, \beta)$ be a financial system and assume that*

$$r_i \geq \beta.$$

Then any compression is immediately beneficial for any pair (i, j) , for $j \in N$.

Proof. Let i, j be such that $\Delta\pi_{ij} > 0$ and i defaults in X under $p \vee c$. We have $r_i = a'_i(p)/l_i \geq \beta$ by assumption and thus $a'_i(p \vee c) \geq a'_i(p) \geq \beta l_i$. We now receive:

$$\begin{aligned} \Delta\pi_{ij}a'_i(p \vee c) + c_{ij} &\geq \beta\Delta\pi_{ij}l_i + c_{ij} \\ &= \beta\pi_{ij}^c c_i - \beta c_{ij} + c_{ij} \geq \beta\pi_{ij}^c c_i \end{aligned}$$

as required, where the second line is by the identity $c_{ij} = \pi_{ij}^c c_i - \Delta\pi_{ij}l_i$. \square

The previous theorem is only meaningful when $\beta < \alpha$. This is because, if $\beta \geq \alpha$, it is easy to see that if we have $r_i \geq \beta = \max(\alpha, \beta)$, then already $r_i = 1$, i.e., i does not default under p and in particular not under $p \vee c$.¹⁸ A default cost regime where $\beta < \alpha$ is one in which external assets are relatively liquid, but interbank payments are subject to high costs such as delays or legal uncertainty (see Section 2.2). Under this regime, the theorem states that the creditors of a defaulting bank will still benefit from compression as long as said bank is not too deep in default. Note how this is independent of the compression applied.

Theorem 1 does still not provide a fully local condition because it depends on the recovery rates. We could derive a fully local condition from the theorem using the estimate $r_i \geq \alpha e_i / l_i$. However, the following theorem shows that a weaker version of this condition is still sufficient.

Theorem 2. *Let c be a compression for $X = (N, e, l, \alpha, \beta)$ and assume that $\alpha > 0$ and*

$$l_i = c_i \quad \text{or} \quad \frac{e_i}{l_i - c_i} \geq \frac{\beta}{\alpha}.$$

Then c is immediately beneficial for any pair (i, j) for $j \in N$.

Proof. The condition implies that $\alpha e_i \geq \beta(l_i - c_i)$ and thus $a'_i(p \vee c) \geq a'_i(c) = \alpha e_i + \beta c_i \geq \beta l_i$. The statement now follows like in the proof of Theorem 1. \square

The term $\frac{e_i}{l_i - c_i} = \frac{e_i}{l_i^c}$ is the reciprocal of the *loan-to-value ratio* $\frac{l_i^c}{e_i}$, which is a common measure used to gauge the riskiness of giving a loan of l_i^c to a would-be

¹⁸This holds for any bank under any clearing vector. We provide a proof of this “gap lemma” in Schuldenzucker, Seuken and Battiston (2019, Lemma 1). There, we consider an extension of the model from the present paper to other contract types, but the argument is the same.

debtor that has own capital of value e_i . This is used especially in the context of mortgages.

The precondition of the theorem implies that $\alpha e_i \geq \beta l_i^c$ and thus $r_i^c \geq \beta$. It can thus be seen as a counterpart to Theorem 1 in X^c . If $\beta \geq \alpha$, then like in Theorem 1, the precondition implies that i does not default in X^c . It might, however, default in X .

Our final sufficient condition enables us to relax the condition that we need to consider all pairs (i, j) for which $\Delta\pi_{ij} > 0$. Recall from Remark 4 that

$$\Delta\pi_{ij} \leq 0 \Leftrightarrow c_{ij} \geq \pi_{ij} c_i \Leftrightarrow \frac{c_{ij}/l_{ij}}{c_i/l_i} \geq 1.$$

That is, compression is immediately beneficial if the relative decrease in the individual liability l_{ij} is at least as big as the relative decrease in the total liabilities l_i of i . The following theorem allows us to relax this condition from “at least as big” to “not too much smaller.”

Theorem 3. *Let c be a compression for a financial system $X = (N, e, l, \alpha, \beta)$. Let*

$$\eta : [0, 1) \rightarrow [0, \infty)$$

$$\eta(x) = \begin{cases} 0 & \text{if } x = 0 \\ \frac{x}{1-x} = \frac{1}{\frac{1}{x}-1} & \text{otherwise.} \end{cases}$$

Assume that

$$\frac{\eta(c_{ij}/l_{ij})}{\eta(c_i/l_i)} \geq \beta.$$

Then c is immediately beneficial for the pair (i, j) .

Proof. Whenever $\Delta\pi_{ij} > 0$, we must have $c_{ij} < l_{ij}$ and in particular $c_i < l_i$, otherwise $\pi_{ij}^c = 0$ and thus $\Delta\pi_{ij} \leq 0$. We must further have $c_i > 0$, otherwise $\Delta\pi_{ij} = 0$. Hence, all fractions above are well-defined. We now have:

$$\frac{\eta(c_{ij}/l_{ij})}{\eta(c_i/l_i)} = \frac{\frac{c_{ij}}{l_{ij}-c_{ij}}}{\frac{c_i}{l_i-c_i}} = \frac{c_{ij}}{\frac{l_{ij}-c_{ij}}{l_i-c_i} c_i} = \frac{c_{ij}}{\pi_{ij}^c c_i}$$

Thus, $c_{ij} \geq \beta \pi_{ij}^c c_i$. This implies the precondition of Lemma 1. \square

The function η is a continuous, monotonic, non-linear, and convex. We have $\eta(0) = 0$ and $\lim_{x \rightarrow \infty} \eta(x) = \infty$. Figure 8 provides an illustration. It is easy to see that $x \geq y \Rightarrow \eta(x)/\eta(y) \geq 1$. Thus, the theorem provides a true relaxation over the condition $\Delta\pi_{ij} \geq 0$. Note that if $c_{ij} = 0 < l_{ij}$, then $\eta(c_{ij}/l_{ij}) = \eta(0) = 0$ and so the precondition is only satisfied in the trivial case where $\beta = 0$. This holds in particular if $c = (C, \mu)$ is a cycle, $i \in C$, and $(i, j) \notin C$.

Figure 8 Function η from Theorem 3

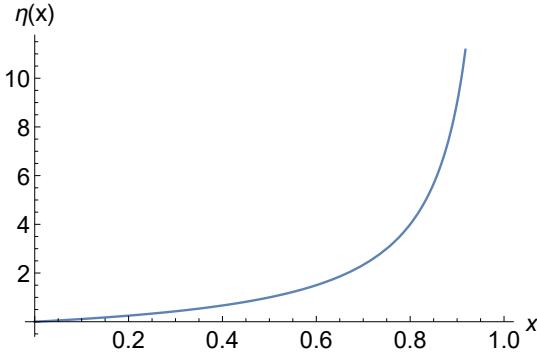
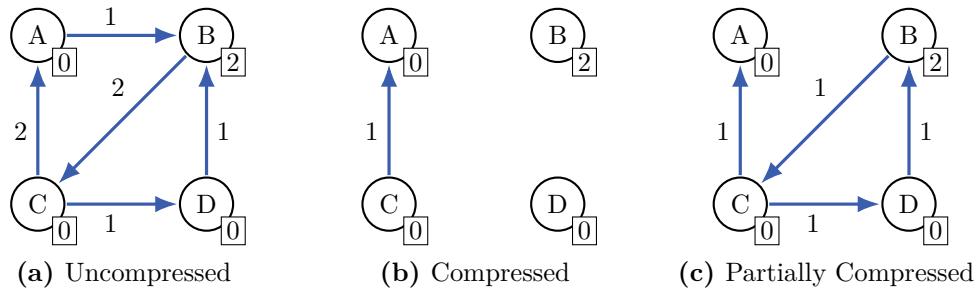


Figure 9 Financial System with a more complex Pareto-improving compression. Let $\beta = 0.5$. α is irrelevant.



Note how Theorem 3 does not depend on e , but only on c and l , i.e., on the network structure and the modifications we make to it.

For an example where the above theorem can be useful, consider Figure 9a and the compression $c = (A-B-C, 1) + (B-C-D, 1)$. This is clearly not a single cycle like in the examples before. Figure 9b shows the compressed network. Let $\beta = 0.5$. Whenever $(i, j) \neq (C, A)$ we have $\Delta\pi_{ij} \leq 0$ because $\pi_{ij}^c = 0$. However, we also have $\Delta\pi_{CA} = 1 - 2/3 = 1/3 > 0$. At the same time,

$$\frac{\eta(c_{CA}/l_{CA})}{\eta(c_C/l_C)} = \frac{\eta(1/2)}{\eta(2/3)} = \frac{1}{2}.$$

Thus, for $\beta = 0.5$, this compression will always be a Pareto improvement for any choice of the external assets.

We can use this example to illustrate another phenomenon. Let $b = (A-B-C, 1)$. Note that $b \leq c$, i.e., b corresponds to executing only part of the compression c . Figure 9c shows the result of compressing by b . Calculation shows that equities correspond to the following table:

i	E_i	E_i^b	E_i^c
A	0	0.25	0
B	0.5	1.125	2
C	0	0	0
D	0	0	0
*	0.5	1.375	2

Note how partial compression first increases the equity of A and then decreases it again. This implies that even though both b and c are Pareto improvements and b is a partial compression of c , this does not imply that c is also a Pareto improvement over b . This provides a case for considering compression by arbitrary circulations, not just by cycles. If we had only considered successive compression by Pareto-improving cycles, we would not have been able to reach c even though it is superior to b with respect to social welfare.

5 Sufficient Conditions for Incentivized Compression

We now turn to non-local conditions on the network structure under which compression is incentivized or a Pareto improvement. Our main result will be that compression is incentivized if there are no *feedback paths*. En-route to this result, we develop a methodology that will enable us to decompose the effects due to compression into two separate phases: an *immediate effect* on only the involved banks and a *feedback effect* due to reverberations of the immediate effect in the rest of the network. When the compression is relatively simple, e.g., just a single cycle, the immediate effect is easy to estimate while the feedback effect depends on the surrounding network. Our results from this and the following section will exploit the fact that we can sometimes know the direction of one of the two effects upfront.

Definition 4. Let c be a compression in a financial system $X = (N, e, l, \alpha, \beta)$. Let $I = \{(i, j) \mid l_{ij} > 0 \wedge c_{ij} > 0\}$ and $O = \{(i, j) \mid l_{ij} > 0 \wedge c_{ij} = 0\}$. Observe that $\Psi_{ij}(q) = 0 \forall q, (i, j) \notin I \cup O$. We can thus ignore pairs $(i, j) \notin I \cup O$, i.e., pairs where $l_{ij} = 0$. To simplify notation, we leave these pairs out. A *feedback path* is a path in O from some bank in $N(c)$ to another or the same bank in $N(c)$. A *chord* is a feedback path of length 1, i.e., an edge $(i, j) \in O$ such that $i, j \in N(c)$. c is called *chord-free* if there are no chords. We call c *normal* for X if $c_{ij} > 0 \Rightarrow \Delta\pi_{ij} \leq 0$ for all $i, j \in N$.

The *immediate-effect payments* \tilde{p} in X due to c are the clearing payments in X^c when all payments in O have been fixed to their values in X . That is, \tilde{p} is the

point-wise maximal solution to the equation¹⁹

$$\tilde{p} = \Psi^c(\tilde{p}_I \cup p_O). \quad (2)$$

Such a maximal solution exists by monotonicity of the function $\Psi^c(\cdot \cup p_O)$. Let $\tilde{E} := E^c(\tilde{p})$ and denote the other values analogously. Define the *adjusted immediate-effect payments* $\tilde{p}' := \tilde{p} + c$ and observe that $E(\tilde{p}') = \tilde{E}$.

The edges in I are those where any compression takes place. However, as we have discussed at the beginning of Section 4, we should only consider those liabilities truly “part of” the compression where $\Delta\pi_{ij} \leq 0$. In a normal compression, the two notions are aligned. Note that any cycle $c = (C, \mu)$ is normal.

Note that if $i \notin N(c)$, then $\tilde{p}_{ij} = p_{ij} \forall j$. This is because $\Psi_{ij}^c(p)$ only depends on the components $(p_{kj})_{k \in N}$ and by assumption we have $(k, i) \in O \forall k$, so $\tilde{p}_{ij} = \Psi_{ij}^c(\tilde{p}_I \cup p_O)$ only depends on p_O and by assumption, $\Psi_{ij}^c = \Psi_{ij}$. Thus, the immediate-effect payments only concern the banks in $N(c)$ and payments within and leaving this set.

The following lemma shows how our decomposition can be useful.

Lemma 2. *Let c be a compression for a financial system X .*

1. *If c is normal for X , then $\tilde{p}'_I \geq p_I$. If c is normal and chord-free for X , then $\tilde{E}_i \geq E_i$ for all $i \in N(c)$.*
2. *If $\tilde{p}_O \geq p_O$, then $p^c \geq \tilde{p}$. In particular, $p_O^c \geq \tilde{p}_O \geq p_O$ and $E_i^c \geq E_i$ for all $i \notin N(c)$. If further, every bank $i \in N(c)$ has a liability in O , then $p' \geq \tilde{p}' \geq p$ and thus c is a Pareto improvement.*
3. *If $\tilde{p}' \geq p$, then $p' \geq \tilde{p}' \geq p$ and thus c is a Pareto improvement.*

The proof can be found in Appendix B.

Part 1 shows that the immediate effects of a (chord-free) compression are always positive to those banks participating in it. Thus, if a compression is not incentivized, this must be due to the feedback effect.

Part 2 and 3 are useful because they tell us that, as far as Pareto improvements (for all banks or the banks not involved in the compression) are concerned, it is sufficient when the immediate effects are beneficial. We will use this fact in Section 6.

Our theorem regarding feedback paths follows from part 1 of the lemma because, if there are no feedback paths, feedback effects do not matter to the banks involved in the compression.

Theorem 4. *If c is normal for X and there are no feedback paths, then c is incentivized.*

¹⁹Recall that we denote restriction of indices to a subset by an index and concatenation along disjoint subsets of indices by a union. Thus, $\tilde{p}_I \cup p_O$ is a matrix with indices in $N \times N$ and $(\tilde{p}_I \cup p_O)_{ij}$ is \tilde{p}_{ij} if $(i, j) \in I$ and p_{ij} if $(i, j) \in O$.

Proof. By Lemma 2 part 1, it is enough to show that $\tilde{E}_i = E_i^c$ for all $i \in N(c)$. This is for the following reason. Let G be the graph with adjacency matrix l . Let

$$\begin{aligned} A &:= \{(i, j) \in O \mid \text{there is a path } j \rightarrow N(c) \text{ in } G\} \\ B &:= \{(i, j) \in O \mid \text{there is a path } N(c) \rightarrow i \text{ in } G\}. \end{aligned}$$

Since there are no feedback paths, A , I , and B are pair-wise disjoint and for all q , $\Psi_A(q)$ is independent of q_I and q_B , and $\Psi_I(q)$ is independent of q_B . The same holds for Ψ^c . This implies that $p_A^c = \tilde{p}_A = p_A$ and $p_I^c = \tilde{p}_I$. And E_i^c for $i \in N(c)$ only depends on p_A^c and p_I^c and likewise for \tilde{E}_i . \square

When we look back at our example from Section 3.1 through the lens of Theorem 1, it is now clear that compression will always be incentivized in this example. Note in particular that this does not depend on the external assets, recovery rates or on which banks default. It also does not depend on the compressed amount (cf. Section 3.5). We also learn from this example that this kind of compression is always incentivized, but need not even increase social welfare, let alone be a Pareto improvement. Section 3.6 provides another example where there are no feedback path and thus compression is always incentivized, but need not be socially beneficial.

Remark 5. It is easy to see that not *all* feedback paths will impair the conclusion of the theorem, but only those where all banks that are part of the feedback path default in X^c . By bounding the assets of these banks from below, perhaps similarly to how we did it in Section 4, one may be able to receive a weaker sufficient condition in Theorem 4.

In the real world, it does not seem that banks consider it a complex strategic decision whether or not they should agree to a compression. Rather, it seems that proposed compressions are agreed to in the vast majority of cases. The lemma and the theorem suggest *local information* or local reasoning as a possible explanation why this might be the case. Feedback paths lie in a part of the network not directly related to the compression. If a bank has limited information about the network structure, it may (incorrectly) assume that feedback paths do not exist. Alternatively, the bank may simply not reason about feedback effects, but only about immediate effects. If it then also assumes that the compression is normal (for example, because it has limited information about the parts of the compression that it is not directly involved in), it may conclude via Theorem 4 of Lemma 2 part 1 that compression is always beneficial. Note that the lemma and the theorem do not depend on the external assets and thus on any shocks, so that this decision can be confidently made ex-ante.

6 Homogeneity

In this section, we consider network structure at an even higher level. Going back to those examples from Section 3 where compression was socially detrimental, we can observe that there was always one bank that was vastly worse off than the others. For example, in Figure 1, bank A was very poorly capitalized while the two other banks in the cycle were very well capitalized. In this section, we show that such a high-level structure is in fact necessary for our examples: if there is a sufficient degree of homogeneity among the balance sheets of involved banks, then compression is always a Pareto improvement.

6.1 Perfect Homogeneity

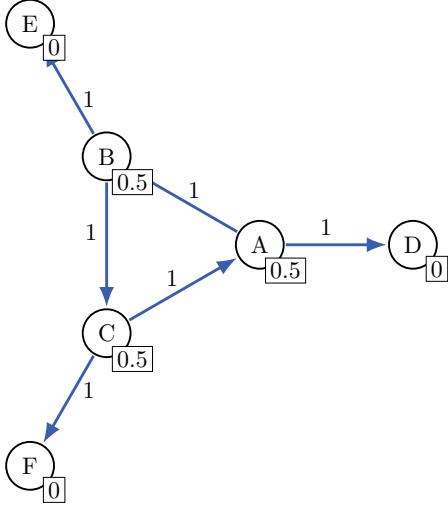
The following theorem proves our statement analytically for the case of *perfect* homogeneity.

Theorem 5. *Let c be a compression for a financial system X and assume that the following values are the same across all $i \in N(c)$: e_i , $p_{oi} := \sum_{j:(j,i) \in O} p_{ji}$, $l_{Ii} := \sum_{j:(j,i) \in I} l_{ji}$, $l_{iO} := \sum_{j:(i,j) \in O} l_{ij}$, $l_{iI} := \sum_{j:(i,j) \in I} l_{ij}$, and c_i . In this case, we call the pair (X, c) homogeneous. Then c is a Pareto improvement.*

The proof of the theorem can be found in Appendix B. Note that, in the previous theorems, none of the *individual* values l_{ij} or c_{ij} need to be equal; they only need to be equal *in aggregate*. This is important for compressions that are more complex than a cycle. There are several simple classes of examples for (X, c) pairs that are *homogeneous* in the sense of Theorem 5:

- Let M be the adjacency matrix of a regular directed graph G , let $\gamma \geq 0$, and let $l = \gamma \cdot M$. Let $e_i := \delta \forall i$, where $\delta \geq 0$ is arbitrary. Let C be any cycle in G , let $\mu \leq \gamma$ and let $c = (C, \mu)$. Then X and c are homogeneous in the sense of Theorem 5.
- Consider the set of homogeneous (X, c) pairs where $N(c)$ has no incoming liabilities from O , i.e., $l_{ji} = 0$ for any $i \in N(c)$ and $(j, i) \in O$. It follows immediately from the definition that this set forms a polytope. In particular, convex combinations of two such homogeneous pairs are homogeneous. Note that Theorem 4 implies that such compressions are always incentivized if they are normal.
- If (X, c) is homogeneous and Y is any financial system, we can consider the disjoint union of the two and connect $N(c)$ to nodes in Y in any way such that $\sum_{j \in Y} l_{ij}$ is the same across $i \in N(c)$. If we call the new financial system Z , (Z, c) is homogeneous.

Figure 10 Simple financial system that is homogeneous together with the compression $c = (A-B-C, 1)$.



6.2 Towards a Sufficient Degree of Homogeneity

The theorem only makes a statement about (X, c) pairs where the different asset and liability values are *exactly* the same. However, intuition seems to suggest that it is not crucial that these values match exactly; there should be some slack. We study this hypothesis using an example of a particularly simple homogeneous (X, c) pair: a cycle where each bank has the same outgoing liabilities and the same external assets (see Figure 10). We then gradually make the pair less homogeneous by changing the external assets of a single node, say A. Given the role of the β parameter we revealed in Section 4, we hypothesize that more homogeneity is required for a Pareto improvement when β is higher. To isolate this effect, we fix $\alpha = 0.5$ and consider the effect of compression on the equities under variation of e_A and β .²⁰

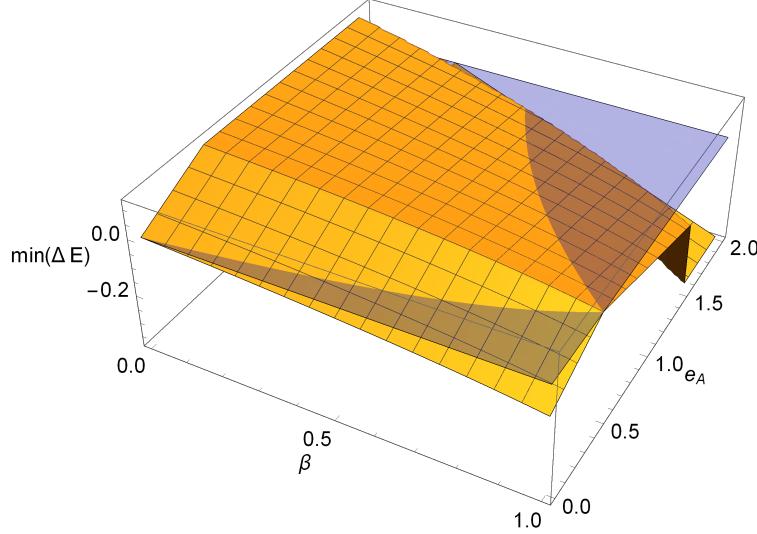
By Theorem 4, c is always incentivized, so the equities of banks A, B, and C will always weakly increase.²¹ Figure 11 depicts the minimum difference in equity across the other three banks D, E, and F. This number is above zero (i.e., the curve is above the blue zero plane in the figure) iff c is a Pareto improvement.

The figure reinforces our hypothesis. Consistent with our findings in Section 4, compression is a Pareto improvement for a wide range of e_A values when β is low. As we increase β , however, this region continuously becomes more narrow until it converges to the point 0.5 for $\beta = 1$: here, we require *exact* homogeneity for compression to be a Pareto improvement.

²⁰Further experiments not documented in the present paper show that the value of $\alpha \in (0, 1]$ affects our finding quantitatively, but not qualitatively. For $\alpha = 0$, banks are indifferent to compression for all but very high values of e_A because all banks default, so all value is destroyed. A choice of $\alpha = \beta$ does not imply a qualitative difference, either. The data is available upon request.

²¹Note, however, that all banks default in this examples unless $e_A \geq 1$, in which case A does not default in X^c . Only for $e_A \geq \approx 1.56$ A does not default in X and compression is detrimental for E.

Figure 11 Minimum difference in equity across the nodes D,E,F in Figure 10 when varying β and e_A . We fix $\alpha = 0.5$. Note that the minimum is *not* taken in absolute value. The blue plane is zero. Resolution: 0.01 steps in both parameters.



Our experiment suggests that one may be able to obtain a quantified version of Theorem 5 where an “amount of homogeneity” inversely dependent on β may be enough to guarantee that compression is a Pareto improvement. An analytical result in this vein would constitute substantial progress towards a deeper understanding of compression.

7 Conclusion

In this paper, we have studied *portfolio compression*, a post-trade mechanism which eliminates cycles in the financial network. We have studied the incentives for banks to engage in portfolio compression and its systemic effects in terms of banks’ equity. We have shown that whether or not compression is socially or individually desirable depends on the parameters of the financial system and the compression in a complex and non-monotonic way and the incentives for participating banks to perform compression may be misaligned with the social good. We have presented sufficient conditions under which compression is a Pareto improvement. We have shown that banks always have an incentive to agree to compression unless there are *feedback paths*. Finally, we have identified *homogeneity* of participating banks’ balance sheets as an important driver of efficiency of compression.

Our results reveal the default costs on interbank payments, encoded by the β parameter in our model, as a central factor for the ex-post desirability of compression. This is intuitive: when part of each payment is lost, it makes sense to have less of them. Our results on homogeneity suggests that portfolio compression may be particularly effective in a financial system where many actors are exposed to the same

(or similar kinds of) risks, so that if a large shock hits certain assets, then it hits all banks at once. Haldane (2009) describes how the pre-crisis financial system, through a convergence of investment and risk management strategies, had indeed been put into this very situation where “financial sector balance sheets became homogenised.” This may be taken as an indication that compression is indeed helpful to protect from another 2008 crisis; of course, further research is needed before any definite conclusions can be drawn.

An important avenue for future work will be to leave the ex-post model considered in this paper and take on an ex-ante perspective. Assuming a probability distribution of shocks and assuming that banks maximize their expected equity, which compressions are efficient and incentivized, respectively, in terms of expected equity? Incorporating uncertainty about the future in this way will be an important step towards an evaluation of compression from a practical point of view. This perspective may provide another explanation for why banks usually agree to compression and it may reveal another disconnect between the incentives of banks and the interests of society: perhaps compression is beneficial in expectation, but has strong detrimental effects in a small part of the probability space.

The insights from an ex-ante study may also help answer the question regarding regulatory incentives for compression. To date, one of the strongest incentives for banks to engage in compression is relief of regulatory capital requirements. Are these capital reliefs justified relative to the systemic impact of compression? Should we further incentivize or disincentivize compression? Or might the issue of portfolio compression point to a fundamental flaw of prescribing capital requirements relative to a single number, namely risk-weighted assets?

Finally, central clearing and portfolio compression in non-centrally cleared markets both serve the same goal: to reduce notional by rewiring the network. Ultimately, we need to be able to judge which of these approaches is more appropriate to reduce systemic risk in the future: central clearing, compression, or perhaps leaving the network as it is while increasing transparency? Answering these questions requires an understanding of the effect of network modifications on systemic risk, to which we hope to have contributed a first step.

Acknowledgments

We would like to thank (in alphabetical order) Stefano Battiston and Vitor Bosshard for helpful comments on this work.

A Effect of Compression on Prior Work

Researchers have previously studied the effect of network structure on systemic risk. We can examine how compression affects network structure at a grand scale and then apply their results. From this, we receive a first indication what effects we might expect from compression on average. Note that all we can expect from this exercise is an a tendency: apart from Veraart (2019), no prior piece of work has studied the effect of compression on a per-network basis.

Glasserman and Young (2015) studied the effect of a random shock to banks' external assets on the aggregate payments in the Eisenberg and Noe (2001) model (i.e., our model without default costs) where every bank has an *outside liability* $b_i \geq 0$ in addition to its interbank liabilities. The authors provided sufficient conditions under which the extent of financial contagion is bounded. These sufficient conditions are based on the following local properties of each bank:

- Its *financial connectivity* $\beta_i = l_i/(l_i + b_i)$, which measures the share of its interbank liabilities relative to all its liabilities.
- Its *leverage of outside assets* $\lambda_i = e_i/E_i(l)$, where $E_i(l) = e_i + \sum_j l_{ji} - l_i$ is the *book value of equity* of i , i.e., the equity of bank i if all its creditors pay in full.
- Its external assets e_i .

Observe that compression reduces financial connectivity and keeps leverage of outside assets and external assets the same. Thus, compression makes the sufficient conditions weaker and thus the overall bounds stronger.²² These results therefore suggest an overall positive effect of compression, which is however very coarse: unless compression transports a financial system from not satisfying the sufficient conditions to satisfying them, we should not expect any implications.

Bardoscia et al. (2017) assessed the stability of a financial system based on the greatest eigenvalue of the *leverage matrix* $\Lambda_{ij} = l_{ij}/E_j(l)$. The lower this eigenvalue, the more stable the system, where 1 is an important threshold value. Since compression reduces l point-wise and keeps $E_i(l)$ the same, it reduces the leverage matrix point-wise and thus also reduces its maximal eigenvalue. Thus, compression should make each individual network more stable in the sense of this paper.

Battiston et al. (2012) studied the spread of financial contagion in a model similar to the SIR (susceptible–infected–resistant) class of models from epidemiology. A level of financial distress of $\psi_i \in [0, 1]$ at bank i causes a level of financial distress of $\Lambda_{ij}\psi_i$ at each creditor j , where Λ is the leverage matrix from above. Since compression reduces Λ point-wise, it reduces contagion in this model. Any SIR-style model where

²²Precisely, this statement should be interpreted as follows: consider any of the sufficient conditions in Glasserman and Young (2015) and fix their parameters. If c is a compression for X and X satisfies the condition, then so does X^c ; there exists an X such that X^c satisfies the condition, but X does not.

infection propagates via liabilities will likely share this property.

Elliott, Golub and Jackson (2014) considered a model of cross-holdings of (essentially) equity cross-holdings and isolate two network measures, both of which must have intermediate levels for contagion to occur: *integration*, which measures the share of an organization held by financial, compared to external, actors. *Diversification* measures how spread out the financial part of the cross-holdings is. While there seem to be several options how one could define compression for a cross-holdership model, it seems that any sensible definition would reduce both integration and diversification. The overall effect would therefore be ambiguous.

Demange (2016) studied aggregate payments in the Eisenberg/Noe model like (Glasserman and Young, 2015), but focused on the partial derivatives of this measure with respect to changes in a banks' external assets. This is what she called the *threat index* μ_i of bank i . The threat index depends on the set of defaulting banks. If the set $D \subseteq N$ of banks default under the maximal clearing matrix of payments, then $\mu_i = 0$ if $i \notin D$ and for the collection $\mu_D = (\mu_i)_{i \in D}$ and $\pi_D := (\pi_{ij})_{i,j \in D}$ we have:

$$\begin{aligned} \mu_D &= \underline{1} + \pi_D \mu_D \\ \Leftrightarrow \mu_D &= (I - \pi_D)^{-1} \underline{1} = \sum_{k=0}^{\infty} \pi_D^k \underline{1}, \end{aligned}$$

where $\underline{1} := (1, \dots, 1)$.²³ Like in the related formula for Katz centrality (cf. the corresponding discussion in Demange (2016)), we thus have for $i \in D$:

$$\mu_i = \sum_{\substack{P \text{ path} \\ \text{starting at } i}} \prod_{(j,k) \in P} \pi_{jk}$$

The risk index of a bank i is the sum over the products of π values of paths starting at i . Compression transforms the π values in a non-linear and non-monotonic way. It may also move a bank into or out of default. This is why compression may increase or decrease a bank's risk index.

B Proofs

Proof of Lemma 1. The key technical step in the proof is that the precondition to the lemma implies the following:

Claim. For any $q \in [0, l - c]$ such that $q + c \geq p$ we have:

$$\Psi^c(q) + c \geq \Psi(q + c) \tag{3}$$

²³For the first line, see Demange (2016, Theorem 1). The second line is by linear algebra and the Neumann series. The series converges because we restricted our attention to the set of defaulting banks.

Proof of the Claim. Consider the inequality for an individual entry (i, j) on both sides. It is easy to see that

$$\begin{aligned} a_i^c(q) &= a_i(q) = a_i(q + c) - c_i \\ a_i'^c(q) &= a_i'(q) = a_i'(q + c) - \beta c_i \end{aligned}$$

and, of course, $l_i^c = l_i - c_i$. Thus, bank i defaults under q in X^c iff it defaults under $q + c$ in X . In case of non-default, (3) for (i, j) is of course equivalent to $l_{ij}^c + c_{ij} \geq l_{ij}$ and trivially true with equality. Thus, assume that i defaults in these cases. Then (3) for (i, j) is equivalent to:

$$\begin{aligned} \pi_{ij}^c a_i'(q) + c_{ij} &\geq \pi_{ij} a_i'(q + c) \\ \Leftrightarrow \pi_{ij}^c a_i'(q + c) - \pi_{ij}^c \beta c_i + c_{ij} &\geq \pi_{ij} a_i'(q + c) \\ \Leftrightarrow \Delta \pi_{ij} a_i'(q + c) + c_{ij} &\geq \beta \pi_{ij}^c c_i \end{aligned} \tag{4}$$

The second line is by the above identity and the third line is by simple algebra. We now distinguish the cases $\Delta \pi_{ij} \leq 0$ and $\Delta \pi_{ij} > 0$.

If $\Delta \pi_{ij} \leq 0$, then we can use the identity $c_{ij} = \pi_{ij}^c c_i - \Delta \pi_{ij} l_i$ to see that (4) is equivalent to:

$$\Delta \pi_{ij} (l_i - a_i'(q + c)) \leq (1 - \beta) \pi_{ij}^c c_i$$

Since i defaults under $q + c$ in X , the parenthesis on the left-hand side is ≥ 0 . Since $\Delta \pi_{ij} \leq 0$, the left-hand side is non-positive and the right-hand side is positive, so the inequality holds trivially.

If $\Delta \pi_{ij} > 0$, we observe that, if q is as specified in the claim, then $q + c \geq p$ and (since $q \geq 0$) also $q + c \geq c$. Thus, $q + c \geq p \vee c$ and this implies that i) i also defaults under $p \vee c$ in X and ii) $\Delta \pi_{ij} a_i'(q + c) \geq \Delta \pi_{ij} a_i'(p \vee c)$. Now (4) follows from the precondition of the lemma.²⁴

$\square_{(\text{Claim})}$

With the claim shown, we can prove the statement of the lemma by induction on the iteration sequence that converges to the clearing payments. More in detail, let:

$$\begin{aligned} p^0 &= l & p^{c,0} &= l - c & p'^n &= p^{c,n} + c \\ p^{n+1} &= \Psi(p^n) & p^{c,n+1} &= \Psi^c(p^{c,n}) \end{aligned}$$

We know that $p^n \rightarrow p$, $p^{c,n} \rightarrow p^c$, and thus $p'^n \rightarrow p'$ for $n \rightarrow \infty$, all from above (see Section 2.2). We now prove by induction that $p'^n \geq p^n \forall n$. This immediately implies $p' \geq p$.

²⁴Note that it is in fact *equivalent* to the precondition of the lemma that (4) hold for all specified q in the $\Delta \pi_{ij} > 0$ case lemma because $\bigwedge \{q + c \mid q \in [0, l - c], q + c \geq p\} = p \vee c$.

For $n = 0$, we trivially have $p'^0 = l = p^0$. Assuming $p'^n \geq p^n$, we have

$$p'^{n+1} = p^{c,n+1} + c = \Psi^c(p^{c,n}) + c \geq \Psi(p^{c,n} + c) = \Psi(p'^n) \geq \Psi(p^n) = p^{n+1}.$$

The first inequality is by the above claim. The claim is applicable because $p^{c,n} + c = p'^n \geq p^n \geq p$ by induction hypothesis. The second inequality is by the induction hypothesis and monotonicity of Ψ . \square

Proof of Lemma 2. 2: By assumption and monotonicity of Ψ^c , we have

$$\Psi^c(\tilde{p}) = \Psi^c(\tilde{p}_I \cup \tilde{p}_O) \geq \Psi^c(\tilde{p}_I \cup p_O) = \tilde{p}.$$

By the Knaster-Tarski fixed point theorem (see, e.g., Granas and Dugundji (2003)), we have

$$p^c = \bigvee \{q \in [0, l - c] \mid \Psi^c(q) \geq q\},$$

where “ \bigvee ” denotes the point-wise supremum. By the above, \tilde{p} is a member of the set on the right-hand side and thus $p^c \geq \tilde{p}$. Now trivially by assumption $p_O^c \geq \tilde{p}_O \geq p_O$.

If $i \notin N(c)$, then $(j, i) \in O$ whenever $l_{ji} > 0$ for all j . Thus, since $p_O^c \geq p_O$, we have $a_i^c(p^c) \geq a_i(p)$. At the same time, since $i \notin N(c)$, $l_i^c = l_i$. This implies $E_i^c \geq E_i$.

For the second sentence, we show that under the preconditions also $\tilde{p}' \geq p'$. Then the statement follows by part 3. Let $(i, j) \in I$. If there is a k such that $(i, k) \in O$, then we have $\tilde{r}_i l_{ik} = \tilde{p}_{ik} \geq p_{ik} = \tilde{r}_i l_{ik}$ and thus $\tilde{r}_i \geq r_i$ and

$$\tilde{p}'_{ij} = \tilde{r}_i(l_{ij} - c_{ij}) + c_{ij} = \tilde{r}_i l_{ij} + (1 - \tilde{r}_i)c_{ij} \geq r_i l_{ij} = p_{ij}.$$

3: The statement implies the precondition to the first sentence of part 2 since $\tilde{p}'_O = \tilde{p}_O$ and $p'_O = p_O$. Therefore, $p' = p^c + c \geq \tilde{p} + c = \tilde{p}'$. And $\tilde{p}' \geq p$ by assumption.

1: Similarly to the proof of Lemma 1, the result is driven by the following statement, which only holds if c is chord-free:

Claim. If c is normal for X and $q_I \in [0, (l - c)_I]$ be arbitrary, then:

$$\Psi_I^c(q_I \cup p_O) + c \geq \Psi_I((q_I + c_I) \cup p_O) \tag{5}$$

Proof of the Claim. Let $(i, j) \in I$ and define for convenience $q'_I := q_I + c_I$. We

distinguish two cases based on the default of i . Note that

$$\begin{aligned} a_i^c(q_I \cup p_O) &= e_i + \sum_{k:(k,i) \in I} q_{I,k,i} + \sum_{k:(k,i) \in O} p_{O,k,i} \\ &= e_i + \sum_{k:(k,i) \in I} (q_{I,k,i} + c_{ki}) + \sum_{k:(k,i) \in O} p_{O,k,i} - c_i = a_i(q'_I \cup p_O) - c_i \\ a_i'^c(q_I \cup p_O) &= a_i'(q'_I \cup p_O) - \beta c_i \\ l_i^c &= l_i - c_i. \end{aligned}$$

Thus, i defaults (i.e., assets are below liabilities) in X^c under $q_I \cup p_O$ iff it defaults in X under $q'_I \cup p_O$. In case of non-default, we have

$$\Psi_{ij}^c(q_I \cup p_O) + c = l_{ij}^c + c = l_{ij} = \Psi_{ij}(q'_I \cup p_O),$$

and in particular (5) holds for (i, j) . In case of default, we have

$$\begin{aligned} \Psi_{ij}^c(q_I \cup p_O) + c &= \pi_{ij}^c a_i'^c(q_I \cup p_O) + c_{ij} = \pi_{ij}^c a_i'(q'_I \cup p_O) - \pi_{ij}^c \beta c_i + c_{ij} \\ \Psi_{ij}(q'_I \cup p_O) &= \pi_{ij} a_i'(q'_I \cup p_O) \end{aligned}$$

Taking the difference, we have $\Psi_I^c(q_I \cup p_O) + c \geq \Psi_I(q'_I \cup p_O)$ iff

$$\Delta \pi_{ij} a_i'(q'_I \cup p_O) + c_{ij} \geq \beta \pi_{ij}^c c_i.$$

Via the elementary identity $c_{ij} = \pi_{ij}^c c_i - \Delta \pi_{ij} l_i$, this is equivalent to

$$\Delta \pi_{ij} (l_i - a_i'(q'_I \cup p_O)) \leq (1 - \beta) \pi_{ij}^c c_i.$$

On the left-hand side, since c is normal, the first factor is $\Delta \pi_{ij} \leq 0$, and since i defaults, the second factor is ≥ 0 . Thus, the left-hand side is non-positive. Since the right-hand side is always non-negative, the inequality holds. \square

With the claim proven, we can show our original statement $\tilde{p}'_I \geq p_I$. Define two sequences of vectors

$$\begin{aligned} \tilde{p}_I^0 &:= l_I - c_I & p_I^0 &:= l_I \\ \tilde{p}_I^{n+1} &:= \Psi_I^c(\tilde{p}_I^n \cup p_O) & p_I^{n+1} &:= \Psi_I(p_I^n \cup p_O). \end{aligned}$$

By definition, \tilde{p}_I is the maximal fixed point of the map $q_I \in [0, (l - c)_I] \mapsto \Psi_I^c(q_I \cup p_O)$. Also, p_I is the maximal fixed point of the map $q_I \in [0, l_I] \mapsto \Psi_I(q_I \cup p_O)$ because p is the maximal fixed point of Ψ .²⁵ Thus, by the well-known extension of

²⁵This follows via the Tarski-Knaster fixed point theorem if we consider Ψ as a function of two variables $\Psi(q_I, q_O)$. See Lemma 3 in Appendix C.

the Tarski-Knaster fixed point theorem (see Remark 1), we have that $\tilde{p}_I^n \rightarrow \tilde{p}_I$ and $p_I^n \rightarrow p_I$ from above for $n \rightarrow \infty$.

Note that $\tilde{p}_I^0 = p_I^0$ and via the claim, $\tilde{p}_I^n + c_I \geq p_I^n \Rightarrow \tilde{p}_I^{n+1} + c_I \geq p_I^{n+1}$. Thus, by induction, $\tilde{p}_I^n + c_I \geq p_I^n$ and by taking the limit, $\tilde{p}'_I = \tilde{p}_I + c_I = \lim_n (\tilde{p}_I^n + c_I) \geq \lim_n p_I^n = p_I$.

Towards the second part of the statement, if c is chord-free and $i \in N(c)$, then for all $j \in N$ we either have $(i, j) \in I$ and thus $\tilde{p}'_{ij} \geq p_{ij}$ or $c_i = c_{ij} = 0$ and thus $\tilde{p}'_{ij} = \tilde{p}_{ij} = p_{ij}$. Therefore, $a_i(\tilde{p}') \geq a_i(p)$ and thus $B_i^c(\tilde{p}) = B_i(\tilde{p}') \geq B_i(p)$. In particular, $\tilde{E}_i \geq E_i$. \square

Proof of Theorem 5. Consider the immediate-effect payments \tilde{p} and consider the recovery rates $\tilde{r}_i = \frac{\tilde{p}_{i*}}{l_i^c}$. We show that $\tilde{r}_i \geq r_i \forall i \in N(c)$. This implies the statement via Lemma 2 because then

$$\tilde{p}'_{ij} = \tilde{r}_i(l_{ij} - c_{ij}) + c_{ij} = \tilde{r}_i l_{ij} + (1 - \tilde{r}_i)c_{ij} \geq \tilde{r}_i l_{ij} = p_{ij} \forall i \in N(c), j \in N.$$

The recovery rates $\tilde{r}_i =: \tilde{\rho}$ and $r_i =: \rho$ are the same across all $i \in N(c)$. To see this, consider the sequences

$$\begin{aligned} \tilde{p}^0 &:= (l - c)_I \cup p_O & p^0 &:= l_I \cup p_O \\ \tilde{p}^{n+1} &:= \Psi^c(\tilde{p}_I^n \cup p_O) & p^{n+1} &:= \Psi(p_I^n \cup p_O). \end{aligned}$$

We know that $\tilde{p}^n \rightarrow \tilde{p}$ and $p^n \rightarrow p$ monotonically decreasing for $n \rightarrow \infty$ (see the proof of Lemma 2). We show by induction that $\tilde{\rho}^n := \tilde{r}_i^n := \frac{\tilde{p}_{i*}^n}{l_i^c}$ and $\rho^n := r_i^n := \frac{p_{i*}^n}{l_i}$ are equal, respectively, across $i \in N(c)$. This implies that also \tilde{r}_i and r_i must be equal, respectively, across $i \in N(c)$, via continuity from above. For $n = 0$, the two statements are trivial because $\tilde{r}_i^0 = r_i^0 = 1$. Given the statement for n , we have:

$$\begin{aligned} a_i^c(\tilde{p}_I^n, p_O) &= e_i + p_{Oi} + \sum_{j:(j,i) \in I} \tilde{p}_{ji}^n = e_i + p_{Oi} + \tilde{\rho}^n (l_{Ii} - c_i) \\ a_i'^c(\tilde{p}_I^n, p_O) &= \alpha e_i + \beta p_{Oi} + \beta \sum_{j:(j,i) \in I} \tilde{p}_{ji}^n = \alpha e_i + \beta p_{Oi} + \beta \tilde{\rho}^n (l_{Ii} - c_i) \\ l_i^c &= l_{iO} + l_{iI} - c_i. \end{aligned}$$

By assumption, all of these values are the same across $i \in N(c)$ and thus,

$$\tilde{r}_i^{n+1} = \begin{cases} 1 & \text{if } a_i^c(\tilde{p}_I^n, p_O) \geq l_i^c \\ \frac{a_i'^c(\tilde{p}_I^n, p_O)}{l_i^c} & \text{if } a_i^c(\tilde{p}_I^n, p_O) < l_i^c \end{cases}$$

are the same across $i \in N(c)$. An analogous argument holds for the r_i^n .

The above-described symmetry implies that either all or no bank in $N(c)$ default

and this is independent of compression. More in detail:

- If, equivalently across $i \in N(c)$, $l_i \leq a_i(p_I^0, p_O) = e_i + p_{Oi} + l_{Ii}$, then $p^1 = p^0$ and thus $p^n = p^0 \forall n$, $p = p^0$, and $\rho = 1$. Further, $l_i^c = l_i - c_i \leq a_i(p_I^0, p_O) - c_i = a_i^c(\tilde{p}_I^0, p_O)$ and thus $\tilde{\rho} = \tilde{p}^0$ and $\tilde{\rho} = 1$.
- If $l_i > e_i + p_{Oi} + l_{Ii}$, then $\rho \leq p^1 < 1$ and $\tilde{\rho} \leq \tilde{p}^1 < 1$ by the same argument.

If $\tilde{\rho} = \rho = 1$, our statement is of course trivially true. If $\tilde{\rho}, \rho < 1$, we have by symmetry and the clearing identity (1) (let $i \in N(c)$ be arbitrary):

$$\begin{aligned}\rho &= \frac{\alpha e_i + \beta p_{Oi} + \beta \rho l_{Ii}}{l_i} \\ \tilde{\rho} &= \frac{\alpha e_i + \beta p_{Oi} + \beta \tilde{\rho} (l_{Ii} - c_i)}{l_i - c_i}\end{aligned}$$

Solving for ρ and $\tilde{\rho}$, respectively, and simplifying yields:

$$\begin{aligned}\rho &= \frac{\alpha e_i + \beta p_{Oi}}{l_i - \beta l_{Ii}} \\ \tilde{\rho} &= \frac{\alpha e_i + \beta p_{Oi}}{l_i - \beta l_{Ii} - (1 - \beta)c_i}\end{aligned}$$

And this obviously implies $\tilde{\rho} \geq \rho$. Note that both fractions are well-defined since, if $l_i - \beta l_{Ii} - (1 - \beta)c_i = 0$, then in particular $l_i - c_i \leq l_{Ii}$ and thus, by the above discussion, i would not default. Likewise for $l_i - \beta l_{Ii}$. \square

C Fixed Points of Monotonic Functions

Lemma 3. *Let K, L be complete lattices and let $F : K \times L \rightarrow K \times L$ be monotonic. Let (x^+, y^+) be the (unique) maximal fixed point of F . Then x^+ is the maximal fixed point of the function $x \mapsto F_1(x, y^+)$.*

Note that the converse of the lemma does not hold. For example, let $K = L = [0, 1]$ and $F(x, y) = (y, x)$. Then 0 is the unique (and therefore maximal) fixed point of $F_1(\cdot, 0)$ and of $F_2(0, \cdot)$, but $(0, 0)$ is not the maximal fixed point of F .

Proof. Let x^* be the maximal fixed point of $x \mapsto F_1(x, y^+)$. We will show that $x^* = x^+$. As x^+ is a fixed point of $F_1(\cdot, y^+)$ by choice of (x^+, y^+) , we have $x^+ \leq x^*$. It remains to show “ \geq ”.

We have $(x^*, y^+) \leq F(x^*, y^+)$. To see this, consider the two components of F separately. We have $x^* = F_1(x^*, y^+)$ by choice of x^* . For the second component, note that $y^+ = F_2(x^+, y^+) \leq F_2(x^*, y^+)$, where the equality is by choice of (x^+, y^+) and the inequality is because $x^+ \leq x^*$. Now, by Tarski’s fixed point theorem, we have

$$(x^+, y^+) = \bigvee \{(x, y) \mid (x, y) \leq F(x, y)\},$$

where “ \vee ” denotes the supremum. By the above, (x^*, y^+) is a member of the set on the right-hand side and thus $(x^+, y^+) \geq (x^*, y^+)$, i.e., $x^+ \geq x^*$. \square

References

- Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi.** 2015. “Systemic Risk and Stability in Financial Networks.” *American Economic Review*, 105(2): 564–608.
- Aldasoro, Iñaki, and Torsten Ehlers.** 2018. “The credit default swap market: what a difference a decade makes.” *BIS Quarterly Review*.
- Amini, Hamed, Damir Filipović, and Andreea Minca.** 2015. “To fully net or not to net: Adverse effects of partial multilateral netting.” *Operations Research*, 64(5): 1135–1142.
- Bardoscia, Marco, Stefano Battiston, Fabio Caccioli, and Guido Caldarelli.** 2017. “Pathways towards instability in financial networks.” *Nature Communications*, 8: 14416.
- Barucca, Paolo, Marco Bardoscia, Fabio Caccioli, Marco D’Errico, Gabriele Visentin, Stefano Battiston, and Guido Caldarelli.** 2016. “Network Valuation in Financial Systems.” Working Paper, <https://ssrn.com/abstract=2795583>.
- Basel Committee on Banking Supervision.** 2011. “Basel III: A global regulatory framework for more resilient banks and banking systems.” <http://www.bis.org/publ/bcbs189.pdf>.
- Battiston, Stefano, Michelangelo Puliga, Rahul Kaushik, Paolo Tasca, and Guido Caldarelli.** 2012. “DebtRank: Too central to fail? Financial networks, the FED and systemic risk.” *Scientific reports*, 2.
- Cui, Zhenyu, Qi Feng, Ruimeng Hu, and Bin Zou.** 2018. “Systemic risk and optimal fee for central clearing counterparty under partial netting.” *Operations Research Letters*, 46(3): 306–311.
- Demange, Gabrielle.** 2016. “Contagion in Financial Networks: A Threat Index.” *Management Science*, 64(2): 955–970.
- D’Errico, Marco, and Tarik Roukny.** 2019. “Compressing Over-the-Counter Markets.” Working Paper, <https://ssrn.com/abstract=2962575>.
- Duffie, Darrell.** 2018. “Compression Auctions With an Application to LIBOR-SOFR Swap Conversion.”
- Duffie, Darrell, and Haoxiang Zhu.** 2011. “Does a central clearing counterparty reduce counterparty risk?” *Review of Asset Pricing Studies*, 1(1): 74–95.
- Duffie, Darrell, Martin Scheicher, and Guillaume Vuillemey.** 2015. “Central clearing and collateral demand.” *Journal of Financial Economics*, 116(2): 237–256.
- Eisenberg, Larry, and Thomas H Noe.** 2001. “Systemic risk in financial systems.” *Management Science*, 47(2): 236–249.
- Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson.** 2014. “Financial Networks and Contagion.” *American Economic Review*, 104(10): 3115–53.
- European Securities and Markets Authority.** 2017. “OTC Derivatives and Clearing Obligation.” Retrieved November 23, 2019, <https://www.esma.europa.eu/regulation/post-trading/otc-derivatives-and-clearing-obligation>.
- Feinstein, Zachary, Weijie Pang, Birgit Rudloff, Eric Schaanning, Stephan Sturm, and Mackenzie Wildman.** 2017. “Sensitivity of the Eisenberg-Noe clearing vector to individual interbank liabilities.” Working Paper, <https://ssrn.com/abstract=3033476>.

- Financial Stability Board.** 2017. “OTC Derivatives Markets Reform — Twelfth Progress Report on Implementation.” <http://www.fsb.org/2017/06/otc-derivatives-market-reforms-twelfth-progress-report-on-implementation/>.
- Glasserman, Paul, and H. Peyton Young.** 2015. “How likely is contagion in financial networks?” *Journal of Banking & Finance*, 50(0): 383–399.
- Granas, Andrzej, and James Dugundji.** 2003. *Fixed point theory*. Springer.
- Haldane, Andrew G.** 2009. “Rethinking the Financial Network.” Speech at the Financial Student Association, Amsterdam, April 28, 2009, <https://www.bankofengland.co.uk/speech/2009/rethinking-the-financial-network>.
- Korte, Bernhard, and Jens Vygen.** 2012. *Combinatorial Optimization: Theory and Algorithms*. . 5th ed., Berlin, Germany:Springer.
- O’Kane, Dominic.** 2017. “Optimising the multilateral netting of fungible OTC derivatives.” *Quantitative Finance*, 17(10): 1523–1534.
- Rogers, LCG, and Luitgard AM Veraart.** 2013. “Failure and rescue in an interbank network.” *Management Science*, 59(4): 882–898.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2018. “Portfolio Compression: Positive and Negative Effects on Systemic Risk.” Research Note, <https://ssrn.com/abstract=3135960>.
- Schuldenzucker, Steffen, Sven Seuken, and Stefano Battiston.** 2019. “Default Ambiguity: Credit Default Swaps Create New Systemic Risks in Financial Networks.” *Management Science*, Published in Articles in Advance.
- Veraart, Luitgard AM.** 2019. “When Does Portfolio Compression Reduce Systemic Risk?” Working Paper, <https://ssrn.com/abstract=3311176>.
- Yellen, Janet L.** 2013. “Interconnectedness and Systemic Risk: Lessons from the Financial Crisis and Policy Implications.” Speech at the American Economic Association / American Finance Association Joint Luncheon, San Diego, California, January 4, 2013, <https://www.federalreserve.gov/newsevents/speech/yellen20130104a.htm>.

Curriculum Vitae

Personal Information

Name	Steffen Schuldenzucker
Date of Birth	September 23, 1988
Place of Birth	Bonn, Germany
Nationality	German

Education

09/2014 – 08/2019	<i>PhD Candidate</i> University of Zurich, Switzerland Department of Informatics Computation and Economics Research Group
02/2012 – 08/2014	<i>MSc in Mathematics</i> University of Bonn, Germany
10/2008 – 01/2012	<i>BSc in Mathematics</i> University of Bonn, Germany

Research Internships

06/2018 – 08/2018	<i>Research Intern</i> Department Financial Stability Deutsche Bundesbank, Frankfurt, Germany
-------------------	---

Other Professional Experience

08/2012 – 08/2014	<i>Student Employee, IT Advisory</i> , KPMG AG, Berlin, Germany
02/2012 – 06/2012	<i>Lecturer</i> , Lernen Bohlscheid, Cologne, Germany
10/2006 – 06/2012	<i>Software Developer</i> , Digital Gecko Ltd., Bonn, Germany
08/2011 – 10/2011	<i>Intern</i> , Tsuru Capital LLC, Tokyo, Japan
02/2010 – 05/2011	<i>Student Employee, Mathematical Software Development</i> , much-net AG, Bonn, Germany
06/2009 – 12/2010	<i>Student Employee, High School Activities</i> , University of Bonn, Germany
06/2008 – 09/2008	<i>Intern IT</i> , BW Fuhrpark, Bonn, Germany