Essays on the New Financial Landscape

A thesis submitted for the degree of
Doctor of Philosophy

Peter Zimmerman

Green Templeton College
University of Oxford
Hilary Term 2019

Abstract

This dissertation consists of three essays, each examining the implications of a major change to the financial landscape over the past decade.

In the first essay, I show that price formation for cryptocurrencies is markedly different from that for other asset classes. I study two distinctive characteristics of cryptocurrencies. First, the blockchain structure of cryptocurrency means that there is a limit on the number of transactions that can be settled in any given time. Second, as a cryptocurrency is a form of money not backed by any asset class, its value depends on the extent to which it is used as a means of payment. Taken together, these characteristics mean that speculative activity can crowd out payment usage. Trading can lower the price of the asset, even when it is buy-side trading. I show that this may explain high price volatility.

In the second essay, we propose a novel channel for credit crunches. A borrower fears that there will be distress in the financial sector, in which case her relationship lender may not be able to finance her project. To insure against this risk, she chooses a smaller project that uninformed outsiders would be willing to finance. Financial distress therefore leads to lower investment, but not via the usual supply or demand side channels. We show that a small increase in distress risk can lead to a large downward jump in investment. This is joint work with Alan Morrison and Joel Shapiro.

Finally, the third essay examines how derivatives trading networks are affected by the introduction of mandatory central clearing. There can be a trade-off: central clearing can increase the mean of net exposures, but reduce the variance. We show that the trade-off can never be in the other direction. This suggests that, the more risk-averse agents are, the beneficial central clearing is likely to be. This is joint work with Rodney Garratt.

Dedication

To Jessica, my C.B.

Acknowledgements

I owe an immense debt of gratitude to my advisers, Joel Shapiro and Alan Morrison, for their guidance and support throughout my studies at Oxford. Their advice kept me focused on what was important. Joel, in particular, has always been ready to discuss ideas and address problems, and for that I owe him special thanks.

I am grateful to my co-author Rodney Garratt for his encouragement over the years. It is he more than anybody else that helped me to realise that research is what I really enjoy. I would also like to thank Itay Goldstein for hosting my visit to the Wharton School and for his support and advice.

Thank you to the ESRC, Green Templeton Partnership, Clarendon Fund, Saïd Business School Foundation, and the Oxford Man Institute for providing financing during my PhD. Parts of this thesis were worked on at the Bank of England, Oxford Man Institute, Saïd Business School, and the Wharton School, all of which I thank for their hospitality.

My family and friends have waited a very long time for me to emerge from my cocoon so that they can see me again. I thank them all for their patience and support.

Most of all, I owe my deepest thanks to my wife Jessica. She has sacrificed a great deal while I have been working on my thesis. She has provided a huge amount of support, even if she has not always realised it. I hope that I have made her proud, and that this enterprise will prove to be worthwhile.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The financial landscape has changed considerably in the wake of the global financial crisis of 2007–8. This dissertation studies the consequences for three of the key functions of banks described by Freixas and Rochet (2008): payment services, credit provision, and risk management.

1.1 Blockchain structure and cryptocurrency prices

Fear of bank failure and the risk of inflation has driven the development and adoption of cryptocurrencies as an alternative means of payment. Satoshi Nakamoto introduced the idea for Bitcoin, the first cryptocurrency, in October 2008 (Nakamoto 2008). On 3 January 2009, he produced the very first block, and encoded the following text into it:

```
The Times 03/Jan/2009 Chancellor on brink of second bailout for banks
```

This text referred to a British newspaper headline that day. As well as providing a time-stamp, Nakamoto was making a comment about the instability he perceived as inherent to the banking system. Reliance on commercial banks and monetary authorities was a key concern of libertarian groups such as the 'Cypherpunk' movement, which were instrumental in the development of cryptocurrencies (Popper 2015). In the wake of the

financial crisis of 2008, increased perception of bank risk helped boost awareness of and interest in this new means of payment that does not rely on banks or central banks.

In my first essay, I show that price formation for cryptocurrencies is markedly different to that for other asset classes. My results depend on two distinctive characteristics of cryptocurrencies. First, the blockchain structure of cryptocurrency means that there is a limit on the number of transactions that can be settled in any given time. For example, the Bitcoin network can handle an average of 7 transactions per second. Second, as a cryptocurrency is fiat money and not backed by any asset class, its value depends on the extent to which it is used as a means of payment. To my knowledge, there is no other asset class that has both of these properties.

The limit on the settlement capacity of cryptocurrencies is a consequence of their decentralised nature. This means that a user need not rely on the competence or honesty of any individual agent in order for a cryptocurrency to function as money. This idea was a key impetus behind the development of cryptocurrencies. Decentralisation is achieved by the ledger of transactions, or blockchain, being maintained by a large number of agents, called miners. Each miner keeps a local copy of the ledger and, so long as about half of the miners can agree on the current state of the ledger, then consensus can be maintained. When a miner wishes to update the blockchain, she creates a list of new transactions — called a block — and proposes that this is appended to the blockchain. To do this, she needs to be able to communicate the update to the other miners, so that decentralised consensus can be maintained. If blocks are large or are created too frequently, then miners may not be able to update their local copies of the blockchain synchronously. Then they will begin to disagree about the current state of the blockchain, and decentralised consensus may break down. To address this risk, cryptocurrency protocols place limits on the size of blocks, and on the rate of block creation.

The limited settlement capacity makes blockchain space scarce, and so there is competition between users who have an urgent need to settle payments. In particular, speculative activity can crowd out payment usage. This seems to have been the case during 2017, when a surge in speculative interest in cryptocurrencies made them more difficult to use as a means of exchange. For example, in January 2018, the North American Bitcoin Conference stopped accepting bitcoin for tickets, because the waiting times were too long

(Choudhury 2018).

Cryptocurrencies have no fundamental value — unlike, for example, commodity-backed money — and so, like fiat money, they only have a positive price to the extent that agents believe that others will accept them as a means of payment. This means that, all else equal, when speculative activity is high and crowds out payment activity on the blockchain, then the price of cryptocurrency should fall.

There is an apparent paradox here: speculators buy cryptocurrency because they believe it will be a popular means of payment but, in doing so, they cause it to become a less useful payments medium, due to the crowding out effect. Interestingly, this implies the demand curve of cryptocurrency can be upward-sloping, since an increase in speculative demand can drive down prices, all else equal. Price formation for cryptocurrencies is therefore very different from that for conventional asset classes. This is because, to my knowledge, no other asset class has both the settlement constraint together with the endogenous valuation property.

I find the price explicitly by solving a model of cryptocurrency usage. Households choose whether to use cryptocurrency or cash to make a payment. They base their decisions on private signals about the usefulness of the underlying blockchain technology. If the cryptocurrency is used for payments, then it is valuable, and speculators buy. I adapt global games techniques — typically used in models of coordination problems, such as currency attacks (Morris and Shin 1998) or bank runs (Goldstein and Pauzner 2005) — to show that there is a unique equilibrium and thus a unique price. I show that the crowding out effect makes the price of the cryptocurrency more volatile than it would otherwise be.

My model has a complication that is not present in other global games models. Users can obtain settlement priority by paying fees to miners, who allocate space to the highest bidders. This means that we do not have global strategic complementarities: if a user believes others are using cryptocurrency and paying a high fee, she might prefer not to use it at all. To deal with this problem, I adapt existing global games techniques, and provide a condition under which there is a unique equilibrium. I make two technical contributions. First, to my knowledge, my model is the first to endogenise both the market for blockchain space and for the cryptocurrency itself. Second, this is the first paper to use global games techniques in a model of money.

1.2 Relationship lending when banks are fragile

Bank lending and real investment both tend to fall following financial crises. Traditionally, this is explained either through 'demand-side' or 'supply-side' channels. Demand-side explanations focus on the choice of borrowers to borrow less and reduce investment due to expectations of lower future economic growth and demand (e.g. Kahle and Stulz 2013). Supply-side explanations emphasise the inability or unwillingness of banks to advance credit to borrowers at the same terms as they did previously, typically due to financial constraints or agency problems (e.g. Duchin, Ozbas, and Sensoy (2010)). There is evidence that both of these channels were active during the recent global financial crisis.

In joint work with Alan Morrison and Joel Shapiro of Oxford University, I postulate a new, novel channel for a credit crunch. In our model, lenders compete to finance a borrower. The borrower's quality is known only to one lender, called the 'insider'. The uninformed lenders are concerned about adverse selection, and so require higher rates than the insider to lend to a good borrower. The borrower can choose the size of her project, which exhibits decreasing returns to scale. Therefore smaller projects tend to attract more competition among lenders, but larger projects are more reliant on the insider for financing. The insider can more easily extract rents from larger projects. The borrower trades off between being able to keep a relatively small piece of a large project (using inside capital), or a bigger proportion of a smaller project (using outside capital).

With some probability, the insider suffers a shock which prevents him from participating in the market for loans. The outcome of this shock is not observed. As the probability of the shock increases, outsiders face less adverse selection, and so offer more generous lending conditions. This makes a smaller project more attractive for the borrower. At the same time, by reducing the size of the project and increasing the expected return, the borrower increases the probability that an outsider offers to finance it. This allows the borrower to insure against the risk that nobody bids for the project and it is not financed at all.

We show that there is a unique equilibrium, whose properties depend on the probability of the insider's shock. When this probability is below a critical value, the borrower chooses larger projects and relies on the insider for funding, at the risk of not being financed at

all. When the shock probability is above the critical value, the borrower chooses a project small enough that it is guaranteed to be financed. At the critical value shock probability, there is a discontinuity in project size and in welfare. This is because the insider is able to extract rent from the larger projects. This discontinuous drop in investment and welfare resembles a credit crunch.

Our credit crunch channel can be described as having both demand- and supply-side aspects. Borrowers respond to supply-side developments by reducing their demand for credit. While the ultimate cause of the crunch is supply-side (increased probability of a shock to informed lending), the proximate cause is demand-side (the borrower chooses to reduce the size of her investment). This channel is driven by informational factors. It does not resemble a standard demand-side channel, because there is no change to the productivity of the borrower's project. Nor does it resemble a standard supply-side channel, because banks do not reduce the amount of available credit. Instead, it is the borrower who decides on a smaller project.

An important friction in our model is that the insider cannot credibly commit ex ante not to exploit his information to the greatest extent possible. We examine a policy intervention that makes such a commitment feasible: a social planner can choose to impose an interest rate cap on the bank. This is equivalent to the insider making a commitment to leave some of the surplus for the borrower, and so inducing the borrower to choose a larger project. This can be welfare-positive, and make projects possible which would be infeasible without such an intervention.

1.3 Centralised netting in financial networks

Poor risk management by financial institutions was an important factor in the recent financial crisis. Market participants who believed they were well-hedged found that they had unexpected exposures, because the counterparty providing the hedge either failed or was unwilling to continue financing their position (Cunliffe 2018). As part of the post-crisis reform effort, policymakers have mandated central clearing in a variety of derivatives markets, particularly high-volume standardised credit default and interest rate swaps. My third essay examines how this reform affects the structure of derivative-trading networks,

and the implications for risk. This is joint work with Rodney Garratt of the University of California, Santa Barbara.

Mandatory central clearing means that all trades in the relevant derivative have to be novated to a central counterparty, or CCP. For example, suppose that dealers A and B enter into a credit default swap agreement with reference entity X, so that A pays B £100 if company X defaults on its debt. After novation to a CCP, this becomes two separate trades: A sells credit protection on X to the CCP, and the CCP sells credit protection on X to B. In this way, B is insulated from counterparty risk with respect to A, and vice versa. Of course, both are now exposed to counterparty risk vis-à-vis the CCP, and the CCP itself is exposed to the failure of A and B. So long as the CCP itself is able to manage its risks more effectively than the market participants, then central clearing can reduce counterparty risks in derivative markets.

The focus of this chapter is on the effect of central clearing on the aggregate level of exposures in the system. By netting offsetting claims, central clearing can reduce systemic risk and collateral requirements. This kind of centralised netting occurs in other markets too, such as triparty repo, and netting in payments networks. Introducing centralised netting in a subset of exposures has the effect of improving netting amongst those exposures, because they are all now novated to the same counterparty. But it disrupts bilateral netting sets amongst those exposures which are not novated. Our essay shows how this trade-off depends on the structure of the network.

Duffie and Zhu (2011) show that, when a single asset class is centrally cleared, bilateral exposures between agents may increase, as a result of reduced netting opportunities across pairs of counterparties. This leads to an overall loss in netting efficiency. They derive conditions, in terms of the number of asset classes and the number of dealers, for whether or not centralised netting in the form of a CCP is beneficial. In Duffie and Zhu's framework, all agents in the network are assumed to trade with one another. But, typically, real-world financial networks are not completely connected. Empirical studies show that there tends to be a number of well-connected counterparties, coupled with a larger number of more poorly-connected counterparties.

We develop a model of centralised netting that extends the analytic framework developed by Duffie and Zhu to more general network models. We develop analytical results that do

not depend on any assumptions about the precise structure of the network. In addition, we examine the effect of centralised netting on the variance, as well as the expectation, of netted exposures. Our rationale is that, if the agents have some aversion to volatility or extreme outcomes, then they are likely to benefit from a lower variance of net exposures, as well as a lower mean. As in Duffie and Zhu, the introduction of centralised netting is more likely to be beneficial when there are more agents, and when there are fewer asset classes. We show that the set of networks for which the expectation of exposures is reduced is a subset of those for which the variance is reduced. This suggests that centralised netting is more likely to be beneficial when agents care more about variance (i.e. when agents are more risk-averse).

We also consider how the results may change if the agents can strategically reform the network in response to the introduction of centralised netting. We show that, in certain circumstances, centralised netting may be more efficient when there are *more* asset classes, rather than fewer. This can happen if the post-clearing network becomes less connected, perhaps as a result of agents raising their bilateral risk limits as a result of the new system infrastructure, and thus needing to form fewer links.

Finally, we discuss how our results can be used in practice, and the implications of our findings for policy. The issue of clearing networks has become an important policy issue in Europe, because the European Commission is considering requiring the clearing of euro-denominated swaps to be carried out within the European Union. Most such clearing is done in the United Kingdom, which has announced its intention to leave the European Union. Our model allows analysts to forecast the disruption to netting in clearing networks that may result.

Chapter 2

Blockchain structure and cryptocurrency prices

2.1 Introduction

Market prices of cryptocurrencies such as bitcoin are much more volatile than those of other assets. Table 2.1 shows that cryptocurrencies have historically had more volatile returns than conventional asset classes. High price volatility and speculative activity contrasts with low current levels of usage of cryptocurrency as a means of payment. Only three of the top 500 e-commerce merchants accept bitcoin, and this number has fallen since 2016.[1] Carney (2018) and Krugman (2018) argue that, because monetary instruments are only valuable to the extent they are used to make payments, speculative activity is symptomatic of a bubble. However, some evidence suggests that prices do respond to news about long-term usage of cryptocurrencies as a monetary instrument, so speculation is, to some degree, rational. For example, the price of bitcoin was sent to new heights in September 2017 by an unsubstantiated rumour that the online retailer Amazon would accept bitcoin (see Greene 2018).[2]

In this paper, I show that high price volatility is a consequence of an unusual feature of

[1] See Faucette et al. (2017).

[2] Biais et al. (2018) find evidence that the price of bitcoin is affected by the ease with which it can be used as money.

Table 2.1: **Comparison of daily price volatility of cryptocurrencies and other asset classes. Volatility is defined as 30-day rolling backward-looking window of standard deviation of daily returns of USD prices, measured in percentage points. Data for cryptocurrencies are from `blockchain.info` and `coinmetrics.io` and run from the earliest available dates to 30 September 2018. Data for other asset classes are from FRED and run from 14 February 1986 to 30 September 2018.**

Asset	Days	Mean	Median	Maximum
Bitcoin	2935	5.03	3.73	35.53
Ether	1119	6.25	5.80	20.36
Litecoin	1950	5.78	4.56	33.09
NASDAQ100	8005	1.27	1.04	6.04
Netflix	4143	3.32	3.10	11.40
Gold	8005	0.79	0.71	3.24
USD-GBP spot rate	8005	0.50	0.47	1.70

cryptocurrencies: transfer of ownership of cryptocurrency between two parties is finalised only when the payment is included in a block and recorded on a ledger, known as a *blockchain*. Each block is created by a miner, who chooses which payments to include in it. Blocks have limited capacity and the average rate of block creation is exogenous and random. This means there is an uncertain waiting time before a transaction can be added to the blockchain. The initiator of the transaction can offer a fee to the miner that creates a block, in exchange for priority. Therefore, fees are a market mechanism that allocates blockchain space to users.

Competition between speculative and transactional users of blockchain space leads to a potential *crowding out effect*. Speculation raises fees and may price out users who would otherwise use the currency to purchase goods. Therefore, speculation hinders the ability of cryptocurrencies to function as a medium of exchange. The crowding out effect is well-documented. The 2017 surge in cryptocurrency trading made merchants more reluctant to accept cryptocurrencies as a means of payment (see, for example, BBC News 2017). Criminal activity has moved away from bitcoin towards other cryptocurrencies with more blockchain capacity (Coleman 2018). In January 2018, the North American Bitcoin Conference had to stop accepting bitcoin payments for tickets, because the process was too slow and costly (Choudhury 2018).

To the extent that the value of money is determined by its usefulness, the crowding out

effect reduces the price of cryptocurrency. This means that buy-side speculation may lower the price of cryptocurrency. That is the opposite of what we would expect with conventional assets. I call this surprising result the *pecuniary effect*.

In my model, there are a large number of potential users, called households. The value of cryptocurrency is an increasing function of the proportion of households that use it. Households consume over two periods, and prefer to consume early. They can choose whether to use cryptocurrency or cash to purchase consumption goods. Using cryptocurrency gives a household a non-pecuniary bonus. This is because cryptocurrency is embedded with a technology that makes it potentially superior to cash as a medium of payment. For example, cryptocurrency may be easy to use and secure relative to cash, or households may value the lack of reliance on authorities such as central or commercial banks.

Each household privately observes a noisy signal about the effectiveness of the payment technology, and uses it to inform her choice of payment method. The non-pecuniary payoff from using crypto is increasing in the number of other households who use it, so there are strategic complementarities between households' decisions. However, there is a trade-off. Cryptocurrency carries settlement risk, because a cryptocurrency payment may take a long time to be recorded on the blockchain, causing consumption to be delayed. This risk can be reduced by offering mining fees.

I use global games techniques to find equilibria. The solution is more complicated than standard global games problems, because there may not be global strategic complementarities in households' actions. Specifically, if one household uses crypto and pays a high mining fee, then another household may prefer not to use crypto at all. I show that the global games solution technique can be modified that there exists a unique threshold equilibrium. Households have a unique strategy, and each pays a fee increasing in her private signal.

I next consider the introduction of a financial market. Speculators receive signals about the effectiveness of the cryptocurrency technology. A speculator can place a buy order for cryptocurrency on an exchange. Once purchased, the cryptocurrency can be moved via the blockchain and traded for consumption goods. Both the households and the speculators therefore require blockchain space. They compete for priority by offering fees to miners. The cost of these fees makes the cryptocurrency less useful to the household as a means

of payment, inducing the crowding out effect.

The price is set by a market maker who only observes the trade order flow, as in Kyle (1985). Crowding out makes the cryptocurrency less useful as a means of payment. The market maker reduces her prior about the value of the cryptocurrency, and so sets a lower market price. This is the pecuniary effect. Since even side buy-side speculation can result in lower prices, the demand curve for cryptocurrency can be locally upward-sloping. Price formation for cryptocurrencies can therefore be very different than for standard assets. For example, one surprising consequence is that a speculator's profits can increase as other informed speculators enter the market and push the price down.

The pecuniary effect gives rise to price volatility. Crowding out lowers the market maker's ex ante belief about the cryptocurrency value and so, if she observes a positive order flow ex post, she raises the price by a relatively large amount. In other words, her price function is more sensitive to the order flow. High price volatility is therefore an intrinsic feature of any monetary asset that settles on a blockchain. A cryptocurrency with a lower blockchain capacity will tend to have higher volatility.

Two distinctive characteristics of cryptocurrencies generate these results. First, because a cryptocurrency is a monetary asset with no fundamental value, its value is governed by its usage as a means of payment. Second, monetary and speculative activities co-exist on the blockchain, causing uncertainty about the timeliness of settlement. Interaction between these two characteristics causes the pecuniary effect and makes prices more volatile.

My model predicts that price volatility is negatively related to blockchain capacity. I discuss how this prediction can be tested empirically, by exploiting a feature of bitcoin in which, by design, the blockchain capacity varies over time. I also predict that price volatility will fall and payments usage increase if, in the future, a greater volume of speculation could be carried out outside the blockchain. Recent developments such as the evolution of cash-settled derivatives markets[3] or the introduction of the Lightning Network could have profound consequences.[4] However, hedging and settlement ultimately

[3] There are two major exchanges for cash-settled bitcoin futures: CBOE and CME. At the time of writing, the total open interest in futures is around $100 million, compared to a total bitcoin market cap of over $100 billion.

[4] The Lightning Network is a secondary layer on top of the existing payments network to facilitate small payments without creating blockchain congestion. Counterparties can open a bilateral channel to make payments off-chain, and then settle net on-chain once the channel is closed. The code has recently been

have to occur on-chain, albeit netted, so blockchain congestion will continue to affect pricing and usage.

This paper is related to a wider literature on the role of cryptocurrencies as a monetary asset, some of which posit alternative explanations for the high level of volatility. Yermack (2013) and Saleh (2018) suggest that, because many cryptocurrencies have a fixed supply schedule, demand shocks cannot be absorbed by adjusting quantity and instead feed through to price. Garratt and Wallace (2018) show that coordination problems among users may mean multiple equilibria exist. Biais et al. (2018) establish a similar result in a dynamic model. In Pagnotta and Buraschi (2018), volatility is due to the sensitivity of the equilibrium outcome to initial conditions. Price volatility could potentially be a result of uncertainty about which equilibrium will emerge, but then what causes beliefs to shift is not clear. My model solves this coordination problem by using global games and showing that there is a unique threshold equilibrium.

Schilling and Uhlig (2018), in a model related to Garratt and Wallace (2018), show that a speculative bubble is a possible equilibrium: if agents expect the price of the cryptocurrency to rise, they will hoard it. Another potential explanation for volatility is market manipulation: Gandal et al. (2018) show that elevated bitcoin price volatility in 2013 may have been due to suspicious trading activity on the Mt. Gox exchange. But these arguments cannot explain why cryptocurrencies are so much more volatile than other asset classes. I argue that volatility is due, at least in part, to the unique characteristics of cryptocurrencies.

This paper relates to a newly developing literature on mining and fees in a blockchain environment. The most closely related papers are Huberman, Leshno, and Moalleni (2017) and Easley, O'Hara, and Basu (2017), which both use queueing theory to assess the effect of blockchain congestion on fees and waiting times. As in my model, agents in both of these papers value having shorter waiting times to process payments. In contrast to my model, neither of these papers examine the trading environment, nor do they incorporate endogenous price formation.

Ciaian, Rajcaniova, and Kancs (2016) use variants of Fisher's equation of exchange (Fisher 1911) to propose supply and demand functions, and so determine a unique ex-

released for the Bitcoin network. For more detail, see the white paper by Poon and Dryja (2016).

change rate between cryptocurrency and fiat money. Athey et al. (2016) obtain a price for cryptocurrency in a model where users are uncertain about whether the technology will fail. Bolt and van Oordt (2016) introduce a speculative element, so that the equilibrium price is affected by speculators' desire to maximise their portfolio values. Fernández-Villaverde and Sanches (2016) study a model in which private entrepreneurs can issue currencies that compete with one another. However, these approaches do not incorporate the technological factors that distinguish cryptocurrencies from other monetary instruments, and so cannot explain why price volatility is higher than for other classes of monetary asset.

A recent paper by Hautsch, Scheuch, and Voigt (2018) studies the limits to arbitrage that arise due to settlement latency, which results from the limited capacity of the blockchain. They use trading data to estimate arbitrage boundaries. In their model, neither prices nor fees are micro-founded. Changes in the price are assumed to follow a Brownian motion. The reduction in latency from a given mining fee is assumed to be fixed, rather than endogenised as in my model.

The paper is arranged as follows. Section 2.2 provides an overview of the key features of cryptocurrency that are useful for understanding the model, including blockchain, the trading environment, and the use of cryptocurrency as money. Section 2.3 introduces a model in which a cryptocurrency is only used to make purchases. Section 2.4 adds a financial market to this model and examines price formation. Section 2.5 shows that the results are robust to the relaxation of many of the key modelling assumptions. Section 2.6 discusses empirical implications of the findings, and Section 2.7 concludes. All proofs are in Appendix B.

2.2 An overview of key blockchain concepts

This Section outlines the technical features of cryptocurrency that are key to my model. In particular, I discuss the limited capacity of the blockchain, the trading environment, and the utility of cryptocurrency as money.

At the time of writing, the website `coinmarketcap.com` lists nearly 1,000 different cryp-

tocurrencies. A description of each of them is beyond the scope of this paper, so I focus here on Bitcoin — the first and largest cryptocurrency — and its offshoots, such as Bitcoin Cash and Litecoin.[5] Together, these three currencies account for more than half of the total cryptocurrency market.[6] With other cryptocurrencies, particularly Ethereum, some details may differ, but all have similar design principles.[7]

2.2.1 Cryptocurrencies and blockchain

A cryptocurrency is a digital asset with no physical form. Ownership is recorded on a decentralised ledger called a *blockchain*, which is maintained by a global network of computer nodes, known as *miners*. The ledger is public, so anybody with a computer and internet access can download the software, become a miner, and join the network. Owners of cryptocurrency store it using software called a *wallet*. A private key is needed to use the contents of a wallet. Only the owner of a wallet has the key and, without it, any attempt to make a transaction will be rejected by miners.

When one person, 'Alice', wants to send currency to another person, 'Bob', she notifies the network.[8] Her payment is then put into the *mempool*, which is a set of pending payments that have not yet been added to the ledger. The transfer of ownership of a coin is finalised only when the corresponding payment is removed from the mempool and added to the blockchain. For this to happen, a miner must include it in a *block*. Every so often, a miner is selected to put transactions from the mempool into a block. A block has finite capacity and can contain only a limited number of payments, so some of the pending payments may remain in the mempool. For Bitcoin, one block contains approximately 1 megabyte of data, which is roughly 4,000 transactions. Because the number of blocks per

[5] In this paper, I follow convention by writing "Bitcoin" to denote the network, and "bitcoin" to denote the currency unit, and similarly for other cryptocurrencies.

[6] As measured by market capitalisation (circulating supply times price) on `coinmarketcap.com` on 30 October 2018. Total market cap was $188bn, of which $110bn was due to Bitcoin, $7bn due to Bitcoin Cash, and $3bn due to Litecoin.

[7] My model examines the usage and trading of a cryptocurrency with its own blockchain, so this Section does not discuss tokens that used a shared platform. For example, most Initial Coin Offering (ICO) tokens use a shared blockchain, typically Ethereum.

[8] If Alice is a node, she can do this directly. Otherwise, she will notify her wallet provider, which is a node that intermediates her access to the network.

unit time is random, total block capacity is uncertain. Appendix A discusses the amount of space used by payments.

Suppose that, among all potential miners, Minnie is selected to mine the next block. Alice can incentivise Minnie to include her payment by attaching a fee to it, denominated in cryptocurrency. If Minnie includes Alice's payment in his block, she receives the fee. The fee system therefore uses a market mechanism to assign priority: those users who are most willing to pay for urgency are settled first.

Once a payment has been submitted to the mempool, the fee cannot be recovered. A user can effectively cancel a transaction in the mempool by submitting another transaction that sends the same coins to herself, but with a higher fee attached. Miners will then prioritise the second transaction. Once it is added to the blockchain, the first transaction becomes invalid because the same coins cannot be spent twice. But Alice cannot recover the mining fee. The fee mechanism can therefore be likened to an all-pay auction, in which the number of units auctioned is random and unknown to the participants, and participation is voluntary.

In addition to Alice's fee, Minnie may receive a certain number of newly minted units of the currency, called the *coinbase reward*. For example, a Bitcoin miner currently receives 12.5 new bitcoins per block mined, in addition to the fees attached to any payments taken from the mempool. The total supply of Bitcoin is fixed at around 21 million, with the last coin expected to be mined around the year 2140. The coinbase reward decreases roughly exponentially as more blocks are created, until all coins have been issued. Thus, as the system matures, coinbase rewards become a less important part of the incentive system for miners, and the fees become correspondingly more important.[9] Figure 2.1 shows how Bitcoin fees and payment confirmation times have evolved since December 2011.[10] To some extent, the recent increase in fees shown in the figure is by design. As the coinbase reward falls, miners demand higher fees from users to compensate them for their mining costs.

[9] Nakamoto (2008), in the original paper that set out the concept of Bitcoin, explains that this predictable and limited supply schedule is designed to eliminate the risk of currency debasement when a monetary authority controls supply. Easley, O'Hara, and Basu (2017) show that fees must necessarily become a more important part of the system as it develops.

[10] Prior to December 2011, confirmation times tended to be negligible.

Figure 2.1: Bitcoin median waiting times in minutes (LHS) and mean fee per transaction, in USD (RHS), 3 December 2011 – 30 September 2018. Data are from `blockchain.info` and are plotted daily over a 30-day backward-looking moving average.

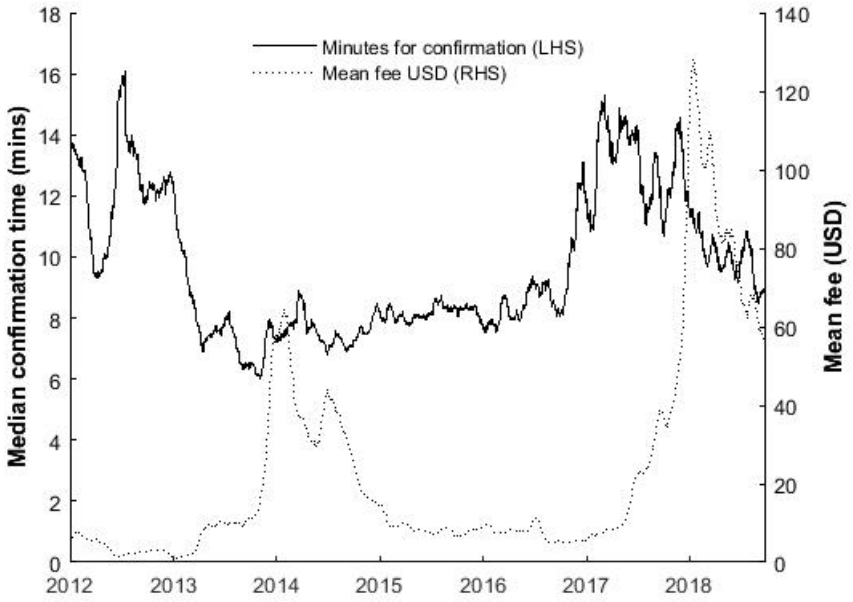

Cryptocurrencies are usually designed so that block creation occurs at a fairly constant rate over time. For example, Bitcoin blocks are created according to a Poisson process, with a block created on average every 10 minutes. Bitcoin employs a 'proof-of-work' algorithm, meaning that miners compete for the right to create the next block by solving a complex computational problem. This problem is different for each proposed block and can only be solved by a 'guess and verify' method, which requires a large number of computations.[11] Other cryptocurrencies use different algorithms to select a miner to create the next block. For example, Ethereum is moving to a 'proof-of-stake' protocol. In this paper, the details of the algorithm are not important. All that matters is that priority in the blockchain is assigned according to fees submitted by users.[12]

[11]The difficulty of the problem is regularly adjusted, depending on total computational power among the miners, in order to keep the mean block creation rate constant. See Section 2.6.3 for more details.

[12]For more details about the proof-of-work mining process, see Harvey (2016). For an analysis of the proof-of-stake algorithm, see Saleh (2017). Although not relevant to this paper, the process for selecting miners is an important design question, and an emerging field called 'cryptoeconomics' is looking at the

Limits on blockchain capacity are the result of a conscious design choice, which trades off security against payment efficiency. Each time a block is added to the blockchain, all nodes must update their own local version of it. If they do not have time to do this before the next block is mined, then there would be disagreement about the true state of the blockchain (Kirkby 2018). A larger block size would also increase the requirements for mining, raising barriers to entry and increasing the risk that one miner could control the blockchain.[13]

Initially, Bitcoin had no block size limit. It was introduced by Bitcoin's creator Satoshi Nakamoto in July 2010 without explanation, most likely because he feared the potential for hostile attack. No technical reason exists why the parameters could not be adjusted to allow a greater capacity for Bitcoin; indeed such an adjustment has been attempted on several occasions, but so far no proposal has been able to achieve the required consensus among the developer community.[14]

2.2.2 Exchanges and the trading environment

Trading in cryptocurrencies typically takes place on exchanges. At the time of writing, at least 200 cryptocurrency exchanges exist (see the website `coinmarketcap.com`). Exchanges differ in several respects such as liquidity, the jurisdictions in which they operate, the types of order permitted, and the trading pairs offered (e.g., trading pairs of cryptocurrencies against each other, or against fiat currency). Typically, an exchange operates double auctions with bids and asks, and charges a commission on trades (see, e.g. Böhme et al. 2015).

Transactions on an exchange are generally settled off-chain; that is, they are settled across the speculators' accounts at the exchange and are not recorded on the blockchain. The exchange simply debits the paying member's account and credits the receiver, without actually moving any cryptocurrency. This process is analogous to settling a payment in

incentives that different designs create. See, for example, Biais et al. (2017).

[13] See Clifford (2017). In a recent paper, Budish (2018) examines miners' trade-off between using their computing power to maintain the blockchain and to attack it.

[14] See Section 2.6.2 for more detail. Morgan (2017) describes the history of attempts to increase Bitcoin blockchain capacity.

commercial bank money rather than central bank money.[15] Off-chain settlement improves delivery times and eliminates fees, at the cost of delaying finality and creating counterparty credit risk between the receiving member and the exchange. This risk is far from trivial. For example, the exchange Mt. Gox failed in 2014 and members have had to wait several years to be fully compensated. Moore and Christin (2013) document a high rate of failure among cryptocurrency exchanges. When a speculator wishes to move his cryptocurrency from the exchange to a wallet, he must do so via the blockchain.

No cryptocurrency exchange allows naked shorts. All require immediate settlement; that is, a seller must have the asset on account at the exchange.[16] It is possible to trade cryptocurrency futures, which provide a speculator with exposure to downside. Until recently, all futures exchanges required physical settlement, so delivery was ultimately still required. However, in December 2017, the Chicago Metal Exchange and the Chicago Board Options Exchange both launched cash-settled bitcoin futures. These contracts do not require physical settlement of bitcoin, so speculation can occur entirely off-chain. As these markets develop, they may allow speculators to adopt speculative positions without being affected by the issues around blockchain capacity that are the focus of this model.[17] However, at the time of writing, the total open interest in cash-settled futures on these exchanges was a small fraction of the size of the cryptocurrency market (see footnote 3).

This paper studies frictions caused by the limited capacity of the blockchain, and assumes the trading environment is efficient. However, evidence shows that cryptocurrency exchanges are beset with frictions such as low liquidity, insider trading and security problems.[18]

[15] Carter and Hasufly (2018) provide a general explanation of how transactions can be batched, how batching saves on blockchain space, and how exchanges have been making increasing use of off-chain settlement.

[16] Betting on downward movements is possible, but it is expensive and risky. See, for example, Wilmoth (2017) and Roberts (2017). I simplify by assuming it is impossible.

[17] Hale et al. (2018) suggest the introduction of futures markets may have led to a downturn in the price of bitcoin.

[18] See, for example, Gandal et al. (2018) and Griffin and Shams (2018). For a recent overview of problems in the trading landscape, see Stecklow et al. (2017).

2.2.3 Cryptocurrency as money

Cryptocurrency has no intrinsic or fundamental value and does not generate a cash flow. The long-term value of cryptocurrency, like other monetary assets without a fundamental value such as token currency, depends on the extent to which it is used as a medium of payment.[19] Cryptocurrency has value to the extent that the technology embedded in it makes it superior to other forms of money as a means of payment. A number of features of cryptocurrency make it a potentially superior form of money:

- no dependence on central authorities (Nakamoto 2009);
- improved anonymity and privacy when carrying out large transactions, especially online (Bech and Garratt 2017; Kahn 2018; Koning 2018);
- superior convenience yield, due to ease of usage and storage (Cochrane 2018);
- immutable record of historic transactions (Luther and Olson 2013);
- markets are more complete as money issuer cannot prohibit access to network (Pagnotta and Buraschi 2018);
- no need to pay fees to financial intermediaries, especially when making international remittances (Athey et al. 2016);
- ability to participate in crypto economy; for example, writing smart contracts, investing in ICOs, and trading on certain cryptocurrency exchanges that do not accept fiat currency (e.g. Sockin and Xiong 2018).

Popper (2015) documents the particular importance of the first two of these factors in the development of cryptocurrency. The early pioneers were motivated by the desire to create a monetary asset that can be used online but does not rely on a central authority, such as a central bank or commercial bank. Such an asset allows for transactions to be made over the internet anonymously and privately, just as cash allows in the offline world. The demand for monetary privacy is often associated with criminal or illicit activity, but legitimate users may require it too (see Kahn 2018 and Koning 2018). For example,

[19] Assuming that speculative bubbles cannot be maintained for an indefinite period, the value of a monetary asset without fundamental value must ultimately depend on how much it is used as a medium of payment.

privacy can protect users from fraud and spam mail.

My model requires that the cryptocurrency technology makes it superior to cash as a means of payment, so that households derive some benefit from paying with it. Any of the features mentioned above could deliver this benefit: the model is agnostic about which particular feature the households find valuable.

2.3 Model with no financial market

I present a model in which the price of a cryptocurrency is endogenously determined by the extent that it is used as a means of payment. There are a large number of technologically adept *households*, who can choose to buy a consumption good using either cryptocurrency, which I call 'crypto' for brevity, or a conventional currency, which I refer to as 'cash'.

Crypto has technological features that make it potentially superior to cash as a means of payment. The technology in my model can be readily interpreted as delivering any of the other potential benefits listed in Section 2.2.3. Each household privately observes a noisy signal about the strength of the crypto technology. A household's payoff from using crypto depends on the strength of the technology, as well as the number of other households that use crypto. However, crypto payments take an uncertain length of time to be settled, leading to the possibility that consumption may be delayed. Each household must trade off the potential technological benefits of crypto against the uncertain settlement time.

2.3.1 Model set-up

I consider a one-period, three-date model. All decisions are taken at a time T_0. Consumption occurs at times T_1 and T_2, after which the game ends. At time T_0, a unit mass of risk-neutral *households*, indexed by $i \in [0,1]$, is born. There are three assets in the economy: crypto, cash, and a consumption good, which acts as the numéraire. Crypto and cash are monetary assets that cannot be consumed but can be exchanged for consumption

goods at T_1 or T_2.

Each household seeks to maximise her utility from consumption over the two periods, T_1 and T_2. Households discount over time, so that consumption at the later period T_2 is subject to a discount factor $\rho \in (0,1)$. Households prefer to stagger consumption over time in the following way: each aims to consume no more than one unit of consumption good at T_1, while any amount can be purchased at T_2. Given consumption of $c_1, c_2 \geq 0$ units of consumption good at times T_1, T_2 respectively, a household's total utility from consumption is:

$$u_H(c_1, c_2) \;=\; \min\{c_1, 1\} + \rho c_2. \tag{2.1}$$

Households have access to an unlimited budget of crypto and cash; for simplicity, this can be thought of as a limitless credit line that is repaid at T_2.

At T_0, each household can choose whether to use crypto or cash to purchase a unit of consumption good for delivery at T_1. Her decision is denoted $y_i \in \{0,1\}$, where $y_i = 0$ if she pays in cash, and $y_i = 1$ if she pays in crypto.[20] Define y to be the total mass of households that use crypto:

$$y \;=\; \int_0^1 y_i \; \mathrm{d}i. \tag{2.2}$$

I make the technical assumption that the set $\{y_i\}$ is Lebesgue-measurable, so that y is well-defined. Total crypto usage y is publicly revealed at T_1.

Definition 2.1 (Crypto payments technology)**.** *Crypto embeds a technology that makes it potentially superior to cash as a means of payment. The strength of this technology is a random variable $R \sim U[0,1]$ that is drawn by nature at the beginning of the game. R is publicly revealed at T_1. A household that uses crypto earns a non-monetary payoff $Rg(y)$, where $g(y)$ is strictly increasing in y, and $g(1) = 1$.*

Using crypto becomes more attractive as the number of other households who use crypto increases. I refer to the function $g(y)$ as the *network effects*. Higher values of $g(y)$ mean stronger strategic complementarities between households.

The strength of the technology R reflects the technological superiority of crypto relative to cash as a means of payment, as described in Section 2.2.3. The households do not

[20] In Section 2.5.1.2, I consider equilibria in mixed strategies.

directly observe R at T_0. Each household i receives a signal R_i, where $R_i - R \sim U[-\sigma, \sigma]$, independently and identically distributed. The parameter $\sigma > 0$ is a measure of the error of the signal. A household's signal R_i tells her not only about the technology R, but also about the signals of the other households.

I interpret this signal structure as capturing disagreement among households about the interpretation of a public signal. Public information about the technology is revealed to all agents at T_0. However, this information is difficult to interpret, so households draw different conclusions about R and hence the signals $\{R_i\}$ are correlated but not identical. The information may, for example, be computer code that is public domain but requires a high degree of technological knowledge to interpret correctly.

Crypto has no underlying cash flows or intrinsic value. Its value at T_1 and T_2 is $v(y)$; that is, it can be used to purchase consumption goods at a rate of $v(y)$ units of crypto per unit of consumption. The value $v(y)$ is endogenously determined by the number of households y who use it at T_1 to purchase the consumption good.[21] The value becomes publicly known at T_1, when y is revealed. The value is strictly increasing in y, with $v(0) = 0$. Let $V = v(1)$. The value of cash is fixed at 1 by a central bank.

Definition 2.2 (Ledger)**.** *A ledger is a record of all transactions settled in a given medium. There are two ledgers, one associated with crypto and one with cash. A transaction must be recorded on the relevant ledger before a consumption good can be delivered to a household. The crypto ledger is called the 'blockchain'.*

The cash ledger is updated instantaneously, so any good purchased at T_1 can be consumed immediately.[22] However, the crypto ledger — that is, the blockchain — has a finite and random capacity at T_1, denoted by N, so that only a measure N of crypto payments can be settled by T_1.

[21] My model is agnostic about the exact nature of that relationship, as long as it is positive. For example, Fisher's equation of exchange implies a linear relationship (Fisher 1911). Network models would tend to predict a convex relationship (Cong, Li, and Wang 2018). The transactions demand for money theory predicts a square-root relationship between money usage and demand (Baumol 1952,Tobin 1956).

[22] For example, the ledger may be maintained by a monetary authority or commercial bank who can verify transactions and update the ledger without delay. Alternatively, if cash is physical tokens (e.g., notes and coins, or commodity money), then the ledger is the physical allocation of tokens across agents, and delivery occurs as soon as the tokens change hands.

Definition 2.3 (Mempool)**.** *The mempool is the set of all crypto transactions that are attempted at T_0.*

If the mempool has Lebesgue measure greater than the blockchain capacity N, then some crypto transactions will have to wait until T_2, so some households' consumption is delayed. At T_2, all remaining transactions can be added to the blockchain. If a household uses crypto and learns at T_1 that her payment will not be delivered, there is not enough time to make a cash payment and have it delivered by T_1. Therefore, once a household i chooses her payment method y_i at T_0, she is committed and cannot change her mind.

Blockchain capacity N is not public information at T_0, so a household faces greater uncertainty about consumption timing when she uses crypto than when she uses cash. She trades this uncertainty off against the potential benefit from the superior technology. The realisation of N does not directly affect the long-term value of crypto $v(y)$, which depends on the total number who use crypto, not on when consumption occurs.

Definition 2.4 (Block rate)**.** *The block rate is a parameter $\lambda \in \mathbb{R}$. The blockchain capacity N available during period T_1 has cumulative distribution function $Z_\lambda(n)$ with full support on $[0,\infty)$, indexed by λ, and has the following properties:*

1. ***First-order stochastic dominance in λ:*** *Given any pair $\lambda_1 < \lambda_2$ then, for all n, $Z_{\lambda_1}(n) > Z_{\lambda_2}(n)$.*
2. ***Monotonicity and continuity:*** *For all λ, $Z_\lambda(n)$ is continuous and strictly increasing in n.*

Blockchain capacity N is realised at T_1.

The block rate λ is the key parameter of interest in my model. The stochastic dominance condition ensures that high values of λ tend to imply high blockchain capacity N. The monotonicity and continuity properties ensure that agents benefit even from slight improvements in their priority. An example of a distribution function that satisfies the definition of $Z_\lambda(n)$ is a normal distribution with mean increasing in λ, with negative realisations mapping to zero. Appendix A discusses the empirical distribution of blockchain space for Bitcoin.

Assumption 2.1 (Upper dominance)**.** *For all* $t \in [0,1]$,

$$(1-\rho)Z_\lambda(t) \;<\; g(t). \tag{2.3}$$

Assumption 2.1 places a lower bound on the network effects $g(\cdot)$ relative to the distribution of blockchain capacity $Z_\lambda(\cdot)$. This assumption simplifies the analysis by ruling out equilibria with no crypto usage. In Section 2.5.1.1, I show that relaxing Assumption 2.1 does not affect the key results in the paper, but merely introduces an additional trivial equilibrium where crypto is never used.

A household i that uses crypto can hedge against the risk of a low realisation of blockchain capacity N by offering a mining fee f_i when she submits her payment. Miners will then prioritise household i's payment over any payment with a fee lower than f_i.

Definition 2.5 (Mining fees)**.** *A fee* $f_i \geq 0$ *is an additional amount of crypto attached by a household* i *to a crypto payment at the time it is submitted to the mempool. The fee must be paid whether or not the transaction is added to the blockchain at* T_1.

Miners always select the N mempool transactions that have the largest fees attached. If two or more payments have equal fees, they are prioritised at random. Fees are denoted in units of crypto and must be paid in crypto. Because agents are assumed to have access to unlimited supplies of crypto, there is no budget constraint. Unconfirmed transactions in the mempool cannot be cancelled, so the fee is forfeit regardless of the outcome.[23]

The events over the three time periods are as follows:

1. At time T_0, the strength of the technology $R \sim U[0,1]$ is determined by nature. Each household $i \in [0,1]$ privately observes a signal R_i, where $R_i - R \sim U[-\sigma,\sigma]$ iid. She chooses a payment medium $y_i \in \{0,1\}$ and, if $y_i = 1$, a fee $f_i \geq 0$.

2. At time T_1, R and y are revealed and the long-term value of crypto $v(y)$ is determined. The amount of block space $N \sim Z_\lambda(n)$ is determined by nature. The N

[23] As explained in Section 2.2.1, a household can effectively cancel her payment but cannot recover the fee. Doing so cannot benefit a household in our model. She would still have to pay the fee, and would still only be able to consume at T_2. I make this modelling choice for the sake of realism, but it is not critical for obtaining my results. A model with an alternative technology that allows fees to be cancellable would have a payoff with different functional forms, but the results would not be qualitatively different.

crypto payments that have the highest fees attached are added to the ledger, as are all cash payments. Any household that had her payment added to the ledger can consume at T_1.

3. At time T_2, all remaining crypto payments are added to the ledger. Consumption occurs and the game ends.

The total expected payoff for household i if she uses crypto, given a signal R_i and a fee f_i, is:

$$\mathbb{E}\Big[Rg(y) + \mathbb{1}\{\text{delivery}\} + \rho\mathbb{1}\{\text{non-delivery}\} - \rho v(y) f_i - \rho \;\Big|\; R_i, f_i\Big]. \tag{2.4}$$

The first term in Equation (2.4) is household i's expected non-pecuniary benefit from using crypto. The second term is the payoff when the payment settles at T_1, in which case the household can consume early. The third term is the payoff when the payment does not settle at T_1, and the household consumes late. The fourth term is the cost of the fee, in terms of foregone consumption at T_2. A unit of crypto paid as a fee to a miner means the household must forego $v(y)$ units of consumption at T_2.

The final term $(-\rho)$ is the cost of making the payment. Because y is public knowledge at T_1, the value $v(y)$ of crypto is publicly known. The household agrees with the seller of the consumption good an exchange rate of $1/v(y)$ with crypto. For each unit of consumption good, the household pays $1/v(y)$ units of crypto, which is worth 1. Her opportunity cost in terms of foregone consumption is thus ρ.[24]

The payoff from using cash is simply $1-\rho$. The household consumes 1 unit at T_1, since cash payments are added to the ledger immediately. She foregoes ρ units of consumption at T_2.

I write $\pi_H(f_i; R_i)$ for difference between household i's total expected payoff from using crypto given in Equation (2.4), and that from using cash, given her signal R_i and her fee

[24] Strictly speaking, this argument is not valid when $v(y) = 0$, so a problem could arise in the event that some households use crypto and $y = 0$. I show that this event can never happen. I find that households employ a switching strategy with threshold $\hat{R}$, so that household i uses crypto if and only if her signal $R_i > \hat{R}$. Then $y = 0$ implies $R \leq \hat{R} - \sigma$, so no household uses crypto.

f_i. Then:

$$\pi_H(f_i;R_i) \;=\; \mathbb{E}\Big[Rg(y)\Big|R_i\Big] \;-\; (1-\rho)\mathbb{P}\Big[\text{non-delivery}\Big|R_i,f_i\Big] \;-\; \rho f_i \mathbb{E}\Big[v(y)\Big|R_i\Big]. \tag{2.5}$$

The first term in Equation (2.5) is the household's expected non-pecuniary benefit from using crypto. The second term is the expected opportunity cost of delayed consumption. The third term in Equation (2.5) is the opportunity cost of paying a fee, in terms of foregone consumption. Note that her choice of f_i affects the probability of delivery, but does not affect y, which depends on other households' signals about the strength of the technology R.

I assume that, if a household earns the same expected payoff from choosing crypto and cash (i.e., $\pi_H = 0$), she will choose cash. Thus, she uses crypto if $\max_{f_i \geq 0} \pi_H(f_i;R_i)$ is strictly greater than zero. Otherwise, she uses cash.

Definition 2.6 (Equilibrium without financial market)**.** *An equilibrium is a pair of decision functions* $\eta : [-\sigma, 1+\sigma] \to \{0,1\}$ *and* $\phi : [-\sigma, 1+\sigma] \to [0,\infty)$*, such that when a household with signal* R_i *believes that all other households play* $y_j(R_j) = \eta(R_j)$ *and* $f_j(R_j) = \phi(R_j)$*, her optimal strategy is to play* $y_i = \eta(R_i)$ *and* $f_i = \phi(R_i)$*.*

The solution concept is Bayesian Nash equilibrium. Because all households are ex ante identical, attention is restricted to symmetric equilibria.

2.3.2 Global games solution

I use global games techniques to obtain an equilibrium. These techniques were pioneered in the finance literature by Morris and Shin (1998) in the study of speculative currency attacks, and by Rochet and Vives (2004) and Goldstein and Pauzner (2005) in the study of bank runs. In these models, agents receive heterogeneous signals about an unknown fundamental value (e.g., the strength of a currency peg, or the value of a bank's assets). Each agent's signal not only provides an estimate of the fundamental value, but also an estimate of the other agents' signals. This information is useful to agents because of strategic complementarities in their payoffs: the benefit to agents from taking a particular action is increasing in the number of other agents who take that same action.

The main result of these models is that, under certain conditions, there is a unique equilibrium. This is a threshold equilibrium; i.e. an equilibrium where agents' strategies are determined by whether their signals exceed a particular threshold. Morris and Shin (2003) provide a clear description of the intuition behind these results. Following the literature, I focus on threshold equilibria.

Definition 2.7 (Threshold equilibrium)**.** *A threshold equilibrium is an equilibrium in which there exists a threshold $\hat{R}$ such that each household i uses crypto if and only if her signal $R_i > \hat{R}$. Each household's fee strategy is a continuously differentiable function of R_i. Both the choice of money and fee are in pure strategies.*

Special cases of threshold equilibria include those where crypto is never used (i.e., $\hat{R} \geq 1$) and where it is always used ($\hat{R} \leq 0$).

In my model, crypto has no fundamental. Instead, each agent receives a signal about the strength of the crypto technology R. A household's signal gives her information not only about the payoff from using crypto, but also about the number of other households y who choose to use crypto. Given that the household's payoff depends on both R and y, it is rational for a household to condition her actions on this signal.

Standard global games models do not have an equivalent of fees: agents in the papers cited above choose from a one-dimensional action space. In my model, agents choose whether to use crypto or cash and, if they choose crypto, a fee. This additional choice makes my model more complex, since any equilibrium must describe households' choices of both payment method and fees. My solution represents a new contribution to the global games literature. In the household's payoff function π_H, there is no direct interaction between her signal R_i and the fee f_i. A household's fee affects her payoff only to the extent that it determines her priority relative to the other households. Therefore, a household will pay a high fee only if she believes others around her do too. Her signal R_i tells her about other households' signals. Her optimal fee depends on her beliefs about how other households set their fees, and not on the actual value of R. This means that fee strategies depend on beliefs about beliefs, and so may not be well-anchored in R.

There are strategic complementarities in households' choices of money, and in their choices of fee. But one household's choice of money is a strategic substitute for an-

other's choice of fees. If household i believes other households pay high fees, she may prefer to use cash. The existence of these strategic substitutes means that a threshold equilibrium may not exist. To understand why, suppose that there is a threshold equilibrium, so that there exists an $\hat{R}$ such that households use crypto if and only if $R_i > \hat{R}$. Choose a $\varepsilon > 2\sigma$. As a household's signal increases from $\hat{R} - \varepsilon$ to $\hat{R} + \varepsilon$, her beliefs about other households' actions change from $y = 0$ to $y = 1$, so her non-pecuniary payoff from using crypto increases by approximately $\hat{R}(1 - g(0))$. However, competition for blockchain space rises too: the expected cost of settlement delay increases from $(1-\rho)Z_\lambda(0)$, up to $(1-\rho)Z_\lambda(1)$; it may be less, if it is optimal to offer a non-zero fee. Therefore, a switching strategy at $\hat{R}$ can only be optimal if the change in the non-pecuniary payoff $g(\cdot)$ is large relative to the change in waiting cost $Z_\lambda(\cdot)$. I impose a restriction that ensures this is the case.

Assumption 2.2 (Existence of threshold equilibrium with no financial market)**.**

$$Z_\lambda(1) \;\leq\; \frac{\int_0^1 Z_\lambda(t)dt}{\int_0^1 g(t)dt} \;\leq\; \frac{Z_\lambda(0)}{g(0)}. \tag{2.6}$$

Assumption 2.2 guarantees that households use crypto just above the threshold, and use cash just below it. It ensures $g(0)$ is small enough relative to $Z_\lambda(0)$ so that, given a signal just below the threshold, the payoff for a household from using crypto is low relative to the risk of not being settled at T_1. Similarly, Assumption 2.2 ensures $g(1) = 1$ is high relative to $Z_\lambda(1)$.

Standard global games models such as Morris and Shin (1998) and Goldstein and Pauzner (2005) have no equivalent to Assumption 2.2, because they have only strategic complementarities among agents, and no strategic substitutes.[25] The assumption is somewhat restrictive: in particular, it requires that $Z_\lambda(0) > 0$, so there is a non-zero probability that there is insufficient blockchain capacity to settle any payments. This is realistic: roughly, blocks arrive according to a Poisson process, so it is possible that no blocks arrive in any given interval of time (see Appendix A). In Section 2.5.1.2, I consider the case in which

[25] Goldstein and Pauzner (2005) in fact have one-sided strategic complementarities, so that strategic complementarities exist only when the fundamentals of the system are above a threshold level. They show that there is still a unique threshold equilibrium in their setup. This finding is not relevant to my model, where strategic substitutes potentially exist everywhere.

Assumption 2.2 does not hold, and show that, if households are allowed to use mixed strategies, then a more general class of equilibrium can exist without this Assumption.

Proposition 2.1 (Unique threshold equilibrium among households)**.** *Suppose that the households' signal accuracy* σ *is small enough to satisfy the following expression:*

$$0 \;<\; \sigma \;<\; \frac{1}{2}\min\Big\{(1-\rho)Z_\lambda(0),\;\; 1-(1-\rho)\max_{y\in[0,1]}\frac{Z_\lambda(y)}{g(y)}\Big\}. \tag{2.7}$$

Then, if Assumption 2.2 holds, there is a threshold equilibrium, in which a household i uses cash if her signal $R_i \leq \hat{R}$*, and uses crypto if* $R_i > \hat{R}$*, where:*

$$\hat{R}\int_0^1 g(t)dt \;\;=\;\; (1-\rho)\int_0^1 Z_\lambda(t)dt. \tag{2.8}$$

This expression is decreasing in λ*. Households employ a deterministic fee function* $f_i = \phi(R_i)$*, where* $\phi(\hat{R}) = 0$ *and* $\phi(R_i)$ *is strictly increasing. If Assumption 2.1 holds, this equilibrium is unique.*

Proof. See Section 2.9.1. □

Proposition 2.1 establishes there exists a unique threshold equilibrium, and each household's optimal fee is increasing in her signal R_i. Assumption 2.2 ensures existence of this equilibrium, and Assumption 2.1 ensures uniqueness.

In the language of Carlsson and Van Damme (1993), we require lower and upper dominance regions to obtain a threshold equilibrium. These are regions where a particular actions dominates regardless of a household's beliefs. There is always a lower dominance region: regardless of beliefs about other households' actions, a household with a low enough signal always prefers to use cash. This anchors households' beliefs. A household with a signal slightly above this region not only believes that the technology is weak, but also believes that almost half of the other households are in the lower dominance region and do not use crypto. She infers the non-pecuniary payoff will be low, and so plays $y_i = 0$. Assumption 2.1 guarantees the existence of an upper dominance region, so that a household with a high enough signal always prefers to use crypto. Therefore, a household

with a signal slightly below the upper dominance region is also likely to prefer crypto.[26] Applying this argument iteratively implies a unique threshold $\hat{R}$ where the household's strategy switches. This is the standard intuition behind global games techniques.

The same intuition can be adapted to explain why households' fees are increasing in their signals in Proposition 2.1. Consider a household who observes a signal just above the threshold $\hat{R}$. She knows it is likely that some of her peers will have a signal below $\hat{R}$ and will use cash, so they do not take up blockchain space. Therefore, she expects competition for blockchain space to be low, and offers a low fee. As R_i increases, the household believes that more of her peers will observe signals about $\hat{R}$, and so there will be more competition for blockchain space. Therefore, she raises her fee. A household with signal $\hat{R}+2\sigma < R_i < \hat{R}+4\sigma$ knows that all other households will use crypto, but does not know whether they know that. Therefore, she believes her peers might not pay high fees, and so sets a moderate fee herself. As R_i increases, the household becomes firmer in her higher order beliefs that others pay fees, and believes these fees are higher. Therefore, she raises her own fee. This argument implies that each household i sets a fee increasing in her signal R_i.

The expression for $\hat{R}$ is derived by considering a marginal household with a signal exactly equal to the threshold. This household adopts a belief that the proportion of households with signals above $\hat{R}$ is uniformly distributed on the unit interval. Morris and Shin (2003) call this uniform prior a 'Laplacian belief'. It is a standard feature of global games models, even when distributions are not uniform. Then Equation (2.8) follows from the fact that the marginal household offers a fee of zero. This logic applies regardless of the value of σ; therefore, as long as the condition in Equation (2.7) is satisfied, the threshold $\hat{R}$ does not depend on σ.

Note that the same expression for $\hat{R}$ would be obtained under any fee mechanism which gives rise to a threshold equilibrium. For example, $\hat{R}$ would be the same if fees are refunded when delivery is not fulfilled by T_1. It would also be the same if we were to replace fees with a lottery, so that available blockchain space is simply allocated at random. Therefore, the design of the blockchain space allocation mechanism may not have much effect on crypto usage, although it will obviously affect total expenditure on fees.

[26] If Assumption 2.1 did not hold, then an additional equilibrium would exist in which all households set $y_i = 0$ regardless of their signals. See Section 2.5.1.1.

The comparative statics of the threshold $\hat{R}$ with respect to each of the parameters are intuitive. The threshold $\hat{R}$ is decreasing in the block rate λ. Higher blockchain capacity makes crypto more appealing as an alternative to cash, and so the threshold required for usage falls. The threshold is also decreasing in $\int_0^1 g(t)\mathrm{d}t$, which reflects the network effects of using crypto. As strategic complementarities between households strengthen, households use crypto more. The threshold is also decreasing in ρ, the discount factor. When households are more patient, they place less weight on the risk of non-delivery, and so are more likely to use crypto.

Proposition 2.2 (Mixed fee strategies in the limit)**.** *Consider the limit as the households' signal error $\sigma \to 0$. There is a unique threshold equilibrium, in which households use the same switching strategy described in Proposition 2.1. The number of crypto users $y \to 1$ if $R > \hat{R}$, $y \to 0$ if $R < \hat{R}$, and $y \to \frac{1}{2}$ if $R = \hat{R}$. In this limit, households set fees using a mixed strategy with complementary cumulative distribution function $\bar{H}(f)$, where:*

$$\bar{H}(f) \;=\; Z_\lambda^{-1}\Big(Z_\lambda(1) - \frac{\rho}{1-\rho} V f\Big) \tag{2.9}$$

for fees over support $f \in [0, \bar{f}]$, where:

$$\bar{f} \;=\; \frac{1-\rho}{\rho V}\Big(Z_\lambda(1) - Z_\lambda(0)\Big), \tag{2.10}$$

and $V = v(1)$ is the value of crypto when all households use it.

Proof. See Section 2.9.2. □

Definition 2.4 implies $Z_\lambda(\cdot)$ is an invertible function, so $\bar{H}(f)$ over the domain $[0, \bar{f}]$. As $\sigma \to 0$, households' signals converge. An equilibrium in pure fee strategies is not feasible because, if a household knows all others use the same fee, she can profitably deviate by offering a slightly higher fee and achieving top priority in the mempool. Therefore, any equilibrium must have mixed fee strategies in the limit.

The mixed fee strategy in Proposition 2.2 does not depend on the value of the household's signal R_i, as long as it is greater than R^*. Given any signal $R_i > R^*$ then, when $2\sigma < R_i - R^*$, a household believes that every other household uses crypto and offers a fee.

Moreover, as $\sigma \to 0$, there is an arbitrarily high-order belief that everybody uses crypto; that is, household i knows that every other household uses crypto, and knows that they know that, and so forth. In this limit, there is no particular reason for the fees to depend on the signals R_i because, given $R > R^*$, there are almost certainly no households with signals below R^* using cash, and so there is nothing to anchor the fee level.[27]

The effect on fees of shocks to the block rate λ is ambiguous and depends on the slope of the blockchain capacity distribution function $Z_\lambda(\cdot)$. A decrease in block capacity can cause households to pay higher fees, as competition for blockchain space becomes more intense. Alternatively, it can lead them to reduce their fees, because even a high fee cannot guarantee timely settlement. In either case, the threshold $\hat{R}$ at which households use crypto falls, so that fewer households use crypto overall.

The fees paid by households are decreasing in ρV. This quantity is the opportunity cost of offering a unit of crypto in fees, when all other households use crypto. Households know that, conditional on $R > \hat{R}$, all households use crypto, and so every unit of crypto paid to a miner represents an opportunity cost of V units of foregone consumption at T_2, which is worth ρV utils.

2.3.3 Welfare analysis

I define welfare to be aggregate utility, which is equal to the aggregate payoff of all the households before fees are paid. The fees are simply a transfer from households to miners and do not affect welfare. Welfare can be interpreted as a measure of how well the households' chosen form of money works. In this Section, I examine the conditions under which crypto adoption is optimal, and whether the equilibrium outcome achieves optimality. I focus on the case with precise signals described in Proposition 2.2; i.e. $\sigma \to 0$.

Given a particular realisation of technology strength R, crypto usage y, and blockchain

[27] Mathematically, in the proof of Proposition 2.1, the inverse of the fee function has a gradient of order $O(1/\sigma)$, so that it becomes arbitrarily steep as $\sigma \to 0$. This suggests that households all choose the same fee, regardless of their signal. But then a pure strategy equilibrium cannot exist, since any individual household can do better by paying a slightly higher fee and achieving top priority. Therefore, there is a mixed strategy equilibrium. The notion that mixed strategies can be interpreted as the limit of pure strategies as payoff uncertainty tends to zero is well-established; see, for example, Harsanyi (1973).

capacity N, total welfare is:

$$\Omega(R,y,N) \;=\; Rg(y)y \;+\; (1-\rho)\max\Big\{\min\Big\{N,y\Big\},0\Big\} \;+\; (1-\rho)(1-y). \quad (2.11)$$

The first term in Equation (2.11) is households' total payoff from using crypto. The second and third terms are the amount of crypto and cash payments, respectively, which settle by T_1. These terms are multiplied by $1-\rho$, the benefit of early consumption. If $N \geq y$, blockchain capacity exceeds demand for space at T_1, and all crypto payments settle early. If $0 < N < y$, then only N crypto payments settle at T_1. If $N \leq 0$, no crypto payments settle at T_1.

I integrate $\Omega(R,y,N)$ with respect to $Z_\lambda(n)$, the distribution of N, to obtain expected welfare given R and y:

$$\Omega(R,y) \;=\; Rg(y)y \;+\; (1-\rho)\Big(1-\int_0^y Z_\lambda(t)\mathrm{d}t\Big). \quad (2.12)$$

The integral term in Equation (2.12) is the expected number of households whose crypto payments fail to settle by T_1. The second term in Equation (2.12) is therefore the number of crypto payments that settle by T_1, multiplied by the relative benefit of early settlement, $1-\rho$.

Welfare is increasing in the strength of the technology R, since a higher R means a higher portfolio value and a higher payoff from using crypto. Welfare is also increasing in the block rate λ, since a higher λ makes crypto a more reliable payment method.

Definition 2.8. *For a given R, the first-best allocation is the number of crypto users* $y \in [0,1]$ *that maximises welfare* $\Omega(R,y)$. *The second-best allocation is the number of crypto users* $y \in \{0,1\}$ *that maximises welfare.*

A benevolent social planner would, upon observing R, choose the first-best allocation. In general, the first-best y may lie in the interior $(0,1)$. Because the only feasible outcomes with a threshold equilibrium in pure strategies are $y \in \{0,1\}$ (except in the special case $R = \hat{R}$, which occurs with probability zero), the global games allocation cannot, in general, deliver first-best. Thus, I use the second-best as a welfare benchmark.

A social planner constrained to the second-best chooses $y = 1$ if:

$$R \;>\; (1-\rho)\int_0^1 Z_\lambda(t)\mathrm{d}t \;:=\; \hat{R}_{2B}. \tag{2.13}$$

The constrained social planner employs a threshold strategy, choosing $y = 1$ if $R > \hat{R}_{2B}$. This threshold is unique, and is decreasing in the block rate λ. As the block rate λ falls, crypto becomes less attractive as a means of payment, and a higher value of R is required for the social planner to have the households use crypto.

Given a threshold equilibrium with threshold A, expected welfare is:

$$\hat{\Omega}(A) \;=\; \int_0^A \Omega(R,0)\mathrm{d}R \;+\; \int_A^1 \Omega(R,1)\mathrm{d}R. \tag{2.14}$$

Lemma 2.1 (Lower block rate means greater under-use of crypto in global game)**.** *Let $\hat{R}_{2B}$ and $\hat{R}$ be the thresholds for crypto usage under the second-best and global games allocations, respectively. Let $\hat{\Omega}_{2B}$ and $\hat{\Omega}$ be expected welfare under the same two allocations. Then:*

1. *the difference between the thresholds $\hat{R} - \hat{R}_{2B}$ is strictly positive, and is decreasing in the block rate λ;*
2. *the difference between the welfares $\hat{\Omega} - \hat{\Omega}_{2B}$ is strictly negative, and is increasing in the block rate λ.*

Proof. See Section 2.9.3. □

Lemma 2.1 shows that the second-best threshold is strictly lower than the global games threshold. In other words, in the global game, households use crypto less than the social planner would like. This under-use occurs because a marginal household with a signal exactly at the threshold $\hat{R}$ is not sure how many of her peers will use crypto, and will apply a uniform prior. The marginal household's uncertainty about the strength of strategic complementarities causes the threshold to be set high, relative to the second-best case where the social planner knows the exact value of y. The difference between the two thresholds — and thus the difference in expected welfare under the two allocations— decreases with

the block rate λ because, for any given realisation of the number of crypto-using households under the uniform prior, they are more likely to be settled on the blockchain at T_1 when λ is higher.

2.4 Financial market

2.4.1 Model setup

I now introduce a financial market to the model. There is a mass $M > 0$ of identical risk-neutral speculators, born at T_0 and indexed by $j \in [0, M]$. These speculators have access to a large supply of cash but no crypto. They have the opportunity to place an order for crypto on an exchange at T_0. At T_1, each speculator attempts to use all of his wealth to purchase consumption goods. Consumption occurs at times T_1 and T_2, and then the game ends.[28]

Like the households, speculators prefer to consume early, so that consumption at time T_2 is discounted by the factor ρ. Unlike the households, speculators do not have a preference for smooth consumption. Nor do they derive any non-pecuniary benefit from paying with crypto. Thus, a speculator's utility is simply equal to the amount of consumption goods he is able to acquire, discounted over time. That is, given consumption of $c_1, c_2 \geq 0$ at times T_1, T_2 respectively, a speculator's payoff is:

$$u_S(c_1, c_2) \;=\; c_1 + \rho c_2. \tag{2.15}$$

Sell orders must be covered by crypto positioned on-exchange. In other words, uncovered short-sell orders are not possible. As the speculator has no endowment, he cannot sell. This restriction is not key to the model. It is imposed because the cases of buying and selling are not symmetric. In Section 2.5.2.2, I describe this asymmetry in more detail, and show that the model can be extended to allow sell orders without substantively affecting the results.

[28] The mass of speculators M is taken to be an exogenous parameter. In Section 2.5.2.3, I consider endogenising the number of speculators.

At T_0, before making any decisions, each speculator observes a signal Y_j about the strength of the crypto technology R, where $Y_j - R \sim U[-\sigma, \sigma]$ iid. Thus, the speculators have signals of the same accuracy as the households.[29]

Each speculator faces a trade-off. If he can buy crypto for a lower price than he believes it is worth, he can potentially make a trading gain. However, this gain can only be crystallised once the newly purchased crypto is used to buy consumption goods. To make such a purchase, the speculator must move the crypto via the blockchain, just like a household. Uncertainty over when the crypto will be delivered implies a risk that his consumption will be delayed to T_2, in which case his payoff is discounted by a factor ρ.

A speculator j who buys crypto can offer a fee to miners in order to increase the probability of settlement. Because the speculator has no other source of crypto, he must finance the fee from his order.[30] Suppose that he buys one unit of crypto. If he then offers a mining fee of f_j, he will only have $1 - f_j$ remaining to purchase consumption goods. A speculator j who places a buy order chooses a mining fee $f_j \in [0,1]$ to maximise his expected consumption $\hat{K}(f_j; Y_j)$, where:

$$\hat{K}(f_j; Y_j) \;=\; (1-f_j)\mathbb{E}\Big[v(y)\Big(\mathbb{1}\{\text{delivered}\} + \rho\,\mathbb{1}\{\text{not delivered}\}\Big) \;\Big|\; f_j, Y_j\Big]. \quad (2.16)$$

I refer to $\hat{K}$ as the *on-chain value* of crypto. It is the expected utility associated with holding one unit of crypto at time T_1. This utility is, in general, less than $v(y)$, because the probability that consumption goods purchased with crypto are not delivered until T_2 is non-zero.

His expected payoff is then:

$$\pi_S(f_j; Y_j) \;=\; \hat{K}(f_j; Y_j) \;-\; \mathbb{E}\Big[\text{price} \;\Big|\; Y_j\Big], \quad (2.17)$$

where the market price is set by a market maker, described in the next subsection. The payoff π_S is relative to the benchmark of submitting no order; that is, the payoff from not

[29] Nothing substantial changes if the accuracy of the speculators' signals differs from those of the households. See Section 2.5.2.1.

[30] If a speculator were to submit a fee greater than the amount of crypto held, no miner would ever accept the payment into a block. A miner is easily able to assess whether adequate funds are available, because the speculator must specify 'unspent transaction outputs' (UTXOs) in his transaction. These UTXOs are coins that have not yet been spent, and are therefore available to be offered to a miner.

submitting an order is normalised to zero.

2.4.1.1 The trading environment

Trade takes place on an *exchange*. There are three types of participant on the exchange: the *speculators*, *liquidity traders*, and a *market maker*. The exchange permits only market orders, and order sizes are restricted to one unit of crypto, as in Glosten and Milgrom (1985). After speculators and liquidity traders place orders, the market maker sets a price, and trading takes place. If a speculator achieves the same payoff from trading and not trading, he prefers not to trade. To justify this assumption, we can suppose there is some infinitesimal non-pecuniary cost of trading, such as an effort cost. Nothing substantial hinges on this.

Trading on the exchange is off-chain. Crypto is settled by the exchange operator across accounts on the exchange, and is not recorded on the blockchain. Blockchain capacity therefore does not affect the ability of any of the agents to trade. This assumption is realistic; to my knowledge, all major cryptocurrency exchanges have off-chain trading, in order to expedite fast settlement and eliminate mining fees. Limited blockchain capacity does, however, restrict the ability of the speculators to use their newly purchased crypto to consume at T_1.

Each speculator j chooses whether to place a buy order of one unit, or not to trade; that is, he chooses an order $x_j \in \{0,1\}$. Total informed order flow is defined as x, where:

$$x = \int_0^M x_j \, \mathrm{d}j, \tag{2.18}$$

and I assume that the set $\{x_j\}$ is Lebesgue-measurable, so that x is well-defined.

Definition 2.9 (Liquidity traders)**.** *There is a mass M of liquidity traders. Each is born at T_0 with an endowment of at least one unit of crypto. With probability $\frac{1}{2}$, they all place a sell order at T_1, and otherwise they place no order. Define the total liquidity order as u, where:*

$$u = \begin{cases} -M, & \textit{with probability } \frac{1}{2}, \\ 0, & \textit{with probability } \frac{1}{2}. \end{cases} \tag{2.19}$$

The liquidity traders are non-strategic agents whose trading behaviour is solely determined by the realisation of a common liquidity shock. They do not attempt to speculate on the public signal about the technology. I interpret them as miners who are structurally long crypto. On occasion, they are hit by a common shock, which creates a liability that must be satisfied in cash. For example, the shock may be a tax demand or a regulatory action that affects the entire industry and requires them to raise cash.

The price on the exchange is set by a market maker. As in Kyle (1985), the market maker is subject to competitive pressure, due to a free entry condition, that ensures she makes zero expected profits. The market maker observes the order flow z but is not able to interpret the public signal about crypto technology, and so her beliefs about R are defined by the posterior given the prior $R \sim U[0,1]$, updated by the information contained in z. She sets a price equal to her posterior belief about the value that a speculator derives from holding a unit of crypto at T_1.[31]

Definition 2.10 (Market maker)**.** *There is a single market maker, born at T_0 with a large endowment of crypto and cash. She balances the market by absorbing excess supply or demand. She observes total order flow $z = x + u$ and sets a price $p(z)$ equal to the expected consumption value of a unit of crypto given z:*

$$p(z) \;=\; \mathbb{E}\Big[\max_{f\geq 0}\Big\{\hat{K}(f;Y_j) \;\Big|\; Y_j, f\Big\} \;\Big|\; z\Big]. \tag{2.20}$$

The price given by Equation (2.20) is not necessarily equal to the market maker's expected long-term value of crypto $v(y)$. This is because there is a non-zero probability that the crypto cannot be used to purchase consumption at T_1, and so the value of crypto to a speculator is, in general, less than $v(y)$. The market maker sets the price so that she expects excess demand in the market to be equal to zero, which is the same logic as in Kyle (1985).

As is common in the market microstructure literature, the presence of the liquidity traders means the market maker cannot perfectly observe the actions of the informed speculators. This potentially allows the speculators to profitably trade on their private information

[31] In principle, the market maker could pay a household or speculator for information about R. However, competitive pressure implies that, if she were to do so, she would not adjust her price, because a competitor would always match it.

about R. Because the focus of this paper is on the effect of the blockchain structure and mining fees, rather than market microstructure, the trading environment is kept simple. There already exists an extensive literature around order sizes and their effect on pricing and liquidity; see Foucault, Pagano, and Roell (2013). Section 2.5.2 discusses the robustness of the results in this model to changes in the trading environment.

2.4.1.2 Payoffs and equilibrium concept

The events over the three time periods are as follows:

1. At time T_0:
   (a) The strength of the technology $R \sim U[0,1]$ is determined by nature. Each household $i \in [0,1]$ privately observes a signal R_i, where $R_i - R \sim U[-\sigma,\sigma]$ iid. She chooses a payment medium $y_i \in \{0,1\}$ and, if $y_i = 1$, a fee $f_i \geq 0$.
   (b) Each speculator $j \in [0,M]$ privately observes his signal Y_j, where $Y_j - R \sim U[-\sigma,\sigma]$ iid. He chooses an order $x_j \in \{0,1\}$ and a fee $f_j \in [0,1]$.
   (c) Nature chooses an order for the liquidity traders, $u \in \{-M,0\}$ with equal probabilities.
   (d) The market maker observes total order flow $z = x + u$ and sets a price $p(z)$ using Equation (2.20).
2. At time T_1, R and y are revealed and the long-term value of crypto $v(y)$ is determined. The amount of block space $N \sim Z_\lambda(n)$ is determined by nature. The N crypto payments that have the highest fees attached are added to the ledger, as are all cash payments. Any household that had her payment added to the ledger can consume. Any speculator that bought crypto and had his payment added to the ledger can consume.
3. At time T_2, all remaining crypto payments are added to the ledger. Consumption occurs and the game ends.

Using Equation (2.17), a speculator j who places a buy order $x_j = 1$ and offers a mining

fee $f_j \in [0,1]$ earns the following expected payoff:

$$\pi_S(f_j;R) \;=\; \mathbb{E}\Big[(1-f_j)v(y)\Big(1-(1-\rho)\mathbb{1}\{\text{not delivered}\}\Big)-p(1+u) \;\Big|\; f_j,R\Big]. \quad (2.21)$$

The speculator's payoff is zero if he does not place an order. The household's payoff is again denoted π_H and is given by Equation (2.5), as before.

Definition 2.11 (Equilibrium with financial market)**.** *An equilibrium is defined as a set of the following functions:*

- *a household order function* $\eta : [-\sigma, 1+\sigma] \to \{0,1\}$*, so that* $y_i(R_i) = \eta(R_i)$ *for all* i*;*
- *a household fee function* $\phi_H : [-\sigma, 1+\sigma] \to [0,\infty)$*, so that* $f_i(R_i) = \phi_H(R_i)$ *for all* i*;*
- *a speculator order function* $\chi : [-\sigma, 1+\sigma] \to \{0,1\}$*, so that* $x_j(Y_j) = \chi(Y_j)$ *for all* j*;*
- *a speculator fee function* $\phi_S : [-\sigma, 1+\sigma] \to [0,1]$*, so that* $f_j(Y_j) = \phi_S(Y_j)$ *for all* j*;*
- *a pricing function* $p : \{-M, 0, M\} \to [0,\infty)$*, so that* $p(z)$ *satisfies Equation* (2.20)*;*

such that the actions of the households and speculators are optimal given beliefs consistent with the equilibrium actions of the other agents.

As before, the solution concept is Bayesian Nash equilibrium. All households and all speculators are assumed to play the same strategies as their peers.

2.4.2 Equilibrium with a financial market

As before, I follow the global games literature and focus on the case $\sigma \to 0$, when households' and speculators' signals become arbitrarily accurate. As in Proposition 2.2, households employ a threshold strategy for their monetary choice, and employ a mixed strategy for fees. Speculators do the same. An equilibrium where crypto is never used cannot exist. To see why, note that if crypto is never used, then $y = 0$ for sure. The speculators do not trade, because it is common knowledge that $v(y) = 0$ regardless of the value of R,

and so they cannot profit from trading. But I show in the proof of Proposition 2.1 that households with sufficiently high signals will still use crypto, even if they believe that $y = 0$. This result is ensured by Assumption 2.1. Therefore, we can focus on threshold equilibria where crypto is used for some values of R.

Given that total potential demand for blockchain space is now $1 + M$ rather than 1, I amend Assumption 2.2 in the following logical way.

Assumption 2.3 (Existence of threshold equilibrium when there is a financial market)**.**

$$Z_\lambda(1+M) \;\leq\; \frac{\int_0^1 Z_\lambda(t(1+M))dt}{\int_0^1 g(t)dt} \;\leq\; \frac{Z_\lambda(0)}{g(0)}. \tag{2.22}$$

Assumption 2.3 guarantees the existence of a threshold equilibrium by ensuring that, as a household's signal R_i increases from just below the threshold to just above it, the increase in the non-pecuniary payoff exceeds the loss in utility from reduced blockchain space. It performs the same role as Assumption 2.2 in the previous section.

Lemma 2.2 (Equilibrium with financial market)**.** *Suppose that Assumptions 2.1 and 2.3 both hold. Then, in the limit as $\sigma \to 0$, there is a unique threshold equilibrium. All agents play a switching strategy with a threshold R^*, where:*

$$R^* \int_0^1 g(t)dt \;=\; (1-\rho)\int_0^1 Z_\lambda\big(t(1+M)\big)dt. \tag{2.23}$$

When $R_i \leq R^*$, household i uses cash. When $R_i > R^*$, household i uses crypto. Each speculator j places a buy order if he observes a signal $Y_j > R^*$, and places no order if his signal $Y_j \leq R^*$.

The structure of fee strategies falls into one of three mutually exclusive categories:

1. *If $\rho V \leq 1 - (1-\rho)Z_\lambda(1)$, speculators choose lower fees than the households.*
2. *If $\rho V \geq 1 + (1-\rho)(Z_\lambda(1+M) - Z_\lambda(M) - Z_\lambda(0))$, speculators choose higher fees than the households.*
3. *Otherwise, the support of the speculators' fees is nested within the support of the households' fees.*

The market maker sets price $p(z)$ where:

$$p(z) = \begin{cases} 0, & \text{if } z = -M, \\ (1-R^*)K, & \text{if } z = 0, \\ K, & \text{if } z = M, \\ \int_0^1 v(t)(1-(1-\rho)Z_\lambda(t(1+M)))dt, & \text{if } z = -\frac{M}{2} \text{ or } \frac{M}{2}, \end{cases} \tag{2.24}$$

where K is the expected value of crypto to a speculator who observes a signal above R^* and follows the equilibrium strategy; that is, $K = \hat{K}(f_j; Y_j)$, where $Y_j > R^*$ and f_j is the optimal fee.

Proof. See Section 2.9.4. Explicit solutions for the fee functions are contained in the proof; for brevity they are omitted from the statement of the Proposition. □

The expression for the threshold R^* given in Equation (2.23) is obtained by considering a marginal household who observes a signal exactly at the threshold. She has a posterior belief that the distributions of households and speculators who observe signals above the threshold are uniform, and so she sets a fee of zero. Assumption 2.3 ensures that this is indeed a threshold.

The speculators employ the same threshold as the households. This is because the speculators' payoffs do not depend directly on R, but only on the households' actions. Therefore, they buy when they believe the households use crypto, and do not buy when they believe the households use cash.

Exactly one equilibrium exists. This equilibrium can fall into one of three categories of equilibrium, depending on the value of ρV. These categories are described fully in the proof. When ρV is low, households tend to choose high fees relative to the speculators and, when it is high, households tend to choose relatively low fees. The reason is that ρV is a measure of the opportunity cost to a household from paying a fee. When $y = 1$, the opportunity cost to a household per unit fee is a reduction in T_2 consumption of V units, which is worth ρV utils. If ρV is low, households are willing to offer higher nominal fees, because the benefit of early consumption outweighs the cost of the higher fee. If ρV is high, households prefer to set lower fees. For speculators, both the trading gain and the

price are linear in V, so changes in the value of V do not directly affect the fees offered.

Proposition 2.3 (Crowding out)**.** *R^* is decreasing in λ and increasing in M.*

Proof. The right-hand side of Equation (2.23) is decreasing in λ and increasing in M, so the result follows immediately. □

Proposition 2.3 shows that a lower block rate λ means a lower probability of settlement for a given fee, and so crypto is less useful as a means of payment. As before, the threshold is also decreasing in the discount factor ρ and the magnitude of the network effects between households $\int_0^1 g(t)\mathrm{d}t$. These comparative statics are the same as in the case without a financial market, examined in Section 2.3.2.

There is now a *crowding out* effect: as the number of speculators M increases, the threshold for crypto usage R^* also rises. This effect is due to competition for blockchain space from speculators. As M increases, demand for blockchain space from speculators goes up, meaning that a higher fee is required to achieve a given probability of settlement for crypto. This increased demand for blockchain space makes crypto less useful as a means of payment, and so households raise the threshold R^* for crypto usage. Note that, when $M = 0$, then $R^* = \hat{R}$, and we have the situation described in Proposition 2.2.

The effects of block rate λ and mass of speculators M on the price are ambiguous. Lemma 2.2 tells us the expected price $\mathbb{E}[p(z)]$ is equal to $(1-R^*)K$, where R^* is the threshold for crypto usage and K is the on-chain consumption value of a unit of crypto at T_1.

Proposition 2.4 (Pecuniary effect)**.** *Consider the random variable $p(z)/K$, which is the market maker's price normalised by the value of crypto. This quantity has first order stochastic dominance as λ increases and M decreases. That is, for all q, $\mathbb{P}[p(z)/K > q]$ is weakly increasing in λ, and there is at least one point q for which it is strictly increasing. And, for all q, $\mathbb{P}[p(z)/K > q]$ is weakly decreasing in M, and there is at least one point q for which it is strictly decreasing. When $V \leq 1$, there is a stronger result: the price $p(z)$ has first order stochastic dominance as λ increases and M decreases.*

Proof. The normalised price of crypto $p(z)/K$ is 0 with probability $\frac{1}{2}R^*$, $1-R^*$ with probability $\frac{1}{2}$, and 1 with probability $\frac{1}{2}(1-R^*)$. We know from Proposition 2.3 that R^*

is monotonically strictly decreasing in λ and increasing in M, so the results for $p(z)/K$ follow immediately.

A speculator's value of holding crypto K is, in general, not monotonic in λ or M. A higher value of λ (or lower value of M) can increase the value of K, because it implies that the probability of settlement is higher, which makes crypto more valuable at T_1. But it may also mean that the households, and other speculators, pay higher fees. This means that, in general, the comparative statics of price $p(z)$ with respect to λ and M are not monotonic. However, the case $V \leq 1$ implies $\rho V \leq 1-(1-\rho)Z_\lambda(1)$, so any threshold equilibrium must be of the first type described in Lemma 2.2. The proof of Lemma 2.2 shows that the on-chain value of crypto at T_1 is $K = 1-(1-\rho)Z_\lambda(M)$, which is increasing in λ and decreasing in M. Therefore, we have first-order stochastic dominance in the price as λ increases and M decreases. More detailed proof in Section 2.9.5. □

To obtain tractable conclusions, I focus on the first type of equilibrium described in Lemma 2.2, which has the simplest comparative statics.

Proposition 2.4 describes a *pecuniary effect*. As the block rate λ increases, crypto becomes less useful as a means of payment, which reduces the price set by the market maker. The same effect occurs as the number of speculators M increases. Therefore, a greater amount of informed buy-side speculation can reduce the price of a cryptocurrency. That is a rather surprising result. In models such as Kyle (1985), buy-side trading tends to increase the price of an asset set by a market maker. The pecuniary effect arises due to the interaction of the crowding out effect with the endogenous value of crypto. It is therefore a novel effect that is peculiar to cryptocurrencies.

The pecuniary effect relies on the crowding out effect. Endogenous valuation without the blockchain friction is not sufficient to generate the pecuniary effect. Lemma 2.3 shows that, without the crowding out effect, a change in the amount of speculation M would have no effect on the price of the asset (recall that M is not only the quantity of informed speculators, but is also the amount of liquidity trading). This result explains why there is no pecuniary effect for token-based money that does not settle on a blockchain, such as commodity-backed money.

Lemma 2.3 (Pecuniary effect relies on crowding out effect)**.** *Suppose that the probability*

of early settlement is fixed and exogenously given, so that it is not affected by λ, or by the activity of other households or speculators. Then the price is not dependent on the number of speculators.

Proof. If the probability of settlement $Z_\lambda(n)$ depends on neither λ nor n, then the distribution of the price will not depend on these parameters either. The number of speculators M does not affect the probability of settlement and, because the number of liquidity traders is equal to M, changes to M do not affect the informativeness of the order flow. Full proof in Section 2.9.6. □

The pecuniary effect is counter to the usual intuition in the microstructure literature. Standard wisdom is that more informed buy-side trading should not decrease the price — see, for example, Kyle (1985) or Glosten and Milgrom (1985). It is surprising that the presence of buy-side speculators can *decrease* the price. A model with a more general distribution of liquidity trading can still exhibit the pecuniary effect, as long as blockchain capacity is low relative to the magnitude of liquidity trading. The pecuniary effect implies price formation for cryptocurrencies may be qualitatively different from that for standard asset classes, and different even from that for other forms of money.

The pecuniary effect can cause speculative profits to be increasing in the number of other speculators. The reason is that, as M increases and causes R^* to rise, the market maker attaches a lower prior probability to the event that $R > R^*$, so a speculator can purchase crypto on the exchange at a lower price. This effect imposes losses on the liquidity traders, who must now sell at a lower price, so there is a wealth transfer from liquidity traders to speculators. However, it is not always the case that profits increase with M: higher M also means that the on-chain value K of a unit of crypto at T_1 is lower. Therefore, the overall effect of higher M on the speculator's profits is ambiguous. Although a speculator can buy crypto more cheaply when M is higher, it is harder to consume. Lemma 2.4 gives a sufficient condition for profits to be increasing in the number of speculators.

Lemma 2.4 (Speculative profits are increasing in the number of informed speculators)**.** *When $V \leq 1$, a speculator's expected payoff is increasing in M, conditional on $R > R^*$. As $\lambda \to \infty$, the speculator's payoff tends to zero.*

Proof. See Section 2.9.7. □

Lemma 2.4 shows that an increase in the number of speculators can actually raise the trading profit of each individual speculator, due to the pecuniary effect. When $V \leq 1$, speculators pay lower fees than households according to Lemma 2.2, so the crowding out effect has a greater effect on households' payoffs than speculators'.

As λ tends to infinity, there is enough blockchain space for everyone, and the crowding out effect disappears. All households use crypto. Knowing this, the market maker sets the price to V, and the order flow carries no additional information. The speculators' private information becomes worthless and they make zero trading profit.

2.4.3 Price volatility

I now study how the volatility of the crypto price is affected by block rate λ and the number of speculators M. By Lemma 2.2, the distribution of the price is:

$$p(z) = \begin{cases} 0, & \text{when } z = -M, \text{ which occurs with probability } \frac{1}{2}R^*, \\ (1-R^*)K, & \text{when } z = 0, \text{ which occurs with probability } \frac{1}{2}, \\ K, & \text{when } z = M, \text{ which occurs with probability } \frac{1}{2}(1-R^*). \end{cases} \tag{2.25}$$

The mean and variance of the price are, respectively:

$$\mathbb{E}[p(z)] = (1-R^*)K, \qquad \mathbb{V}[p(z)] = \frac{1}{2}R^*(1-R^*)K^2. \tag{2.26}$$

The mean price is equal to $(1-R^*)K$, which is her prior expectation of the value of a unit of crypto to a speculator at T_1. This is a consequence of the law of iterated expectations.

The asset pricing literature typically makes cross-sectional comparisons between different assets using volatility of returns; that is, the standard deviation of the change in price over a given time period. For example, computations of the Sharpe ratio — which allows for comparison of risk-reward across portfolios — are done on the basis of returns, rather than nominal prices. Table 2.1, which demonstrates the high volatility of cryptocurrencies,

also measures the volatility of returns. In my model, since there is only one period, the appropriate denominator is the ex ante mean price $\mathbb{E}[p(z)]$. We can think of the market maker setting a prior price of $\mathbb{E}[p(z)]$ before observing the order flow z, and a posterior price of $p(z)$ after observing it. Then $p(z)/\mathbb{E}[p(z)]$ is the return on crypto.

I define *price volatility* as the standard deviation of the price normalised by its mean. This metric is also known as the coefficient of variation (Pinches and Kinney Jr. 1971), or relative standard deviation.

Definition 2.12 (Price volatility)**.** *Unconditional price volatility is defined as the standard deviation of the price divided by its mean:*

$$\Gamma_U = \frac{\sqrt{\mathbb{V}[p(z)]}}{\mathbb{E}[p(z)]}. \tag{2.27}$$

Conditional price volatility is the standard deviation of the price divided by its mean, conditional on a realisation of R:

$$\Gamma_C(R) = \frac{\sqrt{\mathbb{V}[p(z)|R]}}{\mathbb{E}[p(z)|R]}. \tag{2.28}$$

Alternatively, this definition can be thought of as time-series volatility. Suppose that the market maker posts a price before T_0, when she has no information about R other than its prior distribution. She will set a price equal to the ex ante of her prior of the crypto value, which is $(1-R^*)K$. Then volatility, as defined in Definition 2.12, is the standard deviation of the return obtained from buying a unit of crypto at this prior price, and selling it at the ex post price $p(z)$.

Using this definition of price volatility, I establish the following results.

Proposition 2.5 (Price volatility in the threshold equilibrium)**.** *Unconditional price volatility is:*

$$\Gamma_U = \sqrt{\frac{\frac{1}{2}R^*}{1-R^*}}, \tag{2.29}$$

which is increasing in R^*.

Price volatility conditional on $R < R^*$ is equal to 1. Price volatility conditional on $R > R^*$ is:

$$\frac{R^*}{2-R^*}, \tag{2.30}$$

which is increasing in R^*, decreasing in λ, and increasing in M. Price volatility conditional on $R = R^*$ is equal to zero.

Proof. See Section 2.9.8. □

Unconditional price volatility Γ_U is increasing in the threshold R^*, and thus decreasing in the block rate λ. When the blockchain capacity is expected to be low, price volatility increases. This is a result of the crowding out effect. As R^* increases, the standard deviation becomes larger, relative to the mean.

The pecuniary effect explains the effect of changes to R^* on conditional price volatility. A higher value of R^* reduces the market maker's prior about the value of crypto, so she sets a lower price. The order flow is then more informative: if the market maker observes $z = 1$, she infers that crypto will be used as a means of payment, and sets the price to 1. Because the market maker's price is more sensitive to order flow, the volatility is higher. As R^* is decreasing in λ, an immediate implication is that price volatility is decreasing in blockchain capacity.

This crowding out effect on unconditional price volatility occurs even in the absence of informational frictions in the market. Suppose the market maker can perfectly observe R. Then she sets the price equal to K if $R > R^*$, and 0 otherwise. There is still a crowding out effect, because higher blockchain capacity makes crypto more attractive and reduces R^*. But there is no pecuniary effect, because the market maker does not use the order flow in setting the price. The price is simply a Bernoulli random variable, equal to K with probability $1 - R^*$, and zero otherwise. The unconditional price volatility is $\sqrt{R^*/(1-R^*)}$, which is increasing in R^*, but the conditional price volatility is zero. Thus, the crowding out effect alone can explain the effect of λ on *unconditional* price volatility described in Proposition 2.5, but not its effect on *conditional* price volatility. To explain the behaviour of conditional price volatility, we need the pecuniary effect.

Two features are important here. First, the monetary nature of a cryptocurrency implies that its value is endogenously determined by the amount of usage. Second, the blockchain

infrastructure implies that speculative activity has a negative impact on the value. Together, these features generate an interaction between speculative activity and price that is absent from other asset classes. My results suggest that high volatility is the result of the unique features of blockchain-based cryptocurrencies.

2.4.4 Welfare

As before, welfare is aggregate utility, which is equal to the aggregate payoff of all the households before fees are paid. Trading is a zero-sum game, because any profits made by the speculators are exactly matched by losses borne by the liquidity traders.

Proposition 2.6 (Welfare with a financial market)**.** *Welfare under the financial market allocation is lower than that without a financial market. Welfare under the financial market allocation is increasing in* λ *and decreasing in* M*.*

Proof. See Section 2.9.9. □

Speculators take up blockchain space and push up fees. The fees have no direct effect on welfare, because they represent a transfer from households and speculators to miners. But fees do reduce the usefulness of crypto as a means of payment, due to the crowding out effect. Therefore, the presence of speculators may cause households to use cash when it would be otherwise optimal to use crypto. The more speculators there are, the greater the welfare loss relative to the case with no speculators, and thus relative to the second-best benchmark (as we know from Lemma 2.1 that the allocation without speculators is worse than the second-best benchmark). As in Lemma 2.1, a higher block rate λ always makes society better off. As the rate of block creation increases, the crowding out effect is reduced, and crypto becomes a better means of payment.

Households in my model are risk-neutral, so price volatility does not directly affect welfare. If agents were risk-averse, they may be reluctant to use crypto when volatility is high, because they would be exposed to price movements while they hold it.[32] This reluctance

[32] Saleh (2018) presents a model in which risk-averse households prefer payment methods that are less volatile. However, in his model, volatility arises due to the fixed supply schedule of crypto, rather than for speculative reasons.

would raise the threshold R^*, magnifying the crowding out effect.

This analysis is predicated on trading being a zero-sum game, so speculation never has any social value. In this setting, if a social planner could prohibit or tax speculative trading, society would be better off. But, in a model where households are imperfectly informed about the value of the technology R, the public signal contained in prices may have some social value. There is a large literature on feedback between financial markets and the real economy (see Bond, Edmans, and Goldstein 2012 for a review). For example, consider a setup where households' signals are much less accurate than speculators', and where households can observe the price set by the market maker before choosing their payment method. The price can help the households to form more accurate posteriors for R, which may be efficient. In global games models with public information, agents tend to react more strongly to public information than to private signals (see Morris and Shin 2003), so the existence of accurate public information can raise volatility even higher and lead to multiple equilibria (Angeletos and Werning 2006).

2.5 Variations to the model

In this section, I discuss some of the modelling choices I have made, and examine whether my key result — that a lower block creation rate increases price volatility — is robust to generalisation. I divide these modelling choices into two broad groups. First, I look at relaxing constraints that were used to guarantee the existence or uniqueness of equilibria. Second, I look at changes to the trading environment.

2.5.1 Alternative equilibria

2.5.1.1 Uniqueness of threshold equilibrium: relaxing Assumption 2.1

Assumption 2.1 guarantees the uniqueness of the threshold equilibrium in 2.2 (and Proposition 2.1) by ensuring that, if a household receives a sufficiently high signal, she will use crypto regardless of her beliefs. This condition implies that there is an upper dominance region, in the language of Carlsson and Van Damme (1993), meaning that any equilibrium

must have some crypto usage. Furthermore, the Assumption guarantees that the threshold R^* is less than 1, meaning a threshold exists.

If this Assumption is relaxed, we have two potential equilibria: the threshold equilibrium found in Lemma 2.2, and an additional equilibrium in which crypto is never used. In this new equilibrium, every household believes that $y = 0$ for sure, and so using cash is always better. Then $v = 0$ for sure, so the speculators never buy and the market maker sets a price of zero. There is no financial market activity and no price volatility.

Lemma 2.5 (Relaxing Assumption 2.1)**.** *Suppose that there is a financial market and Assumption 2.1 does not hold. Let* $\sigma \to 0$*. Two equilibria are now possible: the financial market equilibrium described in Lemma 2.2, and a 'cash-only' equilibrium where all households use cash and there is no trading. If* $g(0) < (1-\rho)Z_\lambda(0)$*, the cash-only equilibrium exists. If* $R^* < 1$*, the financial market equilibrium exists. It is possible to have two, one, or no equilibria.*

Proof. See Section 2.9.10. □

Relaxing Assumption 2.1 has two effects. First, it makes possible the existence of a cash-only equilibrium, as long as $g(0) < (1-\rho)Z_\lambda(0)$. This condition ensures that the benefits of using crypto are low relative to the probability of failed settlement. Second, relaxing the Assumption introduces the possibility that a threshold equilibrium may not exist. As a consequence, we can end up with both equilibria, exactly one, or none. In Section 2.5.1.2, I discuss the situation where the threshold equilibrium does not exist.

The cash-only equilibrium is trivial and implies a price of zero. It cannot help us to understand cryptocurrency price formation or why cryptocurrency prices have been so volatile. The fact that prices and price volatility are non-zero in the real world suggests that markets place a non-zero probability on crypto adoption and coordinate on the equilibrium with a financial market.

2.5.1.2 Existence of threshold equilibrium: relaxing Assumptions 2.2 and 2.3

Given that Assumption 2.1 holds, Assumption 2.2 guarantees the existence of a threshold equilibrium for households in the absence of a financial market. In other words, As-

sumption 2.2 ensures that, given a threshold $\hat{R}$, households with signals just above $\hat{R}$ are willing to use crypto, whereas those with signals just below $\hat{R}$ are not. The Assumption says the marginal network benefits of crypto are high relative to the risk of non-delivery. Therefore, the expected increase in y that occurs as R_i increases from below $\hat{R}$ to above $\hat{R}$ causes a household to switch her choice of money. Similarly, Assumption 2.3 guarantees existence when there is a financial market.

Let us focus on the case without a financial market. In a threshold equilibrium, $\hat{R}$ as given in Equation (2.8) is the only viable threshold so, if Assumption 2.2 does not hold, then we cannot have a threshold equilibrium in pure strategies. I introduce a more general equilibrium concept that allows mixed strategies, as long as those strategies are continuous functions of the signal.

Definition 2.13 (Mixed equilibrium)**.** *A mixed equilibrium is an equilibrium in which each household's strategy is to use crypto with probability* $\eta(R_i)$ *and cash with probability* $1-\eta(R_i)$*, where* $\eta(\cdot)$ *is a continuous function. Each household uses a fee strategy that is a continuously differentiable function* $\phi(R_i)$*.*

In a threshold equilibrium, a household must pay with either crypto or cash. Now I allow a household to adopt mixed strategies, so she can choose crypto with probability $\eta(R_i)$, and cash with probability $1-\eta(R_i)$. Alternatively, we can think of this strategy as a household splitting her purchase, so that she buys a portion $\eta(R_i)$ of a unit of the consumption good with crypto, and $1-\eta(R_i)$ with cash. As before, the household adopts a pure strategy for the fees when $\sigma > 0$. The decision function $\eta(R_i)$ is restricted to be continuous in R_i. A threshold equilibrium is therefore a limiting case of a mixed equilibrium.

Lemma 2.6 (Crowding out in mixed equilibria)**.** *Consider a household i who receives a signal* R_i *such that* $0 < \eta(R_i) < 1$*. Then, holding fees constant, a negative shock to the block rate* λ *reduces aggregate crypto usage by households.*

Proof. See Section 2.9.11. □

Lemma 2.6 shows that the crowding out effect still exists in a mixed equilibrium: a reduction in the block rate tends to reduce crypto usage. The financial market equilibrium and pricing function are rather complicated in this version of the model, because crypto

can take values other than 0 and V. But the economic intuition suggests a pecuniary effect will still exist.

2.5.2 Changes to the trading environment

2.5.2.1 Accuracy of speculators' information

The model assumes that both households and speculators receive equally informative signals about R, with their private signal identically and independently distributed $\sim U[-\sigma,\sigma]$. The main results still hold if households and speculators have signals of different accuracies. For example, the households could be technical experts who are better at understanding public information about R at T_0. They may therefore be able to obtain more accurate signals about R than the speculators can, and trade on the basis of their expertise. Equally, the opposite could be true.

Consider a generalised setup in which each speculator j has signal Y_j such that $Y_j - R \sim U[-\varepsilon,\varepsilon]$ iid. The speculator's signal error parameter $\varepsilon(\sigma,\tau)$ is parametrised by σ, the error of the household's signals, and by τ, a parameter independent of σ. It is continuous in both parameters. As σ and τ both tend to zero, $\varepsilon(\sigma,\tau)\to 0$. This general structure allows for some correlation between the errors of the signals of the households and speculators as they tend to zero.

Suppose σ,τ satisfy $\varepsilon(\sigma,\tau)\geq\sigma$. The threshold for crypto usage R^* is then given by:

$$R^*\int_{R^*-\sigma}^{R^*+\sigma}\frac{1}{2\sigma}g\Big(\frac{R-R^*+\sigma}{2\sigma}\Big)\mathrm{d}R \;=\; (1-\rho)\int_{R^*-\sigma}^{R^*+\sigma}Z_\lambda\Big(\frac{R-R^*+\sigma}{2\sigma}+M\big(\frac{R-R^*+\varepsilon}{2\varepsilon}\big)\Big)\mathrm{d}R, \tag{2.31}$$

since a household with signal exactly equal to R^* offers a fee of zero. We can make a substitution $t=(R-R^*+\sigma)/2\sigma$ to obtain:

$$R^*\int_0^1\frac{1}{2\sigma}g(t)\mathrm{d}t \;=\; (1-\rho)\int_0^1 Z_\lambda\Big(t+M\Big(\frac{1}{2}+\frac{\sigma}{\varepsilon}\big(t-\frac{1}{2}\big)\Big)\Big)\mathrm{d}t, \tag{2.32}$$

so R^* is decreasing in λ and increasing in M. In the case $\varepsilon(\sigma,\tau)<\sigma$, the right-hand side of Equation (2.31) is a little more complicated, but the signs of the derivatives of R^* with

respect to λ and M do not change. Therefore, the crowding out and pecuniary effects occur in more general information structures.

2.5.2.2 Allowing sell orders

The setup restricts the speculators to buy orders, because exchanges do not allow uncovered short-selling, and the speculators have zero endowment of crypto. This restriction is imposed because buying and selling are not symmetric. However, it is not key to the results. In this Section, I consider the problem faced by speculators who can place sell orders, and show that the crowding out and pecuniary effects still occur.

Suppose each speculator has access to unlimited units of crypto, held in a wallet off-exchange. He can, if he wishes, sell his coins on the exchange. To do so, he must transfer them via the blockchain. Once the crypto is on the exchange, he can sell it at the price set by the market maker. He can then withdraw the cash and use it to buy consumption goods without any further frictions.[33]

When placing a buy order, the speculator trades first, and then takes on delivery risk when he attempts to use crypto for consumption. When placing a sell order, he takes on delivery risk first and then trades only if delivery is successful. Therefore, there is a chance that he will not be able to trade at all. The buying and selling cases are not symmetric. This point is important, because the market maker will now need to take delivery risk into account when setting prices. She will find it useful to observe the realisations of both order flow z and blockchain capacity N, and use these observations to improve her estimate of y.

This delivery risk requires a tweak to the timing of the model. Previously, the market maker observed order flow z and set her price at T_0. In a model with selling, speculators who wish to sell can only do so if delivery is successful, so the order flow only crystallises after blockchain capacity is realised. Therefore, the market maker now observes the order

[33] Equivalently, we could consider the problem of a speculator with no endowment, but who is able to borrow a unit of crypto from a lender. This crypto must be moved by the lender to the speculator's exchange account via the blockchain before it can be sold. We must also assume the lender is ignorant and non-strategic: she has no information about R and does not update her beliefs after observing the speculator's request to borrow. For example, the lender may be an uninformed miner who has an urgent demand for cash. This lender charges no interest, but must be repaid in kind at T_2 and compensated for any fee paid.

flow at T_1. She also observes blockchain capacity at T_1, and so sets the price given all of the information contained in z and N. Timing is as follows:

1. At T_0, the strength of the technology R is determined. Each household i observes a signal R_i and chooses her means of payment y_i and, if she uses crypto, a fee f_i. Each speculator j observes a signal Y_j and places an order x_j and, if relevant, chooses a fee f_j. Nature chooses an order for the liquidity traders.

2. At T_1, R and y are revealed. Blockchain capacity N is realised and observed by the market maker. The market maker also observes order flow z and sets a price $p(z,N)$. Ledgers are updated and consumption can occur. Speculators who placed a sell order are able to consume only at T_1 if their order has been added to the blockchain.

3. At T_2, all remaining payments are added to the ledger and consumption occurs.

As before, order sizes are restricted to one unit, so a speculator's order is $x_j \in \{-1,0,1\}$. A speculator can offer any fee. As before, speculators weakly prefer not trading to trading, if they yield the same monetary payoff.

I make an adjustment to the structure of the liquidity trading, so that total liquidity trades $u \in \{-M,M\}$, with probability $\frac{1}{2}$ attached to each outcome. This adjustment makes it difficult for the market maker to distinguish between the speculators placing a buy order and the speculators successfully selling. With the previous liquidity trading structure $u \in \{-M,0\}$, the market maker would only face uncertainty if the sellers were unsuccessful in placing their trades, in which case selling would never be a profitable activity.

Proposition 2.7 (Buy and sell orders)**.** *Suppose that each speculator can place both buy and sell orders. Suppose also that household signal error $\sigma \to 0$, and that liquidity trades are $u \in \{-M,M\}$ with equal probability. Then there is a unique equilibrium where each household i uses crypto if her signal $R_i > R^*$, where R^* is the financial market threshold identified in Lemma 2.2, and uses cash otherwise. Each speculator j places a buy order if his signal $Y_j > R^*$, a sell order if $Y_j < R^*$, and no order if $Y_j = R^*$.*

When a household observes $R_i > R^*$, or a speculator observes $Y_j > R^*$, he or she adopts the same mixed strategy for fees as in Lemma 2.2. When a speculator observes $Y_j < R^*$, he is indifferent between all fee strategies.

The market maker sets a pricing function $p(z,N)$, where:

$$p(z,N) = \begin{cases} K, & \text{if either } z = 2M, \text{ or if } z = 0 \text{ and } N < M, \\ (1-R^*)K, & \text{if } z = 0 \text{ and } N \geq M, \\ 0, & \text{otherwise}, \end{cases} \tag{2.33}$$

where K is the on-chain consumption value of a unit of crypto at T_1.

Unconditional price volatility is:

$$\Gamma_U = \sqrt{\frac{\frac{1}{2}R^*}{1-R^*}\Big(1+Z_\lambda(M)\Big)}. \tag{2.34}$$

Conditional price volatility is:

$$\Gamma_C(R) = \begin{cases} \frac{1+Z_\lambda(M)}{1-Z_\lambda(M)}, & \text{if } R < R^*, \\ \frac{R^*}{2-R^*(1-Z_\lambda(M))}\sqrt{1-Z_\lambda(M)^2}, & \text{if } R > R^*. \end{cases} \tag{2.35}$$

Proof. See Section 2.9.12. If a speculator observes $Y_j < R^*$, he knows crypto is worthless and places a sell order. There is no benefit to holding onto crypto after T_0, so he is willing to pay as high a fee as possible to miners. However, if $N < M$, then only N speculators are able to position their crypto on the exchange by T_1 and sell. Given N and $z = -N \pm M$, the market maker can tell immediately that speculators are selling, and so sets a price of zero. Thus, selling speculators can only make a profit if $N \geq M$, in which case there is enough blockchain space for them all to sell regardless of fees. Therefore, speculators are willing to set any fee. □

Volatility is not only driven by the crowding out and pecuniary effects, but also by uncertainty about whether sell orders can be placed. The crowding out and pecuniary effects occur only when speculators buy, but not when they sell. This is because speculators sell only when they believe that the households do not use crypto; that is, when $R < R^*$. Therefore, in a selling environment, households never require blockchain space, and so they cannot be crowded out by sellers.

Unconditional price volatility is increasing in M and decreasing in λ. The same is true

of conditional price volatility when $R < R^*$. However, when $R > R^*$, conditional price volatility can be locally decreasing in M or increasing in λ. The reason for this non-monotonicity is that, when λ is low or M is high, then $Z_\lambda(M)$ is high, and there is a high chance that $N < M$ and some sell orders cannot be placed. When $R < R^*$ and $N < M$, then the market maker can easily discern that the value of y is 0. Therefore, a high probability that $N < M$ implies low price volatility. The magnitude of this effect depends on the sensitivity of $Z_\lambda(M)$ to M and λ.

In general, the crowding out and pecuniary effects exist in this setup, and influence price volatility. However, the way that limited blockchain space affects sell orders means that, in some cases, a decrease in blockchain space actually leads to lower volatility, because it reduces the market maker's uncertainty about the value of crypto $v(y)$.

2.5.2.3 Endogenous number of speculators

So far, I have assumed that the mass M of speculators is exogenously determined, so that we can think of a change to M as an exogenous shock. It may be more realistic to consider a model with an endogenous number of speculators, where new speculators enter the market until the marginal benefit of doing so equals the marginal cost (i.e. a free entry condition). We have already derived a relationship between M and the profitability of speculation in the proof of Lemma 2.4, so it is straightforward to consider M as a function of the cost of entry.

Suppose now that there is an unlimited mass of potential speculators. Let $c > 0$ be the one-time fixed cost to speculation. This may represent the cost, for example, of learning enough about blockchain technology to be able to understand the source code and thus acquire a private informative signal.

A potential speculator who does not pay c receives no signal and is thus no more informed than the market maker or liquidity traders. In a Kyle (1985)-type model, an uninformed speculator would make zero profit. In this model, he will do worse, because the blockchain constraint means there is a non-zero probability that he cannot use the crypto to buy consumption goods at T_1. As his return from trading is negative, a potential speculator who chooses not to buy the signal will certainly not trade. This means that we can

interpret M as the number of potential speculators who buy the signal.

Definition 2.14 (Free entry of speculators)**.** *Let $\hat{\pi}_S(M)$ be the expected profit from becoming an informed speculator when there are M other informed speculators, and let c be the cost of obtaining a signal. Then M satisfies the free entry condition if $\hat{\pi}_S(M) = c$.*

The free entry condition in Definition 2.14 simply says that speculators enter the market until their marginal revenue is equal to the marginal cost.

Lemma 2.7 (Endogenous mass of speculators)**.** *Suppose that $V \leq 1$ and M speculators buy the signal. The speculator's expected payoff from trading, prior to observing a signal, is:*

$$\hat{\pi}_S(M) \;=\; \frac{1}{2}R^*(1-R^*)\Big(1-(1-\rho)Z_\lambda(1+M)\Big), \tag{2.36}$$

where R^* is as given in Equation (2.23)*. There may be as many as three solutions to the free entry condition, and as few as zero.*

Proof. As R^* and $Z_\lambda(1+M)$ are both increasing in M, the expression $\hat{\pi}_S(M) - c$ has up to three roots in M. It is easy to show that existence is not assured; since $\hat{\pi}_S(M) > 0$ for all M, the case $c = 0$ has no solution. For more detail, see Section 2.9.13. □

Lemma 2.7 demonstrates that the free entry condition does not uniquely determine the number of speculators. This result can be generalised to other values of V. This result is in contrast to standard market entry models, where monotonicity of the profit function ensures uniqueness, and existence can be guaranteed by a coverage condition on the range of the profit function.

In my model, when $\lambda \to \infty$, $R^* \to 0$ and so $\hat{\pi}_S(M) \to 0 < c$ for all M, and so the free entry condition cannot be satisfied. There is zero speculation. Therefore, the blockchain capacity constraint can lead to more speculative trading than might otherwise be expected, although the solution M is not, in general, monotonic in λ.

When there is more than one solution, it is useful to comment on their stability. Any solution M with $\hat{\pi}'_S(M) < 0$ is stable. No uninformed potential speculator has an incentive to purchase the signal. Speculators who do purchase the signal are indifferent between doing so and not. However, we can also have solutions with $\hat{\pi}'_S(M) > 0$. In this case,

if a non-zero mass of potential speculators decide to join, they push up trading profits for everyone. Then it will be rational for more speculators to enter. Ruling out unstable equilibria can cut down the number of solutions, but there can still be 0, 1 or 2. This kind of unstable equilibrium does not occur in standard models of entry, because in those models profits are monotonically decreasing in the number of entrants.

2.5.2.4 Speculator order size

I impose a restriction that only one unit of crypto can be traded. This assumption is made for simplicity, because the aim of this paper is to examine the effect of blockchain capacity and fees, rather than to retread a well-studied issue in the market microstructure literature. Restricting trades to one unit prevents the market maker from learning any information about R from the size of the speculator's order (as she can, for example, in Kyle 1985).

In fact, in the limit as $\sigma \to 0$, allowing different order sizes would not make any substantive difference to the results. Conditional on $R > R^*$, neither the speculator's order strategy nor his fee strategy depend on R. This suggests that, if we allow a speculator to choose any order $x > 0$ then, for all $R > R^*$, he would choose the same x and the same fee strategy. The same is true conditional on $R < R^*$. Therefore, the market maker does not learn anything from the order flow, aside from possibly whether $R > R^*$ or not. The relationship between speculators' trading decisions and the price function is the same as before. Therefore, the outcomes do not change.

2.5.2.5 Structure of liquidity trades

The setup uses a very simple structure for liquidity trades. This assumption is made for tractability. Lemma 2.8 shows that the unconditional price volatility result found in Proposition 2.5 holds with an arbitrary liquidity trading distribution.

Lemma 2.8 (Arbitrary distribution of liquidity trading)**.** *Consider any distribution of liquidity trades u that is independent of R and x. In a threshold equilibrium, unconditional price volatility is increasing in the threshold R^*.*

Proof. See Section 2.9.14. □

Lemma 2.8 tells us that price volatility is increasing in the threshold. As long as the threshold R^* is decreasing in λ, volatility is decreasing in λ. Similarly, R^* increasing in M implies volatility is increasing in M. Therefore, Proposition 2.5 can be generalised to any arbitrary liquidity trading distribution, as long as it is independent of the value of crypto and the speculators' actions.

2.6 Empirical implications

The model makes three predictions:

1. speculative trading crowds out adoption of the cryptocurrency as a medium of payment,
2. an increase in the amount of speculative trading leads to higher price volatility, and
3. a reduction in the block creation rate leads to higher price volatility.

This section discusses how these results may be verified empirically.

The crowding out effect is well-understood in the cryptocurrency community. The year 2017 saw a surge in speculative interest in cryptocurrencies, which led to higher fees and longer waiting times for settlement. It appears that speculation may have had an economically significant impact on the use of cryptocurrencies as a payment instrument (e.g. BBC News 2017). Therefore, there is plenty of anecdotal evidence for the first prediction. Empirical validation would require quantitative data on the speculative and monetary uses of cryptocurrency. Unfortunately, no reliable data on this are available, because the pseudonymous nature of blockchain transactions means that there is no easy way to discern the purpose of a given transaction. Some attempts have been made to identify and quantify the purpose of transactions on the blockchain. These studies typically rely on assumed heuristics about the nature of different types of transactions. For example, Meiklejohn et al. (2013) develop clustering techniques to group together Bitcoin wallets that appear to have common operators, and examine changes over time in the Bitcoin market. Coinmetrics Team (2018) discusses some of the issues involved in this kind of analysis. However, verifying whether these proposed metrics are reliable measures of payments usage is currently impossible.

Because speculative activity cannot be easily measured, I focus on the third prediction, linking the block rate to volatility. Daily data on cryptocurrency prices are available at websites such as `blockchain.info` and `coinmarketcap.com`. Each of the major cryptocurrencies trades on several exchanges, so price is generally measured as an average across exchanges, weighted by volume. These data can be used to construct a measure of daily price volatility.

To test the third prediction, I need to find a plausibly exogenous change in the block rate λ that is anticipated by agents. In my model, agents' optimal actions depend on the value of λ, and all actions are made before the blockchain capacity N is realised. Therefore, their behaviour is affected by changes in the block rate only if they observe these changes. There are three candidates for such an observable exogenous change in λ: cross-sectional differences between cryptocurrencies, changes to protocols that affect blockchain capacity, and the periods between difficulty adjustments. I discuss each of these candidates below.

2.6.1 Cross-sectional variation in the block creation rate

Different cryptocurrencies vary in blockchain capacity. For example, Bitcoin has one of the slowest block creation rates, at an average of one block every 10 minutes. Other cryptocurrencies tend to be faster: Litecoin has an average block inter-arrival time of 2.5 minutes; Monero, two minutes; and Ethereum, around 15 seconds. In addition, most other cryptocurrencies have larger blocks than Bitcoin. Thus, we might expect, all else being equal, that the bitcoin price should be more volatile than the prices of these other cryptocurrencies.

Any such analysis would need to take account of differences in liquidity between the markets for different cryptocurrencies. As the largest and most established cryptocurrency, Bitcoin is by far the most traded, which is a confounding factor when doing cross-sectional analysis. Liquidity between these markets is likely to be correlated. In addition, most other cryptocurrencies have not existed for very long — for example, Ethereum, the second largest by market cap, was created in 2015 — and so any time-series dimension will be fairly short.

A literature has recently developed that examines cryptocurrency prices and the factors that drive them. For example, Liu and Tsyvinski (2018) find that the prices of Bitcoin, Ripple and Ethereum are driven by different factors to standard asset prices. Biais et al. (2018) show that the usefulness of Bitcoin as a means of payment affects its price. To my knowledge, no empirical studies of the block creation rate as a pricing factor have yet been conducted.

2.6.2 Protocol changes

Occasionally, changes are made to the protocol that governs the blockchain of a cryptocurrency. These changes are typically enacted in order to make an improvement, such as a change to the block size or the block creation rate, a bug fix, or a reversal of the effects of a hack. The miners decide whether to use the updated protocol, in a process known as a *fork*. If a majority of miners agree to use it, the new protocol is adopted. If there is disagreement, the blockchain can split in two. Some cryptocurrencies have been created this way — for example, Bitcoin Cash and Litecoin were created as forks from Bitcoin. Biais et al. (2017) discuss forks in more detail.

Several attempts have been made to create a fork to increase the block creation rate of Bitcoin (Morgan 2017). The only attempt that can be described as a success was the activation of Segregated Witness (SegWit) in August 2017. SegWit does not actually change the block creation rate or the block size; rather, it changes the way data are stored on blocks, so that each transaction takes up less space. It effectively increases the capacity of each block by a factor of up to four.[34]

Bitcoin miners signal their acceptance of the SegWit protocol on their blocks. Parsing the blockchain data provides a numerical proxy for SegWit adoption.[35] This series can then be tested for an effect on price volatility. My model predicts that greater SegWit adoption should reduce volatility, as long as it is anticipated. Any regression model would need to control for other possible causes of volatility.

[34] See Song (2017) for more information.

[35] See `https://transactionfee.info/charts/payments/segwit`.

2.6.3 Periods between difficulty adjustments

The effect of the block rate on price volatility can be analysed by exploiting the fact that the block rate varies predictably in the short run. In the model, I used a constant block creation rate λ. This assumption of a constant block rate is a simplification. It is true in the long run but, in the short run, the block rate can vary. In a proof-of-work system such as Bitcoin, the protocol sets the *difficulty* of the computational problem that miners have to solve in order to earn the right to create a block. The difficulty determines the expected amount of computing power — called *hash rate* — required to solve the problem. If the aggregate hash rate changes, the protocol adjusts the difficulty accordingly, thus keeping the supply of new blocks broadly constant over time.

In the case of Bitcoin, the protocol aims for an average rate of one new block every 10 minutes. The difficulty is adjusted every 2016 blocks, which usually takes around two weeks. Between these adjustments, difficulty is fixed, so changes in the hash rate will affect the supply schedule of Bitcoin. Suppose, for example, that the average hash rate over blocks $\{n, \cdots, n+2015\}$ is h. After block $n+2015$, the protocol resets the difficulty to $h\omega$, where ω is a normalising constant to target the desired block creation rate. Now suppose that, immediately after the adjustment, the miners double their computing power to $2h$. Then blocks will be created twice as quickly — that is, every five minutes, on average — until block number $n+4031$, about one week later. Then the protocol will double the difficulty to $2h\omega$, in order to try to maintain a stable block creation rate. Between readjustments, the block creation rate λ, which is related to the hash rate divided by the difficulty, is variable.

Figure 2.2 illustrates this process. It shows an example in which the hash rate is exponentially increasing. Every 2016 blocks, the difficulty adjusts to be proportional to the average hash rate since the last adjustment. Therefore, changes in difficulty always lag behind changes in the hash rate. The block rate λ is proportional to the hash rate divided by difficulty. It tends to be relatively high just before an adjustment, and relatively low just after.

Table 2.2 shows that the rate of block creation is different immediately before and after difficulty adjustment dates. On the first full day after an adjustment takes place, 151.3 blocks are created on average. On the last full day before an adjustment, 159.6 blocks are

Figure 2.2: Illustration of difficulty adjustment process in Bitcoin. The hash rate increases exponentially. Difficulty adjusts every 2016 blocks based on the average hash rate since the last difficulty adjustment. It is always lagging. The ratio of hash rate to difficulty varies predictably.

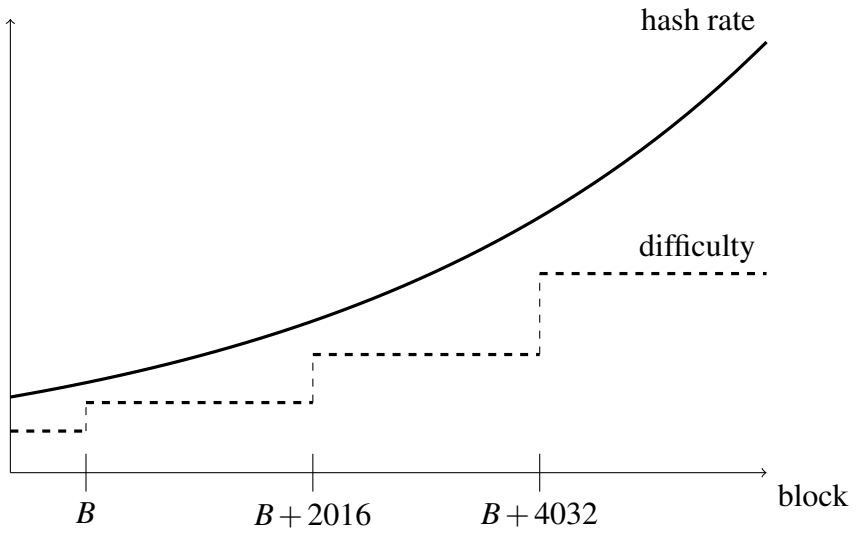

created on average. This difference is statistically significant. This result arises because of a secular rise in the hash rate over time, driven by the entry of new miners, and by existing miners investing in new, more efficient mining technology. It also explains why the mean rate of block creation is higher than the protocol's target of 144 per day.

Table 2.2: **Blocks created in Bitcoin and price volatility around difficulty adjustments, 28 April 2013 – 30 September 2018. Price volatility is measured as the difference between the daily high and low price, divided by the mean price. Block data are from `blockchain.info` and price data from `coinmarketcap.com`.**

	Adjustment yesterday	**Adjustment tomorrow**	**Other days**
Observations	156	156	1514
Mean blocks created	151.3	159.6	156.9
Std dev of blocks created	15.0	25.8	19.0
Mean price volatility	0.054	0.049	0.052
Std dev of price volatility	0.052	0.043	0.052

The secular increase in the hash rate is predictable, and so is the timing of difficulty

adjustments. Households and speculators are able to anticipate that, as the time since the last difficulty reset increases, the block creation rate λ rises. By Proposition 2.5, a higher block creation rate should be associated with lower price volatility. Table 2.2 appears to back up that hypothesis, but the daily measure of price volatility is itself very volatile, which may make it difficult to establish significance.

Any empirical analysis will need to identify and control for other factors that drive volatility, as such factors could confound the analysis. An example would be an omitted factor that causes high price volatility and creates inventory risk for miners, leading them to reduce their hash rate. A causality analysis can help disentangle these effects. In any case, as long as these confounding factors are not plausibly related to the timing of the difficulty adjustments, then evidence that difficulty adjustments affect volatility after controlling for the hash rate would be compelling evidence in favour of my theory.

2.7 Conclusion

I study the effect of restricted blockchain capacity on cryptocurrency price volatility. When a cryptocurrency has low blockchain capacity, it is less useful as a form of money and has lower value. Speculators require blockchain space in order to trade, so they reduce the utility of cryptocurrency as a payments medium, which in turn reduces its value. The cryptocurrency becomes cheaper for speculators to buy, and that makes the order flow more informative to a market maker, resulting in increased price volatility.

My model makes three important contributions. First, it identifies a reason for the price volatility of cryptocurrency that is unique to this asset class, thus explaining why volatility is so high relative to other asset classes. Second, it demonstrates that the price impact of cryptocurrency trading may differ in sign from that of other asset classes. Third, I make a methodological contribution by solving a new type of global games model where users can pay fees to prioritise their payments.

Roy Amara's celebrated law (Amara 2016) states that the impact of technological innovations are often overestimated in the short term and underestimated in the long term. This adage is particularly true of cryptocurrencies. Hype about future adoption leads to specu-

lation, which makes adoption more expensive and thus less likely. My model cannot say whether cryptocurrencies will eventually be adopted as money, but it does suggest that, despite the slow rate of adoption so far, adoption can still occur in the future.

2.8 Appendix A: Distribution of blockchain space

This section contains a discussion of the distribution of blockchain space, and in particular examines the empirical distribution for $Z_\lambda(n)$.

In a proof-of-work cryptocurrency such as Bitcoin, miners earn the rate to create a new block by solving a computationally difficult problem set by the protocol; Harvey (2016) provides a good explanation. These problems can be likened to a 'guess-the-number' game in which the range of possible answers is very high. Once one problem is solved and a block created, miners immediately move on to the next problem. The Poisson distribution is therefore an appropriate model for the number of blocks created in a given period of time (Nakamoto 2008). For example, Bitcoin is designed to generate an average of one block every 10 minutes.

In reality, Bitcoin block generation is faster than the target rate of one block every 10 minutes, because adjustments to the difficulty of the problem lag behind changes in the miners' computing power, or hash rate, as explained in Section 2.6.3. Bowden et al. (2018) show that changes in the hash rate and lagged difficulty adjustments mean blocks do not follow a homogeneous Poisson distribution in the short term. Other cryptocurrencies that have faster difficulty adjustments are better approximated by a homogeneous Poisson distribution.

Each block can contain a certain amount of data, measured in bytes. For example, a Bitcoin block can currently contain up to 1 megabyte of data. Transactions are not homogeneous in size. The amount of memory required for each transaction can vary widely because a user can encode additional information along with her transaction. Most famously, block zero (the 'genesis block') was encoded with a newspaper headline about a British bank bailout in order to verifiably time-stamp the transaction, and perhaps also to send a political message. Other uses include artwork, signalling support for forks,

marriage proposals, adverts, and, potentially, illicit content; see Sedgwick (2018) and Matzutt et al. (2018). Miners typically charge fees per byte, so transactions that take up more space will require higher fees. The size of the payment or the fees offered do not generally affect the amount of space required.

Figure 2.3 shows a daily measure of blockchain capacity for Bitcoin over its lifespan. Each observation is the total amount of block space generated on a particular day (in bytes), divided by the average size of transactions on that day. Note that this statistic is not the same as the actual block space used, because some blocks may not have been full. The top percentile has been removed due to a very long tail.

Figure 2.3: Daily Bitcoin blockchain capacity 3 January 2009 – 30 September 2018. Top percentile not shown due to long tail.

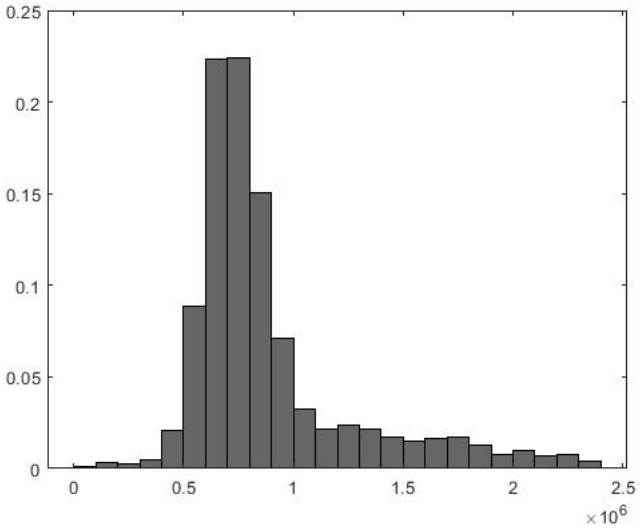

The mean and median total daily capacity over this period are 912,000 and 773,000 transactions respectively. The highest was 7.1 million, and the lowest was 8,157. The long tails are due to a wide variation in the size of transactions. A typical transaction takes up around 440 bytes, but some are as small as 80 bytes, whereas others can be up to half a megabyte. No data on the purpose of individual transactions are available, but it may be reasonable to guess that transactions near the median are straightforward payments. Those transactions that take up more space most likely contain additional information. Smaller transactions may be nuisance entries. Although the model does not capture de-

mand for blockchain space arising from the wish to send messages, we can think of it as a source of random variation in the number of agents' payments that can fit into a block, as captured by the distribution $Z_\lambda(n)$.

2.9 Appendix B: Proofs

2.9.1 Proof of Proposition 2.1

This proof consists of two steps. First, I show that, in any threshold equilibrium, each household sets fees as an increasing function of her signal R_i. I then show that there is a unique fee function that is consistent with equilibrium. Finally, I show that, given this fee function, a unique threshold equilibrium exists. Let $\hat{R}$ be the equilibrium threshold, and let $\phi(R_i)$ be the fee strategy. Note that, by Definition 2.7, $\phi(\cdot)$ is continuously differentiable.

2.9.1.1 Proof that a household's fees are increasing in her signal in any threshold equilibrium

For this part of the proof, I show that, if the fee function has a turning point, then it has other turning points of the same nature spaced at intervals of size 2σ. But the existence of these turning points implies that one exists below $\hat{R}$, which is impossible because the fee is certainly zero there. Thus, the fee function is monotonic. Then I show that, because the fee is zero just below $\hat{R}$ and it cannot be constant over a measurable set, it must be increasing just above $\hat{R}$. Therefore, it must be increasing almost everywhere.

When $R_i > \hat{R}$, household i uses crypto and her payoff function is given by Equation (2.5). Thus, she sets the fee to minimise the following loss function:

$$L(f_i;R_i) \;=\; (1-\rho)\mathbb{P}\Big[\text{non-delivery}\Big|R_i,f_i\Big] \;+\; \rho f_i v(y). \tag{2.37}$$

Because the probability of delivery depends only on the realisation of blockchain capacity N and on other households' fees, a household's signal R_i does not directly affect her optimal fee decision, except to the extent that it affects her beliefs about other households'

fees.

I establish the first part of the proof by establishing a sequence of results.

Result 2.1. *In equilibrium, a household i uses crypto and offers a fee f_i only if she believes that the probability of another household offering a fee of exactly f_i is zero.*

Suppose that household i believes there is a positive probability that other households offer the same fee as she does. Faced with payments with equal fees, miners allocate blockchain space randomly, so a household would always do better to pay a slightly higher fee and obtain a discontinuous increase in priority. Suppose that, given R, i believes $N_0(R)$ other households pay fees strictly higher than f_i, and $N_1(R)$ pay fees exactly equal to f_i, where $N_1(R)$ over a measurable subset of $[R_i-\sigma, R_i+\sigma]$. Then her loss with fee f_i is:

$$L(f_i;R_i) \;=\; \frac{1-\rho}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma}\Big(Z_\lambda(N_0(R))+\int_{N_0(R)}^{N_1(R)}\frac{n}{N_1(R)-N_0(R)}Z'_\lambda(n)\mathrm{d}n\Big)\mathrm{d}R \;+\rho f_i V, \tag{2.38}$$

where $Z_\lambda(n)$ is the first derivative of $Z_\lambda(n)$; that is, the probability density function of the blockchain capacity. If the household deviates to a fee of $f_i+\varepsilon$, her loss is no greater than:

$$L(f_i+\varepsilon;R_i) \;\leq\; \frac{1-\rho}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma}Z_\lambda(N_0(R))\mathrm{d}R \;+\; \rho(f_i+\varepsilon)V, \tag{2.39}$$

which, for small enough ε, is strictly lower than $L(f_i,R_i)$. Thus, there cannot be an equilibrium in which two households choose the same fee with positive probability.

Given this result, household i's loss function can be written as:

$$L(f_i;R_i) \;=\; \frac{1-\rho}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma}Z_\lambda\Big(\frac{1}{2\sigma}\int_{R-\sigma}^{R+\sigma}\mathbb{1}\{\phi(R_j)>f_i\}\mathrm{d}R_j\Big)\mathrm{d}R \;+\; \rho f_i V. \tag{2.40}$$

The probability of non-delivery given a fee f_i is equal to the probability that blockchain capacity N is less than the number of households who offer higher fees than f_i.

Result 2.2. *For $R_i\in[\hat{R},1-\sigma]$, the optimal fee function $\phi(R_i)$ is strictly increasing at R_i if and only if $\phi(R_i-2\sigma)<\phi(R_i+2\sigma)$, and is strictly decreasing if and only if $\phi(R_i-2\sigma)>\phi(R_i+2\sigma)$. The household selects a fee f_i between $\phi(R_i-2\sigma)$ and $\phi(R_i+2\sigma)$.*

The intuition behind this result is as follows. Given the form of household i's loss function

(2.40), her optimal fee depends only on her beliefs about the fees of her peers. The higher i believes those fees to be, the higher her own optimal fee. Suppose that the fee function is strictly increasing at R_i; i.e. there exists a $\delta > 0$ such that, for all $\varepsilon \in (0, \delta)$, $\phi(R_i + \varepsilon) > \phi(R_i)$. As her signal increases from R_i to $R_i + \delta$, household i now has to worry about the fees of households who get signals in $(R_i + 2\sigma, R_i + \delta + 2\sigma]$, but no longer has to worry about fees of households who get signals in $[R_i - 2\sigma, R_i + \delta - 2\sigma)$. Thus, her optimal fee can be increasing only if the former set of households offer higher fees than the latter; i.e. $\phi(R_i + 2\sigma) > \phi(R_i - 2\sigma)$.

In addition, the household's choice of optimal fee must lie between $\phi(R_i - 2\sigma)$ and $\phi(R_i + 2\sigma)$. If it is lower than both, then her probability of delivery does not change, so she is paying a higher fee with no benefit, which cannot be optimal. The same is true if it is higher than both.

An analogous argument follows for the case of a strictly decreasing optimal fee function. When $R_i > 1 - \sigma$, as R_i increases there are no new households included in i's calculation, but some households with lower signals drop out. Therefore, only the inequality relating to lower signals is relevant here.

More formally, household i chooses a fee to satisfy her first-order condition:

$$\frac{\partial}{\partial f_i}\Big[\int_{R_i-\sigma}^{R_i+\sigma} Z_\lambda\Big(\int_{R-\sigma}^{R+\sigma} \mathbb{1}\{\phi(R_j) > f_i\}\mathrm{d}R_j\Big)\mathrm{d}R\Big] \;=\; -2\sigma V\frac{\rho}{1-\rho}. \tag{2.41}$$

Differentiating the first-order condition with respect to R_i yields:

$$\begin{aligned}\frac{\partial^2}{\partial f_i^2}\Big[\int_{R_i-\sigma}^{R_i+\sigma} Z_\lambda\Big(\int_{R-\sigma}^{R+\sigma} \mathbb{1}\{\phi(R_j) > f_i\}\mathrm{d}R_j\Big)\mathrm{d}R\Big]\frac{\partial f_i}{\partial R_i} \\ +\; \frac{\partial^2}{\partial f_i \partial R_i}\Big[\int_{R_i-\sigma}^{R_i+\sigma} Z_\lambda\Big(\int_{R-\sigma}^{R+\sigma} \mathbb{1}\{\phi(R_j) > f_i\}\mathrm{d}R_j\Big)\mathrm{d}R\Big] \;=\; 0.\end{aligned} \tag{2.42}$$

In equilibrium, $\phi'(R_i) = \partial f_i / \partial R_i$. We are interested in the sign of this derivative. After some manipulation, Equation (2.42) implies:

$$sign[\phi'(R_i)] \;=\; sign\Big[\mathbb{1}\{\phi(R_i + 2\sigma) > f_i\} - \mathbb{1}\{\phi(R_i - 2\sigma) > f_i\}\Big], \tag{2.43}$$

Here, $sign[\cdot]$ is a function that returns 1 when the argument is strictly positive, -1 when

it is strictly negative, and 0 when it is exactly zero. Thus, $\phi(R_i)$ is increasing in R_i if and only if $\phi(R_i+2\sigma) > \phi(R_i) > \phi(R_i-2\sigma)$, and decreasing if and only if $\phi(R_i+2\sigma) < \phi(R_i) < \phi(R_i-2\sigma)$.

Result 2.3. *If the fee function $\phi(R_i)$ has a turning point at R', then it has turning points of the same nature at $R'-2\sigma$ and $R'+2\sigma$, as long as those points lie in $[-\sigma, 1+\sigma]$. The fee at each turning point is the same.*

Suppose ϕ achieves a local maximum at R', where $\hat{R} < R' \leq 1-\sigma$. Then there is a $\delta > 0$, such that, for all $\varepsilon \in (0,\delta)$, $\phi(R_i)$ is increasing at $R'-\varepsilon$ and decreasing at $R'+\varepsilon$. Result 2.2 then implies that:

$$\phi(R'-2\sigma-\varepsilon) < \phi(R'-\varepsilon) < \phi(R'+2\sigma-\varepsilon), \tag{2.44}$$

and:

$$\phi(R'-2\sigma+\varepsilon) > \phi(R'+\varepsilon) > \phi(R'+2\sigma+\varepsilon). \tag{2.45}$$

Since these expressions are true for any $\varepsilon \in (0,\delta)$, we must have $\phi(R'-2\sigma) = \phi(R') = \phi(R'+2\sigma)$, by the squeezing lemma.

I now need to show that $R'-2\sigma$ and $R'+2\sigma$ are local maxima. Suppose not, and that $\phi(\cdot)$ is increasing at $R'+2\sigma$. Then there is a $\hat{\delta} > 0$ such that $\phi(R'+2\sigma-\hat{\varepsilon}) < \phi(R'+2\sigma+\hat{\varepsilon})$ for all $\hat{\varepsilon} \in (0,\hat{\delta})$. As there is a local maximum at R', and $\phi(\cdot)$ is continuous, we can find $\varepsilon_1, \varepsilon_2 \in (0,\hat{\delta})$ such that $\phi(R'-\varepsilon_1) > \phi(R'+\varepsilon_2)$, and such that $\phi(\cdot)$ is increasing at $R'-\varepsilon_1$ and decreasing at $R'+\varepsilon_2$. But then we must have $\phi(R'+2\sigma-\varepsilon_1) > \phi(R'+2\sigma+\varepsilon_2)$, which contradicts the definition of $\hat{\delta}$. Therefore, $\phi(\cdot)$ cannot be increasing at $R'+2\sigma$. An analogous argument rules out it being decreasing at this point, so it must be a turning point.

Suppose that $\phi(\cdot)$ has a minimum at $R'+2\sigma$. Since $\phi(R') = \phi(R'+2\sigma)$, the existence of this minimum implies that there is a $\bar{\delta}$ such that, for $0 < \bar{\varepsilon} < \bar{\delta}$, $\phi(R'+2\sigma+\bar{\varepsilon} > \phi(R'+2\sigma)$. But then $\phi(R'+\bar{\varepsilon}) > \phi(R'+2\sigma) = \phi(R')$, which contradicts R' being a local maximum. Therefore, $\phi(\cdot)$ has a local maximum at $R'+2\sigma$.

Analogous arguments establish the result for $R'-2\sigma$, and for when R' is a local minimum.

Result 2.4. *There is a $\zeta > 0$ such that the optimal fee function $\phi(R_i)$ is weakly increasing over $R_i \in (\hat{R}, \hat{R}+\zeta]$.*

Consider a household i with a signal just above $\hat{R}$. She knows that any of her peers who receive signals $R_j \leq \hat{R}$ will use cash, and therefore do not submit payments to the mempool. If there is no neighbourhood $(\hat{R}, \hat{R}+\zeta]$ over which $\phi(\cdot)$ is increasing, then $\phi(\cdot)$ must be constant and equal to zero over a set with positive Lebesgue measure. But then there is a contradiction with Result 2.1, so this cannot be true.

Result 2.5. *The fee function $\phi(R_i)$ is increasing almost everywhere.*

Suppose not. Result 2.4 tells us the fee is increasing close to $\hat{R}$, so there must be a local maximum at some point R'' in $(\hat{R}, 1+\sigma)$. Result 2.3 then implies that there are turning points at $R' - 2\sigma, R' - 4\sigma, \cdots$. Condition (2.7) ensures that $\hat{R} > 2\sigma$, since $\int_0^1 g(t)\mathrm{d}t < 1$. Therefore, there must be a turning point in $[0, \hat{R})$. But that is impossible, since all households with signals below the threshold use cash and pay no fee. We have a contradiction, so the fee must be increasing almost everywhere. The fee function can have local inflexion points where it is weakly increasing, but it cannot decrease anywhere, neither can it be constant over a set with positive Lebesgue measure.

2.9.1.2 Proof that threshold equilibrium is unique

First, I rule out the possibility of a trivial equilibrium where either crypto or cash is always used. Then I solve for the threshold by considering the household who observes a signal exactly at the threshold, and then solve for the fee function of the households with higher signals. Again, I divide this part of the proof into a sequence of results to make it easier to follow.

Result 2.6 (Lower dominance region)**.** *There is a lower dominance region; that is, an $R_L \in (0,1)$ such that a household i always uses cash if her signal $R_i \leq R_L$, regardless of her beliefs about others' actions.*

The most optimistic belief i can have is that all other households use crypto, and pay fees of exactly zero. Then, if she wishes to use crypto, she can pay a negligible fee and achieve priority over all other households. Under this belief, her expected payoff from

using crypto is:

$$\pi_H(f_i;R_i) \;=\; R_i-(1-\rho)Z_\lambda(0)-\rho V f_i. \tag{2.46}$$

Her crypto payment is added to the blockchain by T_1, as long as blockchain capacity N is strictly greater than 0. If $R_i \leq (1-\rho)Z_\lambda(0)$, then this payoff is negative. No household with $R_i \leq (1-\rho)Z_\lambda(0)$ will ever use crypto, regardless of her beliefs. This proves the result.

Result 2.7 (Upper dominance region)**.** *There is an upper dominance region; that is, an $R_U \in (0,1)$ such that a household i always uses crypto if her signal $R_i > R_U$, regardless of her beliefs about others' actions.*

The most pessimistic belief for a household i is not necessarily that $y=0$; she may be worse off if some households use crypto and take up blockchain space. Suppose the most pessimistic belief is that $y=y'$ and all the households that use crypto pay infinitely high fees. Under this belief, her expected payoff from using crypto is:

$$\pi_H(f_i;R_i) \;\geq\; R_i g(y')-(1-\rho)Z_\lambda(y'). \tag{2.47}$$

The right-hand side is her payoff when she sets a fee of zero; no finite fee will be better. For the household to be willing to use crypto regardless of beliefs, we must have $\pi_H > 0$ for all $y' \in [0,1]$. For R_i sufficiently large, this condition is guaranteed by Assumption 2.1. Therefore, an upper dominance region exists.

(A sufficient condition to rule out the threshold equilibrium $\hat{R}=1$ is $g(0) > (1-\rho)Z_\lambda(0)$. I use the stronger Assumption 2.1 because it would be required to obtain a threshold equilibrium.)

Result 2.8. *Suppose the threshold for crypto usage is $\hat{R}$. A household i with signal R_i uses crypto and pays fee $\phi(R_i)$.*

Consider a household with signal $R_i > \hat{R}$. She aims to set a fee to maximise her payoff from using crypto, as given by Equation (2.5). The first term $\mathbb{E}[R_i g(y)]$ is not affected by her choice of fee, but the other two terms are. Therefore, her problem is to set $f \geq 0$ to

minimise a loss function $\kappa(f;R_i)$, where:

$$\kappa(f;R_i) \;=\; (1-\rho)\mathbb{P}[\text{non-delivery}\Big|\; R_i, f] \;+\; \rho v(y) f. \tag{2.48}$$

Let $\phi(R_j)$ be the equilibrium fee strategy that other households use. We know $\phi(R_j)$ must be increasing in R_j, so it has a well-defined inverse over its domain. Let this inverse be ψ, so that $\psi(\phi(R_j)) = R_j$ for all $R_j \in [\hat{R}, 1+\sigma]$. Then the household's loss function is:

$$\kappa(f;R_i) \;=\; \frac{1-\rho}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma} Z_\lambda\Big(\frac{1}{2\sigma}\int_{R-\sigma}^{R+\sigma} \mathbb{1}\{R_j > \psi(f)\}\mathrm{d}R_j\Big)\mathrm{d}R \;+\; \rho f \mathbb{E}[v(y)|R_i]. \tag{2.49}$$

Non-delivery occurs if blockchain capacity N is less than the number of households that make crypto payments with fees higher than f. Because i knows that every other household receives a signal in $[R_i-2\sigma, R_i+2\sigma]$, her optimal fee must either lie in $[\phi(R_i-2\sigma), \phi(R_i+2\sigma)]$, or it is zero. Consider the case where it is non-zero. The loss function is equal to:

$$\kappa(f;R_i) \;=\; \frac{1-\rho}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma} Z_\lambda\Big(\frac{R+\sigma-\psi(f)}{2\sigma}\Big)\mathrm{d}R \;+\; \rho f \mathbb{E}[v(y)|R_i]. \tag{2.50}$$

The first-order condition for f is:

$$0 \;=\; \frac{1-\rho}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma} Z'_\lambda\Big(\frac{R+\sigma-\psi(f)}{2\sigma}\Big)\Big(-\frac{\psi'(f)}{2\sigma}\Big)\mathrm{d}R \;+\; \rho\mathbb{E}[v(y)|R_i]. \tag{2.51}$$

In equilibrium, we must have $\phi(f) = R_i$. I put this into the first-order condition and make a substitution $t = \frac{1}{2\sigma}(R - R_i + \sigma)$ to obtain:

$$\begin{aligned} 0 \;&=\; -\psi'(f)\Big(\frac{1-\rho}{2\sigma}\Big)\int_0^1 Z'_\lambda(t)\mathrm{d}t \;+\; \rho\mathbb{E}[v(y)|R_i], \\ &=\; -\psi'(f)\Big(\frac{1-\rho}{2\sigma}\Big)\Big(Z_\lambda(1) - Z_\lambda(0)\Big) \;+\; \rho\mathbb{E}[v(y)|R_i]. \end{aligned} \tag{2.52}$$

This is a differential equation, which we can solve for $\psi(f)$. To do so, we need to evaluate

$\mathbb{E}[v(y)|R_i]$. For a given value of R, the number of crypto users y is:

$$y(R) = \begin{cases} 0, & \text{if } R \leq \hat{R}-\sigma, \\ \frac{R-\hat{R}+\sigma}{2\sigma}, & \text{if } \hat{R}-\sigma < R < \hat{R}+\sigma, \\ 1, & \text{if } R \geq \hat{R}+\sigma. \end{cases} \tag{2.53}$$

For $R_i > \hat{R}+2\sigma$, we have $\mathbb{E}[v(y)|R_i] = V$. For $\hat{R} < R_i \leq \hat{R}+2\sigma$, we have:

$$\begin{aligned} \mathbb{E}[v(y)|R_i] &= \int_{R_i-\sigma}^{\hat{R}+\sigma} \frac{1}{2\sigma} v\Big(\frac{R-\hat{R}+\sigma}{2\sigma}\Big)\mathrm{d}R \;+\; \int_{\hat{R}+\sigma}^{R_i+\sigma} \frac{1}{2\sigma} V \mathrm{d}R, \\ &= \int_{\frac{R_i-\hat{R}}{2\sigma}}^{1} v(t)\mathrm{d}t \;+\; \Big(\frac{R_i-\hat{R}}{2\sigma}\Big)V \;:=\; \tilde{v}(R_i). \end{aligned} \tag{2.54}$$

We can now rewrite the differential Equation (2.52), using the boundary condition $\psi(0) = \hat{R}$:

$$\begin{aligned} \psi(0) &= \hat{R}, \\ \psi'(f) &= \frac{2\sigma\rho}{(1-\rho)(Z_\lambda(1)-Z_\lambda(0))}\tilde{v}(\psi(f)), \qquad \text{for } \hat{R} < \psi(f) \leq \hat{R}+2\sigma, \\ \text{and } \psi(f) &= \frac{2\sigma\rho V}{(1-\rho)(Z_\lambda(1)-Z_\lambda(0))}(f-\psi^{-1}(\hat{R}+2\sigma))+\hat{R}+2\sigma, \end{aligned} \tag{2.55}$$

where the last condition is true for $\hat{R}+2\sigma < \psi(f) \leq 1+\sigma$. This problem has a unique solution for $\psi(f)$. Then the fee function $\phi(R)$ is given by the inverse of $\psi(f)$ over the domain $[\hat{R}, 1+\sigma]$, and is unique.

2.9.1.3 Proof of existence and uniqueness of threshold equilibrium

Result 2.9. *The threshold for crypto usage $\hat{R}$ satisfies Equation* (2.8).

We know from Results 2.6 and 2.7 that $\hat{R}$ lies in $(0,1)$. Consider a household i who observes a signal R_i exactly equal to $\hat{R}$. This household forms a posterior $R \sim U[R_i-\sigma, R_i+\sigma]$, for any σ that meets the conditions stated in the Proposition. Given R, the mass of households that use crypto is $(R+\sigma-\hat{R})/2\sigma$. By continuity, her optimal fee $\phi(\hat{R}) = 0$.

Therefore, this household's expected payoff is, given a fee f:

$$\pi_H(0;\hat{R}) \;=\; \hat{R}\int_{\hat{R}-\sigma}^{\hat{R}+\sigma}\frac{1}{2\sigma}g\Big(\frac{R+\sigma-\hat{R}}{2\sigma}\Big)\mathrm{d}R \;-\; (1-\rho)\int_{\hat{R}-\sigma}^{\hat{R}+\sigma}\frac{1}{2\sigma}Z_\lambda\Big(\frac{R+\sigma-\hat{R}}{2\sigma}\Big)\mathrm{d}R. \tag{2.56}$$

A substitution $t=(R+\sigma-\hat{R})/2\sigma$ gives:

$$\pi_H(0;\hat{R}) \;=\; \hat{R}\int_0^1 g(t)\mathrm{d}t \;-\; (1-\rho)\int_0^1 Z_\lambda(t)\mathrm{d}t. \tag{2.57}$$

Equation (2.57) illustrates the 'Laplacian prior' of the marginal household, described in Morris and Shin (2003). A household who observes a signal exactly equal to the threshold adopts a posterior in which the number of other households who use crypto is $U[0,1]$.

By continuity, we must have $\pi_H(0;\hat{R})=0$. That gives the required expression for $\hat{R}$, which is unique. Therefore, if a threshold equilibrium exists, it must have a threshold equal to $\hat{R}$. All that remains is to show this equilibrium exists.

Result 2.10. *There is an equilibrium in which households adopt a switching strategy with threshold $\hat{R}$.*

To prove this result, I need to show $\pi_H(\phi(R_i);R_i)>0$ for all $R_i>\hat{R}$, and $\pi_H(0;R_i)<0$ for all $R_i<\hat{R}$. Note that Assumption 2.2 implies $(1-\rho)Z_\lambda(1)\leq\hat{R}\leq(1-\rho)Z_\lambda(0)/g(0)$.

Consider any $R_i>\hat{R}$. If this household offers a fee of zero, her expected payoff is:

$$\begin{aligned}\pi_H(0;R_i) \;=\; &\int_{R_i-\sigma}^{\hat{R}+\sigma}\frac{1}{2\sigma}\Big(R_i g\Big(\frac{R+\sigma-\hat{R}}{2\sigma}\Big)-(1-\rho)Z_\lambda\Big(\frac{R+\sigma-\hat{R}}{2\sigma}\Big)\Big)\mathrm{d}R\\ &+\int_{\hat{R}+\sigma}^{R_i+\sigma}\frac{1}{2\sigma}\Big(R_i-(1-\rho)Z_\lambda(1)\Big)\mathrm{d}R.\end{aligned} \tag{2.58}$$

Given that $R_i>\hat{R}\geq(1-\rho)Z_\lambda(1)$, the final term is strictly positive. Therefore, using a substitution $t=(R+\sigma-\hat{R})/2\sigma$, we have:

$$\pi_H(0;R_i) \;>\; \int_{(R_i-\hat{R})/2\sigma}^1\Big(R_i g(t)-(1-\rho)Z_\lambda(t)\Big)\mathrm{d}t, \tag{2.59}$$

which is strictly positive for $R_i>\hat{R}$. A similar argument shows that the payoff is negative whenever $R_i<\hat{R}$. The proof is complete.

2.9.2 Proof of Proposition 2.2

By Proposition 2.1, we know the households employ a switching strategy with threshold $\hat{R}$ for any $\sigma > 0$, and this is the unique equilibrium. Therefore, they employ the same strategy as $\sigma \to 0$. Taking $\sigma \to 0$ in Equation (2.53), we have $y(R) \to 0$ for all $R < \hat{R}$ and $y(R) \to 1$ for all $R > \hat{R}$.

We now only need to solve for the fee strategy when $R > \hat{R}$. Looking at the solution in the case $\sigma > 0$, it is clear that the pure strategy in fees does not survive into the limiting case $\sigma \to 0$. The differential Equation (2.52) implies $\psi'(f) \to 0$ as $\sigma \to 0$ for all fees f, which would mean all households offer the same fee. But that cannot be an equilibrium, because an individual household can profitably deviate by increasing the fee by a small amount and jumping ahead of all the other households in priority.

The purification theorem (Harsanyi 1973) tells us that a mixed strategy can be interpreted as the probability of playing a pure strategy as uncertainty in payoffs tends to zero. Our problem is continuous so, if there is a unique equilibrium in mixed fee strategies for the households when $\sigma = 0$, the mixed strategy must be the limit of the pure strategies as $\sigma \to 0$. Because the households are identical when $\sigma = 0$ (they all have signal $R_i = R$), they must employ the same strategy.

Let $\bar{H}(f)$ be the complementary cumulative distribution function of the households' fee strategies. Let $\mathcal{F} \subset \mathbb{R}^+$ be the support of their strategies. The households must be indifferent between all fees in this support. As in the proof of Proposition 2.1, we can write a household's loss function in terms of f:

$$(1-\rho)Z_\lambda\big(\bar{H}(f)\big) \;+\; \rho V f \;=\; L(f), \tag{2.60}$$

where $L(f) = \hat{L}$, a constant, for all $f \in \mathcal{F}$, and $L(f) \geq \hat{L}$ for all $f \notin \mathcal{F}$. The rest of her payoff does not depend on the fee and so we do not need to include it here.

First, we show this support has an infimum at zero. Because it is a subset of $\mathbb{R}^+$, it has an infimum, which we call f'. Then $\bar{H}(f') = 1$. By continuity of Z_λ, for any $\varepsilon > 0$, there is a $f'' \in \mathcal{F}$ such that $Z_\lambda(\bar{H}(f'')) > Z_\lambda(1) - \varepsilon$. A household who plays this fee has $L(f'') = (1-\rho)(Z_\lambda(1)-\varepsilon) + \rho V f''$, since almost all of her peers pay a higher fee. But she could pay a fee of zero and earn $L(0) = (1-\rho)Z_\lambda(1)$. For small enough ε, this deviation

must be strictly profitable, unless $f'' = 0$. Thus, we must have $f' = 0$.

Next we show the equilibrium strategy has no atoms; that is, there are no points f where a household places a strictly positive probability mass. Suppose not, and there is some point f^+ where each household places a mass μ. Then the payoff here is:

$$L(f^+) \;=\; (1-\rho)\int_0^{\mu} Z_\lambda(\bar{H}(f^+)+t)\mathrm{d}t \;+\; \rho V f^+, \tag{2.61}$$

because the miners break ties by randomising among households. A household can deviate to $f^+ + \varepsilon$, for some $\varepsilon > 0$. The loss is then $(1-\rho)Z_\lambda(\bar{H}(f^+)) + \rho V(f^+ + \varepsilon)$. For small enough ε, this deviation is profitable.

Next, we show the support is a convex set. Suppose not. Then there is at least one gap in the support: there is a fee f and a number $\varepsilon > 0$ such that $f \in \mathcal{F}$ but no points in $(f-\varepsilon, f)$ are contained in $\mathcal{F}$. But then $\bar{H}(f-\varepsilon) = \bar{H}(f)$, since the strategy distribution has no atoms. Clearly, it is then profitable to deviate to $f-\varepsilon$, so this cannot be an equilibrium.

These results imply the support $\mathcal{F} = [0, \bar{f}]$ for some $\bar{f} > 0$. Putting $f = 0$ into Equation (2.60), we have:

$$\hat{L} \;=\; L(0) \;=\; (1-\rho)Z_\lambda(1). \tag{2.62}$$

Thus, for all $f \in [0, \bar{f}]$, we have:

$$(1-\rho)\Big(1 - Z_\lambda(\bar{H}(f))\Big) + \rho V f \;=\; (1-\rho)\Big(1 - Z_\lambda(1)\Big), \tag{2.63}$$

which can be solved in terms of $\bar{H}(f)$, since $Z_\lambda(\cdot)$ is an invertible function. We use $\bar{H}(\bar{f}) = 0$ to obtain the required expression for $\bar{f}$.

2.9.3 Proof of Lemma 2.1

We first look at the two thresholds. Since $0 < g(t) \leq 1$ for all $t \in [0,1]$, we have:

$$\hat{R} - \hat{R}_{2B} \;=\; (1-\rho)\Big(\int_0^1 Z_\lambda(t)\mathrm{d}t\Big)\Big(\frac{1}{\int_0^1 g(t)\mathrm{d}t} - 1\Big) \;> 0, \tag{2.64}$$

and this expression is decreasing in λ.

Given a threshold A, expected welfare is:

$$\begin{aligned}\hat{\Omega}(A) &= \int_0^A (1-\rho)\mathrm{d}R \;+\; \int_A^1 \Big[R+(1-\rho)\Big(1-\int_0^1 Z_\lambda(t)\mathrm{d}t\Big)\Big]\mathrm{d}R, \\ &= 1-\rho+\frac{1}{2}(1-A^2)-(1-\rho)(1-A)\int_0^1 Z_\lambda(t)\mathrm{d}t.\end{aligned} \tag{2.65}$$

The difference in welfare between the second-best and global games allocations is:

$$\begin{aligned}\hat{\Omega}_{2B}-\hat{\Omega} &= \hat{\Omega}(\hat{R}_{2B})-\hat{\Omega}(\hat{R}), \\ &= \frac{1}{2}\Big(\hat{R}^2-\hat{R}_{2B}^2\Big)+(1-\rho)(\hat{R}_{2B}-\hat{R})\int_0^1 Z_\lambda(t)\mathrm{d}t, \\ &= (\hat{R}-\hat{R}_{2B})\Big((1-\rho)\int_0^1 Z_\lambda(t)\mathrm{d}t+\frac{1}{2}(\hat{R}+\hat{R}_{2B})\Big).\end{aligned} \tag{2.66}$$

Every term in the final line of Equation (2.66) is decreasing in λ. Therefore, $\hat{\Omega}_{2B}-\hat{\Omega}$ is decreasing in λ, and the welfare loss rises as the blockchain capacity goes down, as required.

2.9.4 Proof of Lemma 2.2

2.9.4.1 Threshold and switching strategies

Just as in the case with households only, a household with signal exactly equal to the threshold has an expected payoff of zero and offers a fee of zero. This household has the following expected payoff from using crypto:

$$\pi_H(0;R^*) \;=\; R^*\int_0^1 g(t)\mathrm{d}t \;-\; (1-\rho)\int_0^1 Z_\lambda\big(t(1+M)\big)\mathrm{d}t. \tag{2.67}$$

Note that the $Z_\lambda(t(1+M))$ term comes from the fact that the households and speculators have the same signal accuracy. See Section 2.5.2.1 for a discussion of the case where they have different signal accuracies.

Setting $\pi_H(0;R^*)=0$ gives the required expression for R^*. Assumption 2.1 ensures R^* lies in the range $(0,1)$.

Consider a speculator with signal below R^*. This speculator knows that, as $\sigma \to 0$, the probability that $y > 0$ tends to zero, so the value of crypto is almost surely zero. Thus, his expected payoff from buying is equal to the negative of the price, and he does not buy.

Consider a speculator with signal above R^*. This speculator knows $y = 1$ and thus $v(y) = V$ for sure. The expected consumption value of crypto is K, given that timing of consumption is uncertain. His expected payoff from buying is equal to K minus the average price. He knows all other speculators buy so, given the distribution of liquidity trades, the expected price is equal to $(p(0) + p(M))/2$. As long as $p(M) = K > p(0)$, the trade is profitable and so it is optimal for the speculator to buy. If $p(0) = K$, the expected payoff is zero and the speculator does not trade. I will rule out that case in the following section.

I have shown that, in any threshold equilibrium, both households and speculators employ the same switching strategy as long as $p(0) < K$. This threshold must be equal to R^*.

2.9.4.2 Pricing function

Consider a market maker who believes that all other agents follow the strategies described in the Lemma. If she observes order flow $z = M$, then she knows the speculators all observe a signal greater than R^*. Given that the expected consumption value of crypto for these speculators is K, she sets a price $p(M) = K$.

If the market maker observes $z = -M$, she knows no speculators placed an order. This implies $R < R^*$, and so no households use crypto. Thus, $y = 0$ for sure, and crypto has zero consumption value. She sets a price $p(-M) = 0$.

If the market maker observes $z = 0$, all she can discern is that $u = -x$, and the liquidity trades balance out the informed trades. Because the two liquidity trade outcomes occur with equal probabilities, her posterior is the same as her prior. With probability R^*, $R < R^*$ and crypto has value 0 and, with probability $1 - R^*$, it has value K. Thus, she sets price $p(0) = (1 - R^*)K$.

Finally, suppose the market maker observes $z = M/2$ or $z = -M/2$. She infers $x = M/2$. These order flows can occur only if $R = R^*$ exactly, because then exactly half of speculators observe signals higher than R^* and place buy orders, and exactly half observe lower

signals and do not buy. A marginal speculator who observes a signal of exactly R^* has a Laplacian prior and uses a uniform distribution for the proportion of households and speculators who observe fees above R^*. We also know that, by continuity, this speculator pays a fee of exactly zero. This speculator's expected value of crypto is therefore:

$$\int_0^1 v(t)\Big(1-(1-\rho)Z_\lambda\Big(t(1+M)\Big)\Big)\mathrm{d}t, \tag{2.68}$$

and the market maker sets her price equal to this quantity. I have shown that, in any threshold equilibrium, the pricing function must be as given in Equation (2.24).

This result relies on the maintained assumption that the speculators employ a switching strategy at R^*. I can now show this must be the case. Suppose the speculator does not switch at R^*. In Section 2.9.4.1, I showed the speculator switches at R^* as long as $p(0) < K$, so we must have $p(0) \geq K$. The market maker cannot set a price strictly higher than K, because that would be inconsistent with Equation (2.20), so we must have $p(0) = K$. When $z = M$, the market maker knows the speculator has placed a buy order, so $p(M) = K$. Thus, the speculator knows he always faces a price of K. His expected payoff is then $\pi_S = K - K = 0$, and so he prefers not to trade, regardless of his signal. But then Equation (2.20) implies the market maker should set a price equal to her prior $(1-R^*)K$, regardless of the observed order flow z. We have a contradiction, and so we can conclude that each speculator adopts a switching strategy at R^*.

2.9.4.3 Fee strategies

Because the households are ex ante identical to each other, they must employ the same strategies. As $\sigma \to 0$, the households' signals converge and so do their strategies. Then their fee strategies cannot be pure. If all households offered the same fee, an individual household could profitably deviate from this strategy by offering a slightly higher fee and winning top priority. Thus, all households use the same mixed strategies for fees given R, just as in Proposition 2.2. The same is true for speculators.

No agents employ a fee strategy with point masses. If, say, households place a point mass with measure $\mu > 0$ at some fee f', then an individual household can profitably deviate by offering a slightly higher fee $f' + \varepsilon$ for some $\varepsilon > 0$. This deviation would increase

the probability of settlement by a non-zero amount, at the cost of an arbitrarily small fee. Thus, there can be no point masses.

Define $\mathcal{F}_{\mathcal{H}}$ to be the smallest set that contains the support of the households' fee strategies and can be written as the union of disjoint closed sets. This set exists and is unique. We can define $\mathcal{F}_{\mathcal{S}}$ similarly for the support of the speculators' fee strategies. I shall show that, in a threshold equilibrium, $\mathcal{F}_{\mathcal{H}}$ and $\mathcal{F}_{\mathcal{S}}$ are unique. This implies the threshold equilibrium is unique, up to deviations with Lebesgue measure zero. Such deviations are rather trivial and do not affect the pricing function or expected outcomes.

Consider $\mathcal{F} := \mathcal{F}_{\mathcal{H}} \cup \mathcal{F}_{\mathcal{S}}$, the union of the supports of the households' and speculators' fee strategies. I will show $\mathcal{F}$ is an interval containing 0.

I first claim $\mathcal{F}$ has full Lebesgue measure; that is, given any $f_1, f_2 \in \mathcal{F}$, the set of points $f \in [f_1, f_2]$ that are not contained in $\mathcal{F}$ has Lebesgue measure zero. Suppose not. Then there exists a pair f', f'' such that $f' < f''$ and $f', f'' \in \mathcal{F}$, but almost all points in the open set (f', f'') are not in $\mathcal{F}$. But then there exists a neighbourhood around f'' that cannot be part of an optimal strategy for any agent: offering f' achieves the same probability of settlement and is cheaper. Thus, $\mathcal{F}$ must have full Lebesgue measure.

I claim $0 \in \mathcal{F}$. Suppose not. Let s be the infimum of $\mathcal{F}$. An agent that offers s ranks below everybody else, and so has probability of settlement $Z_\lambda(1+M)$. This agent can do better by offering a fee of zero: such an offer achieves the same probability of settlement and costs less. Thus, zero must be contained in $\mathcal{F}$. Taking these results together, I have shown $\mathcal{F}$ is an interval containing 0, as desired.

Now define $\mathcal{G} := \mathcal{F}_{\mathcal{H}} \cap \mathcal{F}_{\mathcal{S}}$, the intersection of supports of the households' and speculators' fee strategies. I claim that $\mathcal{G}$ contains at most two distinct points. A household is indifferent between all $f \in \mathcal{F}_{\mathcal{H}}$, so:

$$(1-\rho)Z_\lambda\Big(\bar{H}(f) + M\bar{S}(f)\Big) + \rho V f \;:=\; L_H, \qquad \text{for all } f \in \mathcal{F}_{\mathcal{H}}, \tag{2.69}$$

where $\bar{H}(f)$ and $\bar{S}(f)$ are the complementary cumulative distribution functions of the strategies of the households and speculators, respectively. Similarly, because speculators

are indifferent between all fees in their support, we must have:

$$(1-f)\Big(1-(1-\rho)Z_\lambda\Big(\bar{H}(f)+M\bar{S}(f)\Big)\Big) \;:=\; L_S, \qquad \text{for all } f\in\mathcal{F}_\mathcal{S}. \tag{2.70}$$

Putting together Equations (2.69) and (2.70), we have, for all $f\in\mathcal{G}$:

$$0 \;=\; (1-f)(1-L_H+\rho V f)-L_S \;:=\; Q(f). \tag{2.71}$$

Equation (2.71) is a quadratic in f, so it can have at most two real roots. Therefore, $\mathcal{G}$ contains at most two distinct points, given by the solution $Q(f)=0$. When $Q(f)<0$, then $f\in\mathcal{F}_\mathcal{H}$ and $\notin\mathcal{F}_\mathcal{S}$. When $Q(f)>0$, then $f\in\mathcal{F}_\mathcal{S}$ and $\notin\mathcal{F}_\mathcal{H}$.

Suppose the polynomial $Q(f)$ has no roots in $[0,\infty)$. Then $\mathcal{G}$ is empty. $Q(f)$ describes an upside-down parabola, which implies $Q(0)<0$ everywhere, so $\mathcal{F}_\mathcal{S}$ must be empty. Therefore, the speculators never buy. But we know that this cannot be consistent with an equilibrium, since $R^*\in(0,1)$ and it is optimal for the speculators to play a switching strategy around R^*. Thus, $\mathcal{G}$ must contain at least one point, and $Q(f)$ must have at least one non-negative root.

Suppose $Q(f)$ has exactly one root in $[0,\infty)$. The root cannot be at exactly zero, since then the speculator plays a pure strategy at zero, which we know cannot describe an equilibrium. Let the root be $f'>0$. For $f<f'$, $Q(f)>0$ and for $f>f$, $Q(f)<0$. Then there must be fees α,β such that $0<\alpha<\beta$, $\mathcal{F}_\mathcal{S}=[0,\alpha]$, and $\mathcal{F}_\mathcal{H}=[\alpha,\beta]$.

Now suppose $Q(f)$ has exactly two roots in $(0,\infty)$. Then $Q(0)<0$, so zero lies in $\mathcal{F}_\mathcal{H}$, and one of the following must be true:

- there exist fees α,β such that $0<\alpha<\beta$, $\mathcal{F}_\mathcal{H}=[0,\alpha]$, and $\mathcal{F}_\mathcal{S}=[\alpha,\beta]$;
- there exist fees α,β,γ such that $0<\alpha<\beta<\gamma$, $\mathcal{F}_\mathcal{H}=[0,\alpha]\cup[\beta,\gamma]$, and $\mathcal{F}_\mathcal{S}=[\alpha,\beta]$.

The first case can occur because there is a hard limit on the fee that speculators can offer: they only have access to one unit of crypto.

Finally, suppose $Q(f)$ has one root at zero and another root in $(0,\infty)$. Then there are fees α,β such that $0<\alpha<\beta$, $\mathcal{F}_\mathcal{S}=\{0\}\cup[\alpha,\beta]$, and $\mathcal{F}_\mathcal{H}=[0,\alpha]$. This case is equivalent to the case in the first bullet point above, since it differs only over a set of Lebesgue measure zero.

To sum up, in any equilibrium, the supports must take one of the following forms:

1. households pay higher fees than speculators: $\mathcal{F}_{\mathcal{S}} = [0,\alpha]$ and $\mathcal{F}_{\mathcal{H}} = [\alpha,\beta]$;
2. households pay lower fees than speculators: $\mathcal{F}_{\mathcal{H}} = [0,\alpha]$, and $\mathcal{F}_{\mathcal{S}} = [\alpha,\beta]$;
3. household and speculators pay similar fees: $\mathcal{F}_{\mathcal{H}} = [0,\alpha]\cup[\beta,\gamma]$ and $\mathcal{F}_{\mathcal{S}} = [\alpha,\beta]$.

These are the three classes of fee equilibria described in the Lemma. I look at each in turn, and show that each equilibrium occurs if and only if the relevant condition on ρV given in the Lemma holds. For each equilibrium, I write out the fee strategies explicitly.

Equilibrium of type 1: households pay higher fees than speculators

Suppose there are fee levels α,β such that $0<\alpha<\beta$, the speculators' support is $\mathcal{F}_{\mathcal{S}} = [0,\alpha]$, and the households' support is $\mathcal{F}_{\mathcal{H}} = [\alpha,\beta]$. Because the speculators must be indifferent between all fees in their support, we must have, for all $f \in \mathcal{F}_{\mathcal{S}}$, $\pi_S(f) = \pi_S(0)$; that is:

$$(1-f)\Big(1-(1-\rho)Z_\lambda(1+M\bar{S}(f))\Big) \;=\; 1-(1-\rho)Z_\lambda(1+M). \tag{2.72}$$

As $\bar{S}(\alpha)=0$, Equation (2.72) implies:

$$1-\alpha \;=\; \frac{1-(1-\rho)Z_\lambda(1+M)}{1-(1-\rho)Z_\lambda(1)}. \tag{2.73}$$

Note that $0<\alpha<1$, so the speculators' budget constraint does not bind. For all $f \in \mathcal{F}_{\mathcal{H}}$, we must have $\pi_H(f) = \pi_H(\alpha)$, so:

$$(1-\rho)Z_\lambda(\bar{H}(f)) + \rho V f \;=\; (1-\rho)Z_\lambda(1) + \rho V\alpha. \tag{2.74}$$

As $\bar{H}(\beta)=0$, Equation (2.74) implies:

$$\beta \;=\; \alpha + \frac{1-\rho}{\rho V}\Big(Z_\lambda(1)-Z_\lambda(0)\Big). \tag{2.75}$$

Thus, if this is an equilibrium, the fee strategies and their supports are uniquely defined. To verify this is in equilibrium, I must check that no household or speculator has an incentive to deviate to a fee outside of his or her strategy support.

The no-deviation condition for a speculator is, for all $f \in \mathcal{F}_{\mathcal{H}}$, $\pi_S(f) \leq \pi_S(\alpha)$. That condition is equivalent to, for all $f \in \mathcal{F}_{\mathcal{H}}$:

$$\begin{aligned}
& (1-f)\Big(1-(1-\rho)Z_\lambda(\bar{H}(f))\Big) \;\leq\; (1-\alpha)\Big(1-(1-\rho)Z_\lambda(1)\Big), \\
\Leftrightarrow \qquad & (1-f)\Big(1-(1-\rho)Z_\lambda(1)-\rho V\alpha+\rho Vf\Big) \;\leq\; (1-\alpha)\Big(1-(1-\rho)Z_\lambda(1)\Big), \\
\Leftrightarrow \qquad & \rho V(1-f)(f-\alpha) \;\leq\; (f-\alpha)\Big(1-(1-\rho)Z_\lambda(1)\Big), \\
\Leftrightarrow \qquad & \rho V(1-f) \;\leq\; 1-(1-\rho)Z_\lambda(1),
\end{aligned} \tag{2.76}$$

since $f \geq \alpha$ for all $f \in \mathcal{F}_{\mathcal{H}}$. The second line is derived using Equation (2.74).

Note the expression on the left-hand side of the final line of expression (2.76) is decreasing in f. Therefore, for this condition to be true for all $f \in FH$, it is necessary and sufficient for it to be true for $f = \alpha$. The no-deviation condition for the speculator is therefore:

$$\rho V(1-\alpha) \;\leq\; 1-(1-\rho)Z_\lambda(1). \tag{2.77}$$

A similar argument applies for the household. Her no-deviation condition is, for all $f \in \mathcal{F}_S$, $\pi_H(f) \leq \pi_H(\alpha)$. That condition is equivalent to, for all $f \in \mathcal{F}_S$:

$$\begin{aligned}
& (1-\rho)Z_\lambda(1+M\bar{S}(f))+\rho Vf \;\geq\; (1-\rho)Z_\lambda(1)+\rho V\alpha, \\
\Leftrightarrow \qquad & 1-\frac{1-\alpha}{1-f}\Big(1-(1-\rho)Z_\lambda(1)\Big)+\rho Vf \;\geq\; (1-\rho)Z_\lambda(1)+\rho V\alpha, \\
\Leftrightarrow \qquad & (\alpha-f)\Big(1-(1-\rho)Z_\lambda(1)\Big) \;\geq\; \rho V(\alpha-f)(1-f), \\
\Leftrightarrow \qquad & 1-(1-\rho)Z_\lambda(1) \;\geq\; \rho V(1-f).
\end{aligned} \tag{2.78}$$

For this condition to be true for all $f \in \mathcal{F}_S$, it is necessary and sufficient for it to be true at $f = 0$. Thus, the no-deviation condition for the household is:

$$\rho V \;\leq\; 1-(1-\rho)Z_\lambda(1). \tag{2.79}$$

Note condition (2.79) implies condition (2.77), so we have an equilibrium if and only if condition (2.79) holds. This is the required outcome.

The speculator's expected payoff from trading is $K = (1-(1-\rho)Z_\lambda(1+M))V$ minus the price. The market maker uses this value of K to set the price.

Equilibrium of type 2: households pay lower fees than speculators

Now suppose that there are fee levels α, β such that $0 < \alpha < \beta$, and the households' support is $\mathcal{F}_\mathcal{H} = [0, \alpha]$ and the speculators' support is $\mathcal{F}_\mathcal{S} = [\alpha, \beta]$. Because the households must be indifferent between all fees in their support, we must have, for all $f \in \mathcal{F}_\mathcal{H}$, $\pi_H(f) = \pi_H(0)$; that is:

$$(1-\rho)Z_\lambda(\bar{H}(f)+M)+\rho V f \;=\; (1-\rho)Z_\lambda(1+M). \tag{2.80}$$

As $\bar{H}(\alpha) = 0$, Equation (2.80) implies:

$$\alpha \;=\; \frac{1-\rho}{\rho V}\Big(Z_\lambda(1+M) - Z_\lambda(M)\Big). \tag{2.81}$$

For all $f \in \mathcal{F}_\mathcal{S}$, we must have $\pi_S(f) = \pi_S(\alpha)$, so:

$$(1-f)\Big(1-(1-\rho)Z_\lambda(M\bar{S}(f))\Big) \;=\; (1-\alpha)\Big(1-(1-\rho)Z_\lambda(M)\Big). \tag{2.82}$$

As $\bar{S}(\beta) = 0$, Equation (2.82) implies:

$$1-\beta \;=\; (1-\alpha)\frac{(1-(1-\rho)Z_\lambda(M))}{1-(1-\rho)Z_\lambda(0)}. \tag{2.83}$$

Thus, if this is an equilibrium, the fee strategies and their supports are uniquely defined. As before, to verify this is an equilibrium, I must check that no household or speculator has an incentive to deviate to a fee outside of his or her strategy support. In addition, I must also check that $0 < \alpha < \beta < 1$, since a speculator cannot pay a fee greater than 1.

I check the no-deviation conditions using a very similar process to that used for the equilibrium of type 1, so for brevity I do not include all of the stages of algebra. The no-deviation condition for a household is, for all $f \in \mathcal{F}_\mathcal{S}$, $\pi_H(f) \leq \pi_H(\alpha)$. That condition is equivalent to $1-(1-\rho)Z_\lambda(M) \leq \rho V(1-f)$. For this to hold for all $f \in \mathcal{F}_\mathcal{S}$, it is necessary and sufficient for it to hold for $f = \beta$. Thus, the no-deviation condition for a household is $1-(1-\rho)Z_\lambda(M) \leq \rho V(1-\beta)$.

The no-deviation condition for a speculator is, for all $f \in \mathcal{F}_{\mathcal{H}}$, $1-(1-\rho)Z_\lambda(M) \leq \rho V(1-f)$. For this to hold for all $f \in \mathcal{F}_{\mathcal{H}}$, it is necessary and sufficient for it to hold for $f = \alpha$. Thus, the no-deviation condition for a speculator is $1-(1-\rho)Z_\lambda(M) \leq \rho V(1-\alpha)$, which is strictly weaker than the no-deviation condition for a household.

Finally, I need to check that $0 < \alpha < \beta < 1$. From the definitions of α and β, it is clear that $\alpha > 0$ and $\beta > \alpha$. Thus, I need only check that $\alpha < 1$. Suppose $\alpha \geq 1$. Then $\beta > 1$, so $\rho V(1-\beta) < 0$. Then the no-deviation condition for the household cannot hold. Thus, $\alpha < 1$ is implied by the household's no-deviation condition.

I have shown this equilibrium exists if and only if $1-(1-\rho)Z_\lambda(M) \leq \rho V(1-\beta)$, which is equivalent to:

$$\rho V \;\geq\; 1+(1-\rho)\Big(Z_\lambda(1+M)-Z_\lambda(M)-Z_\lambda(0)\Big), \tag{2.84}$$

which is the required condition.

The speculator's expected payoff from trading is $K = (1-\alpha)(1-(1-\rho)Z_\lambda(M))V$ minus the price. The market maker uses this value of K to set the price.

Equilibrium of type 3: households pay similar fees to speculators

Suppose there are fee levels α, β, γ such that $0 < \alpha < \beta < \gamma$, and the households' support is $\mathcal{F}_{\mathcal{H}} = [0,\alpha] \cup [\beta,\gamma]$ and the speculators' support is $\mathcal{F}_{\mathcal{S}} = [\alpha,\beta]$. I proceed using a similar method to the other equilibrium types.

For all $f \in \mathcal{F}_{\mathcal{H}}$, households are indifferent, so:

$$(1-\rho)Z_\lambda\Big(\bar{H}(f)+M\bar{S}(f)\Big)+\rho V f \;=\; (1-\rho)Z_\lambda(1+M). \tag{2.85}$$

As $\bar{H}(\gamma) = \bar{S}(\gamma) = 0$, we have:

$$\gamma \;=\; \frac{1-\rho}{\rho V}\Big(Z_\lambda(1+M)-Z_\lambda(0)\Big). \tag{2.86}$$

For all $f \in \mathcal{F}_S$, speculators are indifferent so:

$$(1-f)\Big(1-(1-\rho)Z_\lambda\Big(\bar{H}(f)+M\bar{S}(f)\Big)\Big) \;=\; (1-\alpha)\Big(1-(1-\rho)Z_\lambda\Big(\bar{H}(\alpha)+M\Big)\Big). \tag{2.87}$$

Since both α and β are members of both $\mathcal{F}_\mathcal{H}$ and $\mathcal{F}_S$, we have:

$$(1-\rho)Z_\lambda\Big(\bar{H}(\alpha)+M\Big)+\rho V\alpha \;=\; (1-\rho)Z_\lambda\Big(\bar{H}(\alpha)\Big)+\rho V\beta, \tag{2.88}$$

and:

$$(1-\alpha)\Big(1-(1-\rho)Z_\lambda\Big(\bar{H}(\alpha)+M\Big)\Big) \;=\; (1-\beta)\Big(1-(1-\rho)Z_\lambda\Big(\bar{H}(\alpha)\Big)\Big). \tag{2.89}$$

We can eliminate α and β from Equations (2.88) and (2.89) to obtain:

$$\rho V-1 \;=\; (1-\rho)\Big(Z_\lambda(1+M)-Z_\lambda(\bar{H}(\alpha)+M)-Z_\lambda(\bar{H}(\alpha))\Big). \tag{2.90}$$

The right-hand side of Equation (2.90) is monotonically decreasing in $\bar{H}(\alpha)$. Thus, it has at most one solution in $\bar{H}(\alpha) \in [0,1]$. It has no solution if $\rho V-1 > (1-\rho)(Z_\lambda(1+M)-Z_\lambda(M)-Z_\lambda(0))$, but then we have an equilibrium of type 2. It also has no solution if $\rho V-1 < -(1-\rho)Z_\lambda(1)$, but then we have an equilibrium of type 1. Thus, the conditions for the three equilibria are mutually exclusive and exhaustive, so one of them exists.

Given the solution for $\bar{H}(\alpha)$, we can solve for α and β using Equations (2.88) and (2.89), and then solve for the strategy functions $\bar{H}(f)$ and $\bar{S}(f)$. There is a unique solution. Given the condition 2.90, the agents' no-deviation conditions always hold, so that condition is necessary and sufficient for an equilibrium of this type. The proof is complete.

The speculator's expected payoff from trading is $K=(1-\alpha)(1-(1-\rho)Z_\lambda(\bar{H}(\alpha)+M))V$ minus the price. The market maker uses this value of K to set the price.

2.9.5 Proof of Proposition 2.4

The pricing function $p(z)$ is given by Lemma 2.2 as

$$p(z) = \begin{cases} 0, & \text{with probability } \frac{1}{2}R^*, \\ (1-R^*)K, & \text{with probability } \frac{1}{2}, \\ K, & \text{with probability } \frac{1}{2}(1-R^*). \end{cases} \tag{2.91}$$

By Equation (2.23), R^* is decreasing in λ and increasing in M. The results for $p(z)/K$ follow immediately.

Suppose $V \leq 1$. Then $\rho V \leq \rho \leq 1-(1-\rho)Z_\lambda(1)$ for all finite λ. Therefore, any threshold equilibrium must be of the first type described in Lemma 2.2. From the proof of Lemma 2.2, we know the support for a speculator's fee strategy contains zero (or, at least, zero is a limit point of the speculator's support). When the speculator sets a fee of zero, he has the lowest priority in the blockchain, and so his payment is settled by T_1 only if $N \geq 1+M$. Thus, the value from holding a unit of crypto is $K = 1-(1-\rho)Z_\lambda(1)$, which is increasing in λ and constant in M.

We know $1-R^*$ and K are both increasing in λ and decreasing in M. This implies that, for any q, $\mathbb{P}[p(z) > q]$ is non-decreasing in λ and non-increasing in M.

2.9.6 Proof of Lemma 2.3

Suppose the probability of settlement by T_1 is fixed, and does not depend on block rate λ or number of speculators M. Let it be equal to Q, an exogenous constant. Because the probability of early delivery is fixed, nobody offers a fee. The problem is reduced to a standard global games model. Equation (2.23) implies the threshold for crypto usage is:

$$R^* = (1-\rho)\frac{Q}{\int_0^1 g(t)\mathrm{d}t}, \tag{2.92}$$

which is independent of λ and M.

The consumption value of holding a unit of crypto at T_1 is:

$$K = 1-(1-\rho)Q, \tag{2.93}$$

which is also independent of λ and M.

The price $p(z)$ is given by:

$$p(z) = \begin{cases} 0, & \text{with probability } \frac{1}{R^*}, \\ (1-R^*)K, & \text{with probability } \frac{1}{2}, \\ K, & \text{with probability } \frac{1}{2}(1-R^*). \end{cases} \tag{2.94}$$

The distribution of the price is therefore independent of λ and M, as required.

2.9.7 Proof of Lemma 2.4

Suppose $R > R^*$ and $V \leq 1$. When $V \leq 1$, we are in an equilibrium of type 1, so the value of crypto to a speculator at T_1 is $K = 1-(1-\rho)Z_\lambda(1+M)$. The expected price of purchasing crypto is $\frac{1}{2}(p(0)+p(M)) = (1-\frac{1}{2}R^*)K$, using the pricing function in Lemma 2.2. Thus, the speculator's expected payoff from trading is $\pi_S = \frac{1}{2}R^*K$. Because R^* is increasing in M and K is constant, the speculator's profit is increasing in M. Note that unconditional profit is $\frac{1}{2}R^*(1-R^*)K$, so the effect of M on unconditional profit is ambiguous. And in equilibria of types 2 and 3, even the conditional effect is ambiguous.

The effect of λ on speculators' profit is also ambiguous. An increase in λ lowers R^* and raises K, so the overall trade-off will depend on the slope of $Z_\lambda(\cdot)$ and the strength of the network effects $\int_0^1 g(t)\mathrm{d}t$.

Now I look at the second part of the Lemma. As $\lambda \to \infty$, $Z_\lambda(n) \to 0$ for all finite n, so $R^* \to 0$ and $K \to 1$. Thus, the speculator's payoff $\pi_S = \frac{1}{2}R^*K \to 0$, as $\lambda \to \infty$. The payoff will tend to zero in the limit regardless of whether we are in an equilibrium of type 1, 2, or 3. The proof is complete.

2.9.8 Proof of Proposition 2.5

The mean and variance of the price are given by Equation (2.26). Dividing the variance by the mean and taking the square root gives:

$$\Gamma_U \;=\; \sqrt{\frac{\frac{1}{2}R^*}{1-R^*}}. \tag{2.95}$$

This gives the required expression for unconditional volatility Γ_U in Equation (2.29).

I now turn to conditional volatility. Conditional on $R < R^*$, the speculators do not trade, so $z = -M$ or 0 with probability $\frac{1}{2}$ each. Thus:

$$\mathbb{E}[p(z)|R<R^*] \;=\; \frac{1}{2}(1-R^*)K, \qquad \mathbb{V}[p(z)|R<R^*] \;=\; \frac{1}{4}(1-R^*)^2K^2, \tag{2.96}$$

so $\Gamma_C(R) = 1$ here.

Conditional on $R > R^*$, the speculators all trade, so $z = M$ or 0 with probability $\frac{1}{2}$ each. Thus:

$$\mathbb{E}[p(z)|R>R^*] \;=\; \Big(1-\frac{1}{2}R^*\Big)K, \qquad \mathbb{V}[p(z)|R>R^*] \;=\; \frac{1}{4}R^{*2}K^2, \tag{2.97}$$

so $\Gamma_C(R) = R^*/(2-R^*)$ here, as required.

Finally, we consider the case $R = R^*$. In this case, half of speculators receive signals above the threshold and buy, while the other half do not. The order flow is thus $z = \pm M/2$ and, in either case, the market maker is able to discern that $R = R^*$. She therefore sets price equal to $\int_0^1 v(t)K\mathrm{d}t = \int_0^1 v(t)(1-(1-\rho)Z_\lambda(t(1+M)))\mathrm{d}t$, as given in Equation (2.24), and conditional volatility is zero.

The proof is complete.

2.9.9 Proof of Proposition 2.6

I reuse the notation from Section 2.3.3. Let $\Omega(R,y,N)$ be total welfare given strength of technology R, crypto usage y, and blockchain capacity N. Then, as before:

$$\begin{aligned}\Omega(R,y,N) &= Rg(y)y + (1-\rho)\mathbb{P}[\text{household has crypto payment settled at } T_1] \\ &\quad + (1-\rho)(1-y).\end{aligned} \tag{2.98}$$

The probability of settlement for a household depends on which type of equilibrium we are in. In a type 1 equilibrium — when the households offer higher fees than the speculators — the probability of settlement is obviously higher than in the other equilibria, when the speculators may offer higher fees than the households. Let us consider each in turn.

Type 1 equilibrium

In this equilibrium, households offer higher fees than the speculators. Given this, miners prioritise households over speculators, so the probability of an individual household achieving early settlement is not directly affected by M. The payment is settled for sure if the amount of blockchain space is sufficient for all of the households' payments; that is, if $N \geq 1$. Otherwise, the probability of settlement is N/y, or zero if $N \leq 0$. This is the same as in the case with no financial market, so $\Omega(R,y,N)$ is given by Equation (2.11). Equation (2.12) gives expected welfare given R and y. Overall expected welfare is then:

$$\begin{aligned}\hat{\Omega}(R^*) &= \int_0^{R^*}(1-\rho)\mathrm{d}R + \int_{R^*}^1 \Big[R + (1-\rho)\Big(1 - \int_0^1 Z_\lambda(t)\mathrm{d}t\Big)\Big]\mathrm{d}R, \\ &= 1-\rho+\frac{1}{2}(1-R^{*2}) - (1-\rho)(1-R^*)\int_0^1 Z_\lambda(t)\mathrm{d}t.\end{aligned} \tag{2.99}$$

Then:

$$\frac{\partial\hat{\Omega}(R^*)}{\partial R^*} = (1-\rho)\int_0^1 Z_\lambda(t)\mathrm{d}t - R^*. \tag{2.100}$$

This partial derivative is negative, because:

$$R^* = (1-\rho)\frac{\int_0^1 Z_\lambda(t(1+M))\mathrm{d}t}{\int_0^1 g(t)\mathrm{d}t} > (1-\rho)\int_0^1 Z_\lambda(t)\mathrm{d}t. \tag{2.101}$$

Welfare in the case without a financial market is the same, with R^* replaced by $\hat{R}$. Because $\hat{R} < R^*$ and welfare is decreasing in R^*, welfare is lower with a financial market than without it.

Note that:

$$\frac{\mathrm{d}\hat{\Omega}(R^*)}{\mathrm{d}\lambda} = \frac{\partial\hat{\Omega}(R^*)}{\partial R^*}\frac{\partial R^*}{\partial\lambda} - (1-\rho)(1-R^*)\int_0^1 \frac{\mathrm{d}}{\mathrm{d}\lambda}Z_\lambda(t)\mathrm{d}t, \tag{2.102}$$

By Proposition 2.3, we know that R^* is decreasing in λ. As $Z_\lambda(t)$ is pointwise decreasing in λ, Equation (2.102) must be positive. Thus, welfare is increasing in λ. Proposition 2.3 also tells us that R^* is increasing in M, so welfare is decreasing in M.

Type 2 equilibrium

In a type 2 equilibrium, the speculators offer higher fees than the households and so are prioritised by the miners. This reduces the probability of a household's crypto payment being settled by T_1, so welfare is lower than in a type 1 equilibrium. Recall that miners are included in the welfare calculation, so lower fees offered by households serve only to reduce welfare, because they make crypto less useful as a means of payment. The probability of a household's payment being settled at T_1 is:

$$\mathbb{P}[\text{household has crypto payment settled at } T_1] = \begin{cases} 0, & \text{if } N \leq M, \\ \frac{N-M}{y}, & \text{if } M < N < y+M, \\ 1, & \text{if } N \geq y+M. \end{cases} \tag{2.103}$$

Welfare given R and y is therefore:

$$\Omega(R,y) = Rg(y)y + (1-\rho)\Big(1 - \int_M^{y+M} Z_\lambda(t)\mathrm{d}t\Big), \tag{2.104}$$

which is decreasing in M. Overall expected welfare is:

$$\hat{\Omega}(R^*) = 1 - \rho + \frac{1}{2}(1-R^{*2}) - (1-\rho)(1-R^*)\int_M^{1+M} Z_\lambda(t)\mathrm{d}t. \tag{2.105}$$

This is lower than welfare in an equilibrium of type 1, and so it is lower than welfare in

the absence of a financial market. Taking derivatives with respect to R^*:

$$\frac{\partial \hat{\Omega}(R^*)}{\partial R^*} = (1-\rho)\int_0^1 Z_\lambda(t+M)\mathrm{d}t - R^*, \tag{2.106}$$

which is negative, because:

$$R^* = (1-\rho)\frac{\int_0^1 Z_\lambda(t(1+M))\mathrm{d}t}{\int_0^1 g(t)\mathrm{d}t} \geq (1-\rho)Z_\lambda(1+M), \tag{2.107}$$

by Assumption 2.3. Because welfare is decreasing in R^*, it must be increasing in λ and decreasing in M, since:

$$\frac{\mathrm{d}\hat{\Omega}(R^*)}{\mathrm{d}R^*} = \frac{\partial \hat{\Omega}(R^*)}{\partial R^*}\frac{\partial R^*}{\partial M} - (1-\rho)(1-R^*)\Big(Z_\lambda(1+M) - Z_\lambda(M)\Big) < 0. \tag{2.108}$$

The proof for an equilibrium of type 3 is very similar, and is omitted.

2.9.10 Proof of Lemma 2.5

Suppose first that there is no financial market. Consider a potential equilibrium where no households use crypto, so $y = 0$ for sure. If an individual household decides to deviate from this equilibrium and use crypto, she earns:

$$\pi_H(f_i; R_i) = R_i g(0) - (1-\rho)Z_\lambda(0) - \rho f_i v(0). \tag{2.109}$$

Because no other households demand blockchain space, her payment is settled as long as blockchain capacity $N > 0$. Thus, her fee does not affect her probability of settlement, and she offers a fee of zero. Her payoff from using crypto is non-positive as long as:

$$R_i \leq (1-\rho)\frac{Z_\lambda(0)}{g(0)}. \tag{2.110}$$

If $(1-\rho)Z_\lambda(0) > g(0)$, then this condition is always true for sufficiently low σ, and so no household deviates. Thus, there is an equilibrium where all households use cash.

Now consider the opening of a financial market. If the speculators believe that $y = 0$ for

sure, then $v(y) = 0$ and crypto offers zero consumption value. They will therefore never place buy orders, and will never take up any blockchain space. The market maker sets a price of zero with probability 1.

The equilibria described in Proposition 2.2 and Lemma 2.2 still exist, as long as the threshold, $\hat{R}$ and R^* respectively, is less than one. Assumption 2.1 guarantees $\hat{R} < 1$ and $R^* < 1$.

Suppose $\hat{R} \geq 1$. Because $R^* > \hat{R}$, the threshold is greater than 1, whether there is a financial market or not. With probability 1, every household uses cash. Therefore, the only possible equilibrium is the cash-only equilibrium.

Now suppose $\hat{R} < 1 \leq R^*$. Then there is no equilibrium where speculators play pure strategies in their trading decision. If speculators trade, then households do not use crypto, and speculators are better off not trading. If speculators do not trade, then households sometimes use crypto (when $R > \hat{R}$), and then it is rational for speculators to place buy orders. However, an equilibrium may exist where speculators play mixed strategies, as explained in Section 2.5.1.2.

2.9.11 Proof of Lemma 2.6

As $0 < \eta(R_i) < 1$, a household with a signal R_i randomly chooses between crypto and cash. She must therefore be indifferent between the two, so $\pi_H = 0$. Therefore, her actions must satisfy the following equation:

$$\begin{aligned}
&\frac{R_i}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma} g\Big(\frac{1}{2\sigma}\int_{R-\sigma}^{R+\sigma}\eta(R_j)\mathrm{d}R_j\Big)\mathrm{d}R \\
&\qquad = \frac{1-\rho}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma} Z_\lambda\Big(\frac{1}{2\sigma}\int_{R-\sigma}^{R+\sigma}\eta(R_j)\mathbb{1}\{\phi(R_j)>\phi(R_i)\}\mathrm{d}R_j\Big)\mathrm{d}R \qquad (2.111)\\
&\qquad\quad + \frac{\rho}{2\sigma}\phi(R_i)\int_{R_i-\sigma}^{R_i+\sigma} v\Big(\frac{1}{2\sigma}\int_{R-\sigma}^{R+\sigma}\eta(R_j)\mathrm{d}R_j\Big)\mathrm{d}R.
\end{aligned}$$

In equilibrium, $\phi(R_j)$ is an increasing function of R_j, so Equation (2.111) becomes:

$$\frac{R_i}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma} g\Big(\frac{1}{2\sigma}\int_{R-\sigma}^{R+\sigma}\eta(R_j)\mathrm{d}R_j\Big)\mathrm{d}R \;=\; \frac{1-\rho}{2\sigma}\int_{R_i-\sigma}^{R_i+\sigma} Z_\lambda\Big(\frac{1}{2\sigma}\int_{R_i}^{R+\sigma}\eta(R_j)\mathrm{d}R_j\Big)\mathrm{d}R + \frac{\rho}{2\sigma}\phi(R_i)\int_{R_i-\sigma}^{R_i+\sigma} v\Big(\frac{1}{2\sigma}\int_{R-\sigma}^{R+\sigma}\eta(R_j)\mathrm{d}R_j\Big)\mathrm{d}R. \tag{2.112}$$

A decrease in λ increases $Z_\lambda(\cdot)$ pointwise. For Equation (2.111) to hold for all R_i with $0<\eta(R_i)<1$, the following quantity must decrease:

$$\int_{R_i-\sigma}^{R_i+\sigma}\frac{1}{2\sigma}\int_{R-\sigma}^{R+\sigma}\eta(R_j)\mathrm{d}R_j. \tag{2.113}$$

But this expression is average crypto usage. The proof is complete.

2.9.12 Proof of Proposition 2.7

Consider a speculator j who observes a signal Y_j. If he does not trade, his payoff is normalised to zero. If he decides to place a buy order, his expected payoff is as given in Equation (2.17). If he decides to sell and offers a fee f_j, his payoff is:

$$\text{selling price}\times\mathbb{1}\{\text{delivery}\} \;+\; \rho v(y)\mathbb{1}\{\text{non-delivery}\} - \;\rho f v(y), \tag{2.114}$$

because the fee is forfeit whether delivery occurs or not.

Let R' be the threshold above which households use crypto. We know such a threshold must exist, due to the existence of upper and lower dominance regions. Consider a speculator's signal Y_j. If $Y_j<R'$, then the speculator believes $y=0$ with probability 1, so the consumption value of crypto is zero, and he will certainly wish to sell. If $Y_j>R^*$, then the speculator believes $y=1$ with probability 1 and he will certainly wish to buy. Therefore, the speculator adopts a switching strategy with the same threshold as the households.

Given the speculators' strategy, a household's problem is exactly the same as in Lemma 2.2. If $R_i>R'$, she competes with all other households and speculators for blockchain

space. If $R_i < R'$, she does not require blockchain space. If $R_i = R'$, the household's payoff is given by Equation (2.23), so $R' = R^*$.

When $R > R^*$, speculators and households adopt the same strategies as in Lemma 2.2. The optimal strategies do not depend on the market maker's pricing function, so the fact that the pricing function may be different when sell orders are possible makes no difference.

When $R < R^*$, each speculator attempts to sell. To do so, he needs to move crypto onto the exchange. If $N \geq M$, they will all be able to do so, because there is no demand for blockchain space from the households. Otherwise, only $\max\{N, 0\}$ of them will be able to do so. When $R = R^*$, equal numbers of speculators attempt to buy and sell, so total informed order flow is zero. Table 2.3 below summarises the informed order flows that may arise, given each realisation of the technological factor R and blockchain space N.

Table 2.3: **Feasible outcomes in financial market equilibrium when speculators can place buy or sell orders and $R \neq R^*$.**

Technological factor R	Blockchain space N	Informed order flow x	Ex ante probability
$R > R^*$	any	M	$1 - R^*$
$R < R^*$	$N \geq M$	$-M$	$R^*(1 - Z_\lambda(M))$
$R < R^*$	$0 < N < M$	$-N$	$R^*(Z_\lambda(M) - Z_\lambda(0))$
$R < R^*$	$N \leq 0$	0	$R^* Z_\lambda(0)$

When $R = R^*$, the possible realisations of x depend on the relative fees offered by sellers and buyers. Because, in equilibrium, sellers are indifferent between all fees, there are a large number of feasible outcomes when $R = R^*$. But $R = R^*$ occurs with prior probability zero, so these outcomes cannot affect the distribution of the price, and so I do not elaborate on them further.

The market maker observes N and $z = x + u$, where u is the total order flow arising from liquidity traders, and then sets her price. The value of a unit of crypto at T_1 for a speculator who buys is K, which is equal to that in Lemma 2.2.

The market maker sets a price of zero whenever she observes $z \leq -2M$ or any z that is not a multiple of M. When the market maker observes $z = -2M$, she knows $R < R^*$ for sure, and so she sets a price of zero. When she observes an order flow z that is not a multiple of M, she knows the only way this could occur is if the speculators attempted to sell but were

not all able to, and so she sets a price of zero. When she observes $z = 2M$, she knows all the speculators placed buy orders, and so sets a price equal to K.

The only other feasible value of z is 0. If $z = 0$ and $N < M$, the market maker knows the speculators cannot have placed sell orders, and so she infers $R > R^*$, and sets a price equal to K. But if $z = 0$ and $N \geq M$ then she cannot be sure whether the speculators placed buy or sell orders, and so sets a price equal to the ex ante mean $(1 - R^*)K$. This is the only case in which the market maker is not certain about the value of y.

Given the market maker's pricing function, a selling speculator knows he faces a price of zero unless $N \geq M$ and $u = M$. In this case, his payment will be settled by T_1 regardless of his fee. At the same time, if $R < R^*$, crypto is worthless at T_2. Therefore, a speculator who observes $Y_j < R^*$ is indifferent about the fee he offers, because it will not affect his payoff. (A speculator who observes $Y_j = R^*$ does place a positive value on crypto at T_2, but as that occurs with prior probability zero, the choice of fee does not affect prior expected outcomes.) A selling speculator therefore earns a payoff of $(1 - R^*)K$ with probability $1 - Z_\lambda(M)$, and zero otherwise.

All that remains is to calculate the unconditional and conditional volatilities. The distribution of the market maker's price is:

$$p(z,N) \quad \begin{cases} 0, & \text{with probability } \frac{1}{2}R^*(1+Z_\lambda(M)) \\ (1-R^*)K, & \text{with probability } \frac{1}{2}(1-Z_\lambda(M)), \\ K, & \text{with probability } \frac{1}{2}(1-R^*)(1+Z_\lambda(M)). \end{cases} \tag{2.115}$$

The mean of the price given the market maker's information is $\mathbb{E}[p(z,N)] = (1 - R^*)K$. As usual, the expected value of the ex post price is equal to the ex ante mean, because of the law of iterated expectations. The variance of the ex post price is:

$$\mathbb{V}[p(z,N)] \;=\; \frac{1}{2}R^*(1-R^*)(1+Z_\lambda(M))K^2, \tag{2.116}$$

and so unconditional volatility is as given in Equation (2.34).

If $R < R^*$, then the price is zero unless $u = M$ and $N \geq M$, so it is zero with probability $\frac{1}{2}(1 + Z_\lambda(M))$ and $(1 - R^*)K$ with probability $\frac{1}{2}(1 - Z_\lambda(M))$. If $R > R^*$, the price is K

with probability $\frac{1}{2}(1+Z_\lambda(M))$ and $(1-R^*)K$ with probability $\frac{1}{2}(1-Z_\lambda(M))$. We have the required expression for conditional volatility. The proof is complete.

2.9.13 Proof of Lemma 2.7

In the proof of Lemma 2.4, I showed that, conditional on purchasing a signal and observing $R > R^*$, the expected value of a unit of crypto to a speculator is:

$$K = 1-(1-\rho)Z_\lambda(1+M), \tag{2.117}$$

and the speculator's expected profit from trading is $\frac{1}{2}R^*K$. If the speculator observes $R < R^*$, he does not trade and so his expected profit is zero. Conditional on purchasing a signal but not yet observing it, his expected profit is therefore $\hat{\pi}(M) = \frac{1}{2}R^*(1-R^*)K$, which is the required expression.

As R^* is increasing in M and K is decreasing, the free entry condition $\hat{\pi}(M) = c$ does not, in general, have a unique solution. Since $\hat{\pi}(M)$ is the product of three monotonic and non-negative functions of M, there cannot be more than three solutions. Since $\hat{\pi}(M) > 0$ everywhere, there can be no solutions, if c is small enough. The proof is complete.

2.9.14 Proof of Lemma 2.8

Consider a threshold equilibrium with threshold R^*. We know that, if the speculators observe $R > R^*$, then the total informed order flow is $x = 1$. If they observe $R < R^*$, then $x = 0$.

Let Z be the set of all possible order flows. We describe an arbitrary liquidity trading distribution as follows: for $z \in Z$, let $k(z)$ be the density function of order flows given that $x = 1$, and let $\ell(z)$ be the density given that when $x = 0$. Then the unconditional density function of z is $k(z)(1-R^*)+\ell(z)R^*$.

I assume the total order flow z is somewhat informative for the market maker, so she tends to set higher prices when $x = 1$ than when $x = 0$. Formally, the distribution of $p(z)$ conditional on $x = 1$ first order stochastically dominates the distribution of $p(z)$

conditional on $x = 0$. Given that $p(z) \geq 0$ with probability 1, this implies:

$$\mathbb{E}[p(z)^2|x=1] \;>\; \mathbb{E}[p(z)^2|x=0]. \tag{2.118}$$

The market maker applies the following pricing function:

$$p(z) \;=\; \frac{k(z)(1-R^*)K}{k(z)(1-R^*)+\ell(z)R^*}. \tag{2.119}$$

Given this pricing function, $\mathbb{E}[p(z)] = (1-R^*)K$ as usual, and:

$$\mathbb{V}[p(z)] \;=\; \int_Z (1-R^*)^2K^2\frac{k(z)^2}{k(z))(1-R^*)+\ell(z)R^*}\mathrm{d}z \;-\; (1-R^*)^2K^2. \tag{2.120}$$

Let Γ_U be unconditional price volatility, as defined in Equation (2.27). Then:

$$\Gamma_U^2 \;=\; \int_Z \frac{k(z)^2}{k(z)(1-R^*)+\ell(z)R^*}\mathrm{d}z \;-\; 1. \tag{2.121}$$

Taking the first derivative with respect to R^* yields:

$$\begin{aligned}\frac{\mathrm{d}(\Gamma_U^2)}{\mathrm{d}R^*} &= \frac{1}{(1-R^*)^2K^2}\int_Z p(z)^2\Big(k(z)-\ell(z)\Big)\mathrm{d}z, \\ &= \frac{1}{(1-R^*)^2K^2}\Big(\mathbb{E}[p(z)^2|x=1]-\mathbb{E}[p(z)^2|x=0]\Big).\end{aligned} \tag{2.122}$$

Expression (2.118) implies that this derivative is positive. Thus, the unconditional volatility is increasing in R^*. This is the required result.

Chapter 3

Relationship lending when banks are fragile[1]

3.1 Introduction

Distress in the financial sector tends to result in lower corporate borrowing and investment, commonly termed a 'credit crunch'. A number of channels have been proposed to explain this phenomenon. Supply-side explanations emphasise the inability or unwillingness of banks to advance credit to borrowers at the same terms as they did previously, while demand-side explanations focus on the choice of borrowers to borrow less and reduce investment due to expectations of lower future economic growth and demand.[2] We identify a new credit crunch channel which does not rely on reduced expectations of future growth, demand or productivity, but is instead driven by informational factors. Borrowers respond to supply-side developments by reducing their demand for credit, so our channel can be described as having both demand- and supply-side aspects.

Lenders compete to finance a borrower. The borrower's quality is known only to one lender, called the 'insider'. The uninformed lenders are concerned about adverse selec-

[1] This chapter is based on joint work with Alan Morrison and Joel Shapiro.

[2] In the context of the recent global financial crisis, Duchin, Ozbas, and Sensoy (2010), Ivashina and Scharfstein (2010) and Campello, Graham, and Harvey (2010) provide evidence for supply-side channels, while Kahle and Stulz (2013) find evidence for a demand-side channel.

tion, and so require higher rates than the insider to lend to a good borrower. The borrower can choose the size of her project, which exhibits decreasing returns to scale. Therefore smaller projects will tend to attract more competition among lenders. Larger projects are more reliant on the insider for financing, and so the insider will capture higher rents. The borrower's trade-off is between being able to keep a relatively small piece of a large project, or a bigger proportion of a smaller project.

With some probability, the insider suffers a shock which prevents him from participating in the market for loans. The outcome of this shock is not observed. As the probability of the shock increases, outsiders face less adverse selection, and so offer more generous lending conditions. This makes choosing the smaller project more attractive for the borrower. At the same time, reducing the size of the project helps the borrower to insure against the risk that nobody bids for the project and it is not financed at all. This is because decreasing returns to scale means that a project above a certain size is not certain to be financed by the outsiders.

We show that there is a threshold equilibrium. When the shock probability is low enough, the borrower chooses large projects and relies on the insider for funding, at the risk of not being financed at all. When the shock probability is above the threshold, the borrower reduces the project size in order to be sure of being financed. There is a discontinuity in project size as the shock probability moves smoothly above the threshold. There is also a discontinuity in welfare, due to an externality imposed on the insider. We call this sudden drop in investment and welfare a 'credit crunch'.

An important friction in our model is that the insider cannot credibly commit not to exploit his information to the greatest extent possible. We examine an extension in which a social planner can impose an interest rate cap on the bank. This is equivalent to the insider making a commitment to leave some of the surplus for the borrower, and thus induces the borrower to choose a larger project. This can be welfare-positive, and make projects possible which are infeasible without such an intervention.

3.1.1 Related literature

Our model is part of a long and burgeoning literature on relationship banking — see Bongini, di Battista, and Nieri (2015) for a recent review. Our key contribution is to examine the effect of relationships and lender distress on project size, which allows us to comment on investment. Most of the theory literature assumes projects of a fixed size, missing out on this interesting aspect. The important exception is Sharpe (1990), as amended by Von Thadden (2004). Our model is close to theirs but we introduce lender distress. Their focus is on how lenders' terms vary across time, using a dynamic model to show that banks may lend too generously to new firms, due to the expectation of extracting profits from them in later periods. This causes low-quality new firms to over-invest. Von Thadden (2004) points out that Sharpe's model has Nash equilibria only in mixed strategies, which applies to our paper too.

Most other theory models in the relationship lending literature assume that projects are of a fixed size, and so these are less suitable for studying investment decisions. Other papers which focus on the rent-seeking behaviour of relationship banks are the classic papers of Rajan (1992) and Von Thadden (1995), which keep project size fixed and investigate the contracts that emerge in equilibrium.

A recent paper which examines the effect of lender distress is Castiglionesi, Li, and Ma (2016). They show that a voluntary information sharing scheme among banks can efficiently reduce the adverse selection that banks face when selling their loans. This is traded-off against a reduced ability to extract rents. The focus of our model is the effect of financial distress on investment choice rather on than incentives to share information.

Aside from the paper just mentioned, there are, to our knowledge, no other papers in the theory literature which examine the effect of financial distress on relationship lending. The focus has been more toward shocks at the *borrower* level. Bolton et al. (2016) examine relationship lending in a model where firms, rather than lenders, encounter distress. They show that, while relationship lending may be costly during periods of economic growth, it can be more efficient during a recession as private information ensures that good projects continue to be financed. In their model, the inefficiency of relationship banking comes from a fixed monitoring cost, rather than rent extraction. Dell'Ariccia and Marquez (2006) associate upturns in the business cycle with reduced adverse selection,

and that this can drive a credit boom.

Several recent empirical papers on relationship banking have exploited datasets of individual bank-borrower relationships. These generally find evidence that bank distress leads to less lending and worse economic outcomes for borrowers, though not necessarily higher interest rates. The effects tend to be more severe for firms which are relatively small in size, and for those which are not publicly listed. This result appears to apply across a range of countries, shocks and financial systems — for example, Khwaja and Mian (2008) examine Pakistani banks exposed to the sudden imposition of currency controls, Beck et al. (2018) use survey data in eastern Europe before and after the global financial crisis, and Chodorow-Reich (2014) studies U.S. lenders in the wake of the financial turmoil of September 2008. These smaller and more opaque firms may be particularly dependent on bank financing, which points to an important role for relationship banking.

3.2 Model

3.2.1 Players and technologies

We consider an economy populated by one entrepreneur and several risk-neutral lenders.[3] For convenience, we use female pronouns for the entrepreneur and male pronouns for the lenders. All agents in the economy are risk-neutral and we normalise the interest rate to 0. The economy has three dates, $t = 0, 1, 2$. The entrepreneur is endowed with a single scaleable project. In exchange for a time $t = 1$ investment of K, the project returns either $Kg(K) > 0$ or 0 at time $t = 2$. We refer to these outcomes as success and failure, respectively. The entrepreneur's outside option yields a payoff of zero. We assume that the per-unit production has a Cobb-Douglas form, so that $g(K) = K^{\gamma}$ for a parameter $\gamma \in (-1, 0)$.[4]

[3] We need at least three lenders: one will be the informed lender, and at least two more are required to ensure that uninformed agents compete for zero expected profit.

[4] The leading coefficient is normalised to 1. This convex functional form is an important difference between our model and those of Sharpe (1990) and Von Thadden (2004), who assume that $g(K)$ is everywhere concave. We find concavity problematic because it implies, for example, that projects above a certain size are guaranteed to be worthless. Moreover, convexity ensures compatibility with the widely-accepted Inada conditions. Our results do actually hold with any concave $g(K)$ which meets the criteria of Sharpe (1990)

The entrepreneur may be one of two types, $\theta \in \{G, B\}$. We denote by p_θ the probability that a type θ entrepreneur's project succeeds, and we assume that $p_G > p_B$. Nature assigns the entrepreneur's type at time $t = 0$. The probability that she has type G is $q \in (0,1)$, which is common knowledge.

We assume that the entrepreneur has an existing relationship with an *inside* lender who is able to observe θ. We refer to the other lenders as the *outsiders*. The outsider is unable to observe θ. We assume too that the entrepreneur is unable to observe θ; this assumption allows us to concentrate upon the effect of information asymmetry between lenders without introducing the added complexity of borrower signalling.[5]

The insider lender may experience a *liquidity shock* with probability $\lambda > 0$. If the insider experiences a liquidity shock then he is unable to fund the entrepreneur's project. A liquidity shock could be a consequence of, for example, deposit withdrawals or a loss of access to the interbank or repo markets.[6]

3.2.2 Timing and strategies

The timing of the game is illustrated in Figure 3.1.

At time $t = 0$, nature selects the entrepreneur's type θ, which is observed only by the insider. The entrepreneur then announces her investment choice K.

At time $t = 1$, the lenders play a bidding game. The insider and the outsider can each choose a bid $r \in [0, \infty)$, or a non-bid, which we denote as bidding ∞. A move for the insider and outsider comprises a probability β of bidding and, conditional upon bidding, a distribution function over the set $[0, \infty)$. For every possible K, the outsider's strategy is therefore a probability β_O and a distribution function F_O. The insider's strategy com-

and Von Thadden (2004); a proof is available upon request.

[5]It is straightforward to show that, if the borrower does know her own type, then the only possible equilibrium is a pooling equilibrium in which lenders respond to any deviation from the equilibrium strategy by assigning very pessimistic beliefs to the borrower's type. There is no separating equilibrium because there is no single-crossing condition: it is always optimal for the B types to mimic the G types.

[6]For simplicity, we rule out that the outsider suffers a similar shock. We can justify this as follows: the market for lending is contestable and a number of banks compete for entry. The one which does not suffer a shock and bids the lowest is labelled the 'outsider', and is the only outside bank of interest to the borrower. As the number of competing banks becomes large, the probability that they all suffer a shock is negligible.

Figure 3.1: Game timing.

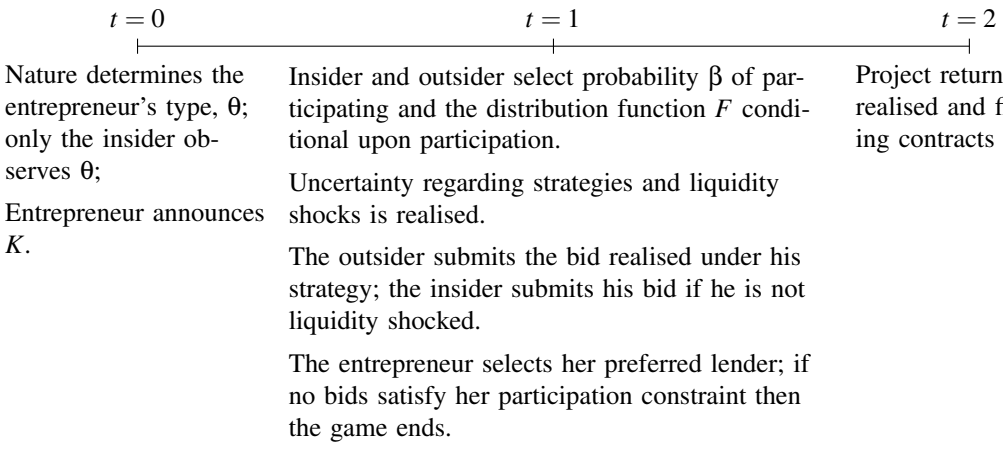

prises probabilities β_I^G, β_I^B and distribution functions F_I^G and F_I^L in the cases where the entrepreneur has a good or a bad project respectively.

The bidding game at $t = 1$ unfolds as follows.

1. The insider and the outsider both select their strategies.
2. The probability of participating β and strategy F are realised for both the insider and the outsider. The insider's liquidity shock is also realised.
3. The insider and the outsider simultaneously bid according to the realisation of their strategies.
4. If no bids satisfy the entrepreneur's participation constraint then the game ends; otherwise, the entrepreneur selects the lowest bid. In the event of a tie, she prefers to be financed by the insider.

At $t = 2$, project returns realise and financing contracts settle.

3.3 Equilibrium

A formal treatment of our game should include a definition of equilibrium that includes strategies for the insider, outsider, and borrower. The equilibrium definition relies upon

profit functions for the insider, outsider and borrower. If the insider bids r and does not experience a liquidity shock, then the outsider bids below him with probability $\beta_O F_O^-(r)$, where $F_O^-(r) = \lim_{x\uparrow r} F_O(x)$. The insider therefore derives expected profit $\pi_I^\theta(r)$ from strategy $(\beta_I^\theta, F_I^\theta)$, where $\theta \in \{G,B\}$ and:

$$\pi_I^\theta(r) = \begin{cases} 0, & \text{if } r = \infty; \\ K(p_\theta(1+r)-1)(1-\beta_O F_O^-(r)), & \text{otherwise.} \end{cases} \tag{3.1}$$

The outsider lends to the entrepreneur if he does not experience a liquidity shock and if he is also the lowest bidder. If he bids r, then the insider beats his bid with probability $(1-\lambda)\beta_I^\theta F_I^\theta(r)$ for $\theta \in \{G,B\}$. The outsider's payoff function $\pi_O(\cdot)$ is therefore given by equation (3.2):

$$\pi_O(r) = \begin{cases} 0, & \text{if } r = \infty; \\ qK\left[\lambda + (1-\lambda)(1-\beta_I^G F_I^G(r))\right](p_G(1+r)-1) & \\ \quad +(1-q)K\left[\lambda+(1-\lambda)(1-\beta_I^B F_I^B(r)\right](p_B(1+r)-1), & \text{otherwise.} \end{cases} \tag{3.2}$$

We seek Perfect Bayesian equilibria of our game. This immediately implies that the lenders' strategies F maximise their expected payoff at time $t=1$ when the scale K of the project is revealed. We therefore begin by considering the $t=1$ continuation game.

Definition 3.1. *A* Perfect Bayesian equilibrium *of the game comprises:*

1. *a time $t=0$ choice of project scale K by the entrepreneur;*
2. *a triple $(\beta_I^G, \beta_I^B, \beta_O)$ of bidding probabilities and a triple (F_I^G, F_I^B, F_O^K) of elements of the set $\mathcal{F}$ of probability distribution functions on $[0,\infty)$,*

such that:

1. *for each K, the insider has no profitable one-shot deviations: for all $\hat{F} \in \mathcal{F}$ and for $\theta \in \{G,B\}$, we have:*

$$\int_0^\infty \pi_I^\theta(r) dF_I^\theta(r) \geq \int_0^\infty \pi_I^\theta(r) d\hat{F}(r); \tag{3.3}$$

2. *for each K, the outsider has no profitable one-shot deviations: for all* $\hat{F} \in \mathcal{F}$*, we have:*

$$\int_0^\infty \pi_O(r)dF_O(r) \geq \int_0^\infty \pi_O(r)d\hat{F}(r); \tag{3.4}$$

3. *the probability* β *of bidding is less than one only if the expected profit from bidding is equal to zero;*

4. *the outsider earns no profits in equilibrium:*

$$\mathbb{E}[\pi_O(r)] = 0. \tag{3.5}$$

Before we solve our game, we establish some basic notation. Suppose that a lender invests K in the entrepreneur's project. If the lender assesses a success probability p for the project then his break-even lending rate r must satisfy equation (3.6):

$$(1+r)Kp = K. \tag{3.6}$$

The break-even rate r is therefore given by equation (3.7):

$$r = \frac{1}{p} - 1. \tag{3.7}$$

An outsider that expects to lend only to bad projects assesses success probability p_B; we denote by r_B the corresponding break-even lending rate. Similarly, if the outsider expects to lend only to good projects then he assesses success probability p_G and expects to break even at lending rate r_G.

If there were no insider, then the outsider would assess entrepreneurial success probability

$$p_M = qp_G + (1-q)p_B. \tag{3.8}$$

We write r_M for the corresponding break-even lending rate.

Suppose that the outsider expects the insider to win every lending contract for which it competes, and that the insider bids only for good projects. Then the outsider's bid will succeed only if the entrepreneur is bad, or if the insider has experienced a liquidity shock.

In this case, the outsider assesses success probability

$$p_N = \frac{\lambda q p_G + (1-q)p_B}{\lambda q + 1 - q} \tag{3.9}$$

for any entrepreneur to whom he lends money. We write r_N for the break-even lending rate that corresponds to p_N.

If a lender demands an interest rate r per unit invested then the lender's participation constraint can be satisfied precisely when $(1+r)K \leq Kg(K)$. The maximum feasible interest rate for an investment K is therefore $g(K) - 1$. Equivalently, because $g(K)$ is decreasing in K, the maximum feasible investment for a given interest rate r is $g^{-1}(1+r)$. Figure 3.2 plots $g(K)$ as a function of investment K, and identifies the maximum investments at which interest rates r_B, r_N, r_M and r_G are feasible; we denote those investment levels as K_B, K_N, K_M, and K_G, respectively.

Figure 3.2: Break-even rates

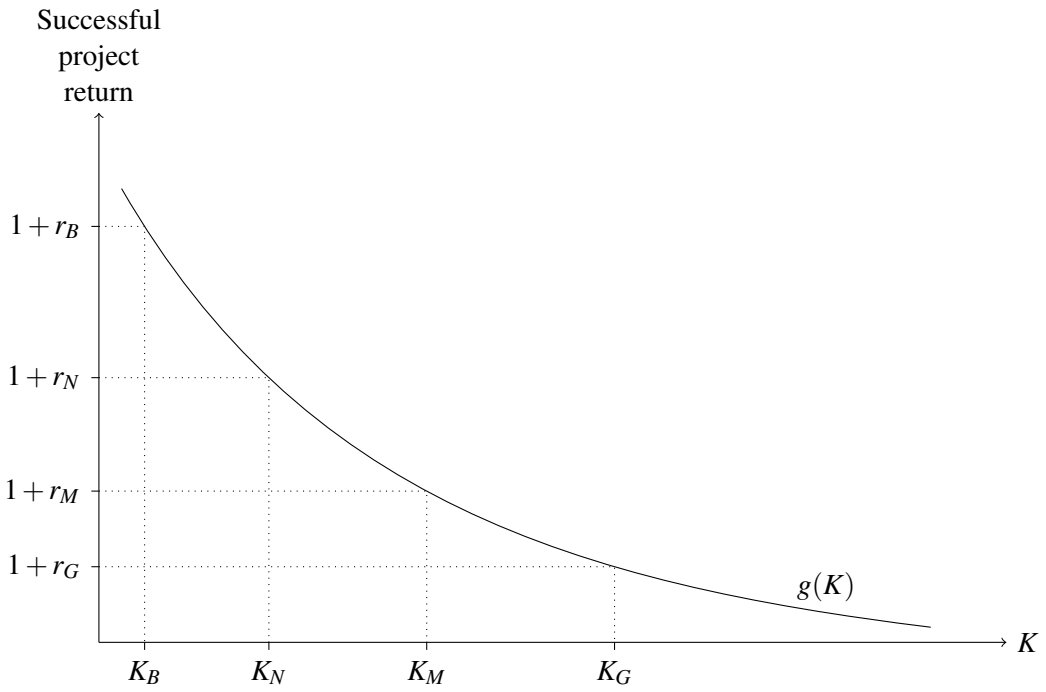

Note that the lowest interest rate a lender would ever accept is r_G, corresponding to the situation where the project is definitely of type G. The highest interest rate that the project

could afford to pay is $g(K)-1$. Hence, the support of any lender's distribution function is

$$S = [r_G, g(K)-1] \cup \{\infty\}. \tag{3.10}$$

We can therefore replace every integral over $[0,\infty]$ in Definition 3.1 with an integral over S.

We start by solving for the optimal lender strategies in the time $t=1$ continuation game; we then determine the entrepreneur's optimal choice of project scale K. The following Lemma shows that insiders aim to pick off G types, but do not compete for B types. The reason for this is that the insider always places a lower value on the B type than the outsiders. As the outsiders make zero expected profit, the insider would have to take a negative profit to win the loan. Thus it is better not to make an offer. It is convenient to establish the following lemma:

Lemma 3.1. *In any Nash equilibrium of the offer subgame, the insider never finances B type borrowers. There is a unique pure strategy Nash equilibrium if and only if* $g(K) \leq 1 + r_M$*, in which case the outsider does not participate. Otherwise, all Nash equilibria are in mixed strategies.*

Proof. See Appendix. □

Proposition 3.1 (Time 1 continuation game)**.** *Lenders adopt the following equilibrium strategies in the time* 1 *continuation game for a given K:*

1. ***Small projects**. If* $K \leq K_N$ *then:*

   - *The outsider always bids, and assigns point probability mass* $(r_M - r_G)/(r_N - r_G)$ *to* r_N *and randomizes over an atomless distribution on* $r \in [r_M, r_N)$ *with probability distribution function*

   $$F(r) = \frac{r - r_M}{r - r_G}. \tag{3.11}$$

   - *If* $\theta = G$ *then the insider always bids, and randomizes over* $r \in [r_M, r_N]$ *with*

probability distribution function

$$\frac{1}{1-\lambda}\frac{p_M}{qp_G}F(r); \tag{3.12}$$

- *If* $\theta = B$ *then the insider does not bid.*

2. ***Large projects***. *If* $K_N < K \leq K_M$ *then:*

   - *The outsider bids with probability* $\beta_O = F(g(K)-1)$. *If he bids, the outsider randomises over an atomless distribution on* $r \in [r_M, g(K)-1)$ *with probability density function*

   $$\frac{F(r)}{F(g(K)-1)}. \tag{3.13}$$

   The outsider assigns zero probability mass to the bid $r = g(K)-1$.

   - *When the entrepreneur's type* θ *is G, the insider randomises over an atomless distribution on* $r \in [r_M, g(K)-1)$ *with probability distribution function*

   $$\frac{1}{1-\lambda}\frac{p_M}{qp_G}F(r) \tag{3.14}$$

   and assigns point mass

   $$1 - \frac{1}{1-\lambda}\frac{p_M}{qp_G}F(g(K)-1) \tag{3.15}$$

   to the bid $r = g(K)-1$.

   - *The insider does not bid when the entrepreneur has type* $\theta = B$.

3. ***Huge projects***. *If* $K_M < K \leq K_G$ *then the outsider does not bid. The insider bids* $g(K)-1$ *for good entrepreneurs and so extracts all of their surplus; he does not bid for bad entrepreneurs.*

4. ***Infeasible projects***. *If* $K > K_G$ *then no lender bids to finance the project.*

Proof. See Appendix. □

The expected cash generated by a small project is sufficient to repay even the most pes-

simistic outsider, since $p_N g(K) > 1$. Therefore outsiders will always be willing to make an offer in this case. However, the expected cash from a large project is not always sufficient to repay the outsider, so with non-zero probability the outsider will choose not to make an offer. When the project is huge, the expected yield is insufficient to repay even the most optimistic outsider ($p_M g(K) < 1$) so the outsider never makes an offer. An infeasible project generates insufficient expected cash to repay even an insider who knows that $\theta = G$, so such a project will never be financed by anyone.

Whereas all small projects are financed, credit rationing may occur for large projects.[7] Large type G projects have positive net present value at any rate offered in equilibrium. However, if the insider suffers a shock and the outsider does not make an offer, then they will not be financed. This occurs with probability $\lambda(1 - F(g(K) - 1))$.

Note that, in the small project case, the lenders' strategies do not depend directly on K. This is because, so long as the boundary condition $K < K_N$ is satisfied, outsiders expect that the project will generate sufficient cash to repay the loan, so the exact value of K does not matter. In contrast, in the large project case the strategies depend explicitly on K.

Interestingly, the outsider's strategy has no direct dependence on λ in the large project case, except to the extent that the boundary condition $K > K_N$ is satisfied.[8] Consider an increase in λ. The outsider faces less adverse selection so, holding his strategy fixed, the outsider expects to finance better loans. The insider must therefore reduce the rate he offers, and bid more aggressively, so that the payoff of the outsider remains at zero.

Now that we have the equilibria from Proposition 3.1, we can calculate the expected payoffs for each party and aggregate welfare for all three parties.

Lemma 3.2 (Expected payoffs)**.** *Expected payoffs depend on the size of the project. There are four cases:*

- *For small projects ($K \leq K_N$), expected payoffs are as follows.*
  - *The outsider earns zero.*
  - *The insider earns zero if $\theta = B$, and $J := p_G(1-\lambda)(r_M - r_G)K$ if $\theta = G$.*

[7] We use the definition of credit rationing in Baltensperger (1978): demand for loans exceeds the supply at the equilibrium price.

[8] Recall that K_N is a function of λ.

  - *Total welfare is* $\Omega_S = (p_M g(K) - 1)K$.
  - *The borrower has payoff* $\Pi_S = \Omega_S - qJ$.

- *For large projects* ($K_N < K \leq K_M$)*, expected payoffs are as follows.*
  - *The outsider earns zero.*
  - *The insider earns zero if* $\theta = B$*, and* J *if* $\theta = G$.
  - *Total welfare is:*

$$\Omega_L = q\Big(1 - \lambda(1 - F(g(K) - 1))\Big)(p_G g(K) - 1)K + (1 - q)F(g(K) - 1)(p_B g(K) - 1)K.$$

  - *The borrower has payoff* $\Pi_L = \Omega_L - qJ$.

- *For huge projects* ($K_M < K \leq K_G$)*, expected payoffs are as follows:*
  - *The outsider earns zero.*
  - *The insider earns zero if* $\theta = B$*, and* $(1 - \lambda)(p_G g(K) - 1)K$ *if* $\theta = G$.
  - *Total welfare is* $\Omega_H = q(1 - \lambda)(p_G g(K) - 1)K$.
  - *The borrower earns zero.*

- *For infeasible projects* ($K > K_G$)*, all parties have zero payoff and total welfare is zero.*

Proof. See Appendix. □

Competition means that the outsider always has expected payoff zero in equilibrium.

The value of K is chosen by the borrower at $t = 0$. From Lemma 3.2, she knows her expected payoff for each possible value of K. Since the borrower does not put up any of her own money, her worst possible outcome is a zero payoff.

Lemma 3.3. *Choosing a small project always earns the borrower a strictly positive payoff. The borrower will never choose K consistent with a huge or infeasible project; i.e. the borrower never chooses* $K > K_M$.

Proof. Small projects are financed with probability 1. Conditional on success, a small project generates cash $g(K)K \geq (1+r_N)K$. The project is financed at a rate no higher than r_N, and with positive probability it is strictly less than r_N. Therefore the expected interest payment is strictly less than $(1+r_N)K$. Thus the borrower's expected payoff is greater than zero.

Lemma 3.2 tells us that the borrower earns zero payoff for sure if she chooses a huge or infeasible project. As she can earn a strictly positive payoff by choosing a small project $K \leq K_N$, this dominates huge or infeasible projects. □

This lemma implies that huge and infeasible projects do not lie on the equilibrium path. From now on we focus on the choice between small and large projects. The borrower faces a trade-off when making this choice. With a small project, she always obtains financing, but the payoff is relatively low. With a large project, she can invest more but there is a non-zero probability that the loan is not financed.

3.4 Comparative statics

We analyse the comparative statics of each equilibrium with respect to the parameter λ. If the borrower chooses a large project, she is relying to some extent on the insider for financing, because it is not certain that the outsider will bid at all. This is because the expected returns for the outsider are too low. Therefore a higher λ makes financing less likely. If the insider fails to make an offer (because he suffers a shock), then the borrower faces the risk that she receives no offers at all. Then the project would not be financed, and the borrower's payoff would be zero. For good-quality borrowers, this unhappy event occurs with a non-zero probability $\lambda(1-F(g(K)-1))$, which is increasing in both λ and K.

In contrast, a higher λ makes the borrower more inclined to choose a small project. This is because the adverse selection problem is reduced, and so outsiders make more generous bids. A small project is NPV positive, so outsiders always bid and the project is always financed. Thus the insider's screening technology gives comparatively little value to the borrower in this case. The insider's borrowing terms are not particularly more generous

than the outsider's, because the insider makes an offer low enough only to win the loan and, conditional on that, retains as much surplus as possible. Moreover, a higher λ reduces the rate offered by the outsider, and therefore by the insider. In this way, a higher λ allows for the 'winner's curse' problem that the outsider faces to be mitigated, because with positive probability the insider cannot exploit his informational advantage. The winner's curse does not entirely disappear unless $\lambda = 1$ (in which event the set of large projects shrinks to a point and a small project is always chosen).

We formally establish these results in the following lemma. Welfare is defined as the sum of the payoffs of the three parties, which is exactly equal to the expected cash flow generated by the project.

Lemma 3.4 (Comparative statics of small and large projects)**.** *Comparative statics differ for small and large projects:*

- *Let K_S^* be the borrower's optimal choice for K under the constraint $K < K_N$. Then:*
  1. *K_S^* is uniquely determined and strictly increasing in λ.*
  2. *At $K = K_S^*$, welfare is increasing over $\lambda \in [0,1]$ (strictly increasing if $\lambda < 1$).*
  3. *At $K = K_S^*$, the borrower's payoff is strictly increasing in $\lambda \in [0,1)$.*
- *Let K_L^* be the borrower's optimal choice for K under the constraint $K_N \leq K < K_M$. Then:*
  1. *K_L^* is uniquely determined and non-decreasing with respect to λ.*
  2. *At $K = K_L^*$, the borrower's payoff is non-increasing with respect to λ.*
  3. *At $K = K_L^*$, welfare is zero when $\lambda = 1$ and strictly positive when $\lambda = 0$.*

Proof. See Appendix. □

In both cases, the borrower's choice of investment is increasing in λ. This is because the mitigation of the winner's curse reduces lending rates and so induces the borrower to invest more. This occurs even though for a large project the probability of not being financed is increasing in both λ and K: Lemma 3.4 shows that this effect is dominated by

the lower lending rates. When λ is high, a large project thus becomes a riskier proposition for the borrower: rewards are high given success (because lending rates are low and investment is high), but the probability of not being financed at all and thus getting zero payoff is also high. The borrower has to consider this trade-off when choosing between the two sizes of project.

For a small project, welfare increases with λ. There are two countervailing effects. While the borrower's payoff increases with λ, the insider's payoff falls with λ as the value of his information monopoly is reduced. Lemma 3.4 shows that the first effect always dominates the second, and so welfare is improved with higher λ. The reason is that the insider's profit is not a welfare-neutral transfer from borrower to lender: it has the effect of reducing incentives to invest, which reduces total social welfare.

For a large project, welfare declines to zero at $\lambda = 1$, though not necessarily monotonically. The insider's screening technology plays a socially important role in channelling funding to *G* type borrowers and away from *B* types. There is a non-zero probability that the outsider does not make an offer at all, and this can be efficient since *B* types have large projects with negative NPV. However, when λ is high this screening technology is less reliable. This is bad for the borrower, because she increasingly has to rely on the outsider, who is not guaranteed to make an offer. It is also bad for the inside lender, because the value of his monopoly decreases. Therefore it is bad for overall welfare too.

The borrower and society have different trade-offs between the two sizes of project. The borrower trades off between a high but risky level of investment (large project) and a safe low level of investment (small project). Society trade-offs between higher investment and better directed credit (large project), against less profit-taking and mitigation of the winner's curse (small project). The action of the borrower — that is, the choice of investment level — imposes an externality on society via its effect on the profit of the inside lender. This suggests that the borrower may not always choose a level of investment which is optimal for society.

We can now analyse the overall payoffs when the borrower chooses between small and large projects, paying particular attention to the point at which the borrower switches between the two.

Proposition 3.2 (Outcome of borrower's choice)**.** *The borrower's choice of project size depends on the probability of the liquidity shock* λ *as follows:*

1. *There exists a* $\lambda^* < 1$ *such that the borrower chooses a small project if* $\lambda > \lambda^*$*, and a large project if* $\lambda < \lambda^*$*.*
2. *At the point when* λ *rises above* λ^**, there is a non-continuous downward jump in investment.*
3. *There exists* $\lambda^\dagger \in (\lambda^*, 1)$ *such that for all* $\lambda < \lambda^\dagger$*, a large project induces higher welfare than a small one. This implies that, in some neighbourhood of* λ^**, the borrower chooses a socially inefficient equilibrium and there is a non-continuous downward jump in welfare when* λ *rises above* λ^**.*

Proof. See Appendix. □

Figure 3.3 illustrates the functions and demonstrate the behaviour at the critical point λ^* when the borrower switches from a large to a small project.[9] The figure confirms the prediction of Proposition 3.2 that, immediately above the point λ^*, both investment and total welfare are lower with a small project than a large one.

The kinks in the charts in Figure 3.3 are due to the constraints on K in each case binding. For small projects, the constraint $K < K_N$ tends to bind for lower values of λ. For large projects, the constraint $K > K_N$ binds for higher values of λ. In the proof of Proposition 3.2, we show that these kinks are not relevant to the borrower's choice. That is, the constraint binds for small projects only when $\lambda < \lambda^*$, and for large projects only when $\lambda > \lambda^*$, so the kinks do not bind on the equilibrium path.

Figure 3.4 shows the overall level of investment chosen by the borrower, and the associated welfare levels. Explicitly, the figure plots K where:

$$K = \begin{cases} K_S^* & \text{when } \lambda > \lambda^* \\ K_L^* & \text{when } \lambda \leq \lambda^* \end{cases}$$

[9] The parameters here are $q = 0.5, p_G = 0.9, p_B = 0.1$ and $g(K) = K^{-0.25}$. These parameters give $\lambda^* = 0.35$.

Figure 3.3: Comparative statics with respect to λ. The left-hand panel shows optimal investment levels under small projects (K_S^*) and large projects (K_L^*) for a given λ, the middle panel shows expected borrower payoff at these investment levels, and the right-hand panel shows total welfare at these investment levels. The dotted vertical line in each panel represents λ^*, the point above which the borrower prefers to choose a small project.

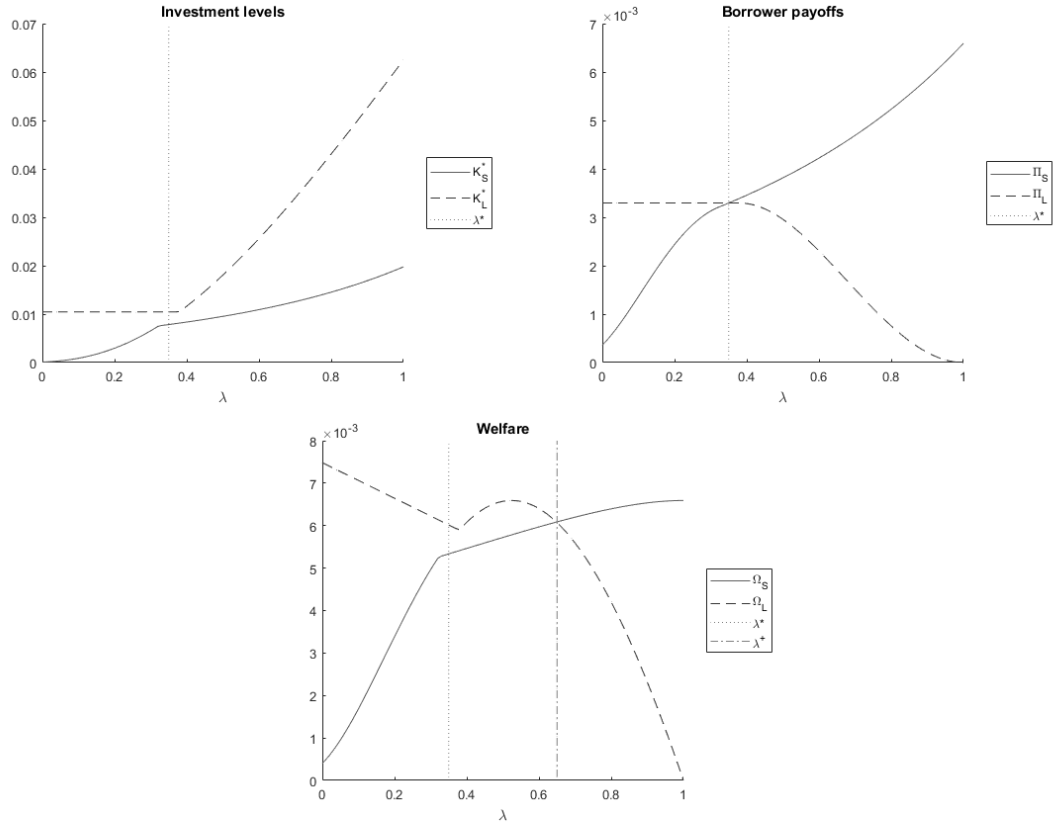

and similarly for Ω. Figure 3.4 clearly shows the drops in investment and welfare that occur at the point when the borrower switches from a large to a small project.

Note that, for high λ, investment and welfare can eventually exceed the levels under lower λ, because the winner's curse effect is mitigated and the outsiders offer low interest rates, incentivising the borrower to invest more as she can keep more of her cash flow. It is an empirical question as to whether such high levels of λ are likely to be observed in practice; for example, in Figure 3.4, λ needs to reach a value of around 0.6 for investment and welfare to exceed the values achieved for the large project when λ is just below λ^*.

Figure 3.4: Investment chosen by borrower for each λ (solid line), and total welfare (dot-dashed line). The vertical line represents λ^*, the point above which the borrower prefers to choose a small project. The parameters are the same as in Figure 3.3.

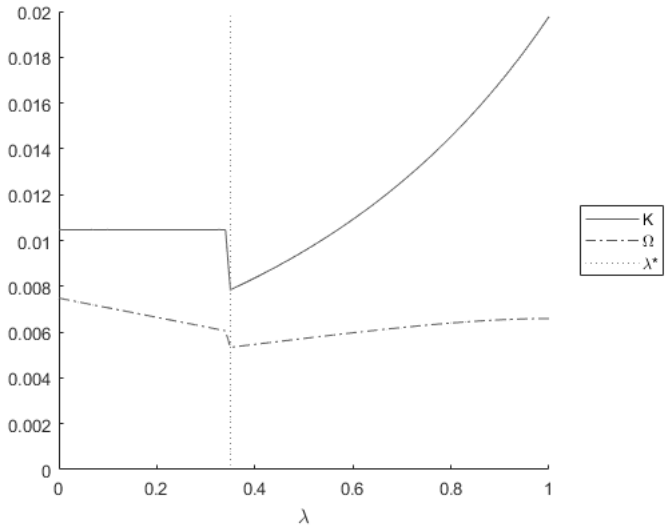

This would mean that a typical bank has a 0.6 probability of defaulting during the tenor of an average loan. Such high levels of bank distress are likely to be associated with bad fundamentals in the economy, which we do not model.

As $\lambda \to 1$, relationship lending fades away and so small project finance can be characterised as transaction or arm's length lending, in the sense of Rajan (1992) or Bolton et al. (2016). Under transaction lending, the lender does not learn the type of the borrower and remains uninformed, so there is no winner's curse. As Ω_S is increasing in λ (by Lemma 3.4), pure transaction lending is associated with higher welfare than any other small project outcome. However, transaction lending to small projects is not necessarily better than relationship lending to large projects, when λ is low. If q is low and $p_G - p_B$ is high (i.e. there is a large degree of information asymmetry) then the social value of the insider's information may make large projects more efficient.

3.4.1 Interpretation as a credit crunch

This model provides a novel reason for the dip in investment, credit demand and real economic activity — that is, a 'credit crunch' — that typically follows banking crises. The proximate cause of our credit crunch is the borrower insuring against the risk of not being financed by choosing a smaller project which the outsider is sure to lend against. We call this a new demand-side channel, because it is the borrower's choice to reduce investment. Unlike existing demand-side explanations of credit crunches, we do not require any change to firm productivity, real economic outcomes, or confidence. All agents believe that projects are just as productive as they were.

The ultimate cause of our credit crunch is the shock to informed capital, so there is a supply-side element to our story as well. Unlike existing supply-side explanations of credit crunches, we do not assume balance sheet constraints, or any changes to the willingness of banks to lend. We need not assume any shock to aggregate capacity to extend credit: our results will still hold if there is an infinite number of outsiders, in which case the relationship bank is an infinitesimally small part of the banking system.

The choice of a smaller project by the outsider is a way of insuring herself against the risk of not being financed at all. Smaller projects are less informationally-sensitive, because they are more likely to be NPV-positive. Therefore smaller projects provide a means for the borrower to continue being financed in an uncertain environment. However, this can be inefficient, because the borrower does not consider the impact of this reduction in investments on the balance sheet of the insider.

Our model predicts that borrowers reduce investment in order to reduce the hold-up problem during times of lender distress. A relevant empirical study in this regard is Almeida et al. (2011), which finds that US firms whose long-term debt matured during the crisis episode of autumn 2007 cut investment by more than comparable firms whose long-term debt matured after 2008. We provide an explanation for this effect: a perceived higher probability of lender stress causes firms to borrow less. Our model does not apply to firms who do not face an immediate to borrow, and so it would predict that such firms reduce investment by less than those which do need financing.

Darmouni (2018) uses data on the U.S. syndicated corporate loan market in the wake of

the default of Lehman Brothers in autumn 2008, and finds evidence that borrowers who seek new lenders are stigmatised. This is because of the lemons problem: new lenders cannot be sure whether the borrower wishes to migrate because the incumbent lender no longer wishes to finance them (implying a negative signal about borrower quality) or because the incumbent lender is suffering liquidity problems (implying a neutral signal about borrower quality). Darmouni (2018) finds that this stigma leads to a substantively lower level of lending, relative to a benchmark in which there is no private information. He argues that rescuing distressed banks can exacerbate this adverse selection problem. This is in line with a prediction of our model: when the level of financial distress is sufficiently high, there is less adverse selection and so there can potentially be more lending. Another study which provides some evidence to support this prediction is Ongena, Smith, and Michalsen (2003), which examines the near-meltdown of the Norwegian banking system during the period 1988-1991 and finds little evidence of a large aggregate impact on borrowing firms.

In addition, our model predicts that interest rates can go down rather than up during a financial crisis, due to reduced adverse selection, even if investment drops. There is some empirical evidence for this in Cingano, Manaresi, and Sette (2016), who show that credit shocks in Italy following the 2007 financial crisis reduced investment expenditure but had no effect on interest rates.

Finally, our shock λ can be interpreted in manners other than a liquidity shock. For example, it can include operational risks such as a cyber attack, technological problem or regulatory changes. Hombert and Matray (2017) argue that deregulation in the US may have contributed to the break-up of lending relationships, as greater competition reduced the benefit of this private information.

3.4.2 Empirical predictions

Our channel can be distinguished empirically from the standard bank lending channel by observing the impact of bank distress on the unconditional probability of a project being financed. In the standard bank lending channel (as described, for example, in Peek and Rosengren 1997), bank distress makes borrowers' projects less likely — or certainly

no more likely — to be financed, because the supply of credit falls. In contrast, our channel predicts that the crunch can make projects more likely to be financed, though these projects are smaller in size than they were prior to the incidence of distress. As borrowers reduce the size of their loans in order to insure against the shock to informed capital, outsiders become willing to lend money with certainty. But outsiders lend to good borrowers at higher rates than the insider, so the lending rate rises.

The following Lemma establishes this formally.

Lemma 3.5 (Probability of financing and average lending rate)**.**

1. *The probability of a project being financed is decreasing with* λ *when* $\lambda \in [0, \lambda^*]$*, and equal to 1 when* $\lambda > \lambda^*$*.*
2. *Conditional on being financed, a borrower's expected loan rate is decreasing with* λ *over* $[0, \lambda^*) \cup (\lambda^*, 1]$*.*
3. *Conditional on being financed, a borrower's expected loan rate is increasing with* λ *at* λ^**.*

Proof. See Appendix. □

Our model predicts the following as λ increases above λ^*:

- **Prediction 1:** Average amount lent per project decreases.
- **Prediction 2:** Average lending rate increases.
- **Prediction 3:** The probability of a project being financed increases.
- **Prediction 4:** The quality of the average project financed decreases.

The last finding arises because when $\lambda < \lambda^*$, good projects are more likely to be financed than bad ones, but when $\lambda > \lambda^*$ all projects are financed.[10]

[10] This is in line with the model of Dell'Ariccia and Marquez (2006), but for a very different reason. In their model, information asymmetry is reduced when firms' productivity improves; i.e. during upturns in the business cycle. In our model, asymmetry is reduced when financial distress is higher.

We are looking for predictions which differ from those predicted by previous models of bank lending channels. Predictions 1 and 2 are fairly standard: distress in the banking system will move the credit supply curve inwards, leading to a lower quantity lent and a higher price.

Prediction 4 is more unusual. In a standard bank lending channel model, we might expect the quality of financed projects to be the same following a bank shock, if banks' monitoring technology does not change. But it could also improve (for example, if banks ration their limited supply of credit and focus on the best borrowers), or to worsen (for example, if banks gamble for resurrection by choosing riskier borrowers[11]). Therefore we cannot conclude that an observation in line with prediction 4 is necessarily evidence for our channel.

Instead, we argue that empirical observations in line with prediction 3 would present evidence for our channel. A standard bank lending channel would imply that the probability of an average project being financed falls — or, at least, does not increase — following a shock to the supply of credit. With our channel, the opposite occurs: projects shrink in size and are thus financed by uninformed lenders. Aggregate investment may go either up or down, depending on the relative magnitude of the effects in predictions 1 and 3. Even so, welfare will always fall when the crunch occurs (by Proposition 3.2), because smaller projects of worse quality are getting funded.

We have been able to find some evidence for this prediction: for example, Duchin, Ozbas, and Sensoy (2010) find evidence that firms which had low debt prior to the crisis cut their investment levels by less than those which had high debt. They interpret this as evidence of the effect of financial constraints, but another explanation is that a larger number of firms are able to obtain credit, because they choose smaller projects than they did prior to the crisis.

However, we have been not been able to find definitive evidence in the empirical literature for prediction 3. This is due, in a large part, to the fact that until now the literature has not focused on this effect. For example, studies which used data sets of matched bank-borrower pairs do not examine borrowers which did not obtain finance prior to the distress

[11] This would not work with the low-quality borrowers in our model, whose returns are first-order stochastically dominated by the returns of the good borrowers. It would require low-quality borrowers to be more risky; i.e. they are second-order stochastically dominated. See, for example, Rochet (2008), chapter 3.4.

event.[12]

This prediction has important implications for policymakers. It may not be possible to assess the impact of a credit crunch, or understand its nature, just by looking at the effect on lending rates and investment. Our model predicts that a crunch may not result in a drop in aggregate investment, but nevertheless welfare may fall due to smaller and worse-quality projects being financed. Thus policymakers should focus on the type of projects being financed, instead of their aggregate size. Where this channel is active, recapitalising the banking system may not be effective because the problem is not with the aggregate supply of credit: it is about the information possessed by the supplier of credit.

3.5 Policy implication: an interest rate cap

In this section, we show that, when financial distress is not too high, a benevolent policy maker can improve welfare by imposing an interest rate cap on banks. Huge projects can be better for society than either small or large projects; for example, when there are a few high quality borrowers with positive NPV projects and a majority of low quality borrowers have negative NPV projects, so that the average projects has negative NPV. Informed capital is thus socially valuable, because it can be directed only towards the good projects. Uninformed capital would not be utilised.

In Lemma 3.3, we showed that the borrower will never choose K consistent with a huge project because that will result in her receiving zero payoff. Therefore these cases will not occur in equilibrium. The huge project is NPV negative for the outsider, but is NPV positive for an insider who knows that he is dealing with a good borrower. However, the insider faces a monopoly and cannot credibly commit not to fully exploit his informational advantage. If a good borrower chooses a huge project, the insider's unique optimal strategy is to demand all of the cash flows. The borrower is thus better off inefficiently choosing a small or large project. An interest rate cap can force the insider to commit to leaving some money for the borrower, inducing the borrower to choose a huge project at $t = 0$, and thus improve welfare.

[12] See, for example, Cingano, Manaresi, and Sette (2016) and Khwaja and Mian (2008).

To illustrate, suppose that $\lambda = 0, p_G = 1, p_B = 0$ and that q is very low. There is a high degree of information asymmetry. Uninformed banks will lend very little because the return r_B they require in order to break even is extremely high, meaning that investment must be very low in order to deliver the required returns. However, $r_G = 0$ and so informed banks are willing to make large loans to the few high-quality firms in the economy. If lending can be directed to such firms, aggregate investment may be higher than under a small or large project. But a huge project cannot occur without intervention, because the informed bank — which has a monopoly over information — cannot commit to taking all of the cash flow once the borrower has announced her choice of K.

Therefore a policy intervention that restricts the ability of the lender to take all of the cash flow, and thus induces the borrower to choose a huge project, can be welfare-enhancing. We show that an interest rate cap is an example of such a policy. We also show that this policy intervention would be justified when both financial distress and corporate productivity are low.

Lemma 3.6. *If the social planner were able to set K, then welfare under a huge project could exceed that under either a small or large project.*

Proof. Suppose that the borrower can be induced to choose a level of investment K consistent with a huge project; that is, $K_M \leq K < K_G$. A huge project is financed if and only if both $\theta = G$ and the insider is not distressed. Therefore total welfare under a huge project is:

$$\Omega_H(K) = q(1-\lambda)(p_G g(K) - 1)K.$$

This is locally maximised when $K = h^{-1}(1+r_G)$. Now, $h(K) < g(K)$ so we certainly have $K < K_G$ at this level, but we cannot guarantee $K \geq K_M$. Thus Ω_H is maximised when:

$$K_H^* = \max\{h^{-1}(1+r_G), K_M)\}.$$

Then the borrower's expected payoff is:

$$\Pi_H = q(1-\lambda)p_G(g(K_H^*) - (1+r_G))K_H^* = \Omega_H(K_H^*).$$

The relevant comparison is with small and large projects under their optimal investment levels. Welfare for small project lending $\Omega_S = (p_M g(K_S^*) - 1)K_S^*$ is maximised when $K_S^* = \min\{h^{-1}(1+r_M), K_N\}$. For large projects, an exact solution is difficult to derive.

Let $\hat{\Omega}_S, \hat{\Omega}_L$ and $\hat{\Omega}_H$ represent the maximised welfare in each case. We need to find an example where $\hat{\Omega}_H$ exceeds both $\hat{\Omega}_S$ and $\hat{\Omega}_L$. Figure 3.5 shows that this can indeed be the case for the same parameter set in Figure 3.3. As one might expect, huge project finance is best when λ is low enough.[13]

Figure 3.5: Total welfare in each case, under an interest rate cap which ensures the socially-optimal investment level. For these parameters, a huge project welfare-dominates a small or large project. The parameters are the same as in Figure 3.3.

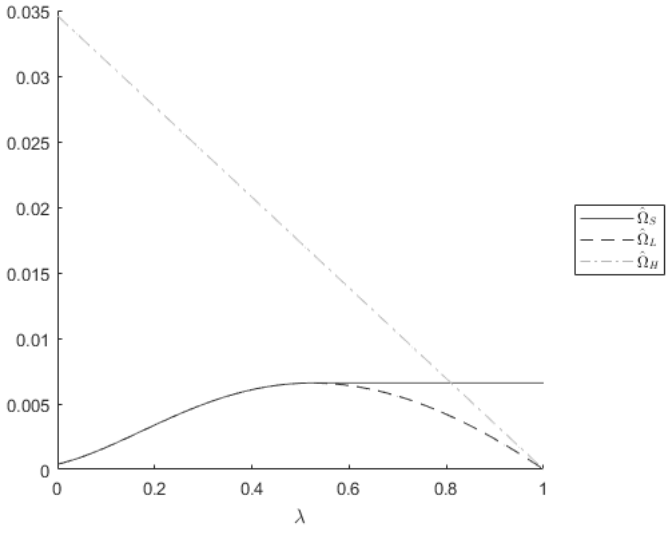

□

Proposition 3.3 (Welfare properties of small and huge projects)**.** *Define* $\Delta := \hat{\Omega}_H - \hat{\Omega}_S$. Δ *is decreasing in both* λ *and* p_B.[14]

Proof. See Appendix. □

[13] In this particular circumstance large projects are never optimal, but that is not true in general.

[14] An equivalent proposition comparing large and huge projects does not in general hold, because welfare for a large project is not monotonic in λ and p_B. It may be possible to establish a weaker version under restrictive assumptions about the parameter values and the curvature of $g(\cdot)$. However, such a result is unlikely to be informative.

Lemma 3.7. *An interest rate cap can induce the borrower to choose a huge project.*

Proof. Suppose that a social planner is able to force the bank not to charge an interest rate above ρ, conditional on lending. The bank is still free not to lend. Then the borrower can certainly make a profit on any funded project of size K where $g(K) > 1 + \rho$. If the social planner chooses $\rho \in [r_G, r_M)$, then the borrower can make a profit with a huge project. The bank will still make a positive profit, and so will lend.

Figure 3.5 shows that the total available surplus available with a huge project can exceed the total surplus with a small or large project. Therefore the borrower may well choose investment consistent with a huge project, if ρ is low enough for her to keep a sufficiently high proportion of the surplus. □

Proposition 3.3 and Lemma 3.7 suggest that an interest rate cap may be socially useful when financial distress is low but aggregate corporate productivity is also poor. For example, it might be used in a recession when there are well-capitalised banks but underinvestment, in order to kick-start investment and overcome information asymmetries. The return on capital cannot meet the interest rates demanded by uninformed lenders. Informed capital — which identifies good projects and so charges lower interest rates — is needed. However, since informed lenders have a monopoly, a policy intervention is required to ensure that firms have some upside and so give them incentives to invest. So long as the bank can take some profit, it is still incentivised to lend. This kind of scenario can expedite a process of Schumpeterian creative destruction, as the capital markets nurture good projects and starve bad ones.

If we allowed the lender and borrower to bargain over the surplus (see, for example, Rajan 1992), then in some cases the borrower may be induced to choose a huge project. For this to work, the borrower would have to be able to credibly threaten to reduce K if the lender did not commit to a lower interest rate.

Other supply-side interventions may work, such as paying corporates directly to invest, incentivising long-term lending contracts, or helping banks to invest in monitoring technology (and thus stimulate entry). However, these interventions all require public funds. An intervention which subsidises or insures credit provision will fail to stimulate investment in this model, because the informed bank will still keep all of the profits.

3.6 Appendix: Proofs

3.6.1 Proof of Lemma 3.1

Suppose that $g(K) < 1 + r_M$. Even in the absence of adverse selection, the highest expected payoff for the outsider is $(p_M g(K) - 1)K < 0$, so the outsider prefers not to make an offer. This is the best response to any strategy of the insider.

As the outsider does not make an offer, the insider knows that he faces no competition in the loans market. He can therefore set his lending rate to $g(K) - 1$ and take all of the cash generated by the project, which is $(p_\theta g(K) - 1)K$. He will do this if the payoff compensates for the risk — that is, if $g(K) \geq 1 + r_\theta$ — and will otherwise not make an offer. As $g(K) < 1 + r_M < 1 + r_B$, in either case the insider never makes an offer for B type loans.

Now suppose that $g(K) > 1 + r_M$. It will be sometimes optimal for the outsider to make an offer. To see why, suppose that the outsider never makes an offer. Then the insider would offer $g(K) - 1$ in order to win the loan with certainty and extract all of the cash. The outsider could then increase his payoff by offering in the range $(r_M, g(K) - 1)$ and winning all of the loans (which have average quality p_M, and so generate a strictly positive expected payoff). But it is not optimal for the outsider to choose pure strategies in this range, because for any given offer the insider can undercut him. Thus any Nash equilibrium must be in mixed strategies.

If $\theta = B$, then the insider will never bid below r_B, because that would entail a strictly negative payoff. Suppose that he makes a bid of r_B of higher. As the worst case for the outsider is a loan of average quality $p_N > p_B$, he will be willing to make an offer in the range $[r_N, r_B)$ and thus will win all of the B loans. Thus the insider never finances type B borrowers.

Finally, consider the case $g(K) = 1 + r_M$. Whatever the action taken by the outsider, it will be optimal for the insider to offer r_M for the G loans (since we wins ties) and not to make an offer for the B loans (since these definitely have negative expected payoff). The only rational responses for the outsider are to offer slightly less than r_M or not to make an offer. But financing the average loan at slightly less than r_M has a negative expected

payoff. Therefore the outsider does not make an offer.

3.6.2 Proof of Proposition 3.1

Let $F_O(r)$ be the density function of the outsider's strategy, conditional on bidding. This has infimum ℓ_O and supremum u_O; in other words:

$$\ell_O = \inf\{r|F_O(r) > 0\}, \qquad u_O = \sup\{r|F_O(r) < 1\}. \tag{3.16}$$

Let β_O be the probability that the outsider bids. Similarly define $F_i^\theta(r), \ell_I^\theta, u_I^\theta, \beta_I^\theta$ for the insider when $\theta \in \{G,B\}$. The payoff functions for the lenders (per unit lent K) are:

$$\begin{aligned}
P_I^G(r) &= \Big(1-\beta_O F_O^-(r)\Big)\Big(p_G(1+r)-1\Big),\\
P_I^B(r) &= \Big(1-\beta_O F_O^-(r)\Big)\Big(p_B(1+r)-1\Big),\\
P_O(r) &= q\Big(\lambda+(1-\lambda)(1-\beta_I^G F_I^G(r))\Big)\Big(p_G(1+r)-1\Big)\\
&\quad +(1-q)\Big(\lambda+(1-\lambda)(1-\beta_I^B F_I^B(r))\Big)\Big(p_B(1+r)-1\Big).
\end{aligned} \tag{3.17}$$

Here $F_O^-(r)$ is defined as the left-limit of $F_O(r)$:

$$F_O^-(r) := \lim_{r'\to r \text{ from below}} F_O(r'). \tag{3.18}$$

We shall require that u_O, u_I^G and u_I^B are all $\le g(K)-1$. Bidding higher than $g(K)-1$ cannot earn the bank any more money (because the borrower cannot pay any more) and also cannot increase the probability of winning the project.[15]

We will prove the Lemma by establishing a series of results.

Result 3.1 (Insider doesn't bid for B types)**.** $\beta_I^B = 0$ *and* $F_I^B(r) = 0$ *for all* $r < r_B$.

[15] Strictly speaking we can obtain multiple equilibria by distributing the mass $1-\beta_O$ over the range (u_O,∞), and similarly for the insider. But the outcomes and payoffs of all of these equilibria would be identical, so it is not interesting to consider them.

Proof. By Lemma 3.1, we knows that the insider doesn't bid for the B types, so this result must hold □

This implies that, for $r < r_B$:

$$\begin{aligned} P_O(r) &= q\Big(\lambda + (1-\lambda)(1-\beta_I^G F_I^G(r))\Big)\Big(p_G(1+r)-1\Big) + (1-q)\Big(p_B(1+r)-1\Big), \\ &= (\lambda q + 1 - q)\Big(p_N(1+r)-1\Big) + q(1-\lambda)\Big(1-\beta_I^G F_I^G(r)\Big)\Big(p_G(1+r)-1\Big). \end{aligned} \tag{3.19}$$

Result 3.2 (Very large and infeasible projects)**.** *When $K \geq K_M$, the outsider never makes an offer.*

Proof. We must have $u_O \leq g(K) - 1 \leq r_M$. As the outsider's average borrower has quality at best p_M, his expected payoff is no more than $(p_M u - 1)K \leq 0$. Thus he clearly does not bid if $K < K_M$. If $K = K_M$, then the only possible strategy will be to bid r_M with certainty. But then the insider will bid r_M with certainty for G types, meaning that the outsider faces a borrower with average quality $p_N < p_M$ and so has strictly negative payoff. Thus $\beta_O = 0$ in this case. □

The statement of the Lemma for huge and infeasible projects follows immediately. With the outsider out of the market, the insider will either bid to take all of the yield from the project (which is worthwhile if $K < K_G$), or not bid at all. For the remainder of this proof, we focus on small and large projects.

Result 3.3 (Lenders don't undercut each other)**.** $\ell_O = \ell_I^G \geq r_M$.

Proof. Suppose that $\ell_O < \ell_I^G$. Choose any $\varepsilon \in (0, \ell_I^G - \ell_O)$. The outsider can do better by moving the mass on $[\ell_O, \ell_I^G - \varepsilon)$ to $\ell_I^G - \varepsilon$. He will win the same set of projects, and finance them at a strictly higher rate.

Similarly, if $\ell_I^G < \ell_O$, then the insider can strictly raise his payoff by deviating.

Finally, the best possible pool that the outsider faces has quality p_M, so he will not lend if his best possible payoff $p_M(1+\ell_O) - 1 < 0$. Thus we must have $\ell_O \geq r_M$. □

Given Result 3.3, we can write $\ell := \ell_O = \ell_I^G$ and drop the indices.

Result 3.4 (Atoms only at the supremum or at ∞)**.** *F_I^G and F_O are continuous on $[\ell, u_I^G)$ and $[\ell, u_O)$, respectively.*

Proof. We prove this by contradiction. Suppose that F_O is discontinuous at some $r \in [\ell, u_O)$. Then $F_O^-(r) < F_O(r)$. Let $\varepsilon = F_O(r) - F_O^-(r) > 0$.

For all $\delta > 0$, $F_O^-(r+\delta) \geq F_O(r)$ and so:

$$\begin{aligned} P_I^G(r+\delta) &= \Big(1 - \beta_O F_O^-(r+\delta)\Big)\Big(p_G(1+r+\delta) - 1)\Big), \\ &\leq \Big(1 - \beta_O F_O(r)\Big)\Big(p_G(1+r+\delta) - 1)\Big), \\ &= \Big(1 - \beta_O (F_O^-(r) + \varepsilon)\Big)\Big(p_G(1+r+\delta) - 1)\Big), \\ &= P_I^G(r) - \varepsilon\beta_O\Big(p_G(1+r) - 1\Big) + \delta p_G\Big(1 - \beta_O(F_O^-(r) + \varepsilon)\Big). \end{aligned} \tag{3.20}$$

Define $\hat{\delta} := \frac{\varepsilon\beta_O(r - r_G)}{1 - \beta_O(F_O^-(r)+\varepsilon)}$. This is strictly positive since $r \geq \ell \geq r_M > r_G$, by Result 3.3.

For any $\delta < \hat{\delta}$, $P_I^G(r+\delta) < P_I^G(r)$. This means that the insider will place zero measure on the interval $(r, r+\hat{\delta})$, and so $F_I^G(\cdot)$ will be constant on this interval.

Consider any r_1, r_2 such that $r < r_1 < r_2 < r + \hat{\delta}$. Then $F_I^G(r_1) = F_I^G(r_2)$, and so:

$$\begin{aligned} P_O(r_2) - P_O(r_1) &= \Big((\lambda q + 1 - q)p_N + q(1-\lambda)(1 - \beta_I^G F_I^G(r_2))p_G\Big)(r_2 - r_1), \\ &> 0. \end{aligned} \tag{3.21}$$

Thus $P_O(\cdot)$ is strictly increasing on $(r, r+\hat{\delta})$, and is continuous everywhere on this interval. This implies that $F_O(\cdot)$ will be constant on this interval, because the point $r + \hat{\delta}$ is strictly preferred to all other points. In particular, this implies that $F_O^-(r) = F_O(r)$, because no mass is placed at r. But that contradicts $F_O^-(r) < F_O(r)$.

The proof of the continuity of F_I^G follows exactly the same logic. Suppose that F_I^G is

discontinuous at some $r \in [\ell, u_I^G)$. Let $\varepsilon = F_I^G(r) - F_I^{G-}(r) > 0$. For any $\delta > 0$,

$$\begin{aligned} P_O(r+\delta) &\leq (\lambda q + 1 - q)\Big(p_N(1+r+\delta) - 1\Big) \\ &\quad + q(1-\lambda)\Big(1 - \beta_I^G(F_I^{G-}(r)+\varepsilon)\Big)\Big(p_G(1+r+\delta) - 1\Big), \\ &= P_O(r) + (\lambda q + 1 - q)p_N\delta + q(1-\lambda)\Big(1 - \beta_I^G(F_I^{G-}(r)+\varepsilon)\Big)p_G\delta \\ &\quad - q(1-\lambda)\beta_I^G\varepsilon\Big(p_G(1+r) - 1\Big), \\ &< 0 \qquad \text{when } \delta < \hat{\delta} := \frac{\varepsilon q(1-\lambda)\beta_I^G(p_G(1+r)-1)}{(\lambda q + 1 - q)p_N + qp_G(1-\lambda)(1-\beta_I^G(F_I^{G-}(r)+\varepsilon))}. \end{aligned} \tag{3.22}$$

$\hat{\delta} > 0$ so long as $p_G(1+r) > 1$; i.e. $r > r_G$, which is always true since $\ell \geq r_M > r_G$. (This explains why the Lemma doesn't apply for F_I^B, because $r < r_B$ and so the equivalent $\hat{\delta}$ would be negative. Thus the density function F_I^B doesn't need to be continuous.)

This means that the outsider places zero measure on the interval $(r, r+\hat{\delta})$, and so $F_O(\cdot)$ is constant on this interval. Then $P_I^G(\cdot)$ is continuous and strictly increasing on this interval, which implies that F_I^G is constant on the interval. This contradicts the assumption of discontinuity at r. □

Result 3.5 (Insider and outsider have same supremum)**.** $u_O = u_I^G$.

Proof. Suppose that $u_O < u_I^G$. Let $\mu = F_I^-(u_I^G) - F(u_O)$. The insider can profitably deviate by moving mass μ from $(u_O, u_I^G]$ to u_I^G; he will win the same loans but earn a higher rate. Thus any equilibrium with $u_O < u_I^H$ must have $\mu = 0$; i.e. the insider does not place any probability mass in the interval (u_O, u_I^G). But then the outsider can profitably deviate by moving mass in a neighbourhood of u_O to just below u_I^G.

By a similar argument, we cannot have $u_I^G < u_O$. Thus we must have $u_I^G = u_O$.

Define $u := u_I^G = u_O$. □

Result 3.6 (Insider always makes an offer for G types)**.** *The insider always bids for G type loans; i.e.* $\beta_I^G = 1$.

Proof. Suppose $\beta_I^G < 1$, so that with positive probability the insider does not make an

offer. When this happens, the insider earns zero payoff for sure. As the insider plays mixed strategies and is risk-neutral, he must therefore earn zero expected payoff on any offer.

Consider an offer of ℓ from the insider. $F_O^-(\ell) = 0$, so:

$$P_I^G(\ell) = p_G(1+\ell) - 1. \tag{3.23}$$

This must equal zero, so $\ell = r_G < r_M$. But that contradicts Result 3.3. □

Result 3.7 (Outsider always bids for small projects)**.** *If $K < K_N$, then $\beta_O = 1$.*

Proof. $K < K_N$ implies that $g(K) > 1 + r_N$. Suppose that $\beta_O < 1$. Then the outsider places non-zero probability on a strategy that yields zero expected payoff. Consider a new strategy for the outsider, which is identical to the current strategy except that, instead of not bidding with probability $1-\beta_O$, he places mass $1-\beta_O$ on bidding $g(K)-1$.

As the insider bids no more than $u_I^G \leq g(K) - 1$ by assumption, the worst case for the outsider is that he only wins at $g(K)-1$ if the insider does not bid or if the borrower is an B type. Thus his gain from the new strategy is no worse than:

$$(1-\beta_O)(\lambda q + 1 - q)(p_N g(K) - 1) > 0. \tag{3.24}$$

The lost gains from deviating from the old strategy are zero, because he earned zero when he did not bid. Thus this is a profitable deviation, and so there cannot be an equilibrium with $\beta_O < 1$ and $r_N < g(K) - 1$. □

Result 3.8 (If outsider always bids, he has an atom)**.** *If $\beta_O = 1$ then the outsider places an atom at u.*

Proof. $F_O^-(\ell) = 0$ so, for all $r \in [\ell, u)$:

$$\begin{aligned} P_I^G(r) &= P_I^G(\ell), \\ &= p_G(1+\ell) - 1, \\ &= p_G(\ell - r_G). \end{aligned} \tag{3.25}$$

Take any $\varepsilon > 0$ and consider $P_I^G(u-\varepsilon)$. Then, since $\beta_O = 1$:

$$\begin{aligned}\Big(1-F_O^-(u-\varepsilon)\Big)p_G(u-\varepsilon-r_G) &= p_G(\ell-r_G),\\ \Rightarrow F_O^-(u-\varepsilon) &= 1-\frac{\ell-r_G}{u-\varepsilon-r_G},\\ &\to \frac{u-\ell}{u-r_G} < 1 \quad \text{as } \varepsilon \to 0.\end{aligned} \tag{3.26}$$

The outsider must place an atom of measure $\frac{\ell-r_G}{u-r_G} > 0$ at u. □

Result 3.9 (When the project is large, insider has an atom and outsider does not)**.** *If $K_N \le K < K_M$, then the insider places an atom at u. The outsider's support does not contain its supremum, and his strategy contains no atoms in $[\ell, u)$. The outsider does not always bid.*

Proof. We have $r_N < g(K) - 1$, and so necessarily $u < r_N$. For any $\varepsilon > 0$:

$$P_O(u-\varepsilon) = (\lambda q + 1 - q)(p_N(1+u-\varepsilon)-1) + q(1-\lambda)\Big(1-F_I^G(u-\varepsilon)\Big)\Big(p_G(1+u-\varepsilon)-1\Big). \tag{3.27}$$

Suppose that the insider does not place an atom at u. Then for any $\delta > 0$, there exists an $\hat{\varepsilon} > 0$ such that $F_I^G(u-\varepsilon) > 1-\delta$ for all $\varepsilon < \hat{\varepsilon}$. Choose $\delta < \frac{(\lambda q+1-q)p_N(r_N-u)}{qp_G(1-\lambda)(u-r_G)}$. For any $\varepsilon < \hat{\varepsilon}$:

$$\begin{aligned}P_O(u-\varepsilon) &< (\lambda q + 1 - q)p_N(u-\varepsilon-r_N) + qp_G(1-\lambda)\delta(u-\varepsilon-r_G),\\ &< (\lambda q + 1 - q)p_N(u-r_N) + qp_G(1-\lambda)\delta(u-r_G),\\ &< 0.\end{aligned} \tag{3.28}$$

Then the outsider places no weight in the interval $(u-\hat{\varepsilon}, u]$. But this contradicts the definition of u as the supremum of the support.

Therefore the insider must place an atom at u. The outsider places no atom here, because:

$$P_O(u) = (\lambda q + 1 - q)(p_N(1+u)-1) < (\lambda q + 1 - q)(p_N(1+r_N)-1) = 0. \tag{3.29}$$

Now suppose $\beta_O = 1$. By Result 3.8 the outsider has an atom at u. But we have just shown that he does not. Thus we must have $\beta_O < 1$. □

Result 3.10 (For small projects, $u = r_N$)**.** *If $K < K_N$ then $u = r_N$.*

Proof. Since the outsider makes zero profit in equilibrium, we must have $P_O(r) = 0$ for all r in the outsider's support. In particular:

$$\begin{aligned} 0 &= P_O(u), \\ &= (\lambda q + 1 - q)(p_N(1+u) - 1), \end{aligned} \tag{3.30}$$

since $\beta_I^G = 1$ (by Result 3.6) and $F_I^G(u) = 1$ (by Result 3.5). Thus we must have $u = r_N$. □

Result 3.11 (For large projects, $u = g(K) - 1$)**.** *If $K_N \leq K < K_M$ then $u = g(K) - 1$.*

Proof. $u > g(K) - 1$ is impossible. Suppose that $u < g(K) - 1$. By Result 3.9, we know that the outsider never bids u but that the insider does. Then the insider wins with a bid of u only if the outsider does not bid at all:

$$P_I^G(u) = (1 - \beta_O)(p_G(1+u) - 1). \tag{3.31}$$

Consider the following deviation for the insider: instead of placing his atom at u, he places it at some $u' \in (u, g(K) - 1]$. Then:

$$P_I^G(u') = (1 - \beta_O)(p_G(1+u') - 1) > P_I^G(u), \tag{3.32}$$

and this is a profitable deviation. Thus the insider sets u as high as possible, and we must have $u = g(K) - 1$. □

Result 3.12 (Lower bound is r_M)**.** *$\ell = r_M$.*

Proof. ℓ is part of the outsider's strategy space. As the outsider has zero expected payoff, we must have $P_O(\ell) = 0$. Thus, since $F_I^G(\ell) = 0$, we have from Equation (3.17):

$$\begin{aligned} 0 &= q(p_G(1+\ell) - 1) + (1-q)(p_B(1+\ell) - 1), \\ &= p_M(1+\ell) - 1, \end{aligned} \tag{3.33}$$

which implies that $\ell = r_M$. □

Result 3.13. *The strategies are as stated in the Proposition.*

Proof. We have already proved the Proposition for huge and infeasible projects in Result 3.2. Thus we focus here only on small and large projects.

$P_O(r) = 0$ for all $r \in [\ell, u)$. That implies that, for all $r \in [\ell, u)$, we have:

$$\begin{aligned} 0 &= q\Big(\lambda + (1-\lambda)(1-F_I^G(r))\Big)\Big(p_G(1+r)-1\Big) + (1-q)\Big(p_B(1+r)-1\Big), \\ \Rightarrow q(1-\lambda)F_I^G(r)\Big(p_G(1+r)-1\Big) &= p_M(1+r)-1, \\ \Rightarrow F_I^G(r) &= \frac{p_M}{qp_G(1-\lambda)} \cdot \frac{r-r_M}{r-r_G}. \end{aligned} \tag{3.34}$$

Note that:

$$\lim_{r \to r_N} F_I^G(r) = \frac{p_M}{qp_G(1-\lambda)} \cdot \frac{r_N - r_M}{r_N - r_G} = 1, \tag{3.35}$$

so the insider indeed has no atom when $u = r_N$ (which is the case when $g(K) > 1 + r_N$, by Result 3.10).

When $K > K_N$, then $u = g(K) - 1 < r_N$ (by Result 3.11) and so $F_I^G(u) < 1$. Therefore the outsider has an atom at u with measure $1 - F_I^G(g(K) - 1)$.

We derive $F_O(r)$ in a similar fashion. The insider's strategy space contains $\ell = r_M$, so for all $r \in [\ell, u)$:

$$\begin{aligned} P_I^G(r_M) &= P_I^G(r), \\ \Rightarrow p_G(1+r_M) - 1 &= \Big(1 - \beta_O F_O^-(r)\Big)\Big(p_G(1+r)-1\Big), \\ \Rightarrow 1 - \beta_O F_O^-(r) &= \frac{r_M - r_G}{r - r_G}, \\ \Rightarrow \beta_O F_O^-(r) &= \frac{r - r_M}{r - r_G}. \end{aligned} \tag{3.36}$$

When $g(K) > 1 + r_N$, then $\beta_O = 1$ (by Result 3.7) and $u = r_N$. $F_O^-(r_N) = \frac{r_N - r_M}{r_N - r_G} < 1$, so the outsider places an atom of measure $1 - F_O^-(r_N) = \frac{r_M - r_G}{r_N - r_G}$ at r_N.

When $g(K) \leq 1 + r_N$, then $u = g(K) - 1$. We know from Result 3.9 that the insider's strategy space contains u, and so Equation (3.36) will hold when $r = u$. We also know that the outsider has no atom at u so $F_O^-(g(K) - 1) = 1$. This means that we must have:

$$\beta_O = \frac{g(K) - 1 - r_M}{g(K) - 1 - r_G}. \tag{3.37}$$

and, for $r \in [r_M, g(K) - 1)$:

$$F_O(r) = F_O^-(r) = \left(\frac{g(K) - 1 - r_G}{g(K) - 1 - r_M}\right) \cdot \frac{r - r_M}{r - r_G}, \tag{3.38}$$

where the equality $F_O(r) = F_O^-(r)$ holds by continuity of $F_O(\cdot)$ over this interval, as shown in Result 3.4. □

3.6.3 Proof of Lemma 3.2

The outsider has zero payoff by assumption of perfect competition. The insider never makes offers when $\theta = B$ (by Lemma 3.1) and so has zero payoff.

As we have mixed strategies with small and large projects, the expected payoff of the insider when $\theta = G$ equals the payoff if he plays any member of the support of his strategy. Thus the expected payoff equals $P_I^G(\ell) = p_G(1 - \lambda)(r_M - r_G)$ in both cases. This is the required value J.

If a project of type θ is financed, the expected welfare is $(p_\theta g(K) - 1)K$. The lending rate is a transfer from borrower to lender and so is irrelevant to total welfare. All small projects are financed with probability 1 and so total welfare is $(p_M g(K) - 1)K$. For large projects, the outsider sometimes does not make an offer and so B projects are financed with probability $F(g(K) - 1)$ and G projects with probability $1 - \lambda(1 - F(g(K) - 1)$. This gives the required values for Ω_S and Ω_L. The borrower's expected payoff is then total welfare minus total expected payoff to the lender.

For a huge project, only type G will be funded and the insider takes all of the cash. No infeasible projects are funded. The payoffs stated in the lemma follow immediately.

3.6.4 Proof of Lemma 3.4

Small projects

1. If the boundary condition $g(K) > 1 + r_N$ does not bind, K_S^* satisfies the first-order condition:

$$0 = p_M h(K_S^{FOC}) - 1 - qp_G(1-\lambda)(r_M - r_G),$$

where $h(K) := g'(K)K + g(K)$ is the derivative of the revenue function $g(K)K$ with respect to K. This means that:

$$K_S^{FOC} = h^{-1}\Big((1+r_M)(1+qp_G(1-\lambda)(r_M-r_G))\Big), \tag{3.39}$$

where we know that $h(K)$ has a well-defined inverse because it is strictly decreasing and continuous (by assumption). Thus K_S^{FOC} is unique. As $h^{-1}(\cdot)$ is a decreasing function, K_S^{FOC} is increasing in λ. Thus the desired result holds in cases where the boundary constraint for the small project does not bind.

Continuity and uniqueness of the first-order condition implies that $K_S^* = K_S^{FOC}$ is the solution iff $K_S^{FOC} < K_N$; otherwise the solution is $K_S^* = K_N$. Thus, the solution is:

$$K_S^* = \min\Big\{h^{-1}\Big((1+r_M)(1+qp_G(1-\lambda)(r_M-r_G))\Big), K_N\Big\}. \tag{3.40}$$

r_N is strictly decreasing in λ and $g^{-1}(\cdot)$ is an increasing function. Thus, both arguments of the right-hand side of (3.40) are strictly increasing in λ, and so K_S^* is strictly increasing in λ, and is unique, as required.

2. Consider any two points $\lambda_1, \lambda_2 \in [0,1]$ with $\lambda_1 < \lambda_2$. From part 1 of this proof,

we know that $K_S^*(\lambda_2) > K_S^*(\lambda_1)$. We have:

$$\begin{aligned}
\Omega_S(\lambda_2) - \Omega_S(\lambda_1) &= (p_M g(K_S^*(\lambda_2)) - 1)K_S^*(\lambda_2) - (p_M g(K_S^*(\lambda_1)) - 1)K_S^*(\lambda_1), \\
&> p_M g(K_S^*(\lambda_1))(K_S^*(\lambda_2) - K_S^*(\lambda_1)) - (K_S^*(\lambda_2) - K_S^*(\lambda_1)), \\
&\qquad \text{as } g(K_S^*(\lambda_2)) < g(K_S^*(\lambda_1)), \\
&> (p_M(1+r_M) - 1)(K_S^*(\lambda_2) - K_S^*(\lambda_1)), \\
&\qquad \text{as } g(K_S^*) > 1 + r_M, \text{ by definition of small project,} \\
&= 0.
\end{aligned} \tag{3.41}$$

Therefore Ω_S is strictly increasing over $\lambda \in [0,1)$.

3. As $\Pi_S - \Omega_S = -qp_G(1-\lambda)(r_M - r_G)K_S^*$ is strictly increasing in $\lambda \in [0,1)$, it follows immediately from part 2 that Π_S is also strictly increasing in $\lambda \in [0,1)$.

Moreover, we can show that Π_S is strictly convex in λ when $K_S^* = K_S^{FOC}$. By the Envelope Theorem,

$$\frac{\mathrm{d}\Pi_S}{\mathrm{d}\lambda}\Big|_{K=K_S^{FOC}} = qp_G(r_M - r_G)K_S^{FOC} > 0,$$

and the second derivative of Π_S is positive, because K_S^{FOC} is increasing in λ.

Large projects

1. From Lemma 3.2, the borrower's payoff for a large project is:

$$\begin{aligned}
\Pi_L &= q(1-\lambda(1-F(g(K)-1)))(p_G g(K) - 1)K \\
&\quad + (1-q)F(g(K)-1)(p_B g(K) - 1)K - qp_G(1-\lambda)(r_M - r_G)K, \\
&= q(p_G g(K) - 1)K + (1-q)F(g(K)-1)(p_B g(K) - 1)K - qp_G(r_M - r_G)K \\
&\quad - \lambda q\Big((1 - F(g(K)-1))(p_G g(K) - 1) - p_G(r_M - r_G)\Big)K.
\end{aligned} \tag{3.42}$$

Using the definition of $F(\cdot)$ in Equation (3.11), the term in λ in (3.42) disappears:

$$\begin{aligned}(1-F(g(K)-1))(p_G g(K)-1)-p_G(r_M-r_G)&\\ &=\Big(\frac{r_M-r_G}{g(K)-1-r_G}\Big)(p_G g(K)-1)-p_G(r_M-r_G),\\ &=\Big(\frac{r_M-r_G}{g(K)-1-r_G}\Big)(p_G g(K)-1-p_G(g(K)-1-r_G)),\\ &=0.\end{aligned} \tag{3.43}$$

Therefore Π_L has no explicit dependence on λ:

$$\begin{aligned}\Pi_L &= q(p_G g(K)-1)K+(1-q)F(g(K)-1)(p_B g(K)-1)K-qp_G(r_M-r_G)K,\\ &= qp_G(g(K)-1-r_M)K+(1-q)\Big(\frac{g(K)-1-r_M}{g(K)-1-r_G}\Big)(p_B g(K)-1)K,\\ &= \frac{p_M(g(K)-1-r_M)^2}{g(K)-1-r_G}K.\end{aligned} \tag{3.44}$$

Π_L depends only on the loan size K and not on λ. As λ increases, any advantage that the borrower has from reduced adverse selection is exactly offset by the risk that the relationship bank is unable to finance the loan. The only way in which λ affects the payoff is when the boundary condition binds: the lower bound of the feasible region for large projects requires $K > K_N$, which is a function of λ.

When this condition binds, the borrower sets $K = K_N$, which is increasing with respect to λ. When it does not, the borrower sets K to locally maximise Π_L, and so K is constant with respect to λ. In either case, K is non-decreasing with respect to λ, and in uniquely determined.

2. We prove that Π_L is non-increasing in λ by establishing a sequence of results. We have already established that Π_L does not depend directly on λ.

Result 3.14. *If Π_L has a unique global maximum in K, then Π_L is non-decreasing in K.*

Proof. The borrower chooses K to maximise Π_L, subject to $K \in [K_N, K_M]$. Suppose that $\Pi_L(K)$ has a unique global maximum K_L^{FOC}. Then, by continuity of Π_L over $[K_N, K_M]$,

the borrower's solution is:

$$K = \begin{cases} K_N, & \text{if } K_L^{FOC} < K_N, \\ K_L^{FOC}, & \text{if } K_N \leq K_L^{FOC} \leq K_M, \\ K_M, & \text{if } K_L^{FOC} > K_M. \end{cases} \tag{3.45}$$

Neither K_L^{FOC} nor K_M depend on λ, so the only way in which λ affects the solution is via its effect on the boundary K_N.

When $K_L^{FOC} \geq K_N$, then K is constant with respect to λ, and so $\Pi_L(K)$ is also constant. Thus, in particular, Π_L is non-increasing with respect to λ.

When $K_L^{FOC} < K_N$, then $K = K_N$ is increasing with respect to λ. As $\Pi_L(K)$ is unimodal and K is above the point which maximises Π_L, we have $\Pi_L'(K) < 0$. Thus Π_L is decreasing with respect to λ. □

Result 3.15. *The borrower never chooses $K = K_M$. In other words, the upper constraint never binds.*

Proof. Suppose this is not true. Then any maximum of $\Pi_L(K)$ must satisfy $K \geq K_M$.

$$\begin{aligned} \Pi_L(K) &= p_M(g(K) - 1 - r_M)K \times \Big(\frac{g(K) - 1 - r_M}{g(K) - 1 - r_G}\Big), \\ \Rightarrow \Pi_L'(K) &= p_M(h(K) - 1 - r_M)\Big(\frac{g(K) - 1 - r_M}{g(K) - 1 - r_G}\Big) \\ &\quad + p_M(g(K) - 1 - r_M)K\Big(\frac{r_M - r_G}{(g(K) - 1 - r_G)^2}\Big)g'(K). \end{aligned} \tag{3.46}$$

When $K > K_M$, then $h(K) < g(K) < 1 + r_M$ and so $\Pi_L'(K) < 0$. $\Pi_L' = 0$ at K_M, but this is a minimum (as can clearly be seen by the fact that $\Pi_L = 0$ at K_M but is strictly positive just beneath it). Thus any maximum in Π_L must satisfy $K < K_M$. □

Result 3.16. *$\Pi_L(K)$ has a unique maximum on $[K_N, K_M]$.*

Proof. Recall that $g(K) = K^\gamma$ for some $-1 < \gamma < 0$. Differentiate Π_L with respect to K to

obtain the following first-order condition in K:

$$0 = \frac{2\gamma K^\gamma (K^\gamma - 1 - r_M)}{K^\gamma - 1 - r_G} - \frac{\gamma K^\gamma (K^\gamma - 1 - r_M)^2}{(K^\gamma - 1 - r_G)^2} + \frac{(K^\gamma - 1 - r_M)^2}{K^\gamma - 1 - r_G}. \tag{3.47}$$

One solution is $K^\gamma = 1 + r_M$, but we know from Result 3.15 that this is a minimum, not a maximum. Let us take out $K^\gamma - 1 - r_M$ as a factor and look for other turning points.

$$0 = \frac{2\gamma K^\gamma}{K^\gamma - 1 - r_G} - \frac{\gamma K^\gamma (K^\gamma - 1 - r_M)}{(K^\gamma - 1 - r_G)^2} + \frac{K^\gamma - 1 - r_M}{K^\gamma - 1 - r_G}. \tag{3.48}$$

This rearranges to the following quadratic equation in K^γ:

$$0 = (\gamma+1)K^{2\gamma} - \Big((1+r_M)(1-\gamma) + (1+r_G)(1+2\gamma)\Big)K^\gamma + (1+r_M)(1+r_G). \tag{3.49}$$

Equation (3.49) has at most two distinct real solutions in K^γ. Therefore we have at most two stationary points. We need to show that at most one of these stationary points is a maximum in the required interval $K \in [K_N, K_M]$.

Note that $\Pi_L(K)$ is continuous everywhere in this interval — its only singularity is at $g(K) = 1 + r_G$, which is well outside of the interval. $\Pi_L(K)$ also has continuous derivatives on this interval. Two consecutive stationary points on a continuous function cannot both be maxima.[16] This completes the proof. □

3. $\Omega_L = \Pi_L + qp_G(1-\lambda)(r_M - r_G)K_L^*$, so when $\lambda = 1$, then $r_N = r_M$ and $\Omega_L = \Pi_L = 0$.

When $\lambda = 0$, $r_N = r_B$ and so:

$$\Omega_L = p_M \frac{(r_B - r_M)^2}{r_B - r_G} K_L^* + qp_G(r_M - r_G)K_L^* > 0.$$

Note that we cannot make a stronger statement about Ω_L than this. During the regime

[16] We can prove this as follows. Let the two consecutive stationary points x_1 and x_2 be maxima, where $x_1 < x_2$. Then $\Pi_L'(x_1) = \Pi_L'(x_2) = 0$ and, for all points in $x \in (x_1, x_2)$, $\Pi_L'(x) \neq 0$. Since both x_1 and x_2 are maxima, we have $\Pi_L''(x_1) \leq 0$ and $\Pi_L''(x_2) \leq 0$. That means that there is a neighbourhood of x_1 where $\Pi_L' \leq 0$ at all points in this set, and a neighbourhood of x_2 where $\Pi_L' \geq 0$ at all points. The intermediate value theorem then implies that $\Pi_L'(x) = 0$ for some $x \in (x_1, x_2)$. But this contradicts the assumption that x_1, x_2 are consecutive stationary points.

when $K_L^* = K_L^{FOC}$, K_L^* and Π_L are both constant and so $\Omega_L = \Pi_L + qp_G(1-\lambda)(r_M - r_G)K_L^*$ is strictly decreasing in λ. However, when $K_L^* = K_N$ (that is, when $\lambda \geq \lambda_B$), then K_L^* is increasing in λ while Π_L is decreasing, so the direction of change of Ω_L depends on the trade-off between these two effects.

3.6.5 Proof of Proposition 3.2

1. By Lemma 3.4, Π_S is strictly increasing in λ and Π_L is weakly decreasing. Therefore there is a unique point λ^* at which they cross, and $\Pi_S > \Pi_L$ iff $\lambda > \lambda^*$. We just need to show that $\lambda^* < 1$. We will do this by showing that $\Pi_S > \Pi_L$ at $\lambda = 1$.

When $\lambda = 1$, $p_N = p_M$ and so $g(K) = 1 + r_M$. Then Equation (3.44) implies that $\Pi_L = 0$. We need to show that $\Pi_S > 0$ when $\lambda = 1$.

When $\lambda = 1$, $K_S^* = \min\{h^{-1}(1+r_M), K_M\}$ from (3.40). For any K, $h(K) = g'(K)K + g(K) < g(K)$, so $h^{-1}(R) < g^{-1}(R)$ for any R. Therefore $K_S^* = h^{-1}(1+r_M)$, and so:

$$\begin{aligned}
\Pi_S &= (p_M g(K_S^*) - 1)K_S^*, \\
&> (p_M h(K_S^*) - 1)K_S^*, \\
&= (p_M(1+r_M) - 1)K_S^*, \\
&= 0,
\end{aligned}$$

and so we must have $\lambda^* < 1$. Note that we cannot rule out that $\lambda^* < 0$. In this event, a large project is never chosen.

2. The idea behind the proof is simple. By definition $K_S^* \leq K_N \leq K_L^*$. Thus they can only be arbitrarily close when the boundary condition binds for both, so that $K_S^* = K_L^* = K_N$. We will show that this can never happen at λ^*. In fact, it can never happen for any λ.

First we prove an auxiliary result, which tells us about the kinks in Figure 3.3:

Result 3.17. *There exist unique λ_S, λ_L such that:*

$$K_S^* = \begin{cases} K_N, & \text{if } \lambda \leq \lambda_S, \\ K_S^{FOC}, & \text{if } \lambda > \lambda_S. \end{cases} \tag{3.50}$$

and:

$$K_L^* = \begin{cases} K_L^{FOC}, & \text{if } \lambda \leq \lambda_L, \\ K_N, & \text{if } \lambda > \lambda_L, \end{cases} \tag{3.51}$$

where K_S^{FOC} and K_L^{FOC} are, respectively, the unconstrained maximisers of $\Pi_S(K)$ and $\Pi_L(K)$. λ_S and λ_L need not lie in $[0,1]$, but $\lambda_S < 1$ always holds.

Proof. Consider the small project first. In the proof of Lemma 3.4, we showed that Π_S has a unique unconstrained maximum at K_S^{FOC}, and that the borrower picks $K_S^* = \min\{K_N, K_S^{FOC}\}$.

Define the set T_S to be the set of $\lambda \in \mathbb{R}$ where $K_S^* = K_S^{FOC}$; i.e.:

$$T_S = \{\lambda \in \mathbb{R} \mid K_N > K_S^{FOC}\}. \tag{3.52}$$

At $\lambda = 1$, $K_S^{FOC} = h^{-1}(1+r_M) < K_M = K_N$, so $1 \in T_S$.

Given the functional form of $g(K)$, we can write $h(K) = g(K) + g'(K)K = (1+\gamma)K^\gamma$. Consider any $\lambda \in T_S$. We have:

$$\begin{aligned} & K_N > K_S^{FOC}, \\ \Leftrightarrow\ & (1+r_N)^{1/\gamma} > \Big(\frac{(1+r_M)(1+qp_G(1-\lambda)(r_M-r_G))}{1+\gamma}\Big)^{1/\gamma}, \\ \Leftrightarrow\ & (1+r_N) < \Big(\frac{(1+r_M)(1+qp_G(1-\lambda)(r_M-r_G))}{1+\gamma}\Big), \quad \text{as } \gamma < 0, \\ \Leftrightarrow\ & (1+\gamma)(1+r_N) < (1+r_M)(1+qp_G(1-\lambda)(r_M-r_G)), \\ \Leftrightarrow\ & p_M(1+\gamma) < p_N(1+qp_G(1-\lambda)(r_M-r_G)), \\ \Leftrightarrow\ & p_M(1+\gamma)(\lambda q+1-q) < (\lambda q p_G + (1-q)p_B)(1+qp_G(1-\lambda)(r_M-r_G)). \end{aligned} \tag{3.53}$$

Thus $\lambda \in T_S$ if and only if $X(\lambda) < 0$, where $X(\cdot)$ is a quadratic function with positive leading coefficient. This means that T_S is some open interval (λ_S, λ_2). We know that

$1 \in T_S$, so $X(1) < 0$ and $\lambda_2 > 1$. We have proved the Lemma for small projects.

Large projects are more straightforward. K_L^{FOC} is constant with respect to λ while K_N is strictly increasing in λ. Define λ_L to be their unique crossing point, and the proof is complete. (Note that λ_L might not lie in the feasible region $[0,1]$.)

□

All that remains is to show that $\lambda_S < \lambda_L$, so that there is no λ where $K_S^* = K_L^* = K_N$.

Result 3.18. $\lambda_S < \lambda_L$.

Proof. Suppose that $\lambda_S \geq \lambda_L$. Choose an $\lambda \in [\lambda_L, \lambda_S]$. At this point:

$$\begin{aligned}
\Pi_L &= p_M \frac{(g(K_N) - 1 - r_M)^2}{g(K) - 1 - r_G} K_N, \\
&= p_M \frac{(r_N - r_M)^2}{r_N - r_G} K_N, \\
&= p_M \Big(1 - \frac{r_M - r_G}{r_N - r_G}\Big)(r_N - r_M) K_N, \\
&= p_M (r_N - r_M) K_N - \frac{p_M (r_N - r_M)}{r_N - r_G} (r_M - r_G) K_N.
\end{aligned} \tag{3.54}$$

Now,

$$\begin{aligned}
\frac{\lambda q p_G + (1-q) p_B}{\lambda q + 1 - q} &= p_N, \\
\Rightarrow (\lambda q p_G + (1-q) p_B)(1 + r_N) &= \lambda q + 1 - q, \\
\Rightarrow p_M (1 + r_N) &= \lambda q + 1 - q + q(1-\lambda) p_G (1 + r_N), \\
\Rightarrow p_M (1 + r_N) - p_M (1 + r_M) &= \lambda q - q + q(1-\lambda) p_G (1 + r_N), \\
\Rightarrow p_M (r_N - r_M) &= q(1-\lambda) p_G (r_N - r_G).
\end{aligned} \tag{3.55}$$

Using Equation (3.55) in (3.54) gives:

$$\begin{aligned}
\Pi_L &= p_M (r_N - r_M) K_N - q p_G (1-\lambda)(r_M - r_G) K_N, \\
&= (p_M g(K_N) - 1) K_N - q p_G (1-\lambda)(r_M - r_G) K_N, \\
&= \Pi_S.
\end{aligned} \tag{3.56}$$

We know from the first part of the Proposition that $\Pi_S = \Pi_L$ if and only if $\lambda = \lambda^*$. Thus $[\lambda_L, \lambda_S]$ can only contain a single point λ^*, and we must have $\lambda_S = \lambda_L = \lambda^*$. All that remains is to show that this leads to a contradiction.

Suppose that $\lambda = \lambda^* = \lambda_S = \lambda_L$. Then $K_L^{FOC} = K_N$ at this point, and so K_N is the unique local maximum for Π_L.

$$\begin{aligned}
\Pi_L &= p_M(g(K_L^*) - 1 - r_M)K_L^* \times \Big(\frac{g(K_L^*) - 1 - r_M}{g(K_L^*) - 1 - r_G}\Big), \\
\Rightarrow \Pi_L' &= p_M(h(K_L^*) - 1 - r_M)\Big(\frac{g(K_L^*) - 1 - r_M}{g(K_L^*) - 1 - r_G}\Big) \\
&\qquad + p_M(g(K_L^*) - 1 - r_M)K_L^*\Big(\frac{r_M - r_G}{(g(K_L^*) - 1 - r_G)^2}\Big)g'(K_L^*), \\
\Rightarrow 0 &= p_M(h(K_N) - 1 - r_M)\Big(\frac{r_N - r_M}{r_N - r_G}\Big) + p_M(r_N - r_M)K_N\Big(\frac{r_M - r_G}{(r_N - r_G)^2}\Big)g'(K_N), \\
\Rightarrow 0 &= (p_M h(K_N) - 1)(r_N - r_G) + p_M(r_M - r_G)\Big(h(K_N) - g(K_N)\Big),
\end{aligned} \tag{3.57}$$

where the last line uses the fact that $g'(K)K = h(K) - g(K)$ for all K.

Now, since $\lambda = \lambda_S$, we also have $K_S^{FOC} = K_N$, and so:

$$\begin{aligned}
h(K_N) &= (1 + r_M)(1 + qp_G(1 - \lambda)(r_M - r_G)), \\
\Rightarrow p_M h(K_N) &= 1 + qp_G(1 - \lambda)(r_M - r_G).
\end{aligned} \tag{3.58}$$

Substitute this into Equation (3.57) to obtain:

$$
\begin{aligned}
&0 = qp_G(1-\lambda)(r_M-r_G)(r_N-r_G)+(r_M-r_G)\Big(1+qp_G(1-\lambda)(r_M-r_G)-p_M(1+r_N)\Big),\\
\Rightarrow &0 = qp_G(1-\lambda)(r_N-r_G)+1+qp_G(1-\lambda)(r_M-r_G)-p_M(1+r_N),\\
\Rightarrow &(p_M-qp_G(1-\lambda))(1+r_N) = 1+qp_G(1-\lambda)(r_M-1-2r_G),\\
\Rightarrow &(\lambda qp_G+(1-q)p_B)(1+r_N) = 1+qp_G(1-\lambda)(r_M-1-2r_G),\\
\Rightarrow &\lambda q+1-q = 1+qp_G(1-\lambda)(r_M-1-2r_G),\\
\Rightarrow &1 = -p_G(r_M-1-2r_G),\\
\Rightarrow &1+r_G = 1+2r_G-r_M,\\
\Rightarrow &r_M = r_G,
\end{aligned}
\tag{3.59}
$$

which is impossible (except in the trivial cases when $q=1$ or $p_G=p_B$ and there is no asymmetric information). Thus we have a contradiction. □

Result 3.18 is sufficient to prove this part of the Proposition, so we can stop here. However, we can establish e a stronger result.

Lemma 3.8. $\lambda_S < \lambda^* < \lambda_L$.

Proof. Suppose that this is not true. We know from Result 3.18 that $\lambda_S < \lambda_L$. Thus we must have either $\lambda^* \leq \lambda_S < \lambda_L$ or $\lambda_S < \lambda_L \leq \lambda^*$. We shall show that each of these leads to a contradiction.

First suppose that $\lambda^* \leq \lambda_S < \lambda_L$. Then, at λ^*, $K_S^* = K_N \leq K_S^{FOC}$ and $K_L^* = K_L^{FOC} > K_N$. That implies:

$$
\begin{aligned}
\Pi_L(K_L^*) > \Pi_L(K_N) \quad &\text{by definition of } K_L^{FOC},\\
= \Pi_S(K_N) \quad &\text{as shown in Equation (3.56).}
\end{aligned}
\tag{3.60}
$$

But, by definition of λ^*, we must have $\Pi_S = \Pi_L$. We have a contradiction.

Now suppose that $\lambda_S < \lambda_L \leq \lambda^*$. Then, at λ^*, $K_S^* = K_S^{FOC} < K_N$ and $K_L^* = K_N \geq K_L^{FOC}$,

so:

$$\begin{aligned} \Pi_S(K_S^*) &> \Pi_S(K_N) \quad \text{by definition of } K_S^{FOC}, \\ &= \Pi_L(K_N) \quad \text{as shown in Equation (3.56),} \end{aligned} \tag{3.61}$$

and once again we have a contradiction. □

3. At $\lambda = \lambda^*$,

$$\begin{aligned} \Omega_S - \Omega_L &= \Pi_S - \Pi_L + qp_G(1-\lambda^*)(r_M - r_G)(K_S^* - K_L^*), \\ &= qp_G(1-\lambda^*)(r_M - r_G)(K_S^* - K_L^*), \qquad \text{as } \Pi_S = \Pi_L \text{ at } \lambda^* \\ &< 0, \end{aligned}$$

since $K_S^* < K_L^*$. Therefore, while it is beneficial for the borrower to switch to a small project at the point λ^*, it would be socially efficient to remain with a large project. The drop in investment is proportional to $K_L^* - K_S^*$, which we have just shown is bounded below by a strictly positive number. Therefore, when the borrower switches from a large to a small project, there is a non-continuous drop in welfare.

3.6.6 Proof of Lemma 3.5

1. Let us write $P(\theta,\Gamma)$ to denote the probability of a type $\theta \in \{G,B\}$ obtaining financing, where Γ denotes whether the investment level K is consistent with a small or large project; i.e. $\Gamma \in \{S,L\}$. We write $P(\Gamma)$ to denote the probability of an average borrower being financed.

When $\lambda > \lambda^*$, the borrower chooses a small project. We know from Proposition 3.1 that the outsider always makes an offer when the borrower chooses a small project, so the borrower always obtains financing and $P(G,S) = P(B,L) = 1$.

When $\lambda \leq \lambda^*$, the borrower chooses a large project and financing does not occur with certainty. According to Proposition 3.1, bad borrowers are financed so long as the outsider makes an offer. Good borrowers will be financed so long as either the outsider makes an

offer or the insider does not suffer a shock. This implies that:

$$P(B,L) = 1 - \Big(\frac{r_M - r_G}{g(K_L^*) - 1 - r_G}\Big), \qquad P(G,L) = 1 - \lambda\Big(\frac{r_M - r_G}{g(K_L^*) - 1 - r_G}\Big), \tag{3.62}$$

The overall probability of being financed is:

$$P = \begin{cases} 1, & \text{if } \lambda > \lambda^*, \\ 1 - (\lambda q + 1 - q)\Big(\frac{r_M - r_G}{g(K_L^*) - 1 - r_G}\Big), & \text{if } \lambda \leq \lambda^*, \end{cases} \tag{3.63}$$

which is decreasing in λ when $\lambda \leq \lambda^*$ and constant and equal to 1 when $\lambda > \lambda^*$. At $\lambda = \lambda^*$, there is a discontinuity and P jumps up to 1.

2. We can solve this part of the Lemma by showing that the distribution of lending rates exhibits first-order stochastic dominance as λ increases. Consider a pair of distress probabilities $\lambda^1 < \lambda^2$ and let $v^1(r), v^2(r)$ be the distribution of rates that the borrower faces in each case. Then we have, for every r, $v^1(r) \leq v^2(r)$, and is strictly greater for some measurable subset of rates. That will ensure that the expected rate that the borrower faces is decreasing in λ.

We will show that the density function of each lender's offered rate exhibits first-order stochastic dominance in this way. That will ensure that the overall lending rate — which is the lower of the two offered rates — also exhibits first-order stochastic dominance as λ increases.

Let us start with the case of a small project. The outsider's strategy function $F_O(r)$ is:

$$F_O(r) = \begin{cases} \frac{r - r_M}{r - r_G}, & \text{if } r \in [r_M, r_N), \\ 1, & \text{if } r \geq r_N. \end{cases} \tag{3.64}$$

$F_O(r)$ depends on λ only via the upper bound of possible actions r_N, which is decreasing in λ. Thus, for any pair $\lambda^1 < \lambda^2$:

$$F_O^2(r) - F_O^1(r) = \begin{cases} 0, & \text{if } r < r_N^2 \text{ or } r \geq r_N^1, \\ 1 - \frac{r - r_M}{r - r_G} & \text{if } r_N^2 \leq r < r_N^1, \end{cases} \tag{3.65}$$

where we use the superscripts 1 and 2 to denote values taken at λ^1 and λ^2 respectively. Thus $F_O^1(r) \leq F_O^2(r)$ for all r, and is strictly greater for a measurable subset of points.

For the insider, the strategy function is:

$$F_I^G(r) = \begin{cases} \frac{p_M(r-r_M)}{(1-\lambda)qp_G(r-r_G)}, & \text{if } r \in [r_M, r_N], \\ 1, & \text{if } r > r_N. \end{cases} \tag{3.66}$$

Therefore:

$$F_I^{G2}(r) - F_I^{G1}(r) = \begin{cases} 0, & \text{if } r < r_M \text{ or } r \geq r_N^1 \\ 1 - \frac{p_M(r-r_M)}{(1-\lambda)qp_G(r-r_G)}, & \text{if } r_N^2 \leq r < r_N^1, \\ \frac{p_M(r-r_M)}{qp_G(r-r_G)}\left(\frac{1}{1-\lambda^2} - \frac{1}{1-\lambda^1}\right), & \text{if } r_M \leq r < r_N^2, \end{cases} \tag{3.67}$$

so we again have the desired property. Finally, the insider never makes offers to a bad borrower, so the desired property weakly holds in that case. As the actual lending rate that the borrower is offered by each lender is a convex combination of these distribution functions, the offered rate also exhibits first-order stochastic dominance in decreasing λ, with strict dominance in some measurable subset of rates. Thus the expected rate that the borrower faces is strictly decreasing in λ.

The argument is similar in the case of a large project. The difference is that the upper bound of the distributions is no longer dependent on λ, because K is constant with respect to λ whenever the large project is chosen in equilibrium (see Results 3.17 and 3.18).

3. Let ρ_S and ρ_L be the expected rates at which the project is financed in the small and large project cases respectively, conditional on obtaining financing. Then, for $\Gamma \in \{S, L\}$:

$$\Pi_\Gamma = \Big(qp_G P(G,\Gamma) + (1-q)p_B P(B,\Gamma)\Big)\Big(g(K_\Gamma^*) - 1 - \rho_\Gamma\Big)K_\Gamma^*. \tag{3.68}$$

The borrower's expected payoff is the probability of being financed times the probability of success times the payoff in the event of success.

Now we know from the first part of the Lemma that $P(G,S) = P(B,S) = 1$, and $P(B,L) <$

$P(G,L) < 1$, so that:

$$\Pi_S = p_M\Big(g(K_S^*) - 1 - \rho_S\Big)K_S^*, \qquad \Pi_L < p_M\Big(g(K_L^*) - 1 - \rho_L\Big)K_L^*. \tag{3.69}$$

We are interested in what happens at $\lambda = \lambda^*$. By definition, we know that at this point $\Pi_S = \Pi_L$, so we must have:

$$\Big(g(K_S^*) - 1 - \rho_S\Big)K_S^* < \Big(g(K_L^*) - 1 - \rho_L\Big)K_L^*. \tag{3.70}$$

By definition of K_S^*, it maximises the borrower's expected payoff, so it maximises the function $(g(K) - 1 - \rho_S)K$. We also know that this function has a unique maximum. Thus:

$$\Big(g(K_L^*) - 1 - \rho_S\Big)K_L^* < \Big(g(K_S^*) - 1 - \rho_S\Big)K_S^*. \tag{3.71}$$

Putting these together implies:

$$\Big(g(K_L^*) - 1 - \rho_S\Big)K_L^* < \Big(g(K_L^*) - 1 - \rho_L\Big)K_L^*, \tag{3.72}$$

which implies that $\rho_S > \rho_L$ at λ^*. This completes the proof.

3.6.7 Proof of Proposition 3.3

We shall show that $\hat{\Omega}_H$ is strictly decreasing in λ and weakly decreasing in p_B, while $\hat{\Omega}_S$ is weakly increasing in λ and strongly increasing in p_B. These results are sufficient to prove the Proposition.

$\hat{\Omega}_H = q(1-\lambda)(p_G g(K_H^*) - 1)K_H^*$ where $K_H^* = \max\{h^{-1}(1+r_G), K_M\}$.

Then:

$$\frac{\mathrm{d}\hat{\Omega}_H}{\mathrm{d}\lambda} = -q(p_G g(K_H^*) - 1)K_H^* < 0,$$

since $p_G g(K_H^*) - 1 > p_G(1+r_G) - 1 = 0$ and K_H^* is not dependent on λ.

$$\frac{\mathrm{d}\hat{\Omega}_H}{\mathrm{d}p_B} = \frac{\partial\hat{\Omega}_H}{\partial p_B} + \frac{\partial\hat{\Omega}_H}{\partial K_H^*}\frac{\mathrm{d}K_H^*}{\mathrm{d}p_B},$$
$$= 0 + q(1-\lambda)(p_G h(K_H^*) - 1)\frac{\partial K_H^*}{\partial p_B}.$$

Now K_H^* is only affected by the value of p_B when it is at the boundary K_M. In that case, it varies positively with p_B because increasing p_B reduces r_M and $g^{-1}(\cdot)$ is a decreasing function.

As $p_G h(K_H^*) \leq 1$, we have $\dfrac{\mathrm{d}\hat{\Omega}_H}{\mathrm{d}p_B} \leq 0$.

$\hat{\Omega}_S = (p_M g(K_S^*) - 1)K_S^*$ where $K_S^* = \min\{h^{-1}(1+r_M), K_N\}$.

Then:

$$\frac{\mathrm{d}\hat{\Omega}_S}{\mathrm{d}\lambda} = \frac{\partial\hat{\Omega}_S}{\partial\lambda} + \frac{\partial\hat{\Omega}_S}{\partial K_S^*}\frac{\partial K_S^*}{\partial\lambda},$$
$$= 0 + (p_M h(K_S^*) - 1)\frac{\partial K_S^*}{\partial\lambda}.$$

K_S^* is only affected by the value of λ when it is at the boundary K_N. In that case it varies positively with λ because increasing λ reduces r_N.

As $p_M h(K_S^*) \geq 1$, we have $\dfrac{\mathrm{d}\hat{\Omega}_S}{\mathrm{d}\lambda} \geq 0$.

Finally,

$$\frac{\mathrm{d}\hat{\Omega}_S}{\mathrm{d}p_B} = \frac{\partial\hat{\Omega}_S}{\partial p_B} + \frac{\partial\hat{\Omega}_S}{\partial K_S^*}\frac{\partial K_S^*}{\partial p_B},$$
$$= (1-q)g(K_S^*)K_S^* + (p_M h(K_S^*) - 1)\frac{\partial K_S^*}{\partial p_B}.$$

Both r_M and r_N are decreasing in p_B, so K_S^* is certainly increasing in p_B. Thus $\dfrac{\mathrm{d}\hat{\Omega}_S}{\mathrm{d}p_B} > 0$.

Chapter 4

Centralised netting in financial networks[1]

4.1 Introduction

Centralised netting in a financial network is the novation of links to a single counterparty. This can reduce the aggregate level of exposures in the system by netting offsetting claims, and so decrease systemic risk and collateral requirements. Examples of centralised netting include the introduction of a central counterparty (CCP) to over-the-counter derivatives markets, triparty repo, and netting in payments networks.[2] Improvements are most likely when all exposures are simultaneously netted.[3] Introducing centralised netting in a subset of exposures has the effect of improving netting amongst those exposures, because they are now novated to the same counterparty. But it disrupts bilateral netting sets amongst those exposures which are not novated. This paper examines how this trade-off depends on the structure of the network.

We focus on the example of the introduction of a central counterparty (CCP) to a deriva-

[1] This chapter is based on joint work with Rodney Garratt, published as Garratt and Zimmerman (2017).

[2] For example, CLS (Continuous Linked Settlement) nets foreign exchange transactions in payments systems.

[3] This may not be the case when the agents strategically reform the network in response to the introduction of centralised netting. See section 4.6.

tives network. Duffie and Zhu (2011) show that, when a single asset class is centrally cleared, bilateral exposures between agents may increase — as a result of reduced netting opportunities across pairs of counterparties — resulting in an overall loss in netting efficiency. They derive conditions, in terms of the number of asset classes and the number of dealers, for whether or not centralised netting in the form of a CCP is beneficial.

Figure 4.1 below illustrates the effect of the introduction of central clearing on netting. Dealer agents A, B and C trade credit default swaps and interest rate swaps with one another. In the left-hand panel, the banks have exposures to each other in CDS, and A is exposed to C in IRS. This may be because, for example, A has written a CDS to B that have moved out-of-the-money, so A is exposed in the event of B's default. A has offsetting exposures to C and so is able to net them out. However, A's CDS exposure to B and B's CDS exposure to C do not have offsetting IRS exposures, and so they are not netted.

The right-hand panel shows the situation after a CCP is introduced that clears credit default swaps. All three CDS exposures are now novated to the CCP. This improves netting efficiency in the CDS market, because all CDS trading is now against the same counterparty, namely the CCP. For example, B now has an exposure both to and from the CCP (previously these were to C and from A, which could not be netted). However, bilateral netting between C and A is disrupted, because C no longer has a CDS exposure to A with which A can net out the IRS exposure. Nonetheless, this is more than compensated for by the improvements in netting elsewhere, so the overall effect is an improvement in netting system-wide. Note that the CCP's exposures to each dealer agent will always be fully netted out.

Figure 4.2 demonstrates a scenario in which the introduction of central clearing can worsen overall netting. Here, there is a greater degree of bilateral netting compared to Figure 4.1. Now B and C benefit from bilateral netting, as well as C and A. Centralised netting disrupts these netting sets, making the overall situation worse.

In Duffie and Zhu's framework, all agents in the network are assumed to trade with one another. By expressing the magnitude of links between agents in each asset class as a random variable, they can then calculate expected exposures with respect to this distribution, and examine how this changes before and after the introduction of centralised

Figure 4.1: Illustration of how centralised netting can reduce network exposures. Each arrow represents an exposure of one unit. Solid lines represent CDS and dashed lines IRS exposures.

Before CDS netting:

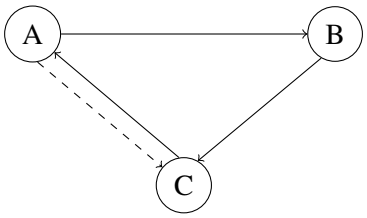

A's net exposure to B = 1
B's net exposure to C = 1
C's net exposure to A = 0
⇒ Total net exposures = 2

After CDS netting:

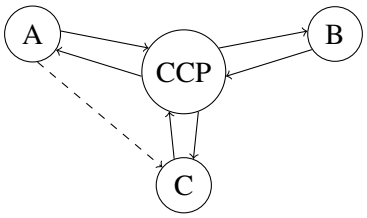

A's net exposure to CCP = 0
B's net exposure to CCP = 0
C's net exposure to CCP = 0
A's net exposure to C = 1
⇒ Total net exposures = 1

Figure 4.2: Illustration of how centralised netting can increase network exposures. Each arrow represents an exposure of one unit. Solid lines represent CDS and dashed lines IRS exposures.

Before CDS netting:

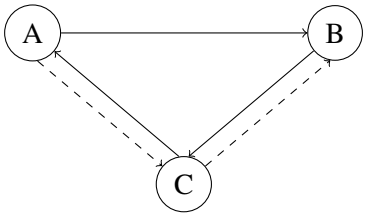

A's net exposure to B = 1
B's net exposure to C = 0
C's net exposure to A = 0
⇒ Total net exposures = 1

After CDS netting:

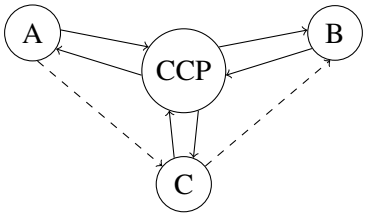

A's net exposure to CCP = 0
B's net exposure to CCP = 0
C's net exposure to CCP = 0
A's net exposure to C = 1
C's net exposure to B = 1
⇒ Total net exposures = 2

netting. A clear advantage of this approach is that it is not necessary to observe the actual network exposures; instead it relates the question of whether or not the introduction of centralised netting in a single asset class is beneficial or not to easily observable parameters. But one drawback to their analysis is the assumption of a completely-connected

network, with all agents trading with one another. Typically, real-world financial networks are not completely connected. Empirical studies show that there is typically a number of well-connected counterparties coupled with a larger number of more poorly-connected counterparties.

In this paper we develop a model of centralised netting that extends the analytic framework developed by Duffie and Zhu to more general network models. We develop analytical results that do not depend on any assumptions about the precise structure of the network. In addition, we examine the effect of centralised netting on the variance, as well as the expectation, of netted exposures. Our rationale is that, if the agents have some aversion to volatility or extreme outcomes, then they are likely to benefit from a lower variance of net exposures, as well as a lower mean. As in Duffie and Zhu, we show that the introduction of centralised netting is more likely to be beneficial when there are more agents, and when there are fewer asset classes.

Our first key result is that, for any given network, if introducing centralised netting reduces expected exposures, then it must also reduce the variance of exposures. However, the converse implication is not true. This means that the set of networks for which the expectation of exposures is reduced is a subset of those for which the variance is reduced.

We derive general expressions to determine cases when centralised netting reduces the expectation and variance of exposures. These expressions are in terms of the network distribution and the number of asset classes. As the benefits of centralised netting are decreasing in the number of asset classes in the network, we can determine the critical number of asset classes above which centralised netting does not deliver benefits. This critical number is higher for the variance than for the expectation of exposures.

These first results require full information about the structure of the network. Our second result weakens this assumption. We establish upper and lower bounds on the maximum number of asset classes required for centralised netting to deliver benefits. These do not depend on knowing the degree distribution of the network. This second result could help a decision maker to judge the benefits of central clearing when the full structure of the network is unknown. We show that assuming the Duffie and Zhu network structure achieves the upper bound, suggesting that their criterion is extreme and will lead to central clearing being introduced too often.

Our additional results relate to specific network structures. The previous literature has demonstrated that real-world financial networks can be accurately modelled as scale-free network models (e.g. Soramäki et al. 2007, Inaoka et al. 2004, Garlaschelli et al. 2005) or by core-periphery models (e.g. Craig and von Peter 2014, Langfield, Liu, and Ota 2014, Markose 2012). For scale-free networks, we find that expected net exposures *always* increase when a single asset is novated to a CCP, regardless of the size of the network. Therefore, for the introduction of central clearing to be a optimal policy, dealer agents or policymakers must have some desire to reduce the variance of net exposures. Moreover, for the asymptotic network we find that, when a single asset is novated to a CCP, expected net exposures always increase and the variance of net exposures always decreases. Hence, the trade-off between mean and variance is a definite feature of the limiting scale-free network. In contrast, for core-periphery networks we find that, for any given number of asset classes, there is a minimum size of the network above which the introduction of a CCP reduces both the mean and variance of net exposures. Thus, for sufficiently large core-periphery networks, centralised netting is unambiguously beneficial.

Our findings can be applied to predict the impact of introducing a CCP on margin requirements, since aggregate margin needs are related to the distribution of netted exposures.[4] Our model can also be applied to any network where the issue of bilateral vs centralised netting is under consideration. For example, it could be used to consider the effect of multilateral netting in the interbank market, the benefits of bilateral vs triparty repo, or the effect on payment system exposures of introducing a netting mechanism (such as CLS or a liquidity-saving mechanism).

The techniques we develop in this paper can be used to evaluate the efficiency gains of novating multiple asset classes to a single CCP or of different configurations of CCPs, which each handle different sets of asset classes. Trivially, if it is efficient to novate one asset class to a CCP, then it will always be efficient to novate another. This follows from the fact that efficiency gains to novation are higher when the number of asset classes is lower and novating an asset effectively reduces the number of asset classes by 1. The logical extension of this idea is that the well-known result that the unconstrained first-best is to put everything through a single CCP. Our final results of the paper show a related

[4] See, for example, Sidanius and Žikeš (2012). More recent papers including Duffie, Scheicher, and Vuillemey (2015) and Campbell (2014) use actual bilateral exposure data to analyse the issue empirically.

finding that (1) merging CCPs will always improve efficiency, again with the ultimate conclusion that a single CCP is best, and (2) it is most efficient to novate asset classes to the largest existing CCP.

We also consider how the results may change if the agents strategically reform the network in response to the introduction of centralised netting. We show that, in certain circumstances, centralised netting may be more efficient when there are *more* asset classes, rather than fewer. This can happen if the post-clearing network is less connected, perhaps as a result of agents raising their bilateral risk limits as a result of the new system infrastructure, and thus needing to form fewer links. Bilateral netting becomes more efficient when there are fewer connections, so it is more likely that the CCP will be beneficial when it disrupts a smaller fraction of the bilateral links.

The final two sections discuss how our results can be used in practice, and the implications of our findings for policy. The issue of clearing networks is likely to come to the fore in Europe in the near future as the European Commission is considering requiring the clearing of euro-denominated swaps to be carried out within the European Union. Most such clearing is done in the United Kingdom, which has announced its intention to leave the European Union in March 2019. Our model allows analysts to forecast the disruption to netting in clearing networks that may result.

4.2 A review of the relevant literature

The framework developed in Duffie and Zhu (2011) has been utilised by other authors in order to investigate specific problems. Heath, Kelly, and Manning (2013) is perhaps the most similar to our paper in that they examine a network other than the completely-connected structure of Duffie and Zhu. They assume a core-periphery structure and use a computational approach to compare the effect of various clearing arrangements on expected netted exposures.

Anderson, Dion, and Pérez Saiz (2013) and Cox, Garvin, and Kelly (2013) apply the Duffie and Zhu framework to explore the policy issue of interoperability between CCPs. They examine whether a regulator can reduce expected netted exposures by mandating

trades to be novated to a local CCP, which can link to a global CCP that clears a range of products. Both papers retain the assumption of a homogeneous link network, though Cox, Garvin and Kelly allow for some heterogeneity between dealer agents in the magnitude of exposures (but not in the existence of links).

Cont and Kokholm (2014) extend the Duffie and Zhu framework by relaxing the assumption of normal exposures between counterparties and show, using a simulation approach, that Duffie and Zhu's conclusions are sensitive to different distributional assumptions. However, they retain the homogeneous network assumption that Duffie and Zhu use. This is in contrast to our paper, which uses more general and realistic network structures.[5] A recent working paper by Menkveld (2015) considers the Duffie and Zhu framework and looks at the mean and variance of the aggregate exposure of the CCP. His goal is to examine the ability of the CCP to survive simultaneous losses in the centrally cleared asset. We, in contrast, look at the impact of introducing a CCP on mean and variance of exposures across all asset classes and our focus is on the aggregate exposures of all counterparties across all asset classes, not just those in the centrally cleared asset class.

Our paper makes two key innovations to the Duffie and Zhu model which, to our knowledge, have not been considered before. First, we provide an *analytical generalisation* of the model so that it can be applied to any network. Second, we look at how the introduction of a CCP affects the *variance* of counterparty exposures, as well as the mean, of net exposures for alternative network structures.

In the broader literature, there are a variety of papers which use a network approach to focus on issues relating to OTC derivatives networks other than netting efficiency. Markose (2012) finds that the empirical OTC derivatives network aggregated over all products can be well-described by a modified core-periphery model, and derives summary statistics to identify institutions which carry the greatest quantity of systemic risk. Heath et al. (2016) use the same data set to show that the empirical network structure has the potential to generate stability risks. Borovkova and Lalaoui El Mouttalibi (2013) use a simulation approach to model the effect of the introduction of a CCP on default cascades in a network. They examine both homogeneous and core-periphery networks, and find that homogeneous networks are more resilient.

[5] While our framework explicitly permits the use of non-normal exposures, some distributional assumption is necessary in order to obtain analytical results.

Jackson and Manning (2007) and Galbiati and Soramäki (2012) use different approaches to examine the desirability of tiering — that is, restricting direct access to the CCP to a limited set of counterparties. Song, Sowers, and Jones (2014) extend the Galbiati and Soramäki framework to study the effect of network structure on the maximum exposure risk of the CCP itself, and use extreme value theory to obtain analytical results.

Our final results discuss the possibility of agents strategically reforming the network in response to the introduction of a CCP. To our knowledge, there is no existing literature which directly addresses this question. However, recent work by Aldasoro, Gatti, and Faia (2016), Chang and Zhang (2015), Hollifield, Neklyudov, and Spatt (2016) and Van der Leij, in't Veld, and Hommes (2016) provides some insight into how networks form under strategic interaction.

4.3 A general network model of exposure netting

We assume that the dealer network is not directly observable but the number of nodes, the degree distribution and the distribution of the magnitude of bilateral exposures is known. This is a realistic assumption for dealer networks, where the regulator and participants often lack exact real-time knowledge of bilateral exposures. This is a generalisation of the assumption made in Duffie and Zhu (2011), in which the exact structure of the network is fixed and known.

Let N be the number of nodes (i.e. market participants) and let S be a random variable denoting the number of links any given node has. S takes values on the non-negative integers. Define J_i to be the set of nodes with which node i has a link. The size of this set is given by a realisation of the random variable S.

Let K denote the number of asset classes. Links are undirected: each link is endowed with a vector reflecting the net exposure in each asset class between the two dealer agents. This may be positive or negative, depending on the direction of the exposure. Define X_{ij} to be a K-vector of values (weights) on the link between nodes i and j, if it exists. The k^{th} element in each vector, denoted X_{ij}^k, represents the net exposure between the two nodes in asset class k. When $X_{ij}^k > 0$ then node i has a net exposure to node j in asset class k,

and when $X_{ij}^k < 0$ then the reverse is true. Each value is generated with the same known distribution, independently of one another and of the link structure of the network.[6]

First consider the situation without a CCP. Consider two linked nodes i and j. Define $Y_{ij}^K \equiv \max\left\{\sum_{k=1}^K X_{ij}^k, 0\right\}$ to be the value of node i's netted exposure to node j. Positive net exposures in one asset class can be partially or wholly offset by negative net exposures in another asset class with the same counterparty. If i and j are not linked, then the net exposure is zero. The total net exposure of node i equals $\sum_{j\in J_i} Y_{ij}^K$.

Now define the function $f(K)$ as the expected net exposure between any two linked nodes, given that there are K asset classes:

$$f(K) \equiv \mathbb{E}\left[Y_{ij}^K\right]. \tag{4.1}$$

The expected total netted exposures for a given node i are:

$$\begin{aligned}\phi_{N,K} &\equiv \mathbb{E}\left[\sum_{j\in J_i} Y_{ij}^K\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\sum_{j\in J_i} Y_{ij}^K \middle| S\right]\right] \\ &= \mathbb{E}\left[Sf(K)\right] \\ &= \mathbb{E}\left[S\right] f(K),\end{aligned} \tag{4.2}$$

where we have used the fact that each Y_{ij}^K is independent from one another, and from S.

Similarly, the variance of the exposure between two nodes after netting is:

$$g(K) \equiv \mathbb{V}\left[Y_{ij}^K\right] \tag{4.3}$$

and the variance of the total netted exposures of the network is:

$$v_{N,K} \equiv \mathbb{V}\left[\sum_{j\in J_i} Y_{ij}^K\right]. \tag{4.4}$$

[6]The constraint $X_{ij}^k = -X_{ji}^k$ does apply, but as our analysis focuses on a representative node, we do not need to apply this constraint.

We can evaluate this expression using the law of total variance:

$$\begin{aligned} v_{N,K} &= \mathbb{E}\left[\mathbb{V}\left[\sum_{j\in J_i} Y_{ij}^K \middle| S\right]\right] + \mathbb{V}\left[\mathbb{E}\left[\sum_{j\in J_i} Y_{ij}^K \middle| S\right]\right] \\ &= \mathbb{E}\left[Sg(K)\right] + \mathbb{V}\left[Sf(K)\right] \\ &= \mathbb{E}\left[S\right] g(K) + \mathbb{V}\left[S\right] f(K)^2. \end{aligned} \tag{4.5}$$

4.3.1 Novating a single asset class to a CCP

Now we introduce a CCP in a single asset class. Without loss of generality, reorder the asset classes so that the centrally cleared asset class is the one labelled K. The net exposure of a given node i becomes:

$$\sum_{j\in J_i} Y_{ij}^{K-1} + \underbrace{\max\left\{\sum_{j\in J_i} X_{ij}^K, 0\right\}}_{Y_{i,CCP}^S} \tag{4.6}$$

where the first term is the sum of a node's exposures to the other nodes, and the second term is its netted exposure to the CCP. We can rewrite the second term as $Y_{i,CCP}^S$, with the S superscript arising because the size of J_i has distribution S.[7]

Note that, for a given realised value of S, the two terms in (4.6) are independent: the first term is determined entirely by exposures arising from the first $K-1$ assets, while the second is determined entirely by exposures arising from asset K.

Now, when there is a CCP, the expected total net exposure of node i is:

$$\begin{aligned} \widetilde{\phi_{N,K}} &= \mathbb{E}\left[\mathbb{E}\left[\left(\sum_{j\in J_i} Y_{ij}^{K-1} + Y_{i,CCP}^S\right)\middle| S\right]\right] \\ &= \mathbb{E}\left[Sf\left(K-1\right) + f(S)\right] \\ &= \mathbb{E}\left[S\right] f\left(K-1\right) + \mathbb{E}\left[f\left(S\right)\right], \end{aligned} \tag{4.7}$$

[7] In section 4.6, we consider the same setting but allow the network structure S to change as a result of the introduction of central clearing.

and the variance of the total net exposure of node i is:

$$\begin{aligned}\widetilde{v_{N,K}} &= \mathbb{E}\left[\mathbb{V}\left[\left(\sum_{j\in J_i} Y_{ij}^{K-1} + Y_{i,CCP}^{S}\right)\Big|S\right]\right] + \mathbb{V}\left[\mathbb{E}\left[\left(\sum_{j\in J_i} Y_{ij}^{K-1} + Y_{i,CCP}^{S}\right)\Big|S\right]\right] \\ &= \mathbb{E}\left[Sg(K-1)+g(S)\right] + \mathbb{V}\left[Sf(K-1)+f(S)\right] \\ &= \mathbb{E}\left[S\right]g(K-1) + \mathbb{E}\left[g(S)\right] + \mathbb{V}\left[Sf(K-1)+f(S)\right].\end{aligned} \tag{4.8}$$

Using (4.2) and (4.7), the change in expected net exposure that results from novating a single asset class to a CCP is:

$$\widetilde{\phi_{N,K}} - \phi_{N,K} = \mathbb{E}\left[f(S)\right] - \mathbb{E}\left[S\right]\left(f(K) - f(K-1)\right). \tag{4.9}$$

Using (4.5) and (4.8), the change in variance that results from novating a single asset class to a CCP is:

$$\begin{aligned}&\widetilde{v_{N,K}} - v_{N,K} = \\ &\mathbb{V}\left[Sf(K-1)+f(S)\right] - \mathbb{V}\left[S\right]f(K)^2 + \mathbb{E}\left[g(S)\right] - \mathbb{E}\left[S\right]\left(g(K)-g(K-1)\right).\end{aligned} \tag{4.10}$$

These results show that introducing centralised netting will change both the mean and variance of total net exposures. Reduction in either the mean or the variance is likely to be positive for users of the system, since it means that counterparty risk — and total margin needs — are lower either in expectation or volatility. Therefore we can say that:

- centralised netting delivers netting benefits when both (4.9) and (4.10) are negative;
- netting is worsened when both expressions are positive; and
- when the expressions have different signs, then there is a trade-off depending on the weight decision-makers place on the mean and the variance of exposures.

The key focus of this paper is to identify the extent to which the effect of centralised netting depends on the underlying network structure.

Note that in the case $K = 1$, the CCP clears all of the assets that the dealers trade with one another. In this case, we would always expect the introduction of a CCP to improve netting, as it does not disrupt any of the existing bilateral netting sets. This is confirmed

by observing that $\widetilde{\phi_{N,1}} \leq \phi_{N,1}$ and $\widetilde{v_{N,1}} \leq v_{N,1}$.[8]

4.3.2 Assigning a distribution to the bilateral exposures

To obtain tractable results for the cases where $K > 1$, we need to assume a distribution for the bilateral exposures between dealers. This will enable us to write down expressions for f and g in equations (4.9) and (4.10). We follow Duffie and Zhu (2011) and assume that each of the bilateral exposures X_{ij}^k is independent and identically distributed normally with mean 0 and variance σ^2.[9] Using the formula for the sum of independent normal random variables, we can write the function f as:

$$\begin{aligned} f(\theta) &= \int_0^\infty \frac{1}{\sqrt{2\pi\theta}\sigma} y e^{-\frac{y^2}{2\theta\sigma^2}} \, dy \\ &= \sigma\sqrt{\frac{\theta}{2\pi}}, \end{aligned} \tag{4.11}$$

and:

$$\begin{aligned} g(\theta) &= \int_0^\infty \frac{1}{\sqrt{2\pi\theta}\sigma} y^2 e^{-\frac{y^2}{2\theta\sigma^2}} \, dy - \left(\sigma\sqrt{\frac{\theta}{2\pi}}\right)^2 \\ &= \frac{1}{2}\int_{-\infty}^\infty \frac{1}{\sqrt{2\pi\theta}\sigma} y^2 e^{-\frac{y^2}{2\theta\sigma^2}} \, dy - \sigma^2\frac{\theta}{2\pi} \\ &= \sigma^2(\pi-1)\frac{\theta}{2\pi}. \end{aligned} \tag{4.12}$$

We substitute these into (4.9) and (4.10) to show that a CCP reduces the expectation of net exposures if and only if:

$$\sqrt{K} + \sqrt{K-1} < \frac{\mathbb{E}[S]}{\mathbb{E}[\sqrt{S}]}, \tag{$\star$}$$

[8] To show this, note that $f(0) = 0 = g(0)$ and that the $\max\{\cdot, 0\}$ function is sub-additive.

[9] We can generalise slightly: the results that follow will also apply asymptotically (i.e. as $N, K \to \infty$) so long as the distribution of the exposures has zero mean and finite variance. This follows from the law of large numbers. In section 4.7.1, we show that using Student's t-distribution instead of the normal does not change the results substantively. This is in line with the findings of Cont and Kokholm (2014).

and introducing a CCP reduces the variance of net exposures if and only if:

$$\mathbb{V}[S\sqrt{K-1}+\sqrt{S}] < K\mathbb{V}[S]. \tag{4.13}$$

Rewriting the first variance term in (4.13) in terms of expectations operators, we can express this condition as:

$$2\sqrt{K-1} < \frac{\mathbb{V}[S]-\mathbb{V}[\sqrt{S}]}{\mathbb{E}[S^{\frac{3}{2}}]-\mathbb{E}[S]\mathbb{E}[\sqrt{S}]}, \tag{†}$$

when the denominator is non-zero.

The denominator of (†) is zero if and only if S is a constant, which implies that the numerator is also zero. In this case, the degree of every node is constant and known with certainty, so the variance of netted exposures does not change upon introduction of centralised netting. We call such a network a "trivial network"; the Duffie and Zhu network is an example of such a network. Trivial networks are not realistic representations of real-world networks, which typically exhibit some heterogeneity in their degree distributions.

4.3.2.1 Effect of changing the number of asset classes K

Proposition 4.1. *As the number of asset classes K increases, there is less benefit from the introduction of centralised netting. This is true whether the benefit is measured in terms of lower expectations or lower variance of netted exposures.*

Proof. We can express this more precisely as follows: for any number of asset classes K, let Σ_K be the set of networks for which centralised netting reduces the mean of netted exposures; i.e. $S \in \Sigma_K$ if and only if S satisfies ($\star$). Then for all $K' > K$, $\Sigma_{K'}$ is a subset of Σ_K. The same is true for the variance. The result follows immediately from the fact that the left-hand sides of both expressions ($\star$) and (†) are increasing in K. □

The intuition behind Proposition 4.1 is that, as K increases, the CCP clears a lower proportion of the dealers' activity with one another, so the benefit of netting with the CCP is reduced.[10] This is consistent with the findings of Duffie and Zhu.

[10] See Section 6 for an examination of the robustness of Proposition 1 to a behavioural response to the

For a given network structure, we can define K^* as the critical value of K below which expected net exposures decrease upon introduction of a CCP, and above which they increase. $K^\dagger$ is the corresponding critical value for variance. Then:

$$K^* \equiv \frac{1}{4}\left(\frac{\mathbb{E}[S]}{\mathbb{E}[\sqrt{S}]} + \frac{\mathbb{E}[\sqrt{S}]}{\mathbb{E}[S]}\right)^2 \tag{4.14}$$

$$K^\dagger \equiv \frac{1}{4}\left(\frac{\mathbb{V}[S] - \mathbb{V}[\sqrt{S}]}{\mathbb{E}[S^{\frac{3}{2}}] - \mathbb{E}[S]\mathbb{E}[\sqrt{S}]}\right)^2 + 1 \tag{4.15}$$

In other words, expressions ($\star$) and ($\dagger$) are equivalent to, respectively, $K < K^*$ and $K < K^\dagger$. These expressions are the focus of the analysis moving forward. They relate the impact of introducing a CCP to the number of asset classes K on the left-hand side, and to the degree distribution S on the right-hand side.

Next we present a general result on the relationship between K^* and $K^\dagger$ that holds for any finite non-trivial network (i.e. a network which has some variance in its degree distribution).

Proposition 4.2. *$K^\dagger > K^*$ for all finite non-trivial networks.*

Proof. See the Appendix. □

An implication of Proposition 4.2 is that the more the decision-maker cares about variance, the wider the range of asset classes K for which the CCP delivers netting benefits. Relative to the existing literature, which only examines conditions for central clearing to reduce *expected* netted exposures, this proposition suggests that central clearing is more likely to be beneficial, so long as the decision-maker places a non-zero weight on variance.[11]

introduction of a CCP by market participants.

[11] In section 4.7.1, we follow Cont and Kokholm (2014) in examining the effect of introducing a CCP on Value-at-Risk (VaR), an alternative target statistic.

4.3.2.2 Bounds on K^* and $K^\dagger$

So far in this paper, we have followed the related literature in treating the degree distribution S as known, even though the exact structure of the network is not. For example, expressions (4.14) and (4.15) relate K^* and $K^\dagger$ to moments of S, which are assumed to be known. However, in reality a decision-maker — such as an association of dealers or a regulator trying to decide whether or not to mandate central clearing in a trading network — may lack knowledge even about the degree distribution of the network. For example, a decision-maker may only know the number of asset classes K and an upper bound on the size of the network N. In such a case, a decision-maker may be interested in bounds on K^* and $K^\dagger$ which can help to inform a decision.

Proposition 4.3 (Bounds on K^* and $K^\dagger$)**.** *Among all networks S with $\mathbb{P}(S \leq N-1) = 1$, K^* lies between 1 and $\frac{N^2}{4(N-1)}$, and both bounds can be achieved. When $K^\dagger$ is defined, it lies between 1 and $N-1$, and both bounds can be achieved.*

Proof. See the Appendix. □

A decision-maker can compare the observed K to the upper bounds, which are both functions of the known quantity N. If K is larger than the upper bound on K^* ($K^\dagger$), then the decision-maker knows that introducing a CCP will increase the mean (variance) of netted exposures. The upper bound for $K^\dagger$ is at least as higher as that for K^* when $N \geq 2$, which means that when the decision-maker is sure that central clearing increases the variance of netted exposures, then she can be sure that it increases the mean too. There may be intermediate cases where a decision-maker is sure that there would be an increase in the mean but cannot be sure about the variance; in such cases, a decision-maker who puts more weight on the effect on variance is more likely to introduce central clearing.

For both K^* and $K^\dagger$, the lower bound is achieved if and only if S takes values only on 0 and 1, which is a trivial case since the network consists only of disconnected pairs and singletons. An immediate corollary to Proposition 4.3 characterises the upper bound on K^*:

Corollary 4.1. *The upper bound on K^* is achieved if and only if we have a complete, homogeneous network.*

As we demonstrate in the next section, this result means that the homogeneous network of Duffie and Zhu (2011) achieves the highest expected netting efficiency under central clearing, compared to all networks of the same size. To put it another way, using the decision criterion in Duffie and Zhu is almost certain to lead to central clearing being accepted too often.

4.4 Examination of different network structures

In this section we consider how the impact of introducing a CCP is affected by the underlying structure of the network. The structure of the network is reflected on the right-hand sides of expressions for K^* and $K^\dagger$ via the distribution of the random variable S.

As a network is a very general object, we restrict our analysis to specific cases identified in the existing literature. First we apply our results to Duffie and Zhu's network and recover their results. Then we apply our analysis to core-periphery and fat-tailed networks, in order to examine the effect on the results of the presence of large well-connected counterparties, which are typically present in real-world financial networks.

4.4.1 Homogeneous network of Duffie and Zhu

Duffie and Zhu (2011) assume the network is completely connected, so every agent has full degree. This means that $S = N - 1$ with certainty, and so $K^* = \frac{N^2}{4(N-1)}$, which corresponds to their equation 6. The Duffie and Zhu network is a trivial network as defined earlier, and introducing a CCP has no effect on variance. The benefit of the CCP can be considered just by examining the effect on the mean.

A slight generalisation of Duffie and Zhu's completely connected network is the Erdős-Rényi random network, where links are formed independently with some probability p. In this network, the nodes are still homogeneous ex ante but there may be random variations in their realised link patterns. As an Erdős-Rényi network grows in size, K^* and $K^\dagger$ grow without limit.[12] However, the Erdős-Rényi homogeneous network is not the focus of our

[12] This is a consequence of Proposition 4.5 — see section 4.4.3.2.

paper, as it is not a good fit to real-world financial networks (as shown, for example, in Markose 2012 and Craig and von Peter 2014). Instead we focus on two models which have been shown to be more realistic: fat-tailed and core-periphery network structures.

4.4.2 Fat-tailed networks

Many real-world financial networks have degree distributions with significant excess kurtosis. There is likely to be a small number of highly-connected nodes (we can think of these as the major dealers), with the majority of nodes having few connections. One popular model for this property is a 'fat-tailed network': that is, a network with a degree distribution which asymptotically obeys a power law $\mathbb{P}(S = s) \sim s^{-\alpha}$, for some real-valued parameter $\alpha > 1$. Fatter tails are associated with lower values of α.

In this section we will focus on scale-free networks, which are a particular class of fat-tailed networks that have been shown to arise in many real-world applications, including in finance.[13] Focusing on this class is instructive because they arise according to a simple and intuitive process, as explained in the following section.

4.4.2.1 How do scale-free networks arise?

Barabási and Albert (1999) show that scale-free networks can be formed via growth and preferential attachment. As time goes on, new nodes join the network and tend to form links with the nodes which are better-connected. This is a realistic model of how a derivatives trading network may develop. Over time we would expect new dealers to enter as the market grows. And there are several reasons why these dealers are more likely to trade with the agents which are already better-connected, such as name recognition, an existing relationship in another market, or economies of scale allowing better-connected agents to offer more attractive terms. Barabási and Albert show that networks formed through this process have fat tails with exponent $\alpha = 3$.

Barabási and Albert's general solution only applies when the number of nodes becomes very large, but the assumption of a large network is not necessarily realistic for our pur-

[13] See, for example, Soramäki et al. (2007), Inaoka et al. (2004) and Garlaschelli et al. (2005).

poses. For example Duffie and Zhu consider a network of size 12, which is the number of entities that, at the time of writing their paper, had partnered with ICE Trust to create a CCP for clearing credit default swaps.[14] In order to model networks of finite size, we need to settle on a specific form of Barabási-Albert network. We use the scale-free network formation process described in Dorogovtsev, Mendes, and Samukhin (2001); henceforth we refer to this as the DMS network. The major advantage of the DMS network is that it has an exact solution for a network of any size. The DMS network generating process is as follows:

1. Begin at time $t = 2$ with 3 nodes. Each has two links connecting to one another.
2. Each time period, a new node is added to the network and connects to two existing nodes. To determine which nodes, choose an existing link at random (each with equal probability). The new node then connects to the two nodes which share that existing link. Repeat.
3. This process generates an undirected DMS network. We now need to determine the value of each link. We assume that if two dealers have a link in one asset class, then they have a link in all K asset classes. For each pair of connected nodes i and j, we generate the net exposures X_{ij}^k, $k = 1, \ldots, K$, as K iid normal random variables with mean 0 and standard deviation σ.

Carried out for t steps, this produces a network of size $N = t + 1$, which tends towards a scale-free network with exponent $\alpha = 3$ as t becomes large. The left-hand panel of Figure 4.3 shows the realisation of a DMS process for $t = 100$.

4.4.2.2 Asymptotic analysis of fat-tailed and scale-free networks

We can use the definition of fat tails to approximate the moments of the degree distribution as the size of the network $N \to \infty$:

$$\mathbb{E}[S^m] \sim \int_Z^{N-1} s^{m-\alpha} ds \tag{4.16}$$

[14] Sizes of other real-world CCP networks can be found in Table 1 in Galbiati and Soramäki (2012) and footnote 12 in Cox, Garvin, and Kelly (2013).

Figure 4.3: Left-hand panel: Scale-free network of size $t = 100$, generated using the DMS process. Right-hand panel: Core-periphery network of size $N = 100$, generated using a Bernoulli distribution with parameters $z = 0.25, p = 0.2, c_0 = 0$.

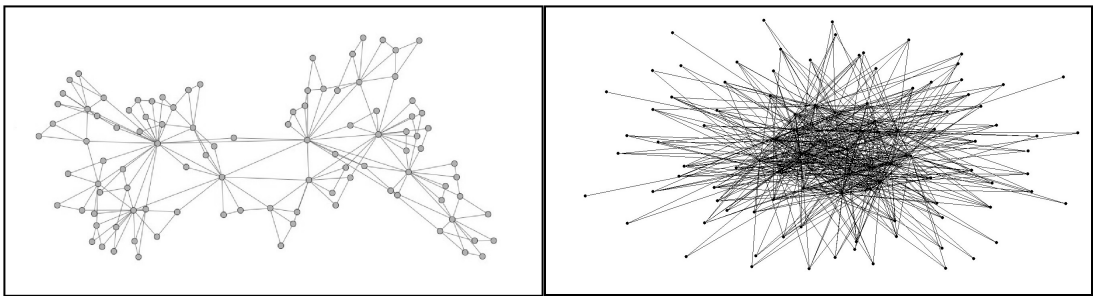

where Z is some constant.

Thus, as $N \to \infty$:

$$\mathbb{E}[S^m] \sim \begin{cases} O((N-1)^{m-\alpha+1}) & \text{if } m > \alpha - 1 \\ O(\log(N-1)) & \text{if } m = \alpha - 1 \\ O(1) & \text{if } m < \alpha - 1 \end{cases} \tag{4.17}$$

and so $\mathbb{E}[S^m]$ has a finite limit if and only if $m < \alpha - 1$.

We can apply this result to our expressions for K^* and $K^\dagger$, (4.14) and (4.15). Table 4.1 below summarises the asymptotics for various values of α.

Table 4.1: **Asymptotic behaviour of K^* and $K^\dagger$ in a fat-tailed network.**

	K^*	$K^\dagger$
$1 < \alpha \leq 2$	tends to positive infinity	tends to positive infinity
$2 < \alpha \leq 3$	tends to finite limit	tends to positive infinity
$\alpha > 3$	tends to finite limit	tends to finite limit

When $\alpha > 3$, the right-hand sides of both expressions (4.14) and (4.15) tend to finite limits. This means that there will be some critical values K^* and $K^\dagger$ beyond which the introduction of a CCP will increase the expectation and variance of net exposures no matter how large the network is. But when $1 < \alpha \leq 2$ the converse is true: for any given number of asset classes K, there will be some critical network size beyond which the expectation and variance of net exposures decline with the introduction of a CCP. This

suggests that a CCP is more likely to deliver benefits to dealers in networks with fatter tails. This makes intuitive sense, because such networks have a large number of major dealers (well-connected nodes), and so the bilateral netting with major dealer is less than efficient in a network where there are fewer of them (i.e. those with thinner tails).

For real-world networks α typically lies between 2 and 3.[15] This corresponds to the intermediate case in Table 4.1, implying a potential trade-off between a higher expected value of net exposures and a lower variance when the number of asset classes is larger than the asymptotic limit of K^*.

Proposition 4.4. *For the infinite DMS network, novating a single asset class to a CCP always increases expected netted exposures, for any number of asset classes $K \geq 2$. But it always decreases the variance of net exposures.*

Proof. See the Appendix. □

The increase in expected exposures that results from the introduction of a CCP ($\widetilde{\phi_{N,K}} - \phi_{N,K}$) has a finite upper bound as $N \to \infty$. However, the reduction in variance ($\widetilde{v_{N,K}} - v_{N,K}$) is without bound. This means that — so long as agents' preference functions put any weight on volatility — then, for any given K, there must be some critical size of the network above which introducing a CCP is beneficial for the dealer agents.

4.4.2.3 Finite analysis of the scale-free network

The distribution of S for a given t is given in Dorogovtsev, Mendes and Samukhin (see their equation 8):

$$\mathbb{P}_t(s) = \frac{t}{t+1}\left[\frac{s-1}{2t-3}\mathbb{P}_{t-1}(s-1) + \left(1 - \frac{s}{2t-3}\right)\mathbb{P}_{t-1}(s)\right] + \frac{1}{t+1}1_{[s=2]} \qquad (4.18)$$

for $t \geq 3$, with initial condition $\mathbb{P}_2(s) = 1_{[s=2]}$. (Here, $1_{[\cdot]}$ denotes the identity function which takes the value 1 if the condition in the subscript is true, and zero otherwise.)

[15] For example, Soramäki et al. (2007) and Inaoka et al. (2004) both estimate $\alpha = 2.1$ for the interbank payment networks in the US and Japan respectively. Garlaschelli et al. (2005) estimate α lies between 2.2 and 3.0 for networks of shareholdings in the US and Italy. For Barabási-Albert scale-free networks, α has an asymptotic limit of 3.

Figure 4.4 shows the effect of introducing a CCP for a range of values of K and t. The distribution of S is determined by equation (4.18) for the corresponding value of t. This has been calculated for DMS networks with up to 250 members. The upper frontier of the black area in the figure corresponds to K^*, while the upper frontier of the grey area corresponds to $K^\dagger$.

Figure 4.4: The effect of the introduction of a CCP in a DMS scale-free network, for a range of values of (K,t). The chart is calculated up to $t = 249$ — that is, $N = 250$. Values of K^* are represented by the frontier between the black and grey areas. Values of $K^\dagger$ are represented by the frontier between the grey and white areas. Illustrative real-world values of N,K are shown by the cross markers (see footnote 16).

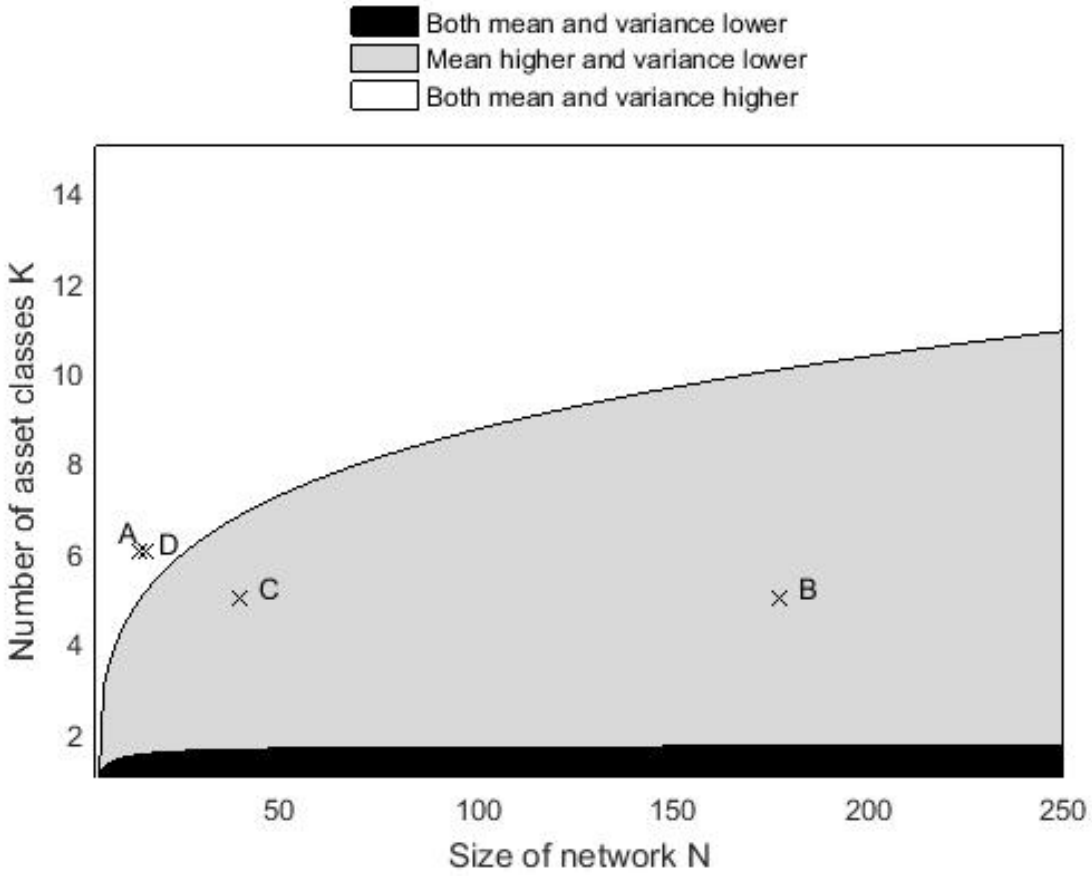

K^* tends monotonically upward with an upper bound of approximately 1.73 (see the proof of Proposition 4.4 in the Appendix). This confirms that the introduction of a CCP

will increase expected net exposures for any non-trivial DMS network (i.e. one where $K \geq 2$). In contrast, $K^{\dagger}$ (shown by the frontier of the grey area) increases without limit as $\sim O((\log(N-1))^2)$, as predicted by the asymptotic analysis. This means that, for sufficiently large networks, the introduction of a CCP will cause a reduction in the variance of exposures.

The cross markers illustrate values of N, K for four real-world networks.[16] In some cases the markers are in the white region, meaning that — assuming the underlying network has a DMS structure — the case for introducing a CCP would need to be motivated by considerations other than netting efficiency. In two cases the markets are in the grey region, suggesting that the introduction of a CCP would be beneficial if sufficient weight is placed on the reduced variance of netted exposures.

4.4.3 Core-periphery networks

Core-periphery networks have been presented in the recent literature as an alternative to scale-free networks as a model of real-world financial linkages. These networks are characterised by a partition of the nodes into two sets: a heavily-connected set of 'core' nodes, along with a sparsely connected set of 'peripheral' nodes.

Borgatti and Everett (1999) present a general model to allow for the detection of core-periphery networks: they assume that in such networks all core nodes are linked to one another, while there are no links between peripheral nodes. They then present a statistic to test for correlations between such an idealised core-periphery network and the actual data. Their model is agnostic about the distribution of links between core and periphery nodes; this is because their specification is aimed at empirical verification of the structure, rather than a general model of a core-periphery network.

Langfield, Liu, and Ota (2014) and Craig and von Peter (2014) use the Borgatti-Everett approach to identify core-periphery structures in the UK and German interbank markets

[16] We have drawn these from published papers where the values of N and K are easiest to deduce. These are as follows:— point A ($N = 12, K = 6$): US derivatives network in June 2010 used by Duffie and Zhu (2011); point B ($N = 176, K = 5$): UK interbank network in 2011 used by Langfield, Liu, and Ota (2014); point C ($N = 38, K = 5$): global OTC derivatives network, used by Markose (2012), restricted to most active participants; point D ($N = 14, K = 6$): CDS network in March 2010 estimated from DTCC Trade Information Warehouse Report.

respectively. Wetherilt, Zimmerman, and Soramäki (2010) and Markose (2012) use alternative methods to show that, respectively, the UK money market and the global network of OTC derivatives can be characterised as having core-periphery structures.

4.4.3.1 How do core-periphery networks arise?

Van der Leij, in't Veld, and Hommes (2016) show that core-periphery structures can arise as the stable outcome of a process of strategic network formation between heterogeneous agents. In their model, there are 'big' (core) banks and 'small' (peripheral) banks, and the payoff from forming a link with a big bank is greater than a link with a small bank. Chang and Zhang (2015) demonstrate that, when agents differ in the amount that they need to trade, they will self-organise into a core-periphery structure, with the agents with the lowest trading need forming the core. Chang and Zhang associate these core nodes with market makers.

Underlying this network generation process is the assumption that, for any given agent, links to core nodes are desirable, while links to peripheral nodes are not. There are plausible reasons why this may be the case for real-world financial networks. Agents may prefer to deal with larger players, with whom they are likely to have existing relationships in other markets. Exposures to larger players may be easier to monitor. And economies of scale may mean that these larger players offer more attractive trading terms.

Abstracting away from consideration of individual nodes' optimal strategies, we can characterise the formation of a core-periphery network using the following simple process:

1. Begin with c_0 core nodes, which are connected to each other.
2. At each step a new node is added. With probability z this new node is labelled 'core'. Otherwise, the node is labelled 'peripheral'.
3. A new core node forms links with all of the existing core nodes with certainty, and forms links with each of the existing peripheral nodes according to some distribution W. A new peripheral node will never form links with existing peripheral nodes, but will form links with existing core nodes according to the distribution W.

This process, carried out for $N - c_0$ steps, will produce a network of size N which meets

the Borgatti-Everett definition of a core-periphery network. The parameters of the model are c_0, z and the distribution W. Borgatti and Everett allow any feasible distribution to determine core-periphery links; for example W could depend on existing links in the network at a given point in time. One natural and simple way to model links between the core and periphery is to assume that each link occurs independently with some fixed probability $p \in (0,1)$ — that is, the link formation process follows a Bernoulli distribution. Under this assumption, the number of links for any randomly chosen node can be expressed as a mix of binomial distributions plus a constant.[17] The right-hand panel of Figure 4.3 shows a realisation of such a network-generating process. We will focus on the Bernoulli core-periphery network in the remainder of this section.

4.4.3.2 Asymptotic analysis of core-periphery networks

In order to make asymptotic inferences about the Bernoulli core-periphery network, we will state and prove a more general Proposition on the asymptotic limit of networks with 'thin tails', which we define as follows.

Proposition 4.5. *Suppose the degree distribution S of the network has the following properties:*

- $\mathbb{V}[S]/\mathbb{E}[S]^2$ *tends to a finite limit as* $N \to \infty$*;*
- *All of the higher moments tend to zero as* $N \to \infty$*.*

Then K^* *and* $K^\dagger$ *are both* $\sim O(\mathbb{E}[S])$ *as* $N \to \infty$*.*

Proof. See the Appendix. □

For a network which meets the conditions in Proposition 4.5, K^* and $K^\dagger$ will increase without bound as the size of the network becomes large. This suggests that, for any given number of asset classes K, there is a minimum size of the network above which the introduction of a CCP would reduce both mean and variance.

[17] Let C be the number of core nodes. Then $C \sim c_0 + Bin(N - c_0, z)$, and the conditional distribution $S|C \sim \frac{C}{N} \cdot [(C-1) + Bin(N-C, p)] + (1 - \frac{C}{N}) \cdot Bin(C, p)$. Note that the case $z = 1$ corresponds to the Duffie and Zhu network with size equal to N.

Examples of network models which meet these 'thin-tailed' conditions are the Duffie and Zhu network, the Erdős-Rényi network described in section 4.4.1, and the Bernoulli core-periphery network described above. In all of these cases, $\mathbb{E}[S] \sim O(N)$.

Proposition 4.5 does not apply to a network with a link formation process W which generates fat tails; such a network would not meet the conditions stated in the Proposition. For example, if the core and peripheral nodes form links according to a Barabási-Albert preferential attachment process, then such a network will exhibit asymptotic results similar to those derived in section 4.4.2.2. Markose (2012) shows that the network of global OTC derivatives exposures can be modelled by a core-periphery network with fat tails in the degree distribution. In this case, the analytics of the fat-tailed network would be more appropriate.

4.4.3.3 Finite analysis of the Bernoulli core-periphery network

While the only parameter in the DMS network is its size N, the Bernoulli core-periphery network has four parameters N, z, p and c_0. In order to derive numerical solutions, we need to choose feasible parameter values, so we turn to the empirical literature. Table 4.2 below summarises parameter estimates from three selected papers: the global OTC derivatives network from Markose (2012), the Dutch interbank market from Van der Leij, in't Veld, and Hommes (2016), and the German interbank market from Craig and von Peter (2014).[18] The value of c_0 is impossible to observe and is likely to make little difference for larger networks, so we assume $c_0 = 0$. We will use $z = 0.10, p = 0.10$ for our numerical solutions — these parameter values are fairly representative of Table 4.2.[19]

Figure 4.5 shows the effect of introducing a CCP for various values of N up to 250, for the parameters $z = 0.10$, $p = 0.10$, $c_0 = 0$. As predicted, K^* and $K^\dagger$ increase approximately linearly with N. There is a grey intermediate area where $K^* < K < K^\dagger$ — in such cases,

[18] For the global OTC network we use the 'inner core' as defined by Markose. Each paper provides the number of nodes N, the size of the core C and the total number of directed links L. We assume that $c_0 = 0$. We can then estimate $z = \frac{C}{N}$ and $p = \frac{L-C(C-1)}{2C(N-C)}$.

[19] The Markose parameter values may be thought to be the best proxy, given that we are interested in central clearing of OTC derivatives and not other forms of interbank lending. However, it should be noted that the Markose network exhibits fat tails and so does not conform very well to a Bernoulli core-periphery structure.

Table 4.2: **Parameter estimates for selected real-world core-periphery networks.**

	Global OTC	Dutch interbank	German interbank
Period	2009 Q4	2008 Q4	2003 Q2
N	204	100	1802
Estimated z	0.10	0.15	0.02
Estimated p	0.01	0.31	0.11

introducing a CCP will increase expected net exposures in the network, but reduce the variance. The cross markers represent the same four real-world networks shown on Figure 4.4.[20]

Comparing Figures 4.4 and 4.5, we can see that network structure has an important effect on the benefits of introducing a CCP. Compare point C in the two networks. This represents the global OTC derivative network explored in Markose (2012). This point lies in the grey region in Figure 4.4 and the white region in Figure 4.5. If we assume a DMS network structure, then introducing a CCP to this network increases the mean of netted exposures, but reduces their variance. Therefore a sufficiently risk-averse policymaker may be willing to introduce a CCP. But, if we assume a core-periphery structure, then introducing a CCP increases both the mean and variance, so even a highly risk-averse policymaker would not wish to introduce a CCP. Clearly, the network structure ought to be an important factor in policymakers' decisions about whether to mandate central clearing.

As Figure 4.5 shows, there are network sizes N where a CCP is unambiguously bad for netting efficiency. For all $N \leq 93$, $K^* < 2$ and so a CCP will never reduce average net exposures for networks of this size. And for all $N \leq 29$, $K^\dagger < 2$ and so introducing a CCP for networks of this size increases both the mean and variance of net exposures.

Testing other parameter values, we also find that K^* and $K^\dagger$ always increase with N and with both z and p. We can interpret higher K^* and $K^\dagger$ as indicating that the introduction of a CCP brings greater netting benefits. These results arise because, when the network

[20] We use different networks in Table 4.2 to those used to illustrate the figures. This is due to differing degrees of confidence in our ability to estimate the relevant parameters, given the data and information provided in the original papers. The networks in Figures 4.4 and 4.5 are those for which we can most confidently estimate N and K. The networks in Table 4.2 are those for which we can most confidently estimate N, z and p.

Figure 4.5: The effect of the introduction of a CCP in a core-periphery network with $z = 0.10, p = 0.10, c_0 = 0$, for a range of values of (K,N) up to $N = 250$. Values of K^* are represented by the frontier between the black and grey areas. Values of $K^\dagger$ are represented by the frontier between the grey and white areas. Illustrative real-world values of N, K are shown by the cross markers (see footnote 16).

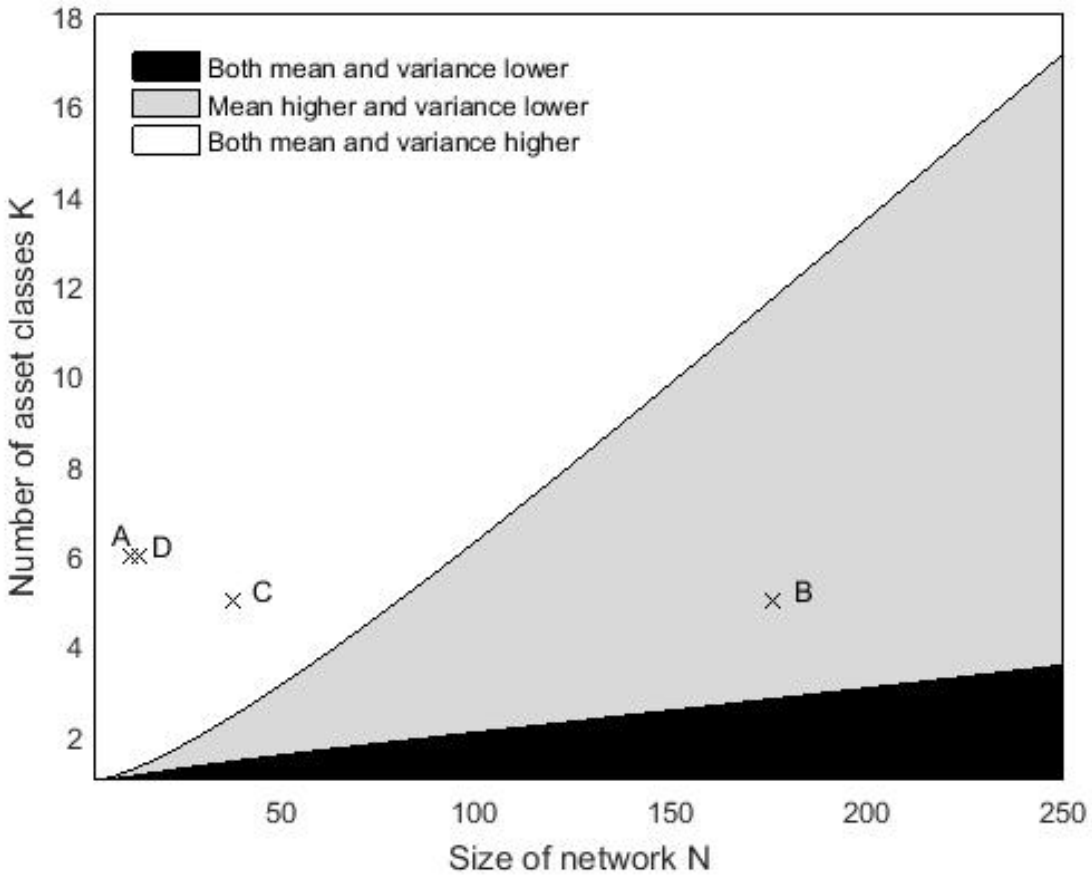

relies on links with a small fraction of key nodes (the scale-free case), bilateral netting substantively reduces net exposures, and central clearing disrupts this. But when links are more spread amongst a larger number of key nodes (the core-periphery case with large z or p), central clearing may bring greater netting benefits.[21]

[21] This is supported by Table 4.1 which suggests that, as α increases in a large fat-tailed network, central clearing tends to bring greater netting benefits. That is because a fatter tail means a greater number of nodes with a substantial number of links, so bilateral netting with individual nodes is relatively less important.

This strongly suggests that knowledge of network structure — as well as the size and the number of asset classes — is important in assessing the possible benefits of central clearing on netting efficiency. Particularly important is the tail of the degree distribution, which determines the connectivity of the largest nodes in the network.

4.5 Generalisation to several CCPs and asset classes

Until now we have only considered the case of the introduction of centralised netting in one asset class. In reality, a single CCP may clear more than one asset class, or there may be several CCPs all clearing different asset classes. This may affect netting efficiency: for example Duffie and Zhu (2011) show that merging multiple CCPs into a single CCP will always improve netting efficiency. In this section, we briefly show how our results can be generalised to a setting with several CCPs.

Suppose that, as before, there are N dealer agents in the OTC derivatives network, and K asset classes. M CCPs are introduced, indexed by $m = 1, \ldots, M$. CCP m clears a number of asset classes a_m. We relabel the asset classes so that asset classes in the set $\{1, \ldots, a_1\}$ are cleared by CCP 1, those in the set $\{a_1 + 1, \ldots, a_1 + a_2\}$ are cleared by CCP 2, and so forth.

The total number of asset classes novated to central clearing is $A := a_1 + \cdots + a_M$. The asset classes with labels $\in \{A+1, \ldots, K\}$ are not centrally cleared.[22]

By the same reasoning as in Section 4.3.1, we can show that the expectation and variance of exposures following centralised netting are:

$$\widetilde{\phi_{N,K}} = \mathbb{E}[S] f(K-A) + \sum_{m=1}^{M} \mathbb{E}[f(a_m S)], \tag{4.19}$$

$$\widetilde{v_{N,K}} = \mathbb{E}[S] g(K-A) + \sum_{m=1}^{M} \mathbb{E}[g(a_m S)] + \mathbb{V}\left[S f(K-A) + \sum_{m=1}^{M} \mathbb{E}[f(a_m S)] \right]. \tag{4.20}$$

[22] We assume here that any given asset class is cleared by at most one CCP. We do not consider the case where an asset class can be cleared by more than one CCP. Our results in this section suggest that such a case is inefficient.

Note that in the case $m = 1,\ a_1 = 1,\ A = 1$, we recover equations (4.7) and (4.8).

These equations provide us with a way of comparing the netting efficiency under different constellations of CCPs, by writing a given constellation in terms of a partition of the set of asset classes. Let us assume again that the underlying bilateral exposures X_{ij}^k are iid normally distributed $\sim N(0,\sigma^2)$ for some parameter σ. We can use the expressions (4.11) and (4.12) to show that introducing a given CCP constellation reduces the expectation and variance of exposures if and only if, respectively:

$$\frac{1}{A}\left(\sqrt{K}+\sqrt{K-A}\right)\sum_{m=1}^{M}\sqrt{a_m} \quad < \quad \frac{\mathbb{E}[S]}{\mathbb{E}[\sqrt{S}]}, \tag{$\star'$}$$

$$2\sqrt{K-A} \quad < \quad \frac{A\mathbb{V}[S]-\left(\sum_{m=1}^{M}\sqrt{a_m}\right)^2\mathbb{V}[\sqrt{S}]}{\left(\sum_{m=1}^{M}\sqrt{a_m}\right)\left(\mathbb{E}[S^{\frac{3}{2}}]-\mathbb{E}[S]\mathbb{E}[\sqrt{S}]\right)}. \tag{$\dagger'$}$$

Expressions ($\star'$) and ($\dagger'$) are generalisations of ($\star$) and ($\dagger$) respectively. As before, all of the K terms are on the left-hand sides of the expressions, which are increasing in K. Therefore a given CCP constellation will tend to be less efficient when the total number of asset classes is higher. The intuition is the same as before: when A is fixed and K increases, then the CCPs collectively clear a smaller proportion of the network, and so the benefits they bring from netting within asset classes are less likely to be greater than the cost from disruption of existing bilateral netting sets across asset classes. We can redefine K^* and $K^\dagger$ as the values of K for which there is equality between the left- and right-hand sides of ($\star'$) and ($\dagger'$), respectively.

We can use these expressions to prove the following general results.

Proposition 4.6.

1. ***Merging CCPs will always improve netting efficiency:*** *Given any constellation $\Psi_1 = \{a_1,\ldots,a_M\}$, consider a new constellation $\Psi_2 = \{a_1+a_2,\ldots,a_M\}$ with one fewer CCP. Then both the expectation and variance of exposures under Ψ_2 will be lower than under Ψ_1.*

2. ***It is most efficient to novate asset classes to the largest existing CCP:*** *Given any constellation $\Psi_1 = \{a_1,\ldots,a_M\}$, consider two alternative constellations $\Psi_2 =$*

$\{a_1+1, a_2, \ldots, a_M\}$ *and* $\Psi_3 = \{a_1, a_2+1, \ldots, a_M\}$. *If* $a_1 > a_2$, *then both the expectation and variance of exposures under* Ψ_2 *will be lower than under* Ψ_3.

Proof. See the Appendix. □

Proposition 4.6 shows that merging CCPs will reduce both the mean and variance of net exposures, and that it is most efficient to clear a new asset class through the largest existing CCP. Taken together, these findings imply that the most efficient arrangement would be to have a single CCP net every asset class. It is easy to see that this is best, because all exposures can be netted against one another. However, such an arrangement may not be optimal once a regulator takes into account systemic risk, because such a CCP would be hugely systemically important. If a regulator does not want too many asset classes being cleared by a single CCP, our expressions ($\star'$) and ($\dagger'$) can help her to compare different potential clearing arrangements, and to assess trade-offs against some other factor of concern, such as systemic risk.

4.6 Strategic behaviour and asset dynamics

Until now we have assumed that the structure of the network is not affected by the introduction of central clearing, except for the trades that are novated. Nor does the distribution of the exposures change. In reality, dealer agents may respond to the introduction of a CCP by adjusting their links with other agents, and changing their bilateral exposures. In this section we allow the introduction of central clearing to change the network structure and the exposures.[23]

We make a distinction between the *trading network* and the *clearing network*. The trading network reflects the structure of derivative trading relationships prior to the introduction of centralised netting. The clearing network reflects the structure after centralised netting. So far, we have assumed that the two networks are identical. Suppose instead that each network reflects the equilibrium result of a strategic network formation process. The dealer agents trade off various effects, such as:

[23] We thank an anonymous referee for this suggestion.

- Bilateral netting benefits: to maximise bilateral netting and thus economise on margin costs, each dealer agent prefers to trade more heavily with each counterparty, and thus have relatively few links.
- Search costs: to minimise search costs, each dealer agent prefers to have fewer links, and consequently to trade more with each counterparty. Links to core nodes entail lower search costs.
- Counterparty risk management: to minimise counterparty credit risk, each dealer agent seeks to diversify its counterparties, forming more links and trading less with each.

As a result of these trade offs, agents form links with several counterparties, but do not distribute their exposures evenly. This gives rise to a particular trading network structure. In fact, we may suppose that the network structures that were studied in the previous sections were the result of optimising behaviour. Now let us suppose, that following the introduction of central clearing, the agents are given the opportunity to re-form the network.

To our knowledge, there is no existing academic literature — either theoretical or empirical — addressing the question of how the clearing and trading networks might differ following the introduction of central clearing. As discussed above, it is likely to depend on agents' preferences and so will be an empirical question. Our approach is to remain agnostic about the answer, and instead consider how our model can be adapted to a given clearing network.

4.6.1 Effect on expected net exposures post-CCP

Let S denote the degree distribution in the trading network, and $\tilde{S}$ the degree distribution in the clearing network. Similarly, let σ^2 and $\tilde{\sigma}^2$ denote the variance of each bilateral exposure in the trading and clearing network, respectively. Let $f(\theta)$ denote the expected net exposure between any two linked nodes in the trading network, given that there are θ asset classes traded between them. Let $\tilde{f}(\theta)$ denote the expected net exposure in the clearing network. This formulation allows not only the structure of the network to change, but also the volatility of the in-the-money positions of the contracts (which may reflect,

for example, changes in the magnitudes of bilateral net positions dealer agents are willing to take against one another).[24]

In the trading network, the expected total netted exposures for a given node i are:

$$\phi_{N,K} = \mathbb{E}[S] f(K) = \sigma \sqrt{\frac{K}{2\pi}} \mathbb{E}[S]. \tag{4.21}$$

In the clearing network, the corresponding quantity is:

$$\begin{aligned} \widetilde{\phi_{N,K}} &= \mathbb{E}[\tilde{S}] \tilde{f}(K-1) + \mathbb{E}[\tilde{f}(\tilde{S})], \\ &= \tilde{\sigma} \sqrt{\frac{K-1}{2\pi}} \mathbb{E}[\tilde{S}] + \tilde{\sigma} \sqrt{\frac{1}{2\pi}} \mathbb{E}[\sqrt{\tilde{S}}]. \end{aligned} \tag{4.22}$$

Define Δ to be the change in expected net exposures after the introduction of central clearing that we defined in (4.9); i.e.:

$$\begin{aligned} \Delta &= \widetilde{\phi_{N,K}} - \phi_{N,K}, \\ &= \frac{1}{\sqrt{2\pi}} \Big(\tilde{\sigma} \Big(\sqrt{K-1} \mathbb{E}[\tilde{S}] + \mathbb{E}[\sqrt{\tilde{S}}] \Big) - \sigma \sqrt{K} \mathbb{E}[S] \Big). \end{aligned} \tag{4.23}$$

Central clearing reduces expected net exposures if and only if $\Delta < 0$. When $S = \tilde{S}$ and $\sigma = \tilde{\sigma}$, this is the same condition $(\star)$ as in the case when the clearing and trading networks were identical.

Proposition 4.7. *Let* $\tilde{\sigma}\mathbb{E}[\tilde{S}] = \beta\sigma\mathbb{E}[S]$. Δ *is increasing with respect to* K *if and only if* $\beta > \left(\frac{K-1}{K}\right)^{\frac{1}{2}}$.

Proof. Substitution of $\beta\sigma\mathbb{E}[S]$ into (4.23) and partially differentiating with respect to K yields the expression $\frac{\sigma\mathbb{E}[S]}{2\sqrt{2\pi}} \left[\beta(K-1)^{-\frac{1}{2}} - K^{-\frac{1}{2}} \right] > 0$. □

Proposition 4.7 is a generalisation of the result in Proposition 4.1, which states that the introduction of centralised netting delivers fewer benefits as the number of asset classes increases. This result is robust to endogenous changes in the trading network so long as the expected number of linkages and the variance of exposures in clearing network are

[24] We continue to assume exposures are determined by the same process as in section 3.2. However, here we are allowing agents to re-optimise in a way that would change the variance of these exposures.

not too much smaller than in the trading network (so that β satisfies the condition given in the statement of Proposition 4.7).

Interestingly, the possibility that traders will reduce the number of linkages or the variance of exposures in response to central clearing opens up the possibility that central clearing may be better when K is high. When the trading and clearing networks are the same, then lower K makes central clearing more likely to be efficient, because the CCP clears a larger proportion of derivatives. But this may no longer apply if the network changes. Suppose that the clearing network is sufficiently less connected than the trading network, so that the condition in Proposition 4.7 is violated. Then, even without considering the impact of the CCP itself, the clearing network has greater bilateral netting efficiency than the trading network, because each agent has her trades concentrated among fewer counterparties. When the CCP is introduced, these bilateral netting sets are disrupted. The higher K is, the less disruption is caused by the introduction of central clearing in one asset class, meaning a more efficient network.

To illustrate this point, suppose that we have a simple two-agent trading network with degree distribution $S \sim Bin(n,p)$ and that the clearing network has degree distribution $\tilde{S} \sim Bin(n,\tilde{p})$. Let $\sigma = \sqrt{2\pi}$ and let $\tilde{\sigma} = \tau\sqrt{2\pi}$. Then:

$$\Delta = \tau n\tilde{p}\sqrt{K-1} - np\sqrt{K} + \tau\mathbb{E}[\sqrt{\tilde{S}}] \tag{4.24}$$

which is increasing in K when:

$$\frac{1}{K} > 1 - \left(\frac{\tilde{p}\tau}{p}\right)^2. \tag{4.25}$$

Consider the case $\tilde{p}\tau < p$. Then there is some $\bar{K}$ such that condition (4.25) fails for all $K > \bar{K}$, and so central clearing is more efficient when there are more asset classes in the market. This occurs either because the clearing network is less connected than the trading network, or it has smaller bilateral exposures (or both). Either way, it means that bilateral netting in the clearing network is more efficient than in the trading network. Introducing a CCP in one asset class disrupts that. But if the one asset class is only a small proportion of the total market, then the impact is less egregious.

We conclude that it is important for policymakers considering the introduction of CCPs

to be aware not only of the trading network structure, but also to try to anticipate how dealer agents' bilateral relationships may change following the introduction of a CCP. For example, suppose that a policymaker believes that a likely consequence of introducing a CCP is that agents increase their bilateral risk positions vis-à-vis one another. Then the clearing network will be less connected than the trading network and have more volatile bilateral exposures across the links that do exist. There is then a risk that netting efficiency may be worse as a consequence of the introduction of central clearing, not better.

The same analysis can be applied to our results in section 4.5 too. In the standard model without strategic network formation, centrally clearing more asset classes is always better than clearing fewer. However, that may no longer be the case if each additional asset class further concentrates the clearing network.

In this section we have focused on the effect of strategic network formation on expected net exposures, but a similar intuition can be applied when considering the variance, or indeed any other feasible summary statistic.

4.7 A computational approach to assessing the benefits of central clearing

In this section, we show how our model can be used by a practitioner — for example, a policy maker who wishes to determine whether or not the introduction of central clearing will improve netting efficiency. We present a computational approach which allows for the estimation of the benefits with knowledge of only a few parameters, rather than the entire network structure. We allow the trading and clearing networks to be different, as described in section 4.6. We focus here on measuring the change in expected netting efficiency, $\widetilde{\phi_{N,K}} - \phi_{N,K}$, but it is straightforward to extend the framework to measuring changes in the variance of netted exposures too, or indeed any other feasible summary statistic.

We make an assumption that the overall willingness of each agent to bear *market risk* does not change following the introduction of central clearing. Market risk is defined as the volatility of an agent's total exposures in each asset class. This measures an agent's

need or desire to trade each asset class. This is based on a premise that participation in the market is not affected by the introduction of the CCP, although the choices of counterparties may be.

This assumption means that each agent bears the same expected market risk in the system in both the trading and clearing network. This is in contrast to *counterparty credit risk*, which a CCP can and generally does affect.

Proposition 4.8. *If agents bear the same expected market risk before and after the introduction of central clearing, then* $\sigma^2\mathbb{E}[S^2] = \tilde{\sigma}^2\mathbb{E}[\tilde{S}^2]$.

Proof. See the Appendix. The idea of the proof is as follows. For each agent and each asset class, market risk is normally distributed with mean zero and variance $\sigma^2 S^2$, conditional on S. The unconditional variance of the market risk is — by assumption — the same in the trading and clearing network, so the desired equality is obtained. □

Proposition 4.8 alone does not allow us to pin down the degree distribution of the clearing network, because it only tells us how the second moment will change. We can make an additional assumption that both the trading and clearing networks have a core-periphery structure, as described in Section 4.4.3. There is plenty of empirical and theoretical support for this assumption. The empirical literature has already been discussed in section 4.4.3. A number of recent theoretical papers have shown that core-periphery structures can arise in equilibrium as a result of heterogeneity between agents.[25] In these papers, agents differ in terms of liquidity needs or urgency of trade. Those agents who have less need to trade take on the role of a market maker and form the core. The agents who have greater need to trade form the periphery.

Assumption 4.1. *The relationship between the trading and clearing networks is governed by the following constraints:*

1. *The trading network is a core-periphery network with parameter set* (N, z, p, σ). *The clearing network is a core-periphery network with parameter set* $(\tilde{N}, \tilde{z}, \tilde{p}, \tilde{\sigma})$. *For simplicity, we make the initial number of core nodes* c_0 *equal to 0.*

[25] See, for example, Aldasoro, Gatti, and Faia (2016), Chang and Zhang (2015), Hollifield, Neklyudov, and Spatt (2016) and Van der Leij, in't Veld, and Hommes (2016).

2. *The number of dealer agents is the same in both networks:* $N = \tilde{N}$.

3. *The proportion of core nodes is the same in both networks:* $z = \tilde{z}$.

4. *Market risk is the same in both networks; that is, Proposition 4.8 holds.*

The number of dealer agents can be assumed to be exogenously determined and not affected by the introduction of central clearing. Similarly, the core nodes are assumed to have intrinsic differences from the peripheral nodes, and so they make up the same proportion of both networks. For example, the core nodes may represent major broker-dealers who have relationships with a large number of smaller financial institutions and non-financial customers in a range of different markets. These existing relationships mean lower search costs for their counterparties, and so they tend to form links with each other whether or not a CCP is introduced.

Given the parameter set of the trading network, the clearing network has four parameters and three constraints. This makes the set of feasible clearing networks a one-dimensional set. We use $\tilde{p}$ to parameterise this set. Clearing networks with higher values of $\tilde{p}$ have more connections. By Proposition 4.8, that implies that each bilateral link will have a lower variance.

Conditional on a number of core nodes $C \sim Bin(N,z)$, the trading network has the following distribution:

$$S|C \sim \frac{C}{N} \cdot \Big(C - 1 + Bin(N-C,p)\Big) + \Big(1 - \frac{C}{N}\Big) \cdot Bin(C,p), \tag{4.26}$$

and similarly for the clearing network. It is straightforward to compute the unconditional mean of S:

$$\mathbb{E}[S] = (N-1)z\Big(z + 2p(1-z)\Big), \tag{4.27}$$

but the other unconditional moments are more complicated. To assess the conditions for a reduction of net exposures we need half-moments, which do not exist in closed form. Therefore it will be necessary to take a computational approach.

To compute the netting efficiency pre- and post-CCP, a planner needs estimates of the parameters (N,z,p,σ), as well as the number of asset classes K. We suggest that these can all be estimated using summary statistics, and thus do not require knowledge of the full

network. Post-crisis, regulators are obtaining better information about trading networks, and so estimation of these parameters should be possible. For further details, see empirical papers such as Craig and von Peter (2014). The results of the previous section suggest that outcomes may be non-monotonic in the parameters, so having precise estimates for the parameters is very important.

As $\tilde{p}$ is only observable ex post, a planner's best approach is to form a prior over the likely connectivity of the clearing network and use that to compute a prior distribution for netted exposures. Future empirical research may help to refine priors over $\tilde{p}$, as regulators can look at the effect of introducing central clearing in other networks.

4.7.1 Robustness checks

We use our computational approach to check whether any of our modelling assumptions drive the results.[26] In particular, we follow Cont and Kokholm (2014) in looking at how the results may change under the following circumstances:

1. exposures are positively correlated;
2. exposures in different asset classes are not identically distributed;
3. a distribution other than the normal is used to generate exposures;
4. the effectiveness of central clearing is assessed by comparing Value-at-Risk (VaR) of netted exposures, rather than the mean or variance.

We examine these separately, using the global OTC derivatives network (identified as point C in Figures 4.4 and 4.5) and assuming a core-periphery structure with $N = 38, K = 5, z = 0.10, p = 0.10, c_0 = 0$. We follow Cont and Kokholm (2014) in implementing each of the robustness checks. For the first, we assume that exposures are normally and identically distributed with unit variance but with correlation coefficient $\rho = 0.1$. For the second, we assume that there is zero correlation but exposures in the centrally cleared asset class K has double the standard deviation of exposures in the other asset classes.For the

[26] We thank an anonymous referee for this suggestion.

third, we use the Student's t-distribution with 3 degrees of freedom and unit variance.[27] For the fourth, we use 99% VaR to report all of the previous results. For each of these, we generate the network one million times and report the statistics in Table 4.3.

Table 4.3: **Robustness checks using a core-periphery network with parameters $N = 38, K = 5, z = 0.10, p = 0.10, c_0 = 0$. The table shows the mean, variance and 99% VaR of netted exposures before and after the introduction of a CCP, under various assumptions. These results are obtained computationally by simulating the network one million times in each case.**

		Base model	Correlated	Different variance	Student's t
Mean	No CCP	0.62	0.74	2.29	0.56
	With CCP	0.75	0.84	0.94	0.65
	Relative change	0.20	0.14	-0.59	0.17
Variance	No CCP	2.58	4.71	35.11	2.40
	With CCP	2.91	4.96	4.37	2.64
	Relative change	0.13	0.05	-0.88	0.10
99% VaR	No CCP	7.79	10.84	28.65	7.40
	With CCP	8.17	11.11	9.92	7.70
	Relative change	0.05	0.02	-0.65	0.04

As shown by point C in Figure 4.5, both mean and variance increase upon introduction of a CCP. Using VaR leads to similar decisions as using the other two metrics. For these parameters, assuming positively correlated exposures or using Student's t makes the CCP mildly more effective (i.e. a smaller increase in the metric), but the direction of the change is the same.

The most striking case is when exposures in asset class K has a higher variance than the other asset classes. Under these parameters, introducing a CCP leads to a substantive improvement in netting under all metrics, meaning that a different decision would be made. The reason is that exposures in asset class K are larger than those in other asset classes, so netting them against each other (in a CCP) is more efficient than netting them against smaller exposures (as with bilateral netting). If instead exposures in asset class K had a lower variance, introducing a CCP would be less effective. All of these results are in line with the findings of Cont and Kokholm (2014).

[27] See Cont and Kokholm (2014) for an explanation of this assumption. Using a distribution with infinite variance or mean will make the results heavily dependent on a few rare extreme observations.

4.8 Implications for policy

In 2009, G-20 Leaders called for central clearing in a variety of derivatives markets (in particular high-volume standardised credit default and interest rate swaps). This has now been introduced into legislation in various member countries; for example in the United States through Title VII of the Dodd-Frank Act of 2010, and in the European Union through the European Market Infrastructure Regulation (EMIR) of 2012. What can our analysis tell us about the need for regulatory intervention? In other words, given that the dealer agents have chosen not to set up a CCP themselves, what market failure does a regulator address by mandating central clearing?

Margin requirements in a derivatives network are related to netted exposures. Let us assume that the dealers are less averse to volatility and the associated extreme outcomes than is socially optimal. This may be the case, for example, if high or volatile margin requirements impose externalities on markets for the collateral assets (Murphy, Vasios, and Vause 2014), or if agency problems mean that dealers take excessive risks. In such cases, the dealers would not wish to introduce centralised netting, even though it may be socially optimal to do so. A social planner may determine that the optimal policy measure is to mandate central clearing, although the dealers themselves may disagree. This explains the need to introduce regulation which mandates that dealers use a CCP.

Following the United Kingdom's vote in June 2016 to leave the European Union, the European Commission is considering changes to the oversight of the clearing of euro-denominated swaps, most of which is done in London.[28] One likely option is a mandate that clearing of such swaps must be done within the European Union. This would mean disruption to existing netting sets among CCPs in the UK and EU (see O'Malia 2017). Or model can help to assess the possible impact of this change on netting efficiency in these CCPs. Proposition 4.6 predicts that the outcome of this will be to make London-based CCPs less efficient and the beneficiaries of this in the EU more efficient. However, as London is by the larger financial centre, the overall effect will be worse aggregate netting efficiency.

[28] See, for example, `https://www.ft.com/content/d3754014-30da-11e7-9555-23ef563ecf9a`.

4.9 Concluding remarks

We show how the introduction of centralised netting in a single asset class affects both the mean and variance of netted dealer exposures, depending on the underlying structure of the network. Centralised netting is more likely to decrease both mean and variance if the network is larger, or if there is a smaller number of asset classes traded in the network. But centralised netting brings fewer netting benefits if the network relies on a small number of key nodes for most of its links.

This has welfare implications because net dealer exposures relate to aggregate counterparty risk in the network and to liquidity needs. For example, if centralised netting reduces both the expectation and variance of exposures, it is likely to be beneficial. If it increases both the expectation and variance, then it is likely to reduce welfare. We show that there will be cases where the decision-makers must trade off lower variance against a higher mean, but not vice versa. This means that some aversion to volatility may be required to favour the introduction of a centralised netting system, such as a CCP. In such cases, whether centralised netting improves welfare or not will depend upon details of risk controls, as well as the preferences of the dealers, policymakers and any other relevant agents.

A criticism of work that analyses changes in financial network structures that result from shocks or policy changes is that nodes in the network (financial institutions) are not assumed to optimally adjust their behaviour following the event studied. This practice arises in part because of difficulties associated with accurately modelling optimal network formation in particular market scenarios. We show that there are trade-offs that can lead to ambiguous conclusions. Our approach is to remain agnostic about the outcome of behavioural effects, but to identify criteria under which such effects are more or less likely to affect conclusions.

4.10 Appendix: Proofs

Proof of Proposition 4.2

As notational shorthand, define $T_j = \mathbb{E}[T^j] = \mathbb{E}[S^{j/2}]$, where S is the degree distribution of the network. Note that, because S can only take non-negative values, T_j is real-valued and non-negative for all $j \geq 0$. In this proof we shall make repeated use of the following result, which is an immediate consequence of Hölder's inequality.

Proposition 4.9. *For any $r > 1$, $a \geq 0$, $b \geq 0$,*

$$T_{a+b} \leq T_{ar}^{\frac{1}{r}} T_{bs}^{\frac{1}{s}}$$

where $\frac{1}{r} + \frac{1}{s} = 1$. Equality holds if and only if $ar = bs$, except in the trivial case when T is a constant.

Let us assume that Proposition 4.2 is false, and that there is some non-trivial network for which $K^* \geq K^\dagger$. We can write this as:

$$\begin{aligned} &\frac{1}{4}\left(\frac{T_2}{T_1} + \frac{T_1}{T_2}\right)^2 \geq \frac{1}{4}\left(\frac{T_4 - T_2^2 - T_2 + T_1^2}{T_3 - T_2 T_1}\right)^2 + 1 \\ \iff &\left(\frac{T_2}{T_1} - \frac{T_1}{T_2}\right)^2 (T_3 - T_2 T_1)^2 \geq (T_4 - T_2^2 - T_2 + T_1^2)^2 \end{aligned}$$

We can take square roots of both sides of this expression without changing the direction of the inequality, so long we are sure that the square root of each side is non-negative. The following statements ensure that it is indeed the case.

$$\begin{aligned} &T_4 - T_2^2 - T_2 + T_1^2 = \mathbb{V}[S] - \mathbb{V}[\sqrt{S}] \geq 0; \\ &\frac{T_2}{T_1} - \frac{T_1}{T_2} = \frac{\mathbb{E}[S]}{\mathbb{E}[\sqrt{S}]} - \frac{\mathbb{E}[\sqrt{S}]}{\mathbb{E}[S]} \geq 0; \text{ and} \\ &T_3 - T_2 T_1 = \mathbb{E}[T^3] - \mathbb{E}[T^2]\mathbb{E}[T] \geq 0. \end{aligned}$$

The first two are obviously true, because $S \geq 1$ with probability 1. The third can be proved using Proposition 4.9 with $a = 2, b = 0, r = \frac{3}{2}, s = 3$.

That means that we can take square roots of both sides of the inequality above to obtain:

$$\begin{aligned}
&\left(\frac{T_2}{T_1}-\frac{T_1}{T_2}\right)(T_3-T_2T_1) \geq T_4-T_2^2-T_2+T_1^2 \\
\iff &\left(\frac{T_2}{T_1}-\frac{T_1}{T_2}\right)T_3-T_2^2+T_1^2 \geq T_4-T_2^2-T_2+T_1^2 \\
\iff &\left(\frac{T_2}{T_1}-\frac{T_1}{T_2}\right)T_3 \geq T_4-T_2 \\
\iff &\ T_2^2T_3+T_1T_2^2 \geq T_1T_2T_4+T_1^2T_3.
\end{aligned}$$

Now we shall use Proposition 4.9 to show that this is not true for any non-trivial network, and thus obtain a contradiction.

Setting $a=\frac{4}{3},b=\frac{2}{3},r=3,s=\frac{3}{2}$, we have $T_2<T_4^{\frac{1}{3}}T_1^{\frac{2}{3}}$. Setting $a=\frac{8}{3},b=\frac{1}{3},r=\frac{3}{2},s=3$, we have $T_3<T_4^{\frac{2}{3}}T_1^{\frac{1}{3}}$. Multiply these together gives $T_2T_3<T_1T_4$, and so $T_2^2T_3<T_1T_2T_4$.

Setting $a=\frac{3}{2},b=\frac{1}{2},r=2,s=2$ in Proposition 4.9, we can show $T_2^2<T_1T_3$. This means $T_1T_2^2<T_1^2T_3$. Adding these up gives $T_2^2T_3+T_1T_2^2<T_1T_2T_4+T_1^2T_3$, and we have our contradiction.

We have shown that $K^*\geq K^\dagger$ can only be true for a trivial network. Thus, the result is proved. □

Proof of Proposition 4.3

Bounds for K^*

For K^*, the upper bound $\frac{N^2}{4(N-1)}$ is achieved if $\mathbb{P}(S=N-1)=1$, that is, if we have a Duffie and Zhu network, and the lower bound of 1 is achieved if S can only take values on 0 or 1.

Let us write $K^*(x)=\frac{1}{4}(x+\frac{1}{x})^2$ where $x=\frac{\mathbb{E}[S]}{\mathbb{E}[\sqrt{S}]}$. In the feasible region $x\geq 1$, K^* has a minimum at $x=1$ and is increasing for $x>1$. Thus the lower bound for K^* is 1. This is attained when $x=1$; i.e. $\mathbb{E}[S]=\mathbb{E}[\sqrt{S}]$, which implies that S only takes values on 0 and 1.

Now we consider the upper bound. Suppose there is some S for which $K^*(x) \geq \frac{N^2}{4(N-1)}$. Then, since $K^*(x)$ is an increasing function, $x \geq \sqrt{N-1}$. That means $\mathbb{E}[S] \geq \sqrt{N-1}\mathbb{E}[\sqrt{S}]$, which can only be true if $\mathbb{P}(S = N-1) = 1$; that is, we have the Duffie and Zhu network. □

Bounds for $K^\dagger$

For $K^\dagger$, the lower bound 1 is achieved when $\mathbb{V}[S] = \mathbb{V}[\sqrt{S}]$; i.e. when S only takes values on 0 or 1. It is easy to see from the definition of $K^\dagger$ that it cannot attain a value lower than 1.

The proof for the upper bound is more complicated, so we first explain heuristically how it proceeds. We are trying to find a candidate probability vector $p \in [0,1)^N$ which maximises $K^\dagger(p)$ subject to $\sum_i p_i = 1$ (where i takes values in $\{0,1,\ldots,N-1\}$). Proposition 4.3 stipulates that we are maximising over the set for which $K^\dagger(p)$ is defined, so we can exclude any vectors p with any element with the value 1. We call these excluded points 'vertices'.

We will show that $K^\dagger(p)$ is maximised on the boundary of this feasible set. That means we reduce our focus from a space with N dimensions to one with $N-1$ dimensions. By the same argument $K^\dagger(p)$ must be maximised on the boundary of that set too which has $N-2$ dimensions, and so on inductively. Eventually we are reduced to one-dimensional line segments; that is, the subspaces where p has at most two non-zero elements. At this point, we cannot employ the inductive argument any more, because the boundaries of these line segments are the vertices at which are not contained in the feasible set for p. Therefore we know $K^\dagger(p)$ is maximised on these line segments, and we can then finish the proof by examining these to find the maximal value.

Write $\Omega(p) = \frac{\mathbb{V}[S]-\mathbb{V}[\sqrt{S}]}{\mathbb{E}[S^{1.5}]-\mathbb{E}[S]\mathbb{E}[\sqrt{S}]}$. Since $\Omega(p) \geq 0$, maximising $K^\dagger = \frac{1}{4}\Omega^2 + 1$ is equivalent

to maximising $\Omega(p)$. We can write down the following Kuhn-Tucker conditions:

$$\frac{\partial \Omega}{\partial p_i} + \lambda_i + \zeta = 0, \quad \forall i \tag{4.28}$$

$$\lambda_i p_i = 0, \quad \forall i \tag{4.29}$$

$$\zeta(\sum_i p_i - 1) = 0, \tag{4.30}$$

where each λ_i is the Lagrange multiplier for the constraint $p_i \geq 0$, and ζ is the Lagrange multiplier for the constraint $\sum_i p_i = 1$. As we are trying to find the maximum of Ω, the Lagrange multipliers are all non-negative.

Computing the partial derivative of Ω, equation (4.28) becomes:

$$\begin{aligned} -(\lambda_i + \zeta)v^2 =& (E_{1.5} - E_1 E_{0.5})(i^2 - 2iE_1 - i + 2\sqrt{i}E_{0.5}) \\ & - (E_2 - E_1^2 - E_1 + E_{0.5}^2)(i^{1.5} - iE_{0.5} - \sqrt{i}E_1), \end{aligned} \tag{4.31}$$

where v is the denominator of Ω, and E_m is shorthand for $\mathbb{E}[S^m]$.

Multiplying (4.31) by p_i and summing over i gives:

$$\begin{aligned} -N\zeta v^2 &= (E_{1.5} - E_1 E_{0.5})(E_2 - 2E_1^2 - E_1 + 2E_{0.5}^2) \\ &\quad - (E_2 - E_1^2 - E_1 + E_{0.5}^2)(E_{1.5} - 2E_1 E_{0.5}) \\ &= E_1(E_{0.5}E_2 - E_1 E_{1.5}) + E_{0.5}(E_{0.5}E_{1.5} - E_1^2). \end{aligned} \tag{4.32}$$

But all the terms on the right-hand side of (4.32) are greater than or equal to zero,[29] while those on the left-hand side are less than or equal to zero. Therefore each term must be exactly equal to zero, and so S is a constant. At such points, Ω is not defined, so we can surmise that $\Omega(p)$ is maximised on the boundary of the feasible set for p.

This boundary is the set of points in $[0,1)^N$ with at least one zero element. Suppose that $\Omega(p)$ is maximised on a point on the boundary where the element in position i' is zero. Then we can try to maximise $\Omega(p)$ on the corresponding subspace, which has $N-1$ dimensions. Repeating the analysis above, we can see that Ω must be maximised on the

[29] We showed this in the proof of Proposition 4.2, where we used Proposition 4.9 to show that $T_1 T_4 \geq T_2 T_3$ and $T_1 T_3 \geq T_2^2$, with equality if and only if T is a constant.

boundary of this subspace, which is the set of points in $[0,1)^N$ with at least two zero elements, one of which is in position i'.

Repeating this argument $N-2$ times we can infer that Ω is maximised on the set of points in $[0,1)^N$ which have exactly two non-zero elements. We cannot go any further, because the boundary of this subspace is the set of points with only one non-zero element; that is, the vertices, which we know do not lie in the feasible set of p.

Consider a general point p in this subspace, with its non-zero elements at i and j, where $i < j$. Suppose $p_i = \delta$ and $p_j = 1-\delta$, for some $0 < \delta < 1$. Then $E_m = \delta i^m + (1-\delta) j^m$ for each m, and so:

$$\Omega(p) = \frac{(j-i)^2 - (\sqrt{j}-\sqrt{i})^2}{(j-i)(\sqrt{j}-\sqrt{i})} \tag{4.33}$$

with the terms in δ cancelling out in the numerator and denominator of (4.33). This means that Ω is constant along any line in the subspace, and we simply need to choose i and j to maximise (4.33).

Noting that $j-i = (\sqrt{j}-\sqrt{i})(\sqrt{j}+\sqrt{i})$, we can cancel terms to obtain:

$$\Omega(p) = \sqrt{j}+\sqrt{i} - \frac{1}{\sqrt{j}+\sqrt{i}}, \tag{4.34}$$

which is maximised when $\sqrt{j}+\sqrt{i}$ is as large as possible. This means that $j = N-1$ and $i = N-2$, and so $\Omega(p) = 2\sqrt{N-2}$ and $K^\dagger = N-1$. □

Proof of Proposition 4.4

We have already shown the second part is true, because $\alpha = 3$ for the DMS network and so $K^\dagger$ grows without bound according to Table 4.1. (In fact this will be true for any scale-free network.) Now we just need to show that the asymptotic value of K^* — which we know tends to a finite limit — is smaller than 2.

Observe that $\mathbb{E}_t(S) = \frac{4t-2}{t+1} \to 4$ in the DMS network, since at each step of the network construction process four additional links are created (each new node has an in-link and an out-link with two existing nodes). Dorogovtsev, Mendes, and Samukhin (2001) show

that:

$$\lim_{t\to\infty}\mathbb{P}_t(s)=\frac{12}{s(s+1)(s+2)},$$

and so the term $\frac{\mathbb{E}_t[S]}{\mathbb{E}_t[\sqrt{S}]}$ is asymptotically equal to:

$$4/\left(12\sum_{s=2}^{\infty}\frac{1}{\sqrt{s}(s+1)(s+2)}\right)=2.17.$$

This gives us an asymptotic value for $K^*=1.73<2$. Therefore in the infinite limit the CCP never reduces expected netted exposures, except in the trivial case where $K=1$; i.e. when the CCP clears the only asset class. □

Proof of Proposition 4.5

Let us use μ and V to denote $\mathbb{E}[S]$ and $\mathbb{V}[S]$ respectively, for a given N. Then, as $N\to\infty$, we can make use of the following approximations:

$$\mathbb{E}[\sqrt{S}]=\mu^{0.5}\mathbb{E}\left[1+\left(\frac{S}{\mu}-1\right)\right]^{0.5}\approx\mu^{0.5}\left[1-\frac{V}{8\mu^2}\right], \tag{4.35}$$

and

$$\mathbb{E}[S^{\frac{3}{2}}]=\mu^{1.5}\mathbb{E}\left[1+\left(\frac{S}{\mu}-1\right)\right]^{1.5}\approx\mu^{1.5}\left[1+\frac{3V}{8\mu^2}\right]. \tag{4.36}$$

In both cases we have expanded the binomial series around 1 and neglected cubic and higher-order terms. Substituting these into our expressions for K^* and $K^\dagger$, we find that both $K^*\sim O(\mu)$ and $K^\dagger\sim O(\mu)$. □

Proof of Proposition 4.6

1. Let K_1^* and K_2^* represent the values of K^* for constellations Ψ_1 and Ψ_2 respectively, as given by equation $(\star')$. We need to show that $K_1^*<K_2^*$. As A and the distribution

S are the same under both constellations, we can show that:

$$\begin{aligned}
& K_1^* < K_2^*, \\
\iff \quad & \frac{1}{A}\Big(\sqrt{K_1^*}+\sqrt{K_1^*-A}\Big) < \frac{1}{A}\Big(\sqrt{K_2^*}+\sqrt{K_2^*-A}\Big), \\
\iff \quad & \Big(\sum_{m=1}^{M}\sqrt{a_m}\Big)^{-1}\cdot\frac{\mathbb{E}[S]}{\mathbb{E}[\sqrt{S}]} < \Big(\sqrt{a_1+a_2}+\sum_{m=3}^{M}\sqrt{a_m}\Big)^{-1}\cdot\frac{\mathbb{E}[S]}{\mathbb{E}[\sqrt{S}]}, \\
\iff \quad & \sqrt{a_1}+\sqrt{a_2} > \sqrt{a_1+a_2},
\end{aligned}$$

which is true for any $a_1, a_2 > 0$.

Now let us define $K_1^\dagger$ and $K_2^\dagger$ similarly. We have:

$$\begin{aligned}
& K_1^\dagger < K_2^\dagger, \\
\iff & \frac{A\mathbb{V}[S]-\Big(\sum_{m=1}^{M}\sqrt{a_m}\Big)^2\mathbb{V}[\sqrt{S}]}{\Big(\sum_{m=1}^{M}\sqrt{a_m}\Big)\Big(\mathbb{E}[S^{\frac{3}{2}}]-\mathbb{E}[S]\mathbb{E}[\sqrt{S}]\Big)} < \\
& \qquad \frac{A\mathbb{V}[S]-\Big(\sqrt{a_1+a_2}+\sum_{m=3}^{M}\sqrt{a_m}\Big)^2\mathbb{V}[\sqrt{S}]}{\Big(\sqrt{a_1+a_2}+\sum_{m=3}^{M}\sqrt{a_m}\Big)\Big(\mathbb{E}[S^{\frac{3}{2}}]-\mathbb{E}[S]\mathbb{E}[\sqrt{S}]\Big)}, \\
\iff & \frac{A\mathbb{V}[S]}{\sum_{m=1}^{M}\sqrt{a_m}} - \Big(\sum_{m=1}^{M}\sqrt{a_m}\Big)\mathbb{V}[\sqrt{S}] \\
& \qquad < \frac{A\mathbb{V}[S]}{\sqrt{a_1+a_2}+\sum_{m=3}^{M}\sqrt{a_m}} - \Big(\sqrt{a_1+a_2}+\sum_{m=3}^{M}\sqrt{a_m}\Big)\mathbb{V}[\sqrt{S}], \\
\iff & 0 < A\mathbb{V}[S]\Big(\frac{1}{\sqrt{a_1+a_2}+\sum_{m=3}^{M}\sqrt{a_m}} - \frac{1}{\sum_{m=1}^{M}\sqrt{a_m}}\Big) \\
& \qquad + \Big(\sqrt{a_1}+\sqrt{a_2}-\sqrt{a_1+a_2}\Big)\mathbb{V}[\sqrt{S}],
\end{aligned} \tag{4.37}$$

which is always true since $\sqrt{a_1}+\sqrt{a_2} > \sqrt{a_1+a_2}$ for any $a_1, a_2 > 0$. Thus the first part of the proposition is proved.

2. Let K_2^* and K_3^* represent the values of K^* for constellations Ψ_2 and Ψ_3 respec-

tively. We need to show that $K_2^* > K_3^*$. We have:

$$\begin{aligned}
& K_2^* > K_3^*, \\
\iff \quad & \frac{1}{A+1}\Big(\sqrt{K_2^*}+\sqrt{K_2^*-(A+1)}\Big) > \frac{1}{A+1}\Big(\sqrt{K_3^*}+\sqrt{K_3^*-(A+1)}\Big), \\
\iff \quad & \Big(\sqrt{a_1+1}+\sum_{m=2}^{M}\sqrt{a_m}\Big)^{-1}\cdot\frac{\mathbb{E}[S]}{\mathbb{E}[\sqrt{S}]} > \Big(\sqrt{a_2+1}+\sum_{m\neq 2}\sqrt{a_m}\Big)^{-1}\cdot\frac{\mathbb{E}[S]}{\mathbb{E}[\sqrt{S}]}, \\
\iff \quad & \sqrt{a_1+1}+\sqrt{a_2} < \sqrt{a_1}+\sqrt{a_2+1}, \\
\iff \quad & \sqrt{a_1+1}-\sqrt{a_1} < \sqrt{a_2+1}-\sqrt{a_2},
\end{aligned} \tag{4.38}$$

which is true since $a_1 > a_2$ and $\sqrt{X+1}-\sqrt{X}$ is a decreasing function.

Now let us define $K_2^\dagger$ and $K_3^\dagger$ similarly. We have:

$$\begin{aligned}
& K_2^\dagger > K_3^\dagger, \\
\iff & \frac{(A+1)\mathbb{V}[S]-\Big(\sqrt{a_1+1}+\sum_{m=2}^{M}\sqrt{a_m}\Big)^2\mathbb{V}[\sqrt{S}]}{\Big(\sqrt{a_1+1}+\sum_{m=2}^{M}\sqrt{a_m}\Big)\Big(\mathbb{E}[S^{\frac{3}{2}}]-\mathbb{E}[S]\mathbb{E}[\sqrt{S}]\Big)} > \\
& \qquad \frac{(A+1)\mathbb{V}[S]-\Big(\sqrt{a_2+1}+\sum_{m\neq 2}\sqrt{a_m}\Big)^2\mathbb{V}[\sqrt{S}]}{\Big(\sqrt{a_2+1}+\sum_{m\neq 2}\sqrt{a_m}\Big)\Big(\mathbb{E}[S^{\frac{3}{2}}]-\mathbb{E}[S]\mathbb{E}[\sqrt{S}]\Big)}, \\
\iff & (A+1)\mathbb{V}[S]\Big(\frac{1}{\sqrt{a_1+1}+\sqrt{a_2}+\sum_{m=3}^{M}\sqrt{a_m}}-\frac{1}{\sqrt{a_1}+\sqrt{a_2+1}+\sum_{m=3}^{M}\sqrt{a_m}}\Big)+ \\
& \qquad \mathbb{V}[\sqrt{S}]\Big(\sqrt{a_1}+\sqrt{a_2+1}-\sqrt{a_1+1}-\sqrt{a_2}\Big) > 0,
\end{aligned} \tag{4.39}$$

which is always true since $a_1 > a_2$. Thus all parts of the proposition are proved. □

Proof of Proposition 4.8

Consider a single agent i and a single asset class k. In the trading network, the agent has S_i links, each distributed $N(0,\sigma^2)$. As these are all exposures to the same asset class, they are

perfectly correlated and so the total market risk conditional on S_i is $M_{ik}|S_i \sim N(0, \sigma^2 S_i^2)$. Then market risk is:

$$\begin{aligned} \mathbb{V}[M_{ik}] &= \mathbb{E}\Big[\mathbb{V}[M_{ik}|S_i]\Big] + \mathbb{V}\Big[\mathbb{E}[M_{ik}|S_i]\Big], \\ &= \mathbb{E}\Big[\sigma^2 S_i^2] + \mathbb{V}[0], \\ &= \sigma^2 \mathbb{E}[S^2]. \end{aligned} \tag{4.40}$$

For the clearing network, exactly the same procedure gives market risk of $\tilde{\sigma}^2 \mathbb{E}[\tilde{S}^2]$ for each agent in each asset class. As the number of agents and asset classes does not change, conservation of market risk means that the two terms must be equal.

The assumption of perfect correlation assumes that the agent has no offsetting derivative positions. We can allow for this, so long as the extent of such hedging is the same in the trading and clearing networks. Let h_{ik} be the proportion of positions in asset class k which do not represent hedging positions. Then overall market risk in the trading network is $h_{ik}\sigma^2 \mathbb{E}[S^2]$ and market risk in the clearing network is $h_{ik}\tilde{\sigma}^2 \mathbb{E}[\tilde{S}^2]$, so Proposition 4.8 still holds.

Chapter 5

Conclusion

In this dissertation, I examine the consequences for the core functions of the banking system of the global financial crisis and subsequent reforms. Freixas and Rochet (2008) describe four main categories of banking functions:

1. liquidity and payment services,
2. transforming assets,
3. processing information and monitoring borrowers,
4. managing risks.

Each of the three essays in this dissertation speaks to one or more of these functions. Cryptocurrencies are designed as a way to avoid relying on banks for payment services. My first essay introduces a model in which agents choose between cryptocurrency or a conventional form of money (e.g. bank deposits) to make payments for consumption goods. I show that price formation in cryptocurrencies is essentially different from that for other asset classes, and demand curves can be upward sloping. I show further that this may explain why cryptocurrency prices are so volatile. My analysis has important implications for the future adoption of cryptocurrencies: it is possible that the low level of adoption so far is a result of speculation taking up blockchain space. The history we have observed up to now is entirely consistent with an outcome where cryptocurrencies are adopted widely as a means of payment.

My second essay speaks to the asset transformation and information processing functions of banks. In the model, a relationship bank provides an important role in monitoring the quality of the borrower, thus allowing credit to be steered to where it is most economically useful. However, concerns about the ongoing health of the lender mean that the borrower cannot rely on the bank to perform asset transformation when it is required. As awareness of bank fragility has increased in the aftermath of the global financial crisis, this channel may become more important. This suggests that firms may choose permanently lower levels of investment than they did before the crisis, even though projects may be no less productive, and individual banks no less willing to lend.

Finally, my third essay assesses a consequence of the regulatory response to banks' risk management failures during the global financial crisis. It extends the earlier literature on central clearing by producing results for a general network configuration, and by considering the impact on the variance, as well as the mean, of netted exposures.

References

Aldasoro, I., D. D. Gatti, and E. Faia (2016). "Bank networks: contagion, systemic risk and prudential policy". *BIS working paper* 597.

Almeida, H., M. Campello, B. Laranjeira, and S. Weisbenner (2011). "Corporate debt maturity and the real effects of the 2007 credit crisis". *Critical Finance Review* 1, pp. 3–58.

Amara, R. (2016). "Oxford Essential Quotations, 4th edition". *Oxford University Press*. URL: http://www.oxfordreference.com/view/10.1093/acref/9780191826719.001.0001/q-oro-ed4-00018679.

Anderson, S., J.-P. Dion, and H. Pérez Saiz (2013). "To link or not to link? Netting and exposures between central counterparties". *Journal of Financial Market Infrastructures* 1.4, pp. 3–29.

Angeletos, G.-M. and I. Werning (2006). "Crises and prices: information aggregation, multiplicity, and volatility". *American Economic Review* 96.5, pp. 1720–1736.

Athey, S., I. Parashkevov, V. Sarukkai, and J. Xia (2016). "Bitcoin pricing, adoption and usage: theory and evidence". *SSRN* 2826674.

Baltensperger, E. (1978). "Credit rationing: issues and questions". *Journal of Money, Credit and Banking* 10.2, pp. 170–183.

Barabási, A. and R. Albert (1999). "Emergence of scaling in random networks". *Science* 286, pp. 509–512.

Baumol, W. J. (1952). "The transactions demand for cash: an inventory theoretic approach". *Quarterly Journal of Economics* 66.4, pp. 545–556.

BBC News (2017). "Steam stops accepting payments in bitcoins". URL: https://www.bbc.co.uk/news/technology-42264622.

Bech, M. and R. Garratt (2017). "Central bank cryptocurrencies". *BIS Quarterly Review*, pp. 55–70.

Beck, T., H. Degryse, R. de Haas, and N. van Horen (2018). "When arm's length is too far: relationship banking over the business cycle". *Journal of Financial Economics* 127.1, pp. 174–196.

Biais, B., C. Bisière, M. Bouvard, and C. Casamatta (2017). "The blockchain folk theorem". *Toulouse School of Economics working paper* 817.

Biais, B., C. Bisière, M. Bouvard, C. Casamatta, and A. Menkveld (2018). "Equilibrium bitcoin pricing". *SSRN* 3261063.

Böhme, R., N. Christin, B. Edelman, and T. Moore (2015). "Bitcoin: economics, technology, and governance". *Journal of Economic Perspectives* 29.2, pp. 213–38.

Bolt, W. and M. R. C. van Oordt (2016). "On the value of virtual currencies". *Bank of Canada staff working paper* 42.

Bolton, P., X. Freixas, L. Gambacorta, and P. E. Mistrulli (2016). "Relationship and transaction lending in a crisis". *Review of Financial Studies* 29.10, pp. 2643–2676.

Bond, P., A. Edmans, and I. Goldstein (2012). "The real effects of financial markets". *Annual Review of Financial Economics* 4.2, pp. 339–360.

Bongini, P., M. L. di Battista, and L. Nieri (2015). "Relationship lending through the cycle: what can we learn from three decades of research?" *SSRN* 2925893.

Borgatti, S. and M. Everett (1999). "Models of core/periphery structures". *Social Networks* 21, pp. 375–395.

Borovkova, S. and H. Lalaoui El Mouttalibi (2013). "Systemic risk and centralized clearing of OTC derivatives: a network approach". *Social Science Research Network* 2334251.

Bowden, R., H. P. Keeler, A. E. Krzesinski, and P. Taylor (2018). "Block arrivals in the Bitcoin blockchain". *arXiv:1801.07447*.

Budish, E. (2018). "The economic limits of bitcoin and the blockchain". *National Bureau of Economic Research working paper* 24717.

Campbell, S. (2014). "Estimating the effect of central clearing on credit derivative exposures". *FEDS Notes*.

Campello, M., J. R. Graham, and C. R. Harvey (2010). "The real effects of financial constraints: evidence from a financial crisis". *Journal of Financial Economics* 97.3, pp. 470–487.

Carlsson, H. and E. Van Damme (1993). "Global games and equilibrium selection". *Econometrica* 61.5, pp. 989–1018.

Carney, M. (2018). "The future of money". *Speech*. URL: https://www.bankofengland.co.uk/speech/2018/mark-carney-speech-to-the-inaugural-scottish-economics-conference.

Carter, N. and Hasufly (2018). "An analysis of batching in Bitcoin". *Coinmetrics*. URL: https://coinmetrics.io/batching.

Castiglionesi, F., Z. Li, and K. Ma (2016). "Bank information sharing and liquidity risk". *Mimeo*.

Chang, B. and S. Zhang (2015). "Endogenous market making and network formation". *Systemic Risk Centre discussion paper* 50.

Chodorow-Reich, G. (2014). "The employment effects of credit market disruptions: firm-level evidence from the 2008–9 financial crisis". *Quarterly Journal of Economics* 129.1, pp. 1–59.

Choudhury, S. R. (2018). "A bitcoin conference has stopped taking bitcoin payments because they don't work well enough". *CNBC*. URL: https://www.cnbc.com/2018/01/10/bitcoin-conference-stops-accepting-cryptocurrency-payments.html.

Ciaian, P., M. Rajcaniova, and D. Kancs (2016). "The economics of bitcoin price formation". *Applied Economics* 48.19, pp. 1799–1815.

Cingano, F., F. Manaresi, and E. Sette (2016). "Does credit crunch investment down? New evidence on the real effects of the bank-lending channel". *Review of Financial Studies* 29.10, pp. 2737–2773.

Clifford, J. (2017). "Understanding the block size debate". URL: https://medium.com/@jcliff/understanding-the-block-size-debate-351bdbaaa38.

Cochrane, J. H. (2018). "The Bitcoin market isn't irrational". *Chicago Booth Review*. URL: http://review.chicagobooth.edu/finance/2018/article/bitcoin-market-isn-t-irrational.

Coinmetrics Team (2018). "On the difficulty of estimating on-chain transaction volume". URL: https://coinmetrics.io/difficulty-estimating-chain-transaction-volume/.

Coleman, L. (2018). "Dark web users ditch Bitcoin for Litecoin due to costly, slow transactions". *CCN*. URL: https://www.ccn.com/dark-web-users-favoring-litecoin-due-to-bitcoins-costly-slow-transactions.

Cong, L. W., Y. Li, and N. Wang (2018). "Tokenomics: dynamic adoption and valuation". *SSRN* 3153860.

Cont, R. and T. Kokholm (2014). "Central clearing of OTC derivatives: bilateral vs multilateral netting". *Statistics and Risk Modeling* 31.1, pp. 3–22.

Cox, N., N. Garvin, and G. Kelly (2013). "Central counterparty links and clearing system exposures". *Research Bank of Australia discussion paper* 12.

Craig, B. and G. von Peter (2014). "Interbank tiering and money center banks". *Journal of Financial Intermediation* 23, pp. 322–47.

Cunliffe, J. (2018). "Central clearing and resolution — learning some of the lessons of Lehmans". *Speech to FIA International Derivatives Expo*.

Darmouni, O. (2018). "Estimating informational frictions in sticky relationships". *SSRN* 2882659.

Dell'Ariccia, G. and R. Marquez (2006). "Lending booms and lending standards". *Journal of Finance* 61.5, pp. 2511–2546.

Dorogovtsev, S., J. Mendes, and A. Samukhin (2001). "Generic scale of the 'scale-free' growing networks". *Physical Review E* 63.062101.

Duchin, R., O. Ozbas, and B. A. Sensoy (2010). "Costly external finance, corporate investment, and the subprime mortgage credit crisis". *Journal of Financial Economics* 97.3, pp. 418–435.

Duffie, D., M. Scheicher, and G. Vuillemey (2015). "Central clearing and collateral demand". *Journal of Financial Economics* 116, pp. 237–56.

Duffie, D. and H. Zhu (2011). "Does a central clearing counterparty reduce counterparty risk?" *Review of Asset Pricing Studies* 1.1, pp. 74–95.

Easley, D., M. O'Hara, and S. Basu (2017). "From mining to markets: the evolution of bitcoin transaction fees". *SSRN* 3055380.

Faucette, J., B. Graseck, J. Moore, S. Shah, and C. Chan (2017). "Bitcoin decrypted: a brief teach-in and implications". *Morgan Stanley Research*.

Fernández-Villaverde, J. and D. Sanches (2016). "Can currency competition work?" *National Bureau of Economic Research working paper* 22157.

Fisher, I. (1911). *The purchasing power of money, its determination and relation to credit, interest and crises, First Edition*. The Macmillan Company.

Foucault, T., M. Pagano, and A. Roell (2013). *Market liquidity: theory, evidence, and policy*. Oxford University Press.

Freixas, X. and J.-C. Rochet (2008). *Microeconomics of Banking*. 2nd ed. MIT Press.

Galbiati, M. and K. Soramäki (2012). "Clearing networks". *Journal of Economic Behavior & Organization* 83, pp. 609–626.

Gandal, N., J. T. Hamrick, T. Moore, and T. Oberman (2018). "Price manipulation in the Bitcoin ecosystem". *Journal of Monetary Economics* 95, pp. 86–96.

Garlaschelli, D., S. Battiston, M. Castri, V. Servedio, and G. Caldarelli (2005). "The scale-free topology of market investments". *Physica A* 350, pp. 491–499.

Garratt, R. J. and N. Wallace (2018). "Bitcoin 1, bitcoin 2, ... : an experiment in privately issued outside monies". *Economic Inquiry* 56.3, pp. 1887–97.

Garratt, R. J. and P. Zimmerman (2017). "Centralized netting in financial networks". *Journal of Banking and Finance* In press.

Glosten, L. R. and P. R. Milgrom (1985). "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders". *Journal of Financial Economics* 14.1, pp. 71–100.

Goldstein, I. and A. Pauzner (2005). "Demand–deposit contracts and the probability of bank runs". *Journal of Finance* 60.3, pp. 1293–1327.

Greene, T. (2018). "Amazon rumor sends Bitcoin surging to record highs". *TNW*. URL: https://thenextweb.com/cryptocurrency/2017/10/13/amazon-rumor-sends-bitcoin-surging-record-highs.

Griffin, J. M. and A. Shams (2018). "Is Bitcoin really un-Tethered?" *SSRN* 3195066.

Hale, G., A. Krishnamurthy, M. Kudlyak, and P. Shultz (2018). "How futures trading changed bitcoin prices". *Federal Reserve Bank of San Francisco Economic Letter* 12.

Harsanyi, J. C. (1973). "Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points". *International Journal of Game Theory* 2.1, pp. 1–23.

Harvey, C. R. (2016). "Cryptofinance". *SSRN* 2438299.

Hautsch, N., C. Scheuch, and S. Voigt (2018). "Limits to arbitrage in markets with stochastic settlement latency". *Mimeo*.

Heath, A., G. Kelly, and M. Manning (2013). "OTC derivatives reform: netting and networks". *Reserve Bank of Australia conference proceedings*.

Heath, A., G. Kelly, M. Manning, S. Markose, and A. Rais Shaghagi (2016). "CCPs and network stability in OTC derivatives markets". *Journal of Financial Stability* 27, 217–233.

Hollifield, B., A. Neklyudov, and C. Spatt (2016). "Bid-ask spreads, trading networks and the pricing of securitizations". *Forthcoming in the Review of Financial Studies*.

Hombert, J. and A. Matray (2017). "The real effects of lending relationships on innovative firms and inventor mobility". *Review of Financial Studies* 30.7, pp. 2413–45.

Huberman, G., J. Leshno, and C. Moalleni (2017). "Monopoly without a monopolist: an economic analysis of the bitcoin payment system". *CEPR discussion paper* 12322.

Inaoka, H., T. Ninomiya, T. Shimizu, H. Takayasu, and K. Taniguchi (2004). "Fractal network derived from banking transaction — an analysis of network structures formed by financial institutions". *Bank of Japan working paper* 04-E-04.

Ivashina, V. and D. Scharfstein (2010). "Bank lending during the financial crisis of 2008". *Journal of Financial Economics* 97.3, pp. 319–338.

Jackson, J. and M. Manning (2007). "Comparing the pre-settlement risk implications of alternative clearing arrangements". *Bank of England working paper* 321.

Kahle, K. M. and R. M. Stulz (2013). "Access to capital, investment, and the financial crisis". *Journal of Financial Economics* 110.2, pp. 280–299.

Kahn, C. M. (2018). "Payment systems and privacy". *Federal Reserve Bank of St. Louis Review*.

Khwaja, A. I. and A. Mian (2008). "Tracing the impact of bank liquidity shocks: evidence from an emerging market". *American Economic Review* 98.4, pp. 1413–1442.

Kirkby, R. (2018). "Cryptocurrencies and digital fiat currencies". *Australian Economic Review* 51.4, pp. 527–539.

Koning, J. P. (2018). "Anonymous digital cash". *American Institute for Economic Research*. URL: https://www.aier.org/article/sound-money-project/anonymous-digital-cash.

Krugman, P. (2018). "Bubble, bubble, fraud and trouble". *New York Times*. URL: https://www.nytimes.com/2018/01/29/opinion/bitcoin-bubble-fraud.html.

Kyle, A. S. (1985). "Continuous auctions and insider trading". *Econometrica* 53.6, pp. 1315–1335.

Langfield, S., Z. Liu, and T. Ota (2014). "Mapping the UK interbank system". *Journal of Banking & Finance* 45, pp. 288–303.

Liu, Y. and A. Tsyvinski (2018). "Risks and returns of cryptocurrency". *National Bureau of Economic Research working paper* 24877.

Luther, W. J. and J. Olson (2013). "Bitcoin is memory". *Journal of Prices & Markets* 3.3, pp. 22–33.

Markose, S. (2012). "Systemic risk from global financial derivatives: a network analysis of contagion and its mitigation with super-spreader tax". *IMF working paper* 12.282.

Matzutt, R., J. Hiller, M. Henze, J. H. Ziegeldorf, D. Müllmann, O. Hohlfeld, and K. Wehrle (2018). "A quantitative analysis of the impact of arbitrary blockchain content on Bitcoin". *International Conference on Financial Cryptography and Data Security*.

Meiklejohn, S., M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage (2013). "A fistful of bitcoins: characterizing payments among men with no names". *Proceedings of the 2013 Conference on Internet Measurement*, pp. 127–140.

Menkveld, A. (2015). "Crowded trades: an overlooked systemic risk for central clearing counterparties". *Social Science Research Network* 2422250.

Moore, T. and N. Christin (2013). "Beware the middleman: empirical analysis of Bitcoin exchange risk". *Financial Cryptography and Data Security* 7859.

Morgan, D. (2017). "The great Bitcoin scaling debate — a timeline". *Hacker Noon*. URL: https://hackernoon.com/the-great-bitcoin-scaling-debate-a-timeline-6108081dbada.

Morris, S. and H. S. Shin (1998). "Unique equilibrium in a model of self-fulfilling currency attacks". *American Economic Review* 88.3, pp. 587–597.

– (2003). "Global games: theory and applications". *Econometric Society Monographs* 1, 56–114.

Murphy, D., M. Vasios, and N. Vause (2014). "An investigation into the procyclicality of risk-based initial margin models". *Bank of England Financial Stability paper* 29.

Nakamoto, S. (2008). "Bitcoin: a peer-to-peer electronic cash system". *Mimeo*.

– (2009). "Bitcoin open source implementation of P2P currency". *P2P Foundation*. URL: http://p2pfoundation.ning.com/forum/topics/bitcoin-open-source.

O'Malia, S. (2017). "CEO opening remarks". *Speech to International Swaps and Derivatives Association*.

Ongena, S., D. C. Smith, and D. Michalsen (2003). "Firms and their distressed banks: lessons from the Norwegian banking crisis". *Journal of Financial Economics* 67.1, pp. 81–112.

Pagnotta, E. and A. Buraschi (2018). "An equilibrium valuation of Bitcoin and decentralized network assets". *SSRN* 3142022.

Peek, J. and E. S. Rosengren (1997). "The international transmission of financial shocks: the case of Japan". *American Economic Review* 87.4, pp. 495–505.

Pinches, G. E. and W. R. Kinney Jr. (1971). "The measurement of the volatility of common stock prices". *Journal of Finance* 26.1, pp. 119–125.

Poon, J. and T. Dryja (2016). "The Bitcoin Lightning Network: scalable off-chain instant payments". *Mimeo*. URL: https://lightning.network/lightning-network-paper.pdf.

Popper, N. (2015). *Digital gold: the untold story of bitcoin*. Allen Lane.

Rajan, R. G. (1992). "Insiders and outsiders: the choice between informed and arm's-length debt". *Journal of Finance* 47.4, pp. 1367–1400.

Roberts, J. J. (2017). "Think bitcoin is a bubble? 5 ways to short it". *Fortune*. URL: http://fortune.com/2017/09/14/bitcoin-bubble-investments-short.

Rochet, J.-C. (2008). *Why are there so many banking crises? The politics and policy of bank regulation*. Princeton University Press.

Rochet, J.-C. and X. Vives (2004). "Coordination failures and the lender of last resort: was Bagehot right after all?" *Journal of the European Economic Association* 2.6, pp. 1116–1147.

Saleh, F. (2017). "Blockchain without waste: proof-of-stake". *SSRN* 3183935.

– (2018). "Volatility and welfare in a crypto economy". *SSRN* 3235467.

Schilling, L. and H. Uhlig (2018). "Some simple bitcoin economics". *National Bureau of Economic Research working paper* 24483.

Sedgwick, K. (2018). "A brief history of hidden messages in the Bitcoin blockchain". *Bitcoin.com*. URL: https://news.bitcoin.com/a-brief-history-of-hidden-messages-in-the-bitcoin-blockchain/.

Sharpe, S. A. (1990). "Asymmetric information, bank lending, and implicit contracts: a stylized model of customer relationships". *Journal of Finance* 45.4, pp. 1069–1087.

Sidanius, C. and F. Žikeš (2012). "OTC derivatives reform and collateral demand impact". *Bank of England Financial Stability paper* 18.

Sockin, M. and W. Xiong (2018). "A model of cryptocurrencies". *Mimeo*.

Song, J. (2017). "Understanding SegWit block size". *Medium*. URL: https://medium.com/@jimmysong/understanding-segwit-block-size-fd901b87c9d4.

Song, R., R. Sowers, and J. Jones (2014). "The topology of central counterparty clearing networks and network stability". *Stochastic Models* 30.1, pp. 16–47.

Soramäki, K., M. Bech, J. Arnold, R. Glass, and W. Beyeler (2007). "The topology of interbank payment flows". *Physica A* 379, pp. 317–33.

Stecklow, S., A. Harney, A. Irrera, and J. Kelly (2017). "Chaos and hackers stalk investors on cryptocurrency exchanges". *Reuters*. URL: https://www.reuters.com/article/us-bitcoin-exchanges-risks-special-repor/chaos-and-hackers-stalk-investors-on-cryptocurrency-exchanges-idUSKCN1C41AU.

Tobin, J. (1956). "The interest-elasticity of transactions demand for cash". *Review of Economics and Statistics* 38.3, pp. 241–247.

Van der Leij, M., D. in't Veld, and C. Hommes (2016). "The formation of a core periphery structure in heterogeneous financial networks". *De Nederlandsche Bank working paper* 528.

Von Thadden, E.-L. (1995). "Long-term contracts, short-term investment and monitoring". *Review of Economic Studies* 62.4, pp. 557–575.

– (2004). "Asymmetric information, bank lending and implicit contracts: the winner's curse". *Finance Research Letters* 1.1, pp. 11–23.

Wetherilt, A., P. Zimmerman, and K. Soramäki (2010). "The sterling unsecured loan market during 2006-08: insights from network theory". *Bank of England working paper* 398.

Wilmoth, J. (2017). "Want to short bitcoin? It's going to cost you". *CCN*. URL: https://www.ccn.com/want-short-bitcoin-going-cost.

Yermack, D. (2013). "Is bitcoin a real currency? An economic appraisal". *National Bureau of Economic Research working paper* 19747.