



The role of financial institutions : limits and perspectives

Louis Bertucci

► To cite this version:

Louis Bertucci. The role of financial institutions : limits and perspectives. Business administration. Université Paris sciences et lettres, 2019. English. NNT : 2019PSLED036 . tel-03222138

HAL Id: tel-03222138

<https://theses.hal.science/tel-03222138>

Submitted on 10 May 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL

Préparée à Université Paris-Dauphine

The Role of Financial Institutions : Limits and Perspectives

Soutenue par

Louis BERTUCCI

Le 19 Septembre 2019

École doctorale n°543

Ecole doctorale de Dauphine

Spécialité

Sciences de gestion

Composition du jury :

Christophe BISIÈRE

Professor, Université de Toulouse-Capitole

Rapporteur

Jérôme DUGAST

Associate Professor, Université Paris-Dauphine

Président du jury

Laurent GERMAIN

Professor, Toulouse Business School

Examineur

Christine PARLOUR

Professor, University of California, Berkeley

Rapporteur

Gilles CHEMLA

Professor, Université Paris-Dauphine

Directeur de thèse

The Role of Financial Institutions : Limits and Perspectives

*Thèse de Doctorat de l'Université PSL - Préparée à l'Université
Paris-Dauphine*

Soutenue par : **Louis BERTUCCI**

Le 19 septembre 2019

Acknowledgments

First and foremost, I would like to express my sincere gratitude to my PhD advisor Professor Gilles CHEMLA. His guidance and comments throughout the years have greatly improved the quality of this thesis. He was always very demanding regarding the quality of my work while supporting me and my ideas even when it was “all over the place”. For the time he dedicated to me, and his patience, I am very grateful.

I would like to thank the whole finance group from the DRM lab at Dauphine University. It has been a real pleasure to write my thesis there. I have benefited a lot from all the PhD courses, from the lab’s partnership with the Essec Business School, and from all the world class professors invited for either teaching top level courses on specific topics and/or for the weekly research seminars. Also being able to talk directly to the invited professors has been a remarkable source of inspiration and improvements on a lot of aspects of my thesis. But above all, I benefited a lot from comments and direct discussions with all the professors of DRM who have proven to be highly available and helpful.

I am also very grateful to Gilles for making me benefit from his network in academia. Upon his instigation, I have had the honor and the privilege to spend the second year of my thesis at the Haas Business School at the University of California, Berkeley. I would like to thank all the faculty members of the finance group at Haas for their warm welcome, but above all I would like to address my sincere thanks to Professor Christine Parlour who sponsored me during my time at Haas. Her comments during the (almost) weekly discussions during the early days have helped shape my thesis in a profound way. This time abroad has been a great source of inspiration and will definitely have an impact on my future in a lot of different aspects. I am also very grateful toward the members of my dissertation committee for the time they take to read my research. Their comments will help me improve my future work.

I would like to address special thanks to my brother Charles. Besides working together on the second chapter of this thesis, our frequent discussions on finance and mathematics in general as well as more direct comments on my thesis and the proofs of each chapter have greatly improved my understanding of some of the most important aspects described in this thesis. I am also very grateful toward my friend

Plamena for taking the time to check and proofread my dissertation.

I thank my fellow labmates for the interesting discussions we had on any kind of topics, especially in the middle of the afternoon while we were all supposed to work, but most of all for the support we were able to provide to each other during all those years.

Last but not least, I would like to thank my friends and my family : my parents and my three brothers for their unconditional support throughout my thesis and for understanding and never complaining about my lack of availability. They made it a lot easier for me to focus on my research.

Contents

1	Introduction	5
1.1	Case study: Central Clearing Party (CCP)	9
1.2	Policy Making with Hidden Information	12
1.3	Blockchain : An Autonomous Institution	15
1.4	The Future of Institutions	21
2	Variation Margins, Liquidity Management and Clearing Structure	25
2.1	Introduction	26
2.2	The Model	32
2.2.1	General Setup	32
2.2.2	The No-Margin Case	33
2.2.3	Bilateral Contracting with Variation Margins	34
2.2.4	Central Clearing through a CCP	36
2.3	Comparison of the Different Cases	39
2.3.1	Global Cash Usage	40
2.3.2	Magnitude and Probability of a Default	42
2.3.3	Some Numerical Calculations	46
2.4	Implications	49
2.4.1	Size of the Market	50
2.4.2	Riskiness of the Market	52
2.5	Conclusion	54
2.6	Appendix	57
2.6.1	Discussion and Comparative Statics on the Optimal Asset Allocation :	57
2.6.2	Proofs of Lemmas	61
2.6.3	Proofs of Propositions	66
3	Policy Making with Hidden Information and Continuous-time Learning	73
3.1	Introduction	74
3.2	The Model	81

3.3	Numerical Calculations	86
3.4	Interpretation	87
3.4.1	General case	87
3.4.2	Erroneously overconfident	92
3.4.3	Impact of the control cost and the signal's noise on the learning process	95
3.5	Exploration vs Exploitation	98
3.6	Conclusion	101
3.7	Appendix	104
3.7.1	Proofs	104
4	Incentives on the Lightning Network : A Blockchain-based Payment Network	107
4.1	Introduction	108
4.2	General Facts about the Lightning Network	114
4.3	Graph Analysis	116
4.3.1	Data	116
4.3.2	Static properties	118
4.3.3	Dynamic properties	123
4.4	The Model : Homogeneous Nodes	127
4.4.1	General setup - Exogenous fee	127
4.4.2	Two stages model - Endogenous routing fee	134
4.5	Heterogeneous Agents and Network Structure	139
4.5.1	Exogenous fee	139
4.5.2	Two stages model - Endogenous routing fee	141
4.5.3	Implications for the network topology	143
4.6	Bridge with the Real Lightning Network	145
4.7	Future Research	147
4.7.1	Theory	147
4.7.2	Empirical	149
4.8	Conclusion	150
4.9	Appendix	153
4.9.1	Proofs of Lemmas	153
4.9.2	Proofs of Propositions	154

Chapter 1

Introduction

Institutions, being the *rules of the game*, are an important determinant of economic activity. However, in the first half the 20th century, as neoclassical economics developed, rather low attention had been devoted to institutional mechanisms. It was not until the work of Williamson (1975, 1985, 1998) and the advent of the New Institutional Economics that proper economic tools were applied to institutional questions. North (1991) defines institutions as “humanly devised constraints that structure political, economic and social interactions”. This literature has its roots in the seminal work of Ronald Coase. In “The Nature of the Firm” (1937), he first defines the *costs of the pricing mechanism* which will later be interpreted as *transaction costs*. This idea, and what will later be called the *Coase Theorem*, was refined in “The Problem of Social Cost” (1960). The Coase theorem states that “with no transaction costs, regardless of the initial distribution of property rights, economic agents will bargain to obtain the pareto efficient allocation”. As explained in Allen (1999), the property rights approach defines transaction costs as “the cost of establishing and maintaining property rights” which includes any direct or indirect costs related to property rights. This provides a direct link between transaction costs, property rights and institutions. Indeed, if there are no transaction costs, property rights are well defined and institutions are not needed. However, the point of Coase, emphasized by Williamson, was to highlight that because there are strictly positive transaction costs in the economy, institutions matter.

Since the establishment of the first societies, the perceived role of institutions is very much in line with this property rights approach. Also, providing efficient “rules of the game” (understand institutions) has always been an important determinant of economic performance, as noted in North (1992). As long as we can trace back in history, we find evidence of property rights enforced by institutions. Ellickson and Thorland (1995) explain how in Mesopotamia, often considered the cradle of humanity, there was already a proper definition of property rights regarding land ownership, and some of the most ancient writings were legal documents specifying

how those property rights were to be enforced by *city-states*¹. Then in ancient Greece and in ancient Rome, philosophers further developed the idea of contracts, that can be viewed as institutional tools to facilitate the definition and transfer of property rights. Otherwise stated, contracts were meant to lower transaction costs. In the Middle-Ages, there was a strong influence of religion but most states continued to provide conditions under which different categories of people could own lands or wealth. Whether property rights were *private* or *public*, it seems that states and social institutions have always endorsed the role of defining how can someone own something, and how can someone dispute or transfer this ownership.

Throughout history, as *inter-state* trades became more and more important, the role of institutions evolved to protect the people within the state from the people outside the state. Around the 16th-18th century, the Mercantilism school of thought became dominant in modernized parts of Europe. At that time, states realized that to increase their wealth, they should try to increase their exports while decreasing their imports, and they did so by increasing transaction costs on international trades. This lasted until the advent of Classical economy lead by Adam Smith. In *The Theory of Moral Sentiment* (1759), he introduced the concept of *invisible hand*. The main idea was that by pursuing self-interested goals, individuals, through their actions, benefit the whole society. This was originally invoked in reference to income distribution but was extended in the late 18th and 19th century by authors like Jean-Baptiste Say and David Ricardo to justify freedom of trade and economic liberalism. With this view, states were meant to provide a legal and operational framework to support international trade. For instance, the Cobden-Chevalier Treaty signed between United Kingdom and France in 1860 was one of the first international trade agreement², and was meant to facilitate trade between the two signing nations.

At the beginning of the 20th century, Neoclassical economists analyzed the micro-foundations of the Classical economic mechanisms. With the seminal work of Alfred Marshall and Vilfredo Pareto (among others), an important paradigm shift occurred and economists started to view economics as a normative science, describing what ought to be. Around the same time, Arthur C. Pigou formalized welfare economics and helped understand what economic outcomes should prevail if agents were to focus only on their utility. Welfare economics introduced some very important changes. This allowed economists to formalize and measure the concept of externality, that is the effect of an agent's decision on other agents' welfare, and subsequently derive implications regarding institutional design. With this view, states then endorsed the role of correcting negative externalities and inefficiencies.

¹We can cite for instance the Laws of Ur-Nammu (c. 2100 b.c.), the Laws of Lipit-Ishtar (c. 1860 b.c.), the Laws of Eshnunna (c. 1930 b.c.), and the Code of Hammurabi (c. 1750 b.c.).

²Grossman (2016) describes it as the first modern trade agreement.

This trend further extends with the Keynesian economics that promotes a high degree of state intervention in the economy. With *The General Theory of Employment, Interest and Money* (1937), John M. Keynes provided a unified framework for analyzing economic growth and population employment, associated with an increased welfare for the whole economy. He identified the main factor influencing those aggregates to be the demand over the production, while pointing out that the state could influence the demand to drive the economy out of recession and crisis. This resulted in an increasing state interventionism in the economy. At this point the perceived role of institutions was not only to correct inefficiencies but also to drive the entire economy in order to increase overall welfare.

Note that we reduced almost all economic history to three main periods, and in each period, institutions endorsed a new role. From the earliest societies in Mesopotamia until the rise of international trade around the 16th century, states' main role was to guarantee and enforce property rights for individuals. Then, institutions were meant to promote and facilitate international trade while keep protecting property rights for individuals. With recent developments of welfare economics, it was recognized that institutions should correct inefficiencies resulting from individual interactions, and even drive the entire economy, while promoting international trade, liberalism and enforcing property rights at the national and international level.

This is not meant as a comprehensive history of economic thought but rather serve as an illustration of how the role of institutions has evolved over time. Each part of the world has had specific economic challenges in each period but this brief history of institutions captures the fundamental view of what was expected, broadly speaking, from the state. For instance, in the 20th century, economists came up with theories of economic regulation to analyze what drives (or what should drive) regulation. Regulatory agencies provide the legal framework in which economic agents can operate and are therefore an important part of the *institutions* as defined by North (1991). The role of those regulatory institutions differs in the two main theories. The Public Interest Theory of Regulation states that regulation is supplied in response to demand of the public for correction of inefficiencies, and is in line with the institution maximizing the welfare of all economic agents. The Capture Theory states that regulation is supplied in response to demand of interest groups failing to maximize the welfare of their members, and was an attempt to understand some recent regulations that were not necessarily correlated with market failures like regulations surrounding taxis, airlines, freight companies, etc³. The Capture Theory

³Note that from a property rights point of view, heavy regulation in the taxi industry for instance could be thought of as a form of property rights granted by the state to a few companies while trying to prevent new entries on the market.

is associated with the work of Georges Stigler and a number of authors have helped formalize this idea (including Bernstein (1955), Stigler (1971), Levine and Forrence (1990) and Laffont and Tirole (1991, 1993))⁴.

It should be noted that from a positive point of view, it is hard to assess the efficiency of the “rules of the game”. Also there are some direct critics about the way institutions implement regulation. Joskow and Rose (1989) propose an overview of the empirical regulation literature and conclude that “the effects of economic regulation often differs considerably from the predictions of ‘public interest’ models which presume that regulation is intended to ameliorate market imperfections and enhance efficiency”. They note that, in general, regulation is surplus-improving but that those results cannot generalize to any kind of *regulation* as this term covers a lot of different ideas. Lucas (1976)’s famous critic of policy evaluation also goes in this direction, stating that the environment is constantly changing and no regulation should be based on aggregate data influenced by regulation.

Broadly speaking, institutions define all aspects of economic and social interactions, however this dissertation focuses on the role of financial institutions. Keeping the same logic as North (1991), we can define *financial institutions* as the set of financial rules as well as organizations that shape the financial landscape. It is common to say that financial institutions should be independent from other state-level institutions but their objective is usually set by governments or anterior institutions. So if the financial policy itself is independent, its goal is not. The recent history of financial institutions is in line with what we described so far. Indeed, the rise of international trade in the 16th-17th century was accompanied with rapid development of those financial institutions. The increase in institutions’ intervention in the economy can relate to the financial history although the latter is somewhat different. Albeit strong regulations were put in place in the aftermath of the 1930’s financial crisis, a wave of deregulation took place in the late 20th century. We can cite for instance, the end of the *Bretton-Woods* system in 1976, the so called *Big-Bang* that dramatically changed the financial rules in the United Kingdom in 1986, or the decay of the *Glass-Steagall Act* throughout the second half of the 20th century. After the global financial crisis of 2008, however, a number of new regulations by financial institutions were enacted with Dodd-Frank act in the United States or EMIR in Europe for instance. Moreover, with mechanisms like low interest rate policy or quantitative easing, the degree of intervention of financial institutions in general remain high.

The goal of this dissertation is to highlight some limits of the role of financial

⁴The document summarizing the work of Jean Tirole when he received the Nobel prize in 2014 is entitled *Market Power and Regulation*, and the 2nd section is a summary of his contribution to regulation economics.

institutions, to understand what makes the process of policy making fundamentally complex and eventually to derive future trends regarding the role of financial institutions. We take a three step process, one presented in each chapter of this dissertation. First, we study the impact of the *central clearing* regulation which is a financial regulation that was introduced as a response to the financial crisis of 2008 through the Dodd-Frank recommendations. Second, we study the issues arising when an institution, like a central bank, has to make continuous policy decisions in presence of hidden information, which yield fundamental limits to the policy maker actions. Lastly, we propose an interpretation of a recent technological advancement, namely the blockchain, as a new kind of financial institution. The third chapter analyzes a payment network built on top of a blockchain, the *Lightning Network*, but this introductory chapter will purposely focus on the blockchain itself. We present the results of those three chapters as well as the link with the theme of financial institutions in the context outlined so far. We end this introduction with a discussion regarding future trends in the role of institutions.

1.1 Case study: Central Clearing Party (CCP)

In the aftermath of the global financial crisis of 2008, financial institutions across the world needed a way to stabilize markets and, above all, restore confidence in the economy. A characteristic of this crisis was the heavy interconnection of the global banking system. When some part of the financial system started to suffer losses, the fear of contagion to the whole system increased. Even worse, the interconnection in the financial system was opaque and largely due to unregulated derivative products. With a notional outstanding of around \$680 trillion by mid-2008, the Over The Counter (OTC) derivative market was seen as a large source of risk for the global financial system. This was viewed as an externality and consistently with the role of financial institutions since the 20th century, new regulations were introduced both in the United States (Dodd-Frank in 2010) and in Europe (EMIR in 2012) in order to mitigate this externality, the contagion risk.

Central clearing of OTC derivative contracts through Central Clearing Platform (CCP) was seen as a solution to prevent the failure of the financial system. Derivative contracts refer to an insurance scheme that a *protection seller* provides to a *protection buyer*. From a practical point of view, the failure of the system may arise from the counterparty exposure agents (banks) have among each others. If one protection seller goes bankrupt, he will not be able to provide insurance anymore in case the protection buyer needs it, which might trigger the default of the protection buyer. This is referred to as *counterparty risk*. Then if this protection buyer happened to have sold insurance to another protection buyer, this could trigger another

default, and so on. Central clearing refers to risk mitigation by a central party, in a way similar to traditional insurance mechanisms. This is essentially like if several agents willing to mitigate this counterparty risk constitute a common *insurance fund* that would be used in case someone defaults. Each agent contributes to the so called *default fund*, held by the CCP, in proportion of the risk they are responsible of. The CCP is a financial institution that effectively acts as the buyer to every seller and the seller to every buyer, removing the previous counterparty risk⁵

It is important to realize that central clearing has been around for a while outside of derivative contracts. All transactions on a market are usually cleared through a CCP. Indeed, when 2 agents contract on a market, the *settlement* is delayed from the actual trade. Therefore, one agent can commit to sell something (a stock, a bond, a commodity, etc.) but might have defaulted by the time the settlement is required. Since the 17th century and the development of international trade, CCPs have been used to insure market participants against the default of their counterparties. Note that the longer the time between the “trade” and the “settlement”, the higher the risk that one counterparty goes bankrupt. On the *spot market*, settlement usually takes place within 2 business days after the trade, therefore the risk is relatively small. However, the very characteristic of derivative contracts is that the settlement can take place up until years after the trade, increasing significantly the probability that one counterparty goes bankrupt. Otherwise stated, it increases significantly the probability that the internal fund of the CCP will be used to cover a member’s default.

It should already be noted that a well functioning CCP relies on good diversification. Indeed, if several members of the CCP were to default at the same time (impacted by a common risk for instance), the CCP itself could go bankrupt, which could trigger the default of the “surviving members” (the members not impacted by the common shock in the first place). Historically, CCP were used in very specific markets, and the risk was shared only among agents active on the same market. Bignon and Vuillemeys (2017) provide the first empirical analysis of the failure of a derivative CCP: the *Caisse de Liquidation* based in Paris. This CCP was particularly active on *sugar futures*. They found that a combination of market characteristics (high volatility of the sugar price), as well as management issues of the CCP itself, led to its default in 1974. This shows that a CCP does not eliminate the risk per se but rather concentrates and modifies it.

Central clearing is often opposed to bilateral clearing as another way to mitigate counterparty risk. In bilateral clearing, the value of the derivative product is computed on a daily basis, and the losing party is required to send the cash-

⁵Note that this introduces a new counterparty risk that is the risk of the CCP’s default, see Bignon and Vuillemeys (2017).

value of the loss (called variation margins) to the other party, at the end of each trading day. This is effectively reducing the risk horizon to one day, because if the counterparty is not able to pay the variation margins at the end of the day it will be declared bankrupt. This system of margin payment is also part of the central clearing scheme, but the variation margins are paid daily to the CCP instead of the bilateral counterparty.

The main motivation of this chapter is to show that clearing, whether bilateral or central, transforms the risk from a one-shot counterparty risk at the maturity of the derivative contract, to a daily liquidity risk. Similar to the traditional risk literature, there are some risks that can be diversified away but there is a part that usually cannot. It can, however, be reshaped and ultimately be borne by someone else. Without any forms of clearing, the protection buyer bears the counterparty risk of the protection seller. With clearing (bilateral or central), the risk is shifted to a liquidity risk borne by the protection seller as he needs to pay variation margins each day (to its counterparty or to the CCP depending on the clearing structure) and the protection buyer is insured by those *margins*. It is not a priori obvious who should bear the risk.

We build a unified framework for analyzing the global cash used by the system, as well as the overall risk, under different clearing structures. In equilibrium, protection sellers constitute a liquidity buffer on their balance sheet in order to face the variation margins payment, that represents a liquidity shock. With central clearing, they also need to lock some cash in the CCP's default fund. We find that the global cash used by the system is higher under central clearing but the global expected loss is lower, so there is a tradeoff when assessing the clearing structure of the economy. We also find that this tradeoff is impacted by the type of market that is considered. In markets with low volatility or with a few clearing members the global risk and cash used is very sensitive to the clearing structure, while the effect is marginal for high volatility markets.

The overall cash used in the economy is an important characteristic of the financial system in the 21st century. In addition to regulations introduced to restore confidence in financial markets, monetary institutions also injected a lot of cash in the system. The Federal Reserve, as well as the European Central Bank, significantly increased the size of their balance sheets by buying assets from banks, through a method referred to as *quantitative easing*. This essentially lowers the opportunity cost of capital and the interest rate, making cash holdings cheap. In the context of clearing this make the cash management of protection seller easier and less consequential but an interesting point is that if and when the opportunity cost of holding cash increases, it could significantly change the tradeoff between cash and risk and therefore the optimal clearing structure.

Institutions around the world have taken the duty to stabilize markets and restore confidence after the global financial crisis. This high level of interventionism should not be judged *per se* but it is interesting to see that from a historical perspective, it is completely in line with the recent trend in financial regulation⁶. Financial institutions came to think that it was their responsibility to drive the financial system and the economy out of the crisis, by correcting externalities like systemic risk and even drive economic growth.

The central clearing case is interesting as the regulator was a bit early compare to academic research regarding the cost and benefit of such measures. The Dodd-Frank act came out in 2010, and the EMIR regulation came out in 2012, but the original decision was taken during the Pittsburgh G20 summit in September 2009. However, a number of academic researchers started to study the real effects and consequences of central clearing after those dates. As a non-exhaustive list we can cite Acharya and Bisin (2014), Amini, Filipovic and Minca (2015), Biais, Heider and Hoerova (2016), Cont and Kokholm (2014), Duffie, Scheicher and Vuillemeys (2015), Duffie and Zhu (2011), Loon and Zhong (2014) and Menkveld (2017). All of them more or less concluded that central clearing introduced some important distortions in the market and that it was not perfectly clear that central clearing was always welfare-improving.

The point here is not to blame financial institutions but rather put in perspective the evolution of their role as perceived by economic agents. Right after the financial crisis, the threat of risk contagion was quite high among market participants and it was somehow expected that institutions should reinsure trust in the system. The system itself was relying on institutions to *do something*. The importance of those institutions seemed to be so high that not only they *can* influence the market, but they also *have to* influence it, and failure to do so can have important consequences on the overall stability. In a way, it is better to do something even if it is a little ahead of time or not studied enough by academic research (like requiring central clearing for most derivative products) than do nothing. The follow up question is therefore “What is the best financial institutions can do when shaping the ‘rules of the game’?” The next section proposes an answer to that by highlighting fundamental issues institutions are facing.

1.2 Policy Making with Hidden Information

This chapter seeks to provide theoretical foundations to how financial institutions can influence markets. Since the formalization of welfare economics, the usual view is an omniscient central planner that can eliminate any kind of friction or externality.

⁶Except for the small deregulation period of the late 20th century.

The underlying assumption is that all the information is available to the regulating institution, which can therefore solve any kind of situation and therefore improve the social welfare. This assumption is relaxed here and the regulation process with hidden information is analyzed.

Economic situations, *games*, with informational asymmetries have been analyzed extensively since the seminal work of Akerlof (1978). The resulting adverse selection and moral hazard are also highly studied subjects in the fields of economics and finance, but not so much from the regulation point of view. Laffont and Tirole (1986) present a model of regulation viewed as a principal-agent relationship, in the context of rent extraction by a group of firms. In their model, the regulator has partial information as he does not observe the true cost function of the firm. The second chapter of this dissertation extends this idea to a much more general framework that can be applied to a wide variety of situations.

The model features an institution that is in charge of regulating a group of agents (referred to as the *economy*), but the institution does not know the private type of the economy⁷. This informational asymmetry is at the core of many economic applications. There are no reason to think that one agent, and a fortiori a regulating institution, would have perfect information about the underlying economy. It seems more practical that the institution would learn agents' characteristics. In this model, the institution is able to learn from an aggregate signal received continuously from the market. The main mechanics is as follows. Given its initial prior belief, the institution sets a control variable (which can be viewed as a macro-variable, or even a signal regarding future policy changes) that influences the payoff of economic agents and therefore their optimal actions (which may be called their *optimal response*). Agents are assumed to be myopic, that is they do not try to anticipate future policy decisions and they just react to the current value of the macro variable. Myopic agents is a restrictive assumption but the goal was to provide theoretical results regarding the implication of hidden information from the regulator point of view, and we expect future work to relax this assumption⁸. Then the agents' responses get aggregated and the institution use Bayes' rule to update its belief and set a new value for the control variable. This is done continuously.

An important result is that the very structure of this model generates an endogenous level of uncertainty from the agents' point of view. The uncertainty solely comes from the learning process of the institution that has no other choice than use try-and-experiment to achieve optimal policy. There are some empirical evidence of policy uncertainty that were proposed by Baker, Bloom and Davis (2016), and oth-

⁷In this model, all agents share the same type. Future work will extend this framework to different private types.

⁸Note that, non-myopic agents would try to learn the internal state of the regulating institution, and that would produce some kind of two-sided learning. See Cisternas (2017) for a similar model.

ers, based on their work. They have built a *policy uncertainty index* out of economic newspapers⁹. They have used this index to explain variations in asset prices or other economic aggregates, taking the index as an exogenous variable from the point of view of the market. The central motivation of this thesis chapter is to explain where does this uncertainty come from, and show that it is in fact endogenous and due to the fundamental asymmetry of information between the financial institution and the market.

The model allows the study of overconfidence of the institution. As in any learning model, the initial belief plays an important role in the equilibrium trajectories¹⁰. When the value of the institution's policy forces the agents to reveal a lot of information, overconfidence makes the institution's belief to converge smoothly toward its long term equilibrium. The institution will learn quickly the value of the agents' type, and the perceived uncertainty will decrease over time. However, when the institution's policy makes it hard for the institution to learn, overconfidence creates inefficiencies and the uncertainty remains high for a longer period.

The running example in this chapter is central bank communication. This is completely similar to any other kind of regulation but it sure represents a better way to think about a continuous signal being sent to the market. Indeed, the central bank usually communicates constantly through the declaration of its representatives and it usually has to decide whether or not to change the main interest rate (the macro variable) on a monthly basis. From the banks point of view, this creates a continuous game in which they react to changes in interest rate (or anticipation of future changes) and the central bank constantly updates its belief regarding the fundamentals of the economy.

This learning game could be analyzed with the traditional *Multi-Armed Bandit* framework. This literature opposes the *exploration* to the *exploitation* when an agent has to learn a value through its actions. In our context, exploration refers to changing the macro variable to try to learn how agents will react to new values, and exploitation refers to keeping the same value for the macro variable once the institution's payoff is good enough. Intuitively, when the cost of exploration is high, the agent that is learning will prefer to stick with the current policy even if this means not learning much about the fundamentals, and when the exploration cost is low it will prefer to learn the best strategy in the first place.

Our results also have implications on the design of regulating institutions. In the

⁹They developed an algorithm that read the newspapers and count the number of occurrence of terms like "economic uncertainty", "uncertain legislation", etc. And they built an index out of this.

¹⁰Note that the initial belief only impacts trajectories because we work with unbounded continuous variables and as long as the "true" value of the unknown parameter is included in the support of the belief distribution, Bayes' rule ensures that the belief will eventually converge to the true value.

model, we are able to impose restrictions on what the institution can do, otherwise stated we can change the cost of exploration. We show that when the institution is heavily constrained (the cost of exploration is high), its learning ability is harmed. The perceived uncertainty is low but the economy will converge *slowly* toward its long term equilibrium. Interestingly, however, when the institution is not constrained at all (low exploration cost), the cost of exploration is so low that the uncertainty perceived by the market is too high. We postulate that there exists an optimal level of constraint that should be applied to the range of institutional actions. It should be emphasized that in order to obtain such an optimal level of constraint, we would need to optimize some criteria, and it is not obvious to us which to choose. It could be the overall perceived uncertainty, but that would assume that uncertainty is bad for the overall welfare of economic agents which might not necessarily be the case.

From a more general point of view on financial institutions, our results shed light on the intrinsic limits of interventionism. The asymmetry of information makes regulation process inherently inefficient in the short run. An institution in charge of regulating a market, will most likely have to learn the market's characteristics and therefore the regulation will not be optimal for the duration of the learning process. It should be noted that, in line with the famous critic of Lucas (1976), the economic environment is constantly changing. Therefore the question of whether the institution will eventually learn the market's characteristic can be asked. Intuitively, if the underlying parameter changes faster than the time it takes the institution to learn it, the institution will keep chasing it without ever reaching optimal regulation. From a technical point of view, the Bayes' rule does not allow the learning of a changing parameter, except if the institution knows when the parameter changed and reset its initial belief, which is unrealistic. The question whether regulation with hidden information could ever be optimal is therefore left open.

1.3 Blockchain : An Autonomous Institution

Let's recall our remarks so far. Historically, the role of financial institutions have gone from *i*) the definition and enforcement of *property rights*, to *ii*) the promotion of international trade, and more recently *iii*) the correction of inefficiencies resulting from individual interactions. Then, we pointed out several limits of active intervention from institutions. First, regulating institutions seem to react strongly and quickly to economic perturbations (like the global financial crisis), which introduces the question of regulation optimality. Second, we emphasized a more fundamental limitation that is the learning process of financial institutions will endogenously produce some level of uncertainty perceived by economic agents.

One of the transaction costs in the economy, according to the *property rights* approach, is the agency problem. First defined by Jensen and Meckling (1976), agency costs refer to issues arising in a principal-agent relationship. As presented in Laffont and Tirole (1986), or in the second chapter of this dissertation, the relationship between an institution (an organization) and the economy can be thought of as an agency situation. This, however, critically depends on the definition of an organization. Even when the objective of a financial institution is clearly defined, as for a central bank for instance, the ultimate decision of the institution is taken by an individual or a group of individuals. As soon as individuals are introduced in the decision process, this introduces some risks like the risk of mistakes, or the risk of corruption for instance. Ultimately, we define the distortion introduced by individuals as *trust*, in the sense that market participants need to *trust* the institution to properly fulfil its goal¹¹. In that sense, trust is embedded in any agency relationship. If everything was perfectly measurable, contractable and enforceable there would not be agency issues, neither trust would be required and transaction costs would be a lot lower.

From a practical point of view, modern forms of *money* (as an institution) entirely rely on trust in financial institutions. Since the end of the Bretton-Woods system, the US dollar is not backed by any physical goods. The intrinsic value solely depends on the belief that future currency holders will keep trusting its value and so on. Another example is that when two agents are contracting, the settlement usually happens through a *trusted third party*, like a CCP, as explained in the first chapter of this dissertation. This is fundamental to any exchange even in the *physical world*. If Alice goes to a shop and buy a cup of coffee from Bob then, either Alice hands over the money to Bob first, or Bob hands over the coffee to Alice first. Regardless of who goes first, the other one trusts that its counterparty will honor his part of the contract. This can be viewed as an institution and this is what North (1991) refers to as informal restraints that are customs, traditions and code of conduct. However, trust is deeply embedded in each one of them.

Several attempts have been made to create a *trustless* institution that would allow economic agents to trade without relying on trust, but all were unsuccessful until Nakamoto (2008) introduced the *blockchain*. Among those attempts we can cite *DigiCash* in 1989, *BitGold* in 1998 or more importantly *Paypal* in 1998. The common issue is that at one point they needed to rely on a *trusted party* somewhere in the system. In other words, they were not fully *trustless*. It has to be realized that the complete absence of trust requires a definition of value that is independent

¹¹In order to explain participation in the stock market Guiso, Sapienza and Zingales (2008) define trust as “the subjective probability individuals attribute to the possibility of being cheated”. The difference is that our definition is not subjective to a particular agent but rather represents an objective assessment of the cheating (or mistake) probability.

of trust, of property rights over this value, and ways to enforce and transfer those property rights, without relying of any particular agent or organization of agents.

Nakamoto (2008) introduced a communication protocol allowing a group of untrusted agents to agree on the content of a distributed database. The term “blockchain” was later coined to refer to the resulting distributed database. In its original form it contains the list of every transactions in the blockchain’s currency, grouped into blocks and chained together. A blockchain needs to have this currency component (called *cryptocurrency*), because it is used as an incentive mechanism for the agents that are maintaining this blockchain (called *miners*) in the form of newly issued coins. Miners create blocks by picking up some transactions from the *pending pool of transactions* up to a *size limit*¹², then they must find a binary value (called a *NONCE*) so that the *hash*¹³ of the selected transactions *and* the NONCE is below a deterministic *target* value. Note that, because hash functions are completely unpredictable, there are no other way to find a correct hash than by brute-forcing it. Brute force refers to randomly computing hashes with different NONCE values until a correct one is found. Once a miner has found the correct hash, the block is effectively “created”, and it is broadcasted to all other miners who then include this hash in the header of the *next* block, and the whole process starts over. There are important considerations here, in particular no one *has to* do anything on a blockchain, so miners could decide to include the hash of a previous block in the current block’s header for instance, but as will be explained miners do not have sufficient incentives to do so.

A blockchain has some very important properties. First, the average *inter-block time* is constant. Since brute force is the only technique possible to find a correct block, the more hashes computed per second, the more likely it is to find a correct hash in a given amount of time. However, the target of the block’s hash is dynamically adjusted so that the average inter-block time is constant. Intuitively, if the inter-block time average over the last few blocks is lower (higher) than the targeted time, the difficulty of the hash-based puzzle increases (decreases) by lowering (increasing) the target value for a correct hash. Second, the circulating supply of coins is deterministic over time. This is a consequence of the previous property. Indeed, when a block is created, a deterministic amount of coins is created and attributed to the miner who solved the block. The issuing schedule can be of different sorts (leading to a maximum supply or not) but is always fully deterministic and

¹²For instance, on the Bitcoin blockchain, the size limit of each block is currently 1 Mb.

¹³A hash is the result of a non reversible algorithm that maps a variable-length input into a fixed-length output so that *i*) the output “looks” random, *ii*) a small change in the input completely changes the output in an unpredictable way, and *iii*) it is impossible to predict the output value without actually running the algorithm.

depends on the block number¹⁴. Third, the miners are incentivized to mine the correct version of the blockchain. Miners need to incur some cost in order to compute hashes (mainly electricity) and are rewarded with new coins when they find a correct hash¹⁵. It can be shown that this provides sufficient incentives for miners to mine the *correct* version of the blockchain. Biais et al. (2019) show that it is a Markov-Perfect-Equilibrium for miners to mine the longest chain¹⁶. Intuitively, if a miner tries to *cheat* by any means, his block, if he finds one, will be rejected by other miner so the effective reward would be zero. The miner would have incurred electricity cost and get no reward. This provides sufficient incentives for all miners not to deviate from the equilibrium behavior.

We define an *autonomous institution* as a *trustless* institution that self-adjusts to achieve a deterministic goal. Let's now show that the blockchain (in the sense of Nakamoto (2008)) is an *autonomous institution*. First of all, everything is deterministic, from the inter-block time to the supply of coins. Second, the system does not rely on any particular agent, no one is more important than anyone else. More precisely, no computing ship is more important than any other, Nakamoto (2008) refers to this as *one-CPU-one-vote*. One agent can be more important in that he has more computing ships than others, but as long as one does not have more than a significant portion of the total computing power, this is not an issue¹⁷. Also, as the blockchain grows and the total computing power of miners increases, it becomes more and more expensive and therefore more unlikely¹⁸. Moreover, it would require a deeper analysis to show that a well functioning blockchain is resistant to small

¹⁴For instance, on the Bitcoin blockchain, the first block created 50 new bitcoins, and the number of new bitcoins created is halved every 210,000 blocks, which with a 10 minutes inter-block time, correspond to roughly 4 years. This decreasing supply will produce a maximum of 21,000,000 coins.

¹⁵In addition to the newly issued coins, they also collect transaction fees associated with the transactions they picked up to form the block. An agent willing to make a transaction *can* give some fees to miners to incentivize them to pick its transaction up.

¹⁶They also show that there are multiple equilibria with persistent forks for instance, but they would arise when there are disagreements on the consensus rules. This would effectively create 2 blockchains, each of them with all the properties of a blockchain.

¹⁷Blockchains are potentially subjected to 51% attack, that is a miner having more than 51% of the computing power could grow his own version of the blockchain faster than the rest of the network. And because of the longest chain rule, his version would become the "truth".

¹⁸For instance, as of May 2019, the Bitcoin total hash rate is estimated to be around 50,000,000 TH/s (1 TH/s = 10^{12} hashes per second). Miners use Application Specific Integrated Circuits (ASIC) to compute hashes. At the same period, one of the most efficient ASIC miner is the Antminer S17 manufactured by Bitmain, each costing USD 1,600. The hashrate of this machine is around 53 TH/sec. With those figures it would require 943,396 Antminer S17 for a total cost of around USD 1.5 billion to have the same power of the current network. This is just the cost of buying the machines, then they would have to be supplied with electricity. Note that there could be some collusion among existing miners, but the biggest miners around that time are in fact mining pools, consisting of shared computational power among several miners and no mining pool accounts for more than 20% of the total computing power. We argue that collusion seems unlikely in such an environment.

perturbations or even attacks from outside users, but as example regarding the Bitcoin blockchain we can cite : the contentious hard-fork of BitcoinCash in August 2017, or perturbations in the hashrate (miners suddenly turning off their machines to lower the hashrate and increase the effective inter-block time) that were common in the early days. None had a long term impact on the Bitcoin blockchain.

Economic agents can use a blockchain in different ways. First, agents can transfer value from one another. If Alice wants to send Bob one unit of cryptocurrency, she will sign a transaction containing an history of unspent transactions, just enough to prove that she indeed has all the property rights over this unit of cryptocurrency, and broadcast this transaction to the network. For this transaction to be effective, it will have to be picked up by a miner, and included in a block that is part of the blockchain. There are two main *transaction costs* associated with this kind of transaction. First, if there are more pending transactions that there is available free space in a block, Alice can incentivized miners to pick her transaction by paying them a fee. If they have the choice, miners will pick the transactions that give them the highest fees. Note that, because the constrained resource is the size of each block, this fee does not depend on the amount being transacted but rather on the actual size of the transaction, in *bytes*. Second, as noted in Biais et al. (2019), there can sometimes be different competing version of the blockchain, this situation is called a *fork*. If Alice really wants to make sure here transaction is included in the “correct” version of the chain she needs to wait for a few blocks to be mined on top of the block that includes her transaction. The probability of having another version of the chain decreases exponentially with the number of child blocks. For instance, on the Bitcoin blockchain, it is common practice to wait for 5/6 confirmation blocks that averages around one hour. This delay is an important part of the cost of transacting.

The second way agents can use a blockchain is to implement what is called *smart-contracts*. A smart-contract is a particular form of transaction that is equivalent to transferring the property rights of some amount of cryptocurrency to a computer program. This program will then have complete ownership and will transfer this value to someone else when a pre-determined set of conditions is satisfied. Those conditions can be time conditions or agreement conditions (like requiring 2 out of 3 agents to agree on whom this value should be transferred to). Those contracts are fully enforced by the blockchain itself, without the ability of anyone to stop a contract from executing as it was designed. We will come back to the importance of this in the last part of this introduction.

There are two remarks that should be discussed. First, the consensus algorithm we described is called *proof-of-work*. The idea is that the miners can be trusted because they have proven their work when broadcasting a valid hash to the network. The rest of the network knows that the only way they would have found this hash

is by incurring a high cost. There are other consensus algorithms, with the most important one being called *proof-of-stake*. Saleh (2018) provides an economic model of this algorithm showing that consensus would arise in equilibrium. However, this algorithm has not been implemented on any large scale blockchain as of May 2019. Second, there is actually an important externality to the proof-of-work chain that is electricity consumption. Indeed, as explained, miners need to incur a lot of electricity which is actually a link to the “real world”, as it could damage the environment for instance.

Once again, it would require a deeper analysis of the situation, but we argue that even this externality will be corrected through time in an autonomous fashion. One of the intuition behind this affirmation is that the main cost of miners is electricity. If a miner has access to a cheaper electricity price, he would earn more (be more profitable). Since the geographic position of miners is not important regarding the functioning of the blockchain, they will be able to move their computing ships where the electricity is the cheapest. With the simple argument of supply and demand, electricity is going to be cheap in region of the world where the supply exceeds the demand. The best example of such a situation regarding the electricity market is where energy is wasted like solar or hydraulic energy in low population parts of the world. This cheap energy is not used by any other industries as it will have to be transported to useful areas but with an internet connection miners can be located anywhere.

This thesis’ third chapter presents a higher order protocol, called the *Lightning Network*, that uses the blockchain’s consensus protocol as base layer. Because of the cost of using a blockchain (the fees and the delay associated with confirmation time), the Lightning Network was proposed by Poon and Dryja (2016) to overcome those issues. The Lightning Network is a decentralized payment network built on top of a blockchain in which intermediary nodes provide a trustless routing service for end users. As such, it is often called a *second-layer* protocol, with the blockchain being the *first-layer* protocol. Using specific types of contract written on the blockchain, nodes open *payment channels* among each other effectively creating a network of payment channels that can be used to route transactions between any two nodes, as long as there exists a route between them. On the Lightning Network, transactions are instantly confirmed and because there is no need for mining, the cost of sending transactions on the Lightning Network will supposedly be very low compare to the blockchain. We show that the strategic interactions between nodes can be described by a Bertrand competition model with capacity constraints. Our results indicate that there is a unique equilibrium in which a centralized network is never optimal. Moreover, we show that when nodes are heterogenous in their opportunity cost of capital only, the equilibrium network structure can match the current state of the

network.

The Lightning Network, in combination with a blockchain, can therefore be thought of as a new form of financial institution, that is *autonomous*, and that can be used to transfer value. This institution does not, however, suffer from the traditional agency problems. Users of a blockchain can trade without the risk of their counterparty defaulting and can always be sure of the security of their funds (as long as the blockchain is well functioning). Traditional interbank settlement usually takes at least one business day when the banks are in the same “banking area” but can take several business days when different jurisdictions are at stake. The promise of blockchains, and furthermore the Lightning Network, is to drastically change value transfer around the world. Note that this is not just a better speed efficiency or (direct) cost reduction, it is a completely new paradigm of transferring value based on the absence of trust. The next section discusses implications of this regarding the future of financial institutions or even institutions in general.

1.4 The Future of Institutions

Let’s recall what we have outlined so far and derive some implications for the future of institutions. We have explained how institutions are part of the economic landscape since the earliest days of humanity. However, the role of national institutions has changed a lot since then. In the beginning, they were in charge of properly defining and enforcing property rights among economic agents. Then, as international trade became more and more important around the 16th century, states endorsed the role of facilitating exchange while protecting people inside the state from people outside the state. Indeed, they realized the nation’s wealth was directly impacted by imports and exports, so they provided a legal and operational framework to increase nation’s wealth. And more recently, during the 20th century, states endorsed a new role of correcting inefficiencies and preventing externalities.

We also provided some limitations of the current role of financial institutions. First, regulation seems to always be behind the current state of the economy. The question raised was if such regulation was ever to be optimal. However, we noted that regulating institutions have become so important that economic agents rely nowadays on those institutions to operate, they expect the regulator to provide them with solutions when the state of the economy is bad. This put regulating institutions in a position in which they need to react quickly to changes in the economy. Second, we explained how such a regulatory framework can raise issues in the presence of hidden information. If the regulating institution needs to learn the internal state of the economy by making decisions, this will endogenously create some level of uncertainty which might not be optimal for economic agents. Overall

this creates a complex learning game for which it is not obvious whether it is better or worse for economic agents.

In this context, blockchains might be helpful. As explained, the fundamental value of a blockchain is the consensus that can arise among untrusted parties. This is a classical problem in the distributed system literature called the *Byzantine General Problem* described by Lamport, Shostak and Pease (1982). It was shown by Fischer, Lynch and Paterson (1982) that it is impossible to reach consensus asynchronously among untrusted process. However, a blockchain, as described by Nakamoto (2008), is an asymptotic solution to this problem¹⁹. Therefore, a blockchain could be used to properly define property rights of individuals without any trusted central party, or organization. As mentioned above, the earliest societies that we know of have been using institutions to define property rights, however with blockchains, and for the first time ever, it is possible to obtain such a proper definition without traditional organizations. From this point of view, the blockchain holds the truth (the consensus), and being completely digital, it has the power to enforce any contract that comply with this truth.

The consensus has the potential to be the basis for a lot of applications. In a way very similar to the Open Systems Interconnection (OSI) model, the blockchain ecosystem could develop in very different aspects of economic activities. The OSI model is the basis of modern computer communications, and it is composed of several layers, each responsible for a specific task. For instance, the well-known *IP* protocol is part of the *network layer* responsible for establishing connections between different sub-network with very few requirements. The *TCP* protocol which is part of the *transport layer* introduces error-checking and data integrity in IP-connections. This example is meant to illustrate that the same layering of roles could be used for economic interactions. As previously explained, if we think of the blockchain as being the consensus layer, the Lightning Network could be the layer used for quick and almost-free transactions. Then, other layers could be introduced in order to managed other types of ownership (of other digital assets for instance), or some other layers could be used on specific markets like international trade or by geographic areas. With the correct set of smart contracts it is possible to establish such a stack for value exchange. The Lightning Network is one of the best example of this layering technique, but a technology called *Sidechain* is also promising²⁰. Then any

¹⁹It is asymptotic in the sense that the probability of reaching consensus increases a lot with time although consensus on very recent events might not be guarantee. As blocks get buried deeper and deeper in the blockchain the probability of this block not to be part of the consensus decreases exponentially.

²⁰Sidechain are “parallel” blockchains that have a one-to-one peg to the main blockchain. The one-to-one peg ensures that any funds on the sidechain corresponds to funds on the main blockchain, in a way very similar to the Lightning Network (using smart contracts). The sidechain uses the consensus of the underlying blockchain so no mining of any sort is required and inter-block time

dispute of ownership in any layer would be automatically resolved by “asking” the underlying layer, until the consensus-blockchain is reached. Because the blockchain owns the truth, it will allocate the value to the corresponding party, very much like a court would do, except no human intervention would be required, and most importantly this will be done with complete absence of trust.

The trustless property of a blockchain is fundamental. Recall that by trustless we mean that the objective probability of being cheated is close to zero. In other words, the power of defining and enforcing property rights and rules is not delegated to a set of individuals but rather to a deterministic piece of software. Therefore economic agents do not perceive any cheating probability as the software will execute according to its design²¹. When economic agents lose trust in international institutions, the consequences can be quite high. For instance Algan, Guriev, Papaioannou and Passari (2017) show that economic distress can produce a decline in trust which can influence the political votes of citizens.

To come back to the Coase theorem (1960), blockchains have important consequences. Coase states that when there are no transaction costs, the initial allocation of property rights does not matter and bargaining among agents will yield a pareto-optimal allocation. However, recall that with the property rights approach, transaction costs should be understood as the cost of defining and enforcing property rights. As mentioned, blockchains are currently costly to use but property rights (over digital value – money) are very well defined and enforceable. It is likely that some cost will remain, like the social cost of energy needed to create consensus, or the actual transfer fees currently needed to incentivized miners, but this means that transaction costs in the coasian sense are very low compared to other institutional designs. Therefore, the Coasian interpretation of the economic environment is likely to become more relevant than ever.

This opens some interesting questions and makes us rethink the role of institutions in general. Indeed, a blockchain does not rely on any nation, state, or frontier so what are going to be the consequences of not needing any nation to define and enforce property rights? Blockchains are still very recent, and it is not clear yet whether this will introduce a real paradigm shift, but if this is the case, nations and institutions around the world will likely have to adapt to this new paradigm. As stated, from a historical perspective, there have been some important changes regarding the role of institutions but blockchains have the potential to introduce an even more dramatic change.

can be quite low for instance.

²¹In a way, trust is moved toward the software developers, as there could be bugs or mistakes in the source code. However, a proper blockchain should be open-sourced (the actual code is available for anyone to read and check) in order to minimize the presence of errors.

Chapter 2

Variation Margins, Liquidity Management and Clearing Structure

Abstract :

In derivatives markets, variation margins require protection sellers to keep a liquidity buffer. This paper builds a unified framework for analyzing protection sellers' cash management incentives. Central clearing requires a larger liquidity buffer than bilateral clearing. However, the expected loss is lower with central clearing so there is a tradeoff when assessing the optimal clearing structure. We find that overall risk and cash used by markets with low volatility or with a few clearing members are very sensitive to the clearing structure, while the effect is marginal for large or high volatility markets.

2.1 Introduction

With a notional outstanding of around \$550 trillion by mid-2016, the Over-The-Counter (OTC) derivatives market, often seen as opaque, represents a large source of risk for the global financial stability¹. Because of its opaqueness it has long remained out of the scope of major world financial regulators. Being held partly responsible for the global financial crisis in 2008, OTC derivatives markets have since attracted a lot of, still growing, interest from both financial regulators and academic researchers. Solutions have been proposed to build a financial environment in which a financial crisis as the last one would not be as damageable. In particular, clearing, either central or bilateral, has largely been viewed as a mean of mitigating risks in such markets.

This chapter focuses on the liquidity aspect of the clearing environment and the use of margins. There are many ways of tackling the inherent counterparty risk in OTC derivatives markets. While lots of effort have been made regarding the general incentives misalignment between a protection seller and a protection buyer, for some reason liquidity concerns remain under-represented in the literature. We propose a framework in which agents must pay uncertain variation margins, and we study how this impacts the global counterparty risk exposure of the market. Variation margins payment creates a liquidity shock for protection sellers and may trigger their default.

When the market is centrally cleared through a Central Clearing Platform (CCP), several layers are responsible for dealing with a particular kind of risks. It is important to distinguish between all those layers and the associated risks which, taken together, constitute the default waterfall. Variation margins are paid at least every day by the losing party, and are meant to cover the variation in the price of the underlying asset. Initial margins are paid once at the inception of any trade by each party. In order not to create additional counterparty risk, they are segregated on the balance sheet of the trade's counterparty. Their computation is based on the intrinsic characteristics of the underlying asset such as variance and correlation with other assets². Initial margins are supposed to cover the loss between two variation margins payments, sometimes referred to as the *Margin Period of Risk*. The default fund is the only resource that is pooled between all market participants and serves as a mutualization mechanism regarding clearing members' default risk. Each participant should contribute to the default fund with respect to their share of the

¹Figures are from the Triennial Central Bank Survey conducted by the Bank of International Settlement (BIS) at the end of June 2016, which captures about 94% of outstanding OTC derivatives positions.

²Initial margins can actually be modified during the life of the derivative product. However, since their computation is based on characteristics such as variance, it would be the case if there was a change in the underlying variance for instance. This issue might be more complicated for complex derivative products, as it could be related to information about the value of the product.

total risk undertaken by the CCP. Last but not least is the CCP equity capital itself, which is often regarded as required skin in the game for the CCP³.

In a bilaterally cleared derivatives market, participants are, or will be, required to use both initial and variation margins⁴. No mutualization structure could be implemented because the market is essentially decentralized, but the initial and variation margins concerns are similar from the point of view of market participants⁵.

The goal of this chapter is to shed light on the liquidity issue arising from the use of day-to-day market valuation of exposure, or variation margins. While liquidity can be important for initial margins and the default fund as well, variation margins are the most relevant because they have to be paid more frequently. If a financial institution were to be illiquid at the time variation margins are required, it will default on the margins payment which constitutes a regular default event under most jurisdictions, and will most likely be followed by a liquidation of the defaulter's open positions.

Our model features homogenous risk-neutral protection sellers, whose only purpose is to maximize their individual profit. After selling derivatives contracts, each protection seller needs to manage its balance sheet in a profitable way. Protection sellers do not care, per se, about remaining solvent in the event of a margin call, they simply do not like those states of the world as they will default and end up with nothing. They do not take into account any bankruptcy costs. Protection sellers have to pay some margins before the maturity of the derivative exposure, which corresponds to variation margins payment.

The key friction in our analysis is a maturity mismatch between some of the protection seller's assets and the derivative exposure itself. Each protection seller faces 2 investment technologies, one that we call *liquid* which is available at any time, and another one that we call *illiquid* which might not be available at the time margins are required. To reflect a liquidity premium, the return on the illiquid assets is higher than that on the liquid one. The illiquid assets can be liquidated before their maturity but in such a case protection sellers only get a fraction of their investment which we call the *collateral value of illiquid assets*.

In our model protection sellers must allocate their initial wealth between these

³CCP's Skin in The Game is actually only part of the CCP's capital. For instance EMIR requires CCP to have at least 25% of operational capital as Skin in The Game. See, e.g., Huang (2016).

⁴The International Organization of Securities Commissions recommend initial and variation margins to be used in bilaterally cleared markets, see CPMI-IOSCO (2012). Most regulations around the world will implement those recommendations in a gradual way, until the end of the 2010's. See EMIR for instance.

⁵There are some differences in practice because the CCP actually acts as a reference point. Indeed, margins calculation on a bank-to-bank basis requires they use similar valuation techniques. For the most complex products, there could be some conflicts but this is out of the scope of this chapter.

two opportunities. They essentially face the tradeoff of a higher solvency probability going along with a higher opportunity cost of not using the illiquid technology as much as possible. They essentially constitute a liquidity buffer to increase their probability of remaining solvent. One of our main results is that this tradeoff which shows up at the individual level is also present at the market level, and might be influenced by the clearing structure.

The main contribution of this chapter is to analyze the protection seller's decision in several settings. We build a general framework which allows us to study the market participants' behavior in a bilaterally and a centrally cleared markets. Our base case scenario is a no margin clearing structure. This trivial setup features no maturity mismatch, therefore protection sellers fully enjoy the illiquid technology. The two frameworks of interest are a bilaterally cleared market with variation margins, and a market centrally cleared through a CCP with variation margins and a default fund.

An important result of this chapter is to characterize the optimal decision of the protection seller. An optimal asset allocation is shown to always exist and to be unique. The protection seller is never tempted to short sell the illiquid asset, but under some circumstances he would be willing to borrow from the liquid asset and invest more than his total wealth in the illiquid asset. This happens when the initial wealth is very low compared to the derivative exposure for instance. In such a case, this can be interpreted as a gambling behavior. If the protection seller is most likely to default even if he invests all his wealth in the liquid asset, he will gamble on the possibility of the exposure to be close to 0, in which case he will actually profit from the illiquid investment.

Regarding the centrally cleared market, this chapter shows that each protection seller invests less in the illiquid technology but also starts with less resources. Indeed, from the protection seller's point of view, the allocation decision is exactly the same as in the bilateral clearing, except he has less initial wealth due the default fund contribution. Our model uses an exogenous default fund rule. To be consistent with industry practices, this rule states that the default fund should be sufficient to cover the default of the two largest members⁶. Because this default fund rule will dictate how much of initial wealth a protection seller must pay upfront to the CCP, it will endogenously determine the quantity of exposure he is willing to undertake, that is the quantity of protection he is willing to sell. However, we show that the quantity of protection sold will not affect the allocation decision as long as the different derivative products are independent.

Our main result is the comparison of all the different setups under a unified

⁶CPMI-IOSCO (2012, 2014, 2016) suggest that a CCP should maintain a default fund satisfying the Cover II principle. In addition stress tests can also be used to refine the exposure estimation.

analysis. To compare all clearing environments, we use the expected loss which is defined as the expectation of losses from the protection buyers' point of view when their counterparty, either a CCP or a protection seller, defaults. This kind of measure is highly relevant when dealing with general policy concerns because it focuses on the loss incurred by the end-users. This is the actual measure used by most regulatory institutions to compute required equity capital for large financial intermediaries⁷. We balance the change in expected loss with the change in total cash used by the system.

First, we find that the number of market participants is an important determinant of the clearing structure's impact on both expected loss and overall cash used by the system. All else being equal, in derivative markets with a few dealers, central clearing decreases a lot the expected loss while requiring much more cash. The effect lies in the exogenous default fund rule which does not take into account the size of the market and always require the default fund to cover the exposure of two members. The individual default fund contribution is then much smaller when the number of market participant increases. Since the individual default fund contribution is essentially the only difference between bilateral and central clearing, as the number of clearing members increases the impact of central clearing becomes smaller.

Second, we show that the riskiness of the underlying derivative product is also an important determinant of the central clearing effect. When the variance of exposure is low, protection sellers have a good idea of what the realization of the variation margin call will be, so they maintain a liquidity buffer just enough to remain solvent and invest the remaining in the, more profitable, illiquid asset. Under central clearing, they must pay upfront the individual default fund contribution, so the wealth available for allocation is lower. However, because the variance of exposure is still low, they will maintain a similar liquidity buffer in absolute terms because it is easy for them to influence their survival probability. So the default fund, which is held in cash by the CCP, is composed of liquid assets that would have been invested in the illiquid technology otherwise. This results in much more cash used overall by the system, but also a great reduction of expected loss. A similar reasoning about markets with high volatility applies. Under high variance regimes, it is hard for protection sellers to influence their survival probabilities so they will pay the default fund contribution with assets that would have been allocated to liquid assets anyway. Their illiquid investment remains similar because it is more profitable for them to benefit from the illiquid investment return than trying to increase their survival probability. So when the variance is high, central clearing has barely any impact on both the expected loss and the global amount of cash needed by the market.

⁷The expected loss is used under Basel 2 to compute required equity capital. See BCBS (2005).

To the best of our knowledge this chapter is the first attempt to study central clearing from the point of view of protection sellers and variation margin payments. Indeed, Biais, Heider and Hoerova (2016) or Thompson (2010) for instance feature models with passive protection sellers, in the sense that they only consider incentive compatibility constraints and in optimum they will be incentivized to make effort. Because we introduce liquidity concerns, active protection sellers maximizing their profit are the prime focus, as their allocation decision is the main risk management tool, which allow us to estimate their demand in terms of liquid versus illiquid assets. While the literature on OTC clearing mechanisms is gaining more and more importance, the liquidity concerns intrinsic to margins remains overlooked. This chapter then contributes to several strands of literature.

First, this model is related to the literature on liquidity shocks. Protection sellers are subject to random liquidity shocks and endogenously maintain a liquidity buffer for precautionary motives (e.g. Alvarez and Lippi, 2009). From a technical standpoint, we use a normally distributed liquidity shock and provide existence and uniqueness theorems, as well as conditions and bounds on optimal solution of the asset allocation problem.

Second, this chapter contributes to the fast growing literature on Central Clearing. Biais, Heider and Hoerova (2016) investigate the risk taking incentives of protection sellers under moral hazard. A welfare-maximizing CCP must set up a contract so that the risk management incentive of protection sellers is preserved under unobservable effort. Thompson (2010) studies the signalling incentives of protection buyers when they have more information about the underlying derivative products. Stephens and Thompson (2013) show that, in an adverse selection framework, good quality insurers are pushed out of the market by bad ones. Pirrong (2009) provides a comprehensive analysis of central clearing under asymmetric information. Acharya and Bisin (2014) show that protection sellers have incentives to undertake too much risk when the composition of their existing derivatives positions is not known by the market. This creates a counterparty risk externality that is absent under central clearing which main function here is trade transparency.

We depart from these models by assuming no information asymmetries. This results in protection sellers having all the decision power regarding their risk mitigation behavior. However, we feature passive protection buyers who are just endowed with some risk they want to unwind because of a higher degree of risk aversion. We assume protection sellers are maximizing their profit and we study their risk mitigation incentives in this context.

Some papers have also studied the implications of having a profit-driven CCP. Murphy (2016) show that such a CCP will reduce equity capital as low as possible but that will drive clearing volume down. Huang (2016) theoretically investigates

the risk management incentives of a for-profit CCP and shows that limited liability prevents the social optimum to be reached. As in this chapter, she shows that higher collateral mitigates potential losses at the expense of forgone investment opportunities.

Unlike those papers, our model features a passive CCP which only purpose is to collect mutualized resources, that is a default fund. We show that the risk reduction versus opportunity costs tradeoff is also present even with a passive CCP, and we account for the collateral buffer that protection sellers have to maintain in order to remain solvent after variation margins calls. The focus of this chapter is to investigate the difference between bilaterally and centrally cleared markets, and the default fund maintained by the CCP is the main difference.

Some empirical work has been done on the effect of central clearing. For instance according to our implications, Loon and Zhong (2013) document a reduction in counterparty risk on voluntary cleared CDS contracts. Duffie and Zhu (2011) show that central clearing may increase overall counterparty risk exposure. Cont and Kokholm (2014) and Anderson et al. (2013) quantify multilateral netting benefits for several market characteristics.

Third, this chapter also contributes to the literature on the collateral demand. Heller and Vause (2012) use simulated data to show that while dealers already have enough collateral to meet initial margins, they might increase their cash holdings to meet recurrent variation margins payments. Sidanius and Zikes (2012) use aggregate data to quantitatively investigate how initial and variation margins needs will affect the system-wide demand for collateral. Duffie, Scheicher and Vuillemeys (2014) use fine-grained, actual data to refine the characteristics of global collateral demand, and in particular are able to investigate the liquidity buffer that market participants have to maintain for variation margin calls.

Our model provides theoretical motivations as well as quantitative estimations for the market-wide collateral demand. This chapter focuses on the part of collateral demand that comes from protection sellers maintaining liquidity buffer on their balance sheet, and we use that to show the differences between bilateral and central clearing. The appendix of this chapter provides complete comparative statics on the size of the cash buffer regarding variation margin calls.

The remaining of the chapter is organized as follows. Section 2 introduces the general formulation of the model as well as all the different frameworks. Section 3 compares the different setups and provides a discussion on which one should be preferred. Section 4 provides some numerical estimations and implications. Section 5 concludes. The appendix presents proofs, as well as numerical estimations and interesting comparative statics on the optimal asset allocation of a protection seller.

2.2 The Model

2.2.1 General Setup

There are 3 dates indexed by $t \in \{0, 1, 2\}$. There is a single, risk neutral, protection seller endowed with A_0 in cash. We study the decision process of this unique protection seller, which we may call the *agent*, and we will add many of them later when introducing a Central Counterparty.

There are two technologies able to transform cash over time, a *liquid* asset paying a return R_l each period, and an *illiquid* asset which yields \tilde{R}_i per unit of investment at date $t = 2$. We could assume that \tilde{R}_i follows a normal distribution with mean μ_R and variance σ_R^2 , but as the protection seller is risk neutral, we will abstract completely from the notion of idiosyncratic risk on assets and simplify it to $\tilde{R}_i = R_i$ always. The illiquid asset can be sold at a premium at the intermediary date. The liquidity premium on this asset is denoted by ℓ , with $0 < \ell < 1$, meaning that the protection seller will get ℓ per unit of assets sold at $t = 1$. The protection seller is only concerned about the liquid versus illiquid feature of assets. To account for the liquidity difference between both assets we assume investing in the illiquid technology gives a higher return over two periods.

$$R_i > R_l^2$$

We do not explicitly model the first contracting stage for the protection seller on the market for protection⁸. In reality, protection sellers may face protection buyers and contract with each other. Because protection sellers are risk neutral, and protection buyers would be risk averse, there exists an optimal contract in which a protection seller provides insurance to a protection buyer against some upfront payment by the protection buyer. We focus on the post-trade exposure management mechanisms and do not focus on the pre-trade risk-sharing behavior.

The protection seller then starts off with one unit of an exogenous exposure on the derivative market, called \tilde{e} , associated with some cash π , corresponding to the upfront payment he gets from a hypothetical protection buyer. By convention if $\tilde{e} > 0$, this means that the protection seller would have to pay to the other party. We assume that the derivative exposure follows a normal distribution with mean μ_e and variance σ_e^2 .

$$\tilde{e} \sim \mathcal{N}(\mu_e ; \sigma_e^2)$$

To simplify the analysis of margins payments we will do further⁹, we assume the

⁸We shall do that in further version but this is not what's going to be important in this story.

⁹Adding a real signal at the intermediary date is likely to complexify the analysis with no obvious additional insights.

derivative exposure is realized at $t = 2$ but the value $\tilde{e} |_{t=2}$ is effectively known by everyone at $t = 1$. All happens as if there were a perfect signal at the intermediary date and the exposure does not change between date $t = 1$ and $t = 2$. In the remainder of the chapter we will refer indifferently to the realization of the derivative exposure at $t = 1$ or $t = 2$, otherwise stated $\tilde{e} |_{t=1} = \tilde{e} |_{t=2}$.

In addition to the cash a protection seller receives from the market for protection, he may have initial resources, denoted by \overline{W} . Let's define by A_0 the total available resources of the protection seller at the first date.

$$A_0 \equiv \overline{W} + \pi$$

The agent must allocate A_0 between the liquid and illiquid technologies. This decision arises at the first date $t = 0$. Let's call θ the dollar amount invested in the illiquid technology so that $(A_0 - \theta)$ is invested in the liquid one. The protection seller will do so by maximizing the mean of its final wealth, which will depend on the setup. We avoid the possibility of short selling any of the assets so that the illiquid investment is constrained to $\theta \in (0 ; A_0)$.

2.2.2 The No-Margin Case

As a benchmark, we analyse the case of bilateral exposure without margins payments. Nothing is occurring at intermediary date $t = 1$ and the protection seller must simply pay $\tilde{e} |_{t=2}$ at the last date. There is no liquidity issue here.

The final wealth of the protection seller, in this case called A_2^{NM} , is a function of the quantity invested in the illiquid asset θ and is equal to the return on investment minus the derivative exposure.

$$A_2^{NM}(\theta) \equiv (A_0 - \theta)R_l^2 + \theta R_i - \tilde{e}$$

The protection seller chooses its illiquid investment θ so that it maximizes the mean of its final wealth. The optimal quantity invested in the illiquid asset is denoted by θ_{NM}^* .

$$\theta_{NM}^* = \arg \max_{\theta} \mathbb{E} [A_2^{NM}(\theta)] = (A_0 - \theta)R_l^2 + \theta R_i - \mu_e$$

The agent solves this program and we have the following proposition.

Proposition 2.1 *In the bilateral exposure environment without variation margins payments, the protection seller always puts all its initial wealth in the illiquid asset.*

$$\theta_{NM}^* = A_0$$

The optimal amount invested in the illiquid asset does not depend on the underlying derivative exposure. This makes sense because in the no-margin case, there is no mismatch between the date at which both assets pay off and the maturity of the derivative exposure \tilde{e} . No liquidity buffer is maintained on the protection seller balance sheet.

2.2.3 Bilateral Contracting with Variation Margins

In order to simulate variation margins payments, all happens as if the derivative exposure was realized at the intermediary date, that is before the illiquid asset pays off. The protection seller is required to pay the realization of \tilde{e} at $t = 1$.

The available resources at the time the derivative exposure is realized, denoted by A_1 , are composed of the liquid investment, available each period, and any cash the protection seller can obtain by selling his illiquid assets. This quantity is a function of the quantity invested in the illiquid assets θ .

$$A_1(\theta) \equiv (A_0 - \theta)R_l + \theta\ell$$

When the derivative exposure is realized, there are 3 possible cases. Either *i*) the liquidity shock is too severe, the protection seller must file for bankruptcy with his counterparty undertaking a loss, or *ii*) the liquid assets are sufficient to face this liquidity shock, or *iii*) he needs to sell some of his illiquid position.

The protection seller will be required to sell some of his illiquid position if the derivative exposure \tilde{e} is higher than his liquid position $(A_0 - \theta)R_l$ and lower than $A_1(\theta)$. In such a case he uses all of his cash, and sell a quantity X of illiquid assets such that it covers exactly the derivative exposure.

$$X = \frac{\tilde{e} |_{t=1} - (A_0 - \theta)R_l}{\ell}$$

In this case, the agent's last period wealth, denoted by $A_2^S(\theta)$, will be the return R_i he gets on whatever quantity of illiquid assets he still has.

$$A_2^S(\theta) = (\theta - X)R_i$$

The agent defaults when $A_1(\theta) < \tilde{e} |_{t=1}$. We do not assume any bankruptcy costs, so the protection seller ends up with zero wealth if he files for bankruptcy. When he does not default, he invests his remaining wealth in the liquid technology. His final wealth, if he does not default and does not need to sell some illiquid assets,

is denoted by $A_2(\theta)$.

$$A_2(\theta) \equiv ((A_0 - \theta)R_l - \tilde{e}) R_l + \theta R_i$$

The protection seller chooses its illiquid investment θ so that it maximizes the mean of its final wealth. The optimal quantity invested in the illiquid assets in this setup is denoted by θ_m^* .

$$\begin{aligned} \theta_m^* = \arg \max_{\theta} \int_{-\infty}^{(A_0 - \theta)R_l} & \left[((A_0 - \theta)R_l - e) R_l + \theta R_i \right] \phi(e) de \\ & + \int_{(A_0 - \theta)R_l}^{A_1(\theta)} \left(\theta - \frac{e - (A_0 - \theta)R_l}{\ell} \right) R_i \phi(e) de \end{aligned}$$

with $\phi(e)$ being the density function of a normal random variable with mean μ_e and variance σ_e^2 .

The agent solves this program and we have the following proposition.

Proposition 2.2 *In the bilateral exposure environment with variation margins payments, there always exists a unique optimal quantity invested in the illiquid asset $\theta_m^* > 0$. The budget constraint is not binding as long as*

$$\frac{\Phi(0)}{\mathcal{H}(\ell)} < \frac{R_i(R_l - \ell)}{R_i - R_l^2}$$

with $\Phi(x)$ being the cumulative distribution function of a normal random variable with mean μ_e and variance σ_e^2 , and $\mathcal{H}(\ell) = \frac{\Phi(A_0\ell) - \Phi(0)}{\ell}$ and $\lim_{\ell \rightarrow 0} \mathcal{H}(\ell) = A_0\Phi'(0)$.

There is no closed-form solution for the optimal quantity invested in the illiquid assets θ_m^* . In the appendix we present a discussion as well as numerical estimations of the solution and show that it behaves as expected.

We do not need to impose a short selling constraint on the illiquid asset because in equilibrium the quantity invested in this asset θ_m^* is positive. Having a negative θ , compared to $\theta = 0$, would increase the agent's survival probability, but in the survival states the agent would have to pay the return on illiquid asset R_i and only receive R_l^2 . So he would actually get less because of the assumption that $R_i > R_l^2$. However, under some values for the parameters the agent can be tempted to borrow against the liquid asset and invest the proceeding in the illiquid one. This will be the case in particular when the difference between the return on both assets is large, and/or when the liquidity value of the illiquid asset is close to 1.

2.2.4 Central Clearing through a CCP

This section introduces some results on the asset allocation and liquidity concern when all bilateral trades are moved to a Central Counterparty. The previous model is slightly adjusted to take into account the CCP. We consider several protection sellers, all identical to the protection seller of the previous section. However, the *quantity* of derivative exposure they have is now endogenous.

The CCP requires some collateral¹⁰ from any protection seller at the inception of any trade. This cash is going to be pooled with all other protection sellers' and is meant to provide an insurance mechanism, in the form of mutualization, for the default of one or several clearing members. To be consistent with current industry practices, we shall refer to these pooled resources as the *default fund*. Each agent contributes to the default fund regarding his contribution to the sum of risk undertaken by the CCP. The return on the default fund paid by the CCP is assumed to be 0. In practice, there could be some return but this is usually very low since the default fund is contributed in high-liquid and low-risk assets. In addition to the default fund the CCP is now in charge of collecting variation margins at the intermediary date.

There are m protection sellers and each is endowed with 0 cash at the first date. We slightly change the notation, in order for the equilibrium quantity to be consistent with those of the previous section. The agent's problem is now a bit more complex because he has to solve for the asset allocation, as before, and also for the quantity of exposure he wants to take. Indeed, the derivative exposure is now endogenous, and the quantity of cash the protection seller needs to post to the CCP will imply a tradeoff for the quantity of exposure the agent is willing to take.

To simplify, we assume that there is a fixed unit of exposure and the agents can take any quantity of this risk. We denote by $n \in \mathbb{N}$ the integer number of derivative contracts taken by an agent, and each unit has the same characteristics as the exposure in the previous section, $\tilde{e} \sim \mathcal{N}(\mu_e; \sigma_e^2)$. Each unit of derivative risk is *i.i.d.*. When undertaking a derivative contract, in addition to the risk, the agent also receives some cash, denoted by π . This cash corresponds to a premium that could be paid by a hypothetical protection buyer to give his risk away. The remaining of the model is similar to before, so overall each agent has now two control variables : θ , the quantity of cash invested in the illiquid technology, and n , the quantity of exposure he is willing to take.

Since the agent is risk neutral, this dual decision can be taken in two steps with two separate programs. Indeed, if we think backward, the allocation decision only depends on the cash available at the beginning of the period. So we can solve for

¹⁰This is usually required to be cash or High Quality and Liquid Assets (HQLA).

the allocation decision, θ , in terms of the available cash, and then plug that in the first decision which is going to be the amount of exposure the agent is willing to take, which will depend on the default fund rule of the CCP.

The default fund rule is exogenous. Consistent with industry practices, the default fund is meant to cover the default of the two largest members¹¹. Since all protection sellers are similar, the default fund will be set to cover the default of any two members. Since in this setup all agents are identical, their decisions will be exactly symmetric, in particular they will all choose the same number of derivative contracts. The default fund is meant to cover, on average, the default of any two members which represent $2n$ contracts, with mean exposure μ_e . The total size of the default fund will be $D = 2n\mu_e$. The m clearing members will have the same individual contribution d to the total default fund D .

$$d(n) = \frac{D}{m} = \frac{2n}{m}\mu_e$$

First the agent needs to solve for the optimal asset allocation decision θ given an amount of initial cash A_0 . To account for a variable risk exposure, we slightly adjust the model of the previous section. The agent undertakes n derivative contracts. We denote by $\tilde{f}(n)$ the random variable corresponding to having n independent derivative exposure contracts \tilde{e} .

$$\tilde{f}(n) \sim \mathcal{N}(n\mu_e; n\sigma_e^2)$$

We denote by $\Pi(A, \theta, n)$ the expected terminal wealth of the protection seller if he chooses an asset allocation θ with A initial wealth and n derivative contracts.

$$\begin{aligned} \Pi(A, \theta, n) = \int_{-\infty}^{(A-\theta)R_l} & \left[((A-\theta)R_l - f) R_l + \theta R_i \right] \phi_{\tilde{f}(n)}(f) df \\ & + \int_{(A-\theta)R_l}^{A_1(A, \theta)} \left(\theta - \frac{f - (A-\theta)R_l}{\ell} \right) R_i \phi_{\tilde{f}(n)}(f) df \end{aligned}$$

with $\phi_{\tilde{f}(n)}(f)$ being the density function of a normal random variable with mean $n\mu_e$ and variance $n\sigma_e^2$, and $A_1(A, \theta)$ the available cash at the intermediary date, which depends on the quantity invested in the illiquid asset θ as well as the initial wealth.

$$A_1(A, \theta) = (A - \theta)R_l$$

The optimal investment in the illiquid assets in this model can then be expressed as a function of n , $\theta_{ccp}^*(n)$. The objective function is given by the following expres-

¹¹Huang (2016) and Armakolla and Laurent (2017) use the same exogenous rule while studying CCP's risk management incentives.

sion.

$$\theta_{ccp}^*(n) = \arg \max_{\theta} \Pi(A, \theta, n)$$

By proposition 2.2, there always exists a unique asset allocation, $\theta_{ccp}^*(n)$, that maximizes the protection seller expected utility. For the remaining of the analysis we will need an additional result which is summarized in the following lemma.

Lemma 2.1 *When a protection seller is maximizing his terminal wealth, if his initial wealth, A , is a multiple of the number of independent derivative contracts he has previously undertaken, n , then his investment in the illiquid asset in proportion of his initial wealth does not depend on the number of contract.*

$$\frac{\partial}{\partial n} \left(\frac{\theta^*(n)}{A(n)} \right) = 0$$

The initial cash available for allocation A is equal to whatever the agent gets by selling insurance contracts minus the individual default fund contribution.

$$A_0(n) = n\pi - d(n) = n\pi - \frac{2n}{m}\mu_e$$

Protection sellers can now solve for the equilibrium number of contracts n^* given their expected profit $\Pi(A, \theta_{ccp}^*(n), n)$.

$$n^* = \arg \max_n \Pi(A_0(n), \theta_{ccp}^*(n), n)$$

The cash available right before the allocation decision $n\pi - 2n\mu_e/m$ is linear in n and equal to zero when $n = 0$. By lemma 2.1, the optimal allocation decision $\theta_{ccp}^*(n)$ in proportion of the initial wealth is then independent of n . This allows us to specify the unit-profit Π_u , which corresponds to the expected wealth for selling one insurance contract. Because of the independence of the allocation decision θ^* in proportion of the wealth with respect to n , we have the resources available for 1 contract, denoted by A_u , equals the total available resources divided by the number of contracts.

$$A_u = \frac{1}{n}A_0(n) = \pi - \frac{2\mu_e}{m}$$

We can rewrite the program with the per unit expected profit.

$$n^* = \arg \max_n n \Pi(A_u, \theta_{ccp}^*(1), 1) = \arg \max_n n \Pi_u$$

Proposition 2.3 *In the central clearing environment, when the individual contribution to the default fund is linear in the number of contracts undertaken, and if protection sellers start with no initial exogenous wealth, they invest a constant proportion of their wealth θ_{ccp}^*/A in the illiquid asset.*

$$\frac{\partial}{\partial n} \left(\frac{\theta^*(n)_{ccp}}{A(n)} \right) = 0$$

Regarding the optimal number of contracts undertaken n^* , if we require it to be positive and finite, $n^* \in [0, n_{max}]$, and if we let Π_u be the expected profit per contract, we have that

- If $\Pi_u > 0$ then

$$n^* = n_{max}$$

- If $\Pi_u < 0$ then

$$n^* = 0$$

The cost of undertaking derivative contracts is essentially the default fund contribution. Because the default fund contribution $d(n)$ is linear in n , there is no interior solution to this program. We require the optimal number of contracts n^* to be positive and finite. Because we did not choose to model the first risk-sharing contracting stage between protection buyers and protection sellers this might not be clear, but both assumptions can easily be justified. Recall that in a richer model, protection sellers will not simply choose the optimal number of contracts but this would come from a contracting stage with protection buyers, before the story in this model takes place. Since a protection buyer would have a higher risk-aversion than protection sellers, the case $n^* < 0$ is not feasible. Regarding the other limit n_{max} , this makes sense if we assume that there is a finite number of protection buyers, with finite insurance needs. Then all protection sellers would be in competition on this insurance market, and because they are all identical, if there are no other frictions, they will evenly share the total number of available insurance contracts. And even in the case they are not strictly identical, they would still get a finite number of contracts in equilibrium.

2.3 Comparison of the Different Cases

Having solved for the protection seller's decision in all setups, we now want to compare them. On the one hand, note that even if the solution of the protection seller's program is the same for a fixed amount of initial cash, since the CCP requires

a default fund contribution, the amount invested in the liquid asset is not the same. On the other hand, with additional resources available to insure the defaulter's counterparty both the probability of default and the size of a default are changed. When assessing the implications and consequences of a default, we do not take into account some negative externalities that could arise from early liquidation of illiquid assets. Given that this effect should affect the bilaterally and centrally cleared setups the same way, it should not be needed for comparison purposes.

This section provides a comprehensive analysis of the cost and benefit of a CCP in terms of cash usage and default characteristics.

2.3.1 Global Cash Usage

Variation margins payments generate a liquidity buffer that would not have been used otherwise. In addition to those resources, the default fund needs also to be accounted for when a CCP is operating on the market. We already computed the size of the default fund, which will depend on the number of derivative contracts undertaken by the agents, so we now need to evaluate the size of the liquidity buffer on the protection seller's balance sheet in different scenarios.

The difference in the quantity invested in the illiquid asset θ resides only in the difference in the initial cash available which is lower with a CCP due to the default fund contribution. With a CCP, some cash corresponding to the default fund contribution is taken away from the allocation decision of the protection seller. If the proportion invested in the illiquid asset were to be constant regarding initial cash, the amount of cash needed in this economy would always be the same. The following lemma helps understanding how the proportion invested in the illiquid asset changes.

Lemma 2.2 *The allocation decision made by a protection seller depends on the initial assets available, A , and we have*

$$0 < \frac{\partial}{\partial A} \theta^*(A) < 1$$

We can also show that the derivative of the proportion invested in the illiquid asset θ is increasing and convex, everything else being equal. See Appendix for a discussion on the derivatives of the optimal allocation.

Intuitively, if the initial assets are very high compared to the derivative exposure, any additional unit of initial asset will be almost entirely allocated to the illiquid asset. The marginal gain in survival probability associated with the next unit invested in the liquid asset is so low that the agent is better off with the higher return on

the illiquid asset. As the proportion of initial assets with respect to the derivative exposure diminishes the margin gain in survival probability increases, so it becomes more and more tempting for the agent to forego the higher return by investing the next unit of cash in the liquid assets. When the initial assets are very low compared to the derivative exposure, the agent chooses to invest almost entirely the next unit of cash in the liquid assets because of the gain realized in survival probability.

We can now assess the cash usage in this economy. The following proposition summarizes the results of this section.

Proposition 2.4 *In a bilateral derivative market, adding a CCP which maintains a default fund increases the global amount of cash needed by the economy.*

It may seem a little surprising since adding a CCP reinforce the resilience of the market as well as providing some mutualization of counterparty risk through the default fund. However, this does not matter for the protection seller. Indeed from his point of view, he still faces a derivative exposure and needs to allocate his wealth in a similar way, the problem has not changed.

By focusing on the protection seller, there might be some effects that are not captured by this model. First, one can argue that mutualization and increased insurance should reduce the need for cash, but here the notion of mutualized insurance is a bit different from the usual one. Indeed, the protection sellers are required to contribute to a default fund at the CCP level but they are not the ones benefiting from this increased insurance. The other side of the initial transactions, the protection buyers are benefiting from this mutualization but if the presence of a default fund is not internalized by the protection buyers and transferred to the protection sellers through the original transaction's price, there is no reason why protection sellers should take that into account and reduce their liquid investment.

Second, in a richer model with initial margins in particular, increasing the default fund might decrease the initial margin requirements. This is, however, out of the scope of this chapter and I would argue that even with initial margins, the effect on variation margin and the liquidity buffer would not change¹².

The way to understand this framework is that if you take away some unconditional cash from the protection sellers' allocation decision, with a constant derivative exposure, the global amount of cash used will increase due to the tradeoff faced by the agent between an increase probability of survival versus the opportunity cost of cash.

¹²See for instance Murphy (2017).

2.3.2 Magnitude and Probability of a Default

We use the *expected loss* to compare the magnitude and probability of default in the different setups. This is especially relevant as this measure is actually used by the regulator, in Basel 2 for instance, to assess the financial resilience of institutions and compute their regulatory capital requirements. The idea behind such measures is that we do not really care about the upsides of the outcome, but are concerned about the downsides. When assessing general policy concerns and regulatory environment, we want to think about the social outcome, and since profits and losses are not shared by the same agents¹³, it makes sense to differentiate the measure as such. We denote the expected loss by EL .

$$EL = \mathbb{E}[Loss] = \mathbb{E} \left[LGD \times \mathbb{1}_{\{default\}} \right]$$

with LGD being the *Loss Given Default*, which is the expected loss in the default area, and $\mathbb{1}_{default}$ being an indicator function that is one when a default occurs.

The *expected loss* is the main quantity that will allow us to derive some real world implications. This should be a sufficient criteria when comparing several frameworks, and policy makers and regulators are heavily relying on measures like expected loss. Even if the welfare effect is not clear, there should be a cost associated with more expected loss. However, the comparison does not take into account any externalities arising from the default of a large institution like a CCP¹⁴. It is likely that the resolution of a defaulting CCP would impact markets in a broad sense, even only through the confidence channel. The analysis will stick to hard measures that we can compute and compare with one another.

We need to define the *default* event. In the bilateral clearing environment, there is no kind of default insurance, the protection seller will default at the intermediary date as long as he does not have enough cash to meet the margins payment. If we denote the default probability for this framework by Δ^m we have

$$\Delta^m \equiv Pr(default) = \mathbb{E} \left[\mathbb{1}_{\{default\}} \right] = Pr(\tilde{e} > A_1(A, \theta_m^*))$$

with $A_1(A, \theta_m^*) = (A - \theta_m^*)R_l + \theta_m^*\ell$. In order to compare it with the central clearing setup, we need to set a number of contracts equal to n , and the per-contract cash received by protection sellers equal to π , as we did in the central clearing case. This allows us to compare two setups in which the protection seller undertakes the same number of contracts as in the central clearing case (where the optimal number of contracts comes from the exogenous default fund rule). If we let $\Phi_{\tilde{e}}()$ be the

¹³This happens mainly because of limited liability and hence moral hazard.

¹⁴A defaulting CCP is likely to have a number of implications even only through its systemically importance. See, e.g., Amini, Filipovic and Minca (2016) and Menkveld (2016).

cumulative distribution function of the random normal variable \tilde{e} , we can rewrite this probability as follows.

$$\Delta^m = 1 - \Phi_{\tilde{e}}(A_1(n\pi, \theta_m^*))$$

A default will occur as soon as a protection seller does not have sufficient funds to pay for the variation margins call at the intermediary date. In such a case they will liquidate everything, give everything to the protection buyer and default. The later then suffers a loss equivalent to the difference between the available fund of the protection seller and the original exposure. In order to properly compare setups, we need to account for m protection sellers (which is needed for the central clearing case). Using the same notation as before, we can write the expected loss in the bilateral clearing setup, EL^m .

$$EL^m = \int_{A_1(n\pi, \theta_m^*)}^{+\infty} m(e - A_1(n\pi, \theta_m^*))\phi(e) de$$

In the central clearing setup, we need to refine the notion of *default*. We are no longer interested in the default of the protection sellers per se, instead we care about the default of the CCP.

The whole purpose of the CCP is to provide insurance against the default of one or several clearing members. If we care about losses incurred by external agents, such as those who bought insurance from protection sellers, it makes sense to consider the default of the CCP. In the previous setup, as soon as a protection seller defaults, losses were incurred by its direct counterparty, which was a protection buyer. Here, if a protection seller defaults, its direct counterparty is the CCP. The CCP has a default fund whose purpose is exactly to cover those losses. So now the buyers of protection will incur losses as soon as the default fund is not sufficient to cover the losses incurred by all protection sellers, otherwise stated when the CCP defaults.

One way to view $A_1(n^*\pi - d, \theta_{ccp}^*)$, which is the available cash of the protection seller at the intermediary date, is as initial margins¹⁵. Recall that in this setup the agent starts with A in cash (obtained from selling insurance contracts) minus its individual default fund contribution $d = 2n^*\mu_e/m$. Initial margins are supposed to be segregated and not used as pooled resources. Then if a protection seller incurs losses, the CCP will draw on its initial margins first, and the remaining, if any, will be drawn from the default fund. To compute how much will be drawn from the default fund, we need to create a new random variable, \tilde{X} , that represents the

¹⁵What could actually be considered as initial margin is only the investment in the liquid assets $(n^*\pi - \theta_{ccp}^*)R_l$. The difference between the two is the value of the illiquid assets in case of early liquidation. So A_1 could still be considered initial margins if the CCP's claim on variation margins has the highest seniority among the protection seller's liability claims.

losses in excess of the available cash at the intermediary date $A_1(n^*\pi - d, \theta_{ccp}^*)$.

$$\tilde{X} = \tilde{e} - A_1(n^*\pi - d, \theta_{ccp}^*)$$

Because the derivative exposure \tilde{e} follows a normal distribution, the excess losses \tilde{X} is also normally distributed.

$$\tilde{X} \sim \mathcal{N}(\mu_e - A_1(n^*\pi - d, \theta_{ccp}^*) ; \sigma_e^2)$$

Since there are m clearing members, the CCP will face losses, denoted by \tilde{Y} , corresponding to the sum of \tilde{X} across m protection sellers. All agents are homogeneous, so the total losses experienced by the CCP is simply m times the excess loss of a protection seller.

$$\tilde{Y} = \sum_{i=1}^m \tilde{X}$$

And because the risk is normally distributed and *i.i.d.* across agents, the total loss experienced by the CCP also follows a normal distribution.

$$\tilde{Y} \sim \mathcal{N}(m(\mu_e - A_1(n^*\pi - d, \theta_{ccp}^*)); m \sigma_e^2)$$

The CCP will default on its obligation as soon as the sum of excess losses from all protection sellers exceeds the size of its default fund. In equilibrium the default fund D^* is a function of the optimal number of derivative contracts undertaken by each protection seller n^* .

$$D^* = 2n^*\mu_e$$

The default probability for this framework, denoted by Δ^{ccp} , can be expressed as follows.

$$\Delta^{ccp} \equiv Pr(\tilde{Y} > D^*) = Pr(\tilde{Y} > 2n^*\mu_e)$$

If we let $\Phi_{\tilde{Y}}()$ be the cumulative distribution function of the random normal variable \tilde{Y} , we can rewrite this probability as follows.

$$\Delta^{ccp} = 1 - \Phi_{\tilde{Y}}(2n^*\mu_e)$$

In the centrally cleared setup, the expected loss is defined as the size of the default of the CCP itself. This is because losses are evaluated from the point of view of the protection buyer and any default from just one protection seller will not affect any protection buyer. Recall that we defined \tilde{Y} to be the total losses experienced by the CCP. The total available resources of the CCP is the default fund contributed by protection sellers, D , so a default will occur as soon as $\tilde{Y} > D$.

In such a case, protection buyers would get the entire default fund and suffer losses equal to the difference between the default fund and the realized exposure. Using the same notation, we can write the expected loss in the centrally cleared setup, EL^{ccp} .

$$EL^{ccp} = \int_D^{+\infty} (Y - D) \phi_{\tilde{Y}}(Y) dY$$

In order to compare the expected loss in those different setups, we have to express them both regarding the same random variable. We are going to make a change of variables so that the expected loss in the centrally cleared environment will be expressed in terms of \tilde{e} the original derivative exposure to which protection sellers are exposed to. The expected loss in the centrally cleared setup becomes the following.

$$EL^{ccp} = \int_D^{+\infty} (Y - D) \phi_Y(Y) dY = \int_Z^{+\infty} m(e - Z) \phi(e) de$$

where $Z = \mu_e + \frac{D - m(\mu_e - A_1(n^* \pi - d, \theta_{ccp}^*))}{m}$, $A_1(n^* \pi - d, \theta_{ccp}^*)$ is the available cash at the intermediary date in the case of central clearing and $\phi()$ the density of a normal random variable with mean μ_e and variance σ_e^2 .

Expressions of expected loss in both setups have very similar functional forms. We can now compare them by comparing the value at the lower bound of the integral, which defines at which point losses start to be positive. Let's recall both expression for expected loss.

$$EL^m = \int_{A_1^m}^{+\infty} m(e - A_1^m) \phi(e) de \quad EL^{ccp} = \int_Z^{+\infty} m(e - Z) \phi(e) de$$

We first notice that the default probability (which can be viewed as the probability of hitting the lower bound in each integral) and the expected loss are intrinsically linked in this framework. Indeed, whenever the probability of default is higher in one setup, the expected loss associated with the same setup is also higher relative to the other one. This is due to the linearity (with the same coefficient) of the loss function inside the integral. The following proposition summarizes the comparison of different setups in terms of default.

Proposition 2.5 *In an OTC derivatives market, introducing a CCP reduces the likelihood of a default as well as the expected loss in the system if the following condition is met*

$$\Delta^m > \Delta^{ccp} \quad \Longleftrightarrow \quad EL^m > EL^{ccp} \quad \Longleftrightarrow \quad m > \frac{D(R_l - 1)}{(R_l - \ell)(\theta_m^* - \theta_{ccp}^*)}$$

where m is the number of participants in the market, D is the size of the default fund, θ_{ccp}^* and θ_m^* are respectively the illiquid asset allocation of the protection seller with and without central clearing.

The specific form of the condition in proposition 2.5 comes directly from the exogenous default fund rule that the CCP is implementing. This rule states that the default fund should be set such as to cover the default of the two largest members. Most regulations around the world have some kind of rules similar to this one. In our model, this is a little artificial, as all the members are identical to each other. And because we allow for an arbitrary number of protection sellers m , it could have been expected that the number of market participants plays a great role in the condition. As tempting as it is, this is not an expression for a threshold in the number of market participants. Indeed, the illiquid allocation when there is a CCP, θ_{ccp}^* , depends on the number of protection seller m .

2.3.3 Some Numerical Calculations

In the previous section, more specifically in proposition 2.5, we showed that introducing a CCP do not always lower the expected loss in a market. In this section we argue that even if in some situation a CCP can actually increase the expected loss, those situations can be qualified as extreme and highly unlikely under usual market conditions.

This section presents some numerical examples whose purpose is to show that under reasonable conditions, a CCP will always decrease the expected loss. The remaining of the section describes the methodology used and presents some arbitrary examples¹⁶.

Recall that we defined the expected loss to be

$$EL = \int_X^{+\infty} f(e)\phi(e) de$$

where $\phi()$ is the density of a normal variable representing the derivative exposure, X is the minimum value for the derivative exposure above which losses are realized, and $f(e)$ a *loss function* so that the expected loss is the mean of the losses. By definition we have $f(X) = 0$. We define the function being integrated as the loss density, and we denote it by $LD(e)$.

$$LD(e) = f(e)\phi(e)$$

¹⁶All python source code and detailed methodology will be made available in the online appendix shortly.

The loss density depends on the realization of the random variable e and gives the product of the loss that would be incurred if e was realized and the probability that e being drawn. This is a loss function taking into account the likelihood of the outcome.

The way to numerically solve the model is standard. A solution for the illiquid asset allocation θ is found and then everything else can be computed. By analyzing the shape of the loss density, we are able to graphically observe the expected loss. Each figure presented here have a caption displaying the parameters used for it, when this is not the base case example. The base case example is represented by the parametrization in table 2.1. This case is not meant to represent a calibration using real world figures but rather serves as an illustration of the mechanisms at stake.

Parameter's name	Variable	Default Value
Initial Wealth	A_0	100
Liquid Asset's Return	R_l	8%
Illiquid Asset's Return	R_i	30%
Mean of Exposure	μ_e	40
Variance of Exposure	σ_e^2	40
Collateral Value of Illiquid Assets	ℓ	0
Number of Market Participants	m	20

Table 2.1: Base Case Example

Figure 2.1 shows the loss density function for both the bilateral and central clearing cases. The hashed area represents the expected loss, that is the area under the curve when the loss density function is positive. When the loss density function is negative this means that a profit is realized. The point at which each curve crosses the x-axis may be used to relatively analyse default probability. Indeed, the moment at which the loss density starts to be positive give a sufficient criteria to rank default probabilities. Moreover, given the functional form of the loss density, the curves will never cross.

We see that in this case the expected loss is higher in the bilateral clearing case, and at the same time the likelihood of losses incurred is higher. If we try to see under which market conditions the reverse would be true, we may use the results of proposition 2.5. This proposition shows that in order to reverse the result we should try to minimize the distance between the allocation in illiquid assets in both setups $(\theta_m^* - \theta_{ccp}^*)$, and increase the collateral value of illiquid assets ℓ . To reduce the distance between both asset allocations we will increase the number of market participants. The effect will be that the individual default fund contribution will be much smaller and then the asset allocation in the central clearing setup will be closer to the one in the bilateral clearing setup.

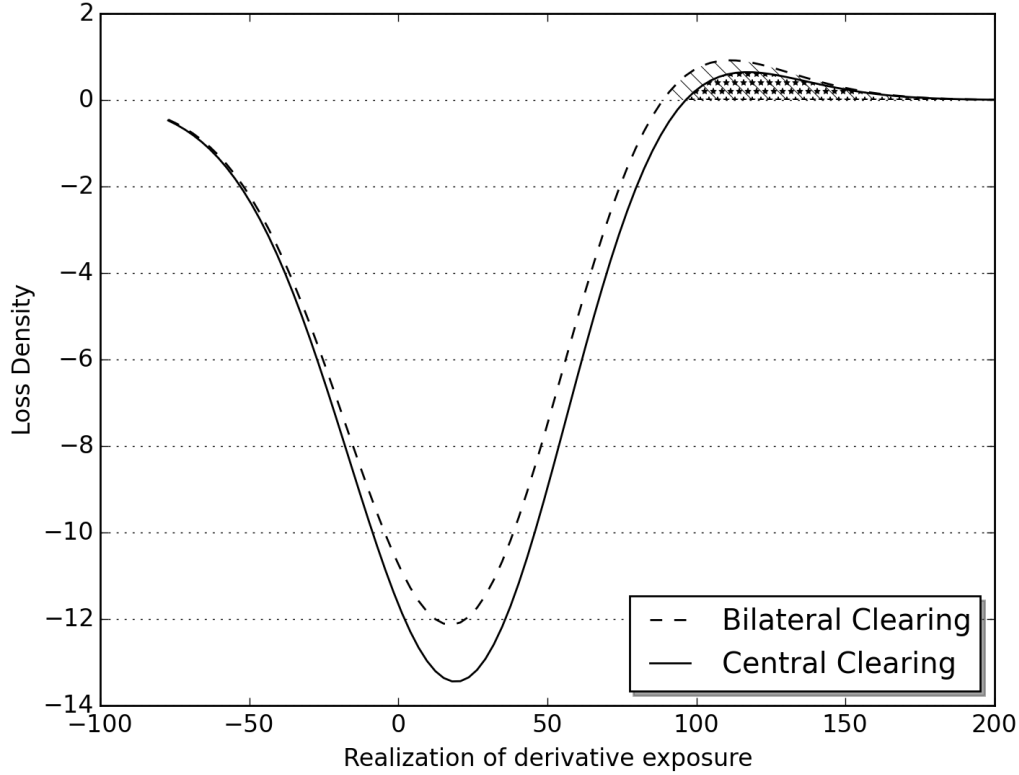


Figure 2.1: Loss Density for the Base Case Example
 $EL^m > EL^{ccp}$

The parameters values used next will be $\mu_e = 64$ and $\ell = 0.75$. They can be considered extreme for two reasons. First the mean of derivative exposure used is $\mu_e = 64$. Given that the protection seller starts with $A_0 = 100$, this represents 64% of her balance sheet exposed to derivative risk. In most markets, clearing members are large banks, and their derivative exposure is likely to be lower than this level.

Regarding the collateral value of illiquid assets set to $\ell = 0.75$, this also can be considered extreme. A recovery rate of 75% on the illiquid assets is likely to be unfeasible in time of market-wide stress. Indeed, a large clearing member defaulting on his position and selling his illiquid investment in the market will probably experience some illiquidity premium which can be especially high in times of crisis.

Figure 2.2 presents the loss density for both setups under the new parameter values. We can see that despite the extreme values, the expected loss in the bilateral clearing case is still a little higher. If we really want to make the expected loss in the centrally cleared market higher, we need to change the parameters value to even more extreme values.

For illustration purposes we propose a set of parameters under which the expected loss is actually higher with a CCP. The next parameters values are $\mu_e = 80$ and $\ell = 0.90$. Figure 2.3 presents both loss density under these new parameters, and




Figure 2.2: Loss Density with $\mu_e = 64$ and $\ell = 0.75$
 $EL^m > EL^{ccp}$

even with those extreme values for parameters, the difference in the two expected loss is still small.

Even though each parameter affects the value of the optimal illiquid asset investment θ , we were able to show some numerical examples in order to investigate when the condition in proposition 2.5 is satisfied. The analysis show that it requires extreme and unlikely parameters values in order for a CCP to increase the expected loss in the market. In the remaining we shall consider that the condition in proposition 2.5 is satisfied under regular values of market parameters.

2.4 Implications

So far the findings are that introducing central clearing in a derivatives market has two impacts. First it increases the overall amount of cash needed by the market, but it also reduces the expected loss to the end users. In this section, we will derive a number of implications resulting from the previous results.

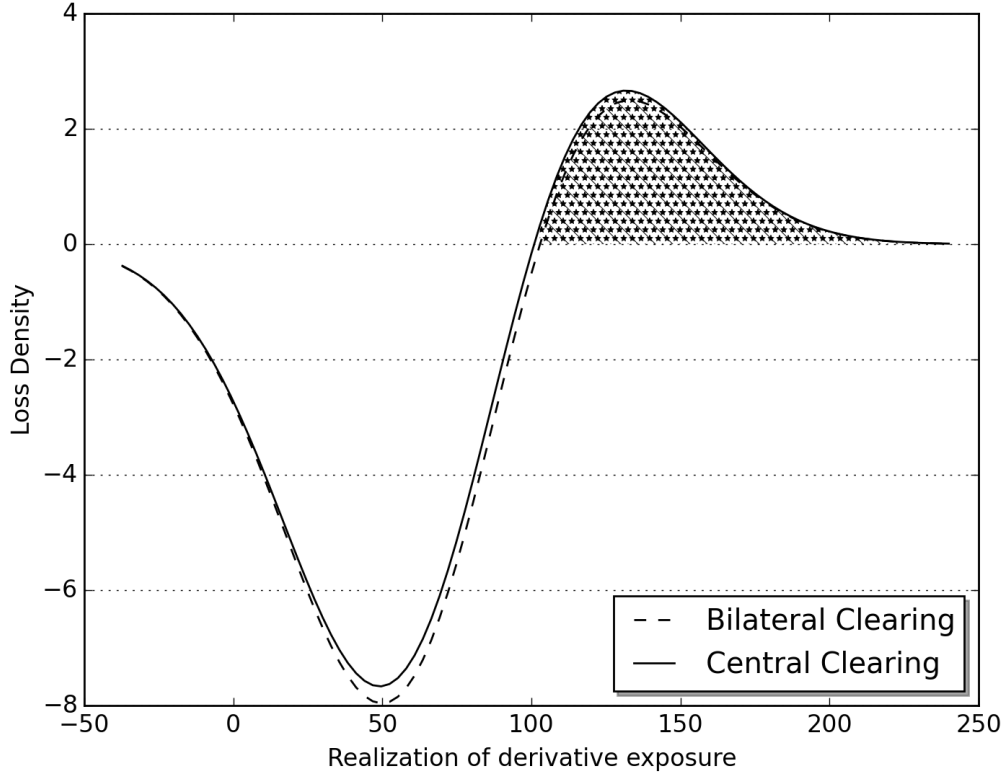


Figure 2.3: Loss Density with $\mu_e = 80$ and $\ell = 0.90$
 $EL^{ccp} > EL^m$

2.4.1 Size of the Market

In this model the size of the market plays an important role. This is due to the exogenous default fund rule that we set up, because this rule does not take the size of the market into account. The rule says that the default fund maintained by the CCP should be sufficient to cover losses made by the two largest members. Then on a market with a few participants, they will each have to contribute a lot to constitute the same default fund, and with more participants the individual default fund contribution will be smaller. A smaller default fund contribution makes the available cash before allocation closer to the one without a CCP, so the illiquid asset allocation with a CCP θ_{ccp}^* is closer the one without a CCP θ_m^* .

We analyse the costs and benefits of introducing central clearing compared to bilateral clearing depending on the size of the market. The way we conduct this analysis is by computing the percentage change in overall cash used in the market and the percentage change in the expected loss.

For the numerical computations we still use the base case example set of value which can be found in table 2.1. Figure 2.4 presents the gains in terms of reduction of expected loss and the costs in terms of quantity of cash used of introducing a

CCP in a derivatives market. This figure represents the cash usage difference (left scale) and expected loss difference (right scale) between the bilateral and the central clearing setup. The way to read it is the following. For a bilaterally cleared market with 25 market participants, introducing a CCP will increase the global cash usage by around 9.5% and on the other hand the expected loss will drop by 30%.




Figure 2.4: Cash Usage and Expected Loss Changes

As the number of market participants increases the marginal effects of central clearing on both the excess cash required and the expected loss diminishes. Again this is due to the specific form of the default fund rule. As the number of market participants increases, the individual default fund contribution becomes smaller and smaller which induces less and less difference between the two setups.

The real effect that is captured by the number of market participant is the size of the individual default fund contribution relative to the initial wealth of protection sellers. When the individual default fund contribution is large (i.e. when there are a few participants) the marginal effect is larger. For instance in Figure 2.4 we can see that with around 10 participants, the introduction of a CCP will use around 16% more cash, but the expected loss will be reduced by around 40%. Recall that the individual default fund contribution d is given by $d = 2n\mu_e/m$. This could correspond to markets with some very large players or markets in which notional

at risk represents a large proportion of protection sellers' balance sheet (μ_e large relative to initial wealth).

2.4.2 Riskiness of the Market

The tradeoff presented in the last section also shows up across another interesting parameters which is the riskiness of the derivatives market (i.e. the variance of derivative exposure). In this section, we analyze the variation in benefits and costs of adding a CCP in a market depending on the riskiness of the market.

As in the last section, figure 2.5 shows the change in cash usage and expected loss from introducing a CCP in a derivatives market.

When the variance of the derivative exposure is low, introducing a CCP will reduce the expected loss by up to 99.99% but the whole system will use up to 40% more cash. In low variance regimes, and when the mean of derivative exposure is low relative to initial wealth, the expected loss is low even without a CCP. What's happening is that the protection seller knows the realization of the exposure will be close to the mean, he can then keep just enough cash to pay for the exposure and invest the remaining in the illiquid asset. A CCP will be able to remove almost all the expected loss because if all protection sellers behave in a way similar to above, the aggregate risk will be quite low.

Moreover, under low variance regimes, the derivative of investment in the illiquid asset with respect to initial wealth is close to 1. Again this is because the protection seller as a good idea of what the realization of the risk will be and then will maintain a liquid investment just enough to cover this exposure and invest the remaining in the illiquid asset. So if he gets one unit less as initial wealth, he will still invest the same amount in the liquid investment and will invest less in the illiquid one. So overall the use of cash will be a lot higher.

In highly risky markets, the effect will be much lower on both the use of cash and the expected loss. For the exact opposite reason, when the variance is high it is very costly for protection sellers to influence their default probabilities by investing more in liquid assets. It is much more profitable to invest in the illiquid asset. It results a lower sensibility of illiquid investment to variation in initial wealth. The amount of cash used in the market will be very similar because with one unit of initial wealth taken away from the protection seller, his illiquid investment will remain the same. So we have taken away one unit of wealth to contribute to the default fund, but this unit would have been invested in the liquid investment by the protection seller himself anyway. This implies a similar use of cash.

The expected loss when the variance of derivative exposure is high does not change much. This is also due to the fact that the liquid investment is less sensitive




Figure 2.5: Cash Usage and Expected Loss Changes

to variation of initial wealth, that is the introduction of a CCP. So the risk is reduced but only by a small amount.

Figure 2.6 presents the absolute value of the cash usage and expected loss for both environments, with and without a CCP. This helps grasping the concepts highlighted above. The non-monotonicity of both cash usage functions is due to the behavior of protection sellers regarding risk in derivatives markets. When the variance is very low, they can maintain precisely the amount of liquid assets they need. As the variance goes up, they need to invest a little bit more in liquid assets because they gain from the increase in survival probability. At some point the variance is so high, that it becomes harder to influence the survival probability and the protection seller benefits more from investing in the illiquid asset. However, as seen in figure 2.5, the difference in cash usage is monotonic in this example.

Analyzing the effect of central clearing in markets with different amount of risk is a central issue. Under most regulations, the riskiest markets are not subject to mandatory central clearing. The reason for that is usually that central clearing requires a good price information because the CCP has to price the derivative contracts in order to call for margins. The analysis conducted here tends to show that for reasons related to protection sellers' risk management incentives the effect




Figure 2.6: Cash Usage and Expected Loss

of central clearing in very risky markets might be limited both in terms of benefits (reduction of expected loss) and costs (increase in cash usage). In our model, the variance of derivative exposure is actually the variance of the variation margins call, perceived as random from an ex-ante perspective. Variation margins risk can come from several sources, either the market is inherently risky or the market fundamentals do not change much but the computation of the variation margin is very complex. If the calculation is perceived as complex, it might result in higher uncertainty surrounding the variation margin call, which will have the exact same effect as a truly riskier market.

2.5 Conclusion

This chapter develops a unified framework for analyzing the protection sellers' cash management incentives under several clearing structures in derivatives markets. Risk neutral protection sellers and risk averse protection buyers contract with each other, then protection sellers allocate their wealth between liquid and illiquid investment. A liquidity shock representing variation margins call hit them at the intermediary date. First, we consider a bilateral clearing structure in which protection buyers

and sellers enter a derivative position and manage variation margins call at the intermediary date by themselves. Second, we consider a centrally cleared market in which a CCP is responsible for managing the variation margin call and also constitutes a default fund, hence providing some risk mutualization to the market. The main difference between the two is the creation of a default fund, which acts as pooled resources effectively insuring the protection buyers against the default of their original counterparty. Protection sellers have to contribute to the default fund unconditionally and before the allocation of their initial wealth.

We analyze both environments in terms of cash used by the whole market and expected loss for the end-users (i.e. protection buyers). We find that, most of the time, introducing central clearing in a derivatives market will reduce the expected loss while increasing the amount of cash needed. It is important to note that the increase in cash used does not come from any specifics of central clearing per se, but rather directly from protection sellers' incentives for cash management. A protection seller does not really take into account who his counterparty is (a protection buyer or a CCP), all he cares about is that there will be a liquidity shock and he wants to maximize the mean of his final wealth which implies remaining solvent at the intermediary date. Those incentives create a situation in which the protection seller has to trade off an increase in survival probability versus the opportunity cost resulting from more liquid investment.

We also find that the magnitude of the effect on cash used and expected loss depends on market characteristics. First, the size of the market modifies the impact of central clearing. The importance of the market size comes directly from the way the default fund is constructed. The CCP sets the default fund, in an exogenous way, so that it covers the default of the two largest clearing members. Because the default fund rule does not take into account the size of the market, protection sellers don't contribute the same amount whether there are a few or a lot of market participants. The asset allocation made by protection sellers depends on the available wealth (i.e. initial wealth minus the individual default fund contribution), so the size of the default fund contribution will be an important determinant for the analysis. When there are a few (a lot of) market participants, the individual default fund contribution is large (small), so the impact on both cash used and expected loss is also large (small).

Second, the riskiness of the market is also an important determinant of the impact of central clearing. We show that when the derivative exposure has a low variance, central clearing can remove almost entirely the expected loss but will also use much more cash. This happens because when the variance is low, protection sellers have a good idea of what the variation margins call will be, then the sensitivity of the illiquid investment to variation in initial wealth is close to 1. Central

clearing takes away some wealth from the protection seller available resources that would have been allocated to illiquid investment, protection sellers' liquid investment does not change much, which results in much more cash used in the system. This makes protection sellers, and the market as a whole, safer and less prone to default. Conversely, in markets with high volatility, for the same reasons, the effect of central clearing is limited. Indeed, in those markets, the individual default fund contribution takes away some wealth that would have been invested in the liquid asset anyway. This implies a smaller reduction of expected loss and a smaller amount of extra cash used by the market.

The effect of central clearing on the cash needed by the market and the expected loss varies a lot. However, the basic tradeoff between more cash used by the market and reduction in expected loss is always significant. The model does not feature protection buyers per se, so we cannot do welfare analysis. This is also a limit of the expected loss measure which represents the amount of losses in the market but does not allow us to derive absolute implications in terms of when it is better to have central or bilateral clearing. The tradeoff depends on the costs associated to more cash used, and also on the level of risk aversion in the market (i.e. the shape of the utility function) to give a utility value for the expected loss.

This chapter does not pretend to give absolute recommendations on the structure of clearing but rather highlights an important tradeoff which should be accounted for when designing the clearing structure in derivatives markets. Even if central clearing provides some counterparty risk diversification, the incentives of the protection buyers are not affected by this risk reduction, and they will hold more cash in proportion of their initial wealth, which results in more cash used by the whole system. Market designers and regulators should also keep in mind that the effect is not identical in all markets and while in some cases it might be welfare improving, in some other cases the effect might not be significant.

The liquidity aspect regarding the clearing structure has generally been overlooked and we hope that more advanced studies will be done on both the theory and empirics of liquidity provision associated with different market designs.

2.6 Appendix

2.6.1 Discussion and Comparative Statics on the Optimal Asset Allocation :

Whenever a protection seller experiences a mismatch between the maturity of his investment and the date at which he will have to pay the margins, he will forego some of the return on the longer maturity assets in order to remain solvent at the intermediary date. Recall that the optimal quantity invested in the liquid asset, θ^* comes from the following maximization program.

$$\theta^* = \arg \max_{\theta} \int_{-\infty}^{(A_0 - \theta)R_l} \left[((A_0 - \theta)R_l - e) R_l + \theta R_i \right] \phi(e) de \\ + \int_{(A_0 - \theta)R_l}^{A_1(\theta)} \left(\theta - \frac{e - (A_0 - \theta)R_l}{\ell} \right) R_i \phi(e) de$$

Proposition 2.2 shows that there always exists a unique θ^* but unfortunately we cannot obtain a closed-form solution. In the remaining of this section, we present some numerical estimations of the solution with respect to all parameters. Whenever the parameter is not the one under investigation, and otherwise specified, its default value is provided by table 2.2.

Parameter's name	Variable	Default Value
Initial Wealth	A_0	100
Liquid Asset's Return	R_l	5%
Illiquid Asset's Return	R_i	25%
Mean of Exposure	μ_e	40
Variance of Exposure	σ_e^2	30
Collateral Value of Illiquid Assets	ℓ	0.5

Table 2.2: Default Parameters

2.6.1.1 Comparative Statics w.r.t. Asset's return

While studying the comparative statics with respect to assets' return, and according to the assumptions made in the model, we need to require the following

$$1 < R_l^2 < R_i$$

Figure 2.7 presents the variation of the quantity invested in the illiquid asset θ as a function of the return on the illiquid asset. This is increasing as one would have expected. As the return on illiquid asset goes up, the opportunity of getting more return, even if it comes with a loss in the survival probability, also increases.




Figure 2.7: Optimal quantity invested in illiquid asset w.r.t. Illiquid asset's return

Regarding the return on the liquid asset the results are less obvious. Indeed, as the liquid asset's return goes up two effects have a role. The first one is similar to the analysis of the illiquid asset's return, namely that as an investment opportunity, the liquid asset becomes more attractive because of a higher return. Hence the quantity invested in the illiquid asset θ should go down.

The second effect is related to the default probability of the agent. Indeed, all else being equal, and in particular with a similar liquidity shock, as the liquid asset's return goes up, the agent needs to invest less in the liquid asset in order to end up with the same amount of available cash at the intermediary date.

Recall that in our model, there is no particular cost of defaulting for protection sellers. When the first effect dominates, if the liquid asset's return is low, they prefer to invest more in the illiquid asset, even if it comes with an increase of the default probability. The relative higher attractiveness of the illiquid asset overcomes the fact that they will default more often. If we were to introduce a convex cost of defaulting, the first effect should never dominate.




Figure 2.8: Optimal quantity invested in illiquid asset w.r.t. Liquid asset's return

2.6.1.2 Comparative Statics w.r.t. Initial Wealth

The variation in illiquid asset investment with respect to initial wealth is discussed in the chapter.

2.6.1.3 Comparative Statics w.r.t. Exposure Parameters

The exposure parameters are especially relevant because they can be used to discriminate between several kinds of derivative exposure. We choose to model the risk exposure as a normal random variable so there are two parameters of importance the mean μ_e and the variance σ_e^2 of the exposure.

Figure 2.9 presents the variation of the optimal illiquid asset investment with respect to the mean of exposure. As expected, an increase in the mean of the derivative exposure increases the probability of default, so the protection seller has more incentives to invest in the liquid asset, and the optimal quantity invested in the illiquid asset goes down.

Regarding the variance of the exposure the results are a little surprising. Figure 2.10 shows that there seems to be two separate effects on the optimal quantity invested in the illiquid asset. When the variance is low, the relation is negative until it reaches a minimum and then an increase in the variance also increases the optimal




Figure 2.9: Optimal quantity invested in illiquid asset w.r.t. Mean of exposure

illiquid investment. This should be related to some gambling behavior. When the variance is very close to zero, the protection seller knows almost for sure what will be the realization of the exposure. In this case he invests the amount required to remain solvent. Then as the variance goes up a little, he needs to invest more in the liquid asset to make sure he remains solvent. At some point, the variance is so large, the protection seller starts to gamble by investing more in the illiquid assets. Figure 2.11 allows us to understand a little more the relation between the variance of the exposure and the optimal illiquid investment as the mean of derivative exposure changes.

2.6.1.4 Comparative Statics w.r.t. Collateral Value

The collateral value of illiquid assets represents the ease at which illiquid assets can be liquidated earlier than their maturity. Figure 2.12 shows how the optimal quantity invested in the illiquid asset θ changes with the liquidity parameter. As expected, when the liquid value of the illiquid asset goes up, the optimal θ increases in a convex relationship.

The protection seller does not benefit from an increase in the liquidation value when it is low as much as when the liquidation is already high.




Figure 2.10: Optimal quantity invested in illiquid asset w.r.t. Variance of exposure

2.6.2 Proofs of Lemmas

Proof of Lemma 2.1 :

To prove this lemma it is easier to rewrite the maximization program of the protection seller by expressing the investment in the illiquid asset in proportion of the initial wealth. So only for the sake of this proof, θ now refers to a proportion of the initial wealth instead of a dollar amount.

With this new notation, the agent is maximizing the following function :

$$G(\theta) = \int_{-\infty}^{n\pi(1-\theta)R_l} \left[(n\pi(1-\theta)R_l - f) R_l + n\pi\theta R_i \right] \phi_{\tilde{f}(n)}(f) df \\ + \int_{n\pi(1-\theta)R_l}^{A_1(\theta,n)} \left(n\pi\theta - \frac{f - n\pi(1-\theta)R_l}{\ell} \right) R_i \phi_{\tilde{f}(n)}(f) df$$

with $A_1(\theta, n) = n\pi[(1-\theta)R_l + \theta\ell]$. This is strictly equivalent to the original program except we are now solving for a proportion of assets invested in the illiquid technology. After some computations, we can show that the first derivative is equal




Figure 2.11: Optimal quantity invested in illiquid asset w.r.t. Mean and variance of exposure

to :

$$G'(\theta) = (R_i - R_l^2) \Phi_{\tilde{f}(n)}(n\pi(1-\theta)R_l) + \left(\frac{\ell - R_l}{\ell} \right) R_i \left[\Phi_{\tilde{f}(n)}(A_1(\theta, n)) - \Phi_{\tilde{f}(n)}(n\pi(1-\theta)R_l) \right]$$

We need to show an intermediate result, that is :

Intermediate result : Let \tilde{f} be a random normal variable with mean μ and variance σ^2 , and $\forall n \in \mathbb{R}$ let \tilde{g} be a random normal variable with mean $n\mu$ and variance $n\sigma^2$. If $\Phi_{\tilde{v}}(x)$ denote the cumulative distribution function of the random variable \tilde{v} at the point x , we have that

$$\forall x \in \mathbb{R} \quad \Phi_{\tilde{g}}(nx) = \Phi_{\tilde{f}}(x)$$




Figure 2.12: Optimal quantity invested in illiquid asset w.r.t. Liquid asset's return

Proof : To show this result let's create a third random normal variable, \tilde{h} , with mean 0 and variance 1. We have the following

$$\Phi_{\tilde{g}}(nx) = \Phi_{\tilde{h}}\left(\frac{nx - n\mu}{n\sigma}\right) = \Phi_{\tilde{h}}\left(\frac{x - \mu}{\sigma}\right)$$

and

$$\Phi_{\tilde{f}}(x) = \Phi_{\tilde{h}}\left(\frac{x - \mu}{\sigma}\right)$$

Q.E.D

Then if we consider applying the implicit functions theorem on the first order condition which is when the first derivative $G'(\theta)$ is equal to 0, in order to determine the derivative of θ with respect to n , we have¹⁷

$$\frac{\partial}{\partial n}\theta = -\frac{\frac{\partial}{\partial n}G'(\theta)}{\frac{\partial}{\partial \theta}G'(\theta)}$$

¹⁷We assume here that the conditions for applying the implicit function theorem are fulfilled.

We can now apply our intermediate result to show that

$$\frac{\partial}{\partial n} G'(\theta) = 0$$

Hence

$$\frac{\partial}{\partial n} \theta = 0$$

And because here θ is expressed in proportion of the initial wealth, if we take the original notation we have that

$$\frac{\partial}{\partial n} \left(\frac{\theta^*(n)}{A(n)} \right) = 0$$

Q.E.D

Proof of Lemma 2.2 :

We are interested in characterizing the derivative of the amount invested in the illiquid asset θ with respect to the initial wealth. Recall the first derivative of the objective function $G'(\theta)$ is given by :

$$G'(\theta) = (R_i - R_l^2) \Phi((A - \theta)R_l) + \left(\frac{\ell - R_l}{\ell} \right) R_i \left[\Phi(A_1(\theta)) - \Phi((A - \theta)R_l) \right]$$

with $A_1(\theta) = (A - \theta)R_l + \theta\ell$. Because at the optimum we have $G'(\theta^*) = 0$, we apply the implicit functions theorem to implicitly get the derivatives.

$$\frac{\partial}{\partial A} \theta^*(A) = - \frac{\partial G'(\theta^*) / \partial A}{\partial G'(\theta^*) / \partial \theta}$$

1) Sign of the derivative

In order to determine the sign of this derivative, we use the sign function $sg(\cdot)$. Because we are interested in the derivative of $G'(\cdot)$ at the optimum we already know that the derivative of $G'(\cdot)$ with respect to θ is negative (the objective function is concave).

$$sg \left(\frac{\partial}{\partial A} \theta^*(A) \right) = sg \left(\frac{\partial G'(\theta^*)}{\partial A} \right)$$

We have

$$\frac{\partial G'(\theta^*)}{\partial A} = (R_i - R_l^2) R_l \phi((A - \theta^*)R_l) + \frac{R_l - \ell}{\ell} R_i R_l \left(\phi((A - \theta^*)R_l) - \phi(A_1(\theta^*)) \right) \quad (2.1)$$

We assume that this derivative is positive and show that we end up with a condition that is always true.

$$\begin{aligned} \frac{\partial G'(\theta^*)}{\partial A} &> 0 \\ \iff \frac{\phi((A - \theta^*)R_l)}{\phi(A_1(\theta^*))} &> \frac{R_i(R_l - \ell)}{R_l(R_i - R_l\ell)} \end{aligned}$$

We then notice that the first order condition, that is $G'(\theta^*) = 0$, can be expressed as

$$\frac{\Phi((A - \theta^*)R_l)}{\Phi(A_1(\theta^*))} = \frac{R_i(R_l - \ell)}{R_l(R_i - R_l\ell)}$$

We can then return to the first derivative of $G'()$ with respect to A .

$$\begin{aligned} \frac{\partial G'(\theta^*)}{\partial A} &> 0 \\ \iff \frac{\phi((A - \theta^*)R_l)}{\phi(A_1(\theta^*))} &> \frac{\Phi((A - \theta^*)R_l)}{\Phi(A_1(\theta^*))} \\ \iff \frac{\phi((A - \theta^*)R_l)}{\Phi((A - \theta^*)R_l)} &> \frac{\phi(A_1(\theta^*))}{\Phi(A_1(\theta^*))} \\ \iff \frac{\partial \ln(\Phi((A - \theta^*)R_l))}{\partial (A - \theta^*)R_l} &> \frac{\partial \ln(\Phi(A_1(\theta^*)))}{\partial A_1(\theta^*)} \end{aligned}$$

This is essentially comparing the growth rate of the same function at two different points. Let's define a function f such that $f(x) = \ln(\Phi(x))$, it can be shown that f is a concave function. We can rewrite the previous inequality as

$$\iff \frac{\partial f((A - \theta^*)R_l)}{\partial (A - \theta^*)R_l} > \frac{\partial f(A_1(\theta^*))}{\partial A_1(\theta^*)} \quad (2.2)$$

Because $A_1(\theta^*) = (A - \theta^*)R_l + \theta^*\ell$, we have

$$(A - \theta^*)R_l < A_1(\theta^*)$$

And because $\Phi(x)$ is strictly increasing, we have that (2.2) is always true. This proves that

$$\frac{\partial}{\partial A} \theta^*(A) > 0$$

2) Value of the derivative

We now want to show some bounds on the value of the derivative. We already computed the derivative of the first order condition with respect to A , let's now

derive it with respect to θ .

$$\frac{\partial G'(\theta^*)}{\partial \theta} = - \left[(R_i - R_l^2) R_l \phi((A - \theta^*) R_l) + \frac{R_l - \ell}{\ell} R_i \left(R_l \phi((A - \theta^*) R_l) - (R_l - \ell) \phi(A_1(\theta^*)) \right) \right]$$

By (2.1) we can rewrite that as

$$\frac{\partial G'(\theta^*)}{\partial \theta} = - \left[\frac{\partial G'(\theta^*)}{\partial A} + (R_l - \ell) R_i \phi(A_1(\theta)) \right]$$

Overall we have

$$\frac{\partial}{\partial A} \theta^*(A) = - \frac{\partial G'(\theta^*) / \partial A}{\partial G'(\theta^*) / \partial \theta} = \frac{\partial G'(\theta^*) / \partial A}{\partial G'(\theta^*) / \partial A + (R_l - \ell) R_i \phi(A_1(\theta))} = \frac{x}{x + y}$$

Because $y > 0$ and from the first part of the proof we know that $x > 0$, this shows that

$$0 < \frac{\partial}{\partial A} \theta^*(A) < 1$$

Q.E.D

2.6.3 Proofs of Propositions

Proof of Proposition 2.1 :

The agent is maximizing the following function :

$$G(\theta) = (A_0 - \theta) R_l^2 + \theta R_i - \mu_e$$

We can see that the objective function is not strictly concave in θ , so no global maximum is reached. The first derivative is given by

$$G'(\theta) = R_i - R_l^2 > 0$$

By assumption this is a positive quantity, so $G(\theta)$ is increasing in θ . The agent will then be able to increase θ until the budget constraint binds, which arises at

$$\theta = A_0$$

Q.E.D

Proof of Proposition 2.2 :1) Existence and uniqueness of a solution

The agent is maximizing the following function :

$$\begin{aligned} \theta^* = \arg \max_{\theta} \int_{-\infty}^{(A_0 - \theta)R_l} \left[((A_0 - \theta)R_l - e) R_l + \theta R_i \right] \phi(e) de \\ + \int_{(A_0 - \theta)R_l}^{A_1(\theta)} \left(\theta - \frac{e - (A_0 - \theta)R_l}{\ell} \right) R_i \phi(e) de \end{aligned}$$

First note that, there is an ℓ at the denominator in the second integral. To allow $\ell = 0$, we could take the limit of that quantity as $\ell \rightarrow 0$ but instead we notice that

$$\ell = 0 \quad \Longleftrightarrow \quad A_1(\theta) = (A_0 - \theta)R_l$$

so the second integral vanishes to 0. We will then separate two cases, one for which $\ell = 0$ and the other for which $\ell > 0$.

We compute the first derivative of the objective function. We define $\Phi(x)$ to be the cumulative distribution function of a normal random variable with mean μ_e and variance σ_e^2 , and $\phi(x) = \Phi'(x)$. After some computations we end up with

- If $\ell = 0$:

$$G'(\theta) = (R_i - R_l^2)\Phi((A_0 - \theta)R_l) - \theta R_i R_l \phi((A_0 - \theta)R_l) \quad (2.3)$$

- If $\ell > 0$:

$$G'(\theta) = (R_i - R_l^2)\Phi((A_0 - \theta)R_l) + \left(\frac{\ell - R_l}{\ell} \right) R_i \left[\Phi(A_1(\theta)) - \Phi((A_0 - \theta)R_l) \right] \quad (2.4)$$

Moreover, we have that

$$\lim_{\ell \rightarrow 0} \frac{\Phi(A_1(\theta)) - \Phi((A_0 - \theta)R_l)}{\ell} = \lim_{\ell \rightarrow 0} \frac{\Phi(A_1(\theta)) - \Phi(A_1(\theta) - \theta\ell)}{\ell} = \theta\phi((A_0 - \theta)R_l)$$

This shows that

$$\lim_{\ell \rightarrow 0} (2.4) = (2.3)$$

So we can simply use (2.4) and still consider the case in which $\ell = 0$ keeping in mind that in such a case, we are actually referring to the limit. We now compute the second

derivative of the objective function. After some computations and simplifications we have that

$$G''(\theta) = \phi(A_1(\theta))R_i \frac{(R_l - \ell)^2}{\ell} - \phi((A_0 - \theta)R_l) \frac{R_l^2}{\ell} (R_i - R_l \ell)$$

We are interested in the sign of the second derivative, because in addition to the limits of the objective function, it will give us some important insights about the shape of the objective function. The limits can be derived directly from the objective function :

$$\lim_{\theta \rightarrow +\infty} G(\theta) = -\infty \quad \text{and} \quad \lim_{\theta \rightarrow -\infty} G(\theta) = \text{cste (finite)}$$

First, let's show a small result which will be needed afterward. If we let X and Y be such that

$$X = R_i \frac{(R_l - \ell)^2}{\ell} \quad \text{and} \quad Y = \frac{R_l^2}{\ell} (R_i - R_l \ell)$$

we can show that

$$0 < X < Y$$

The first inequality is obvious from our assumptions. To prove the second inequality, we assume that this condition is true and simplifies it to show that under our usual assumptions this is actually always true.

$$\begin{aligned} R_i \frac{(R_l - \ell)^2}{\ell} &< \frac{R_l^2}{\ell} (R_i - R_l \ell) \\ R_i \frac{R_l^2}{\ell} - R_l^3 &< R_i \frac{R_l^2}{\ell} + R_l \ell - 2R_i R_l \\ R_l^3 &< R_i (2R_l - \ell) \\ R_l^2 R_l &< R_i (2R_l - \ell) \end{aligned}$$

This is always true because of the two following results we have from the model.

$$R_i > R_l^2 > 0$$

and

$$R_l > \ell \quad \Longleftrightarrow \quad 2R_l - \ell > R_l > 0$$

Now we can rewrite the sign of the second derivative as follow, with $sg(\cdot)$ being the sign function,

$$sg(G''(\theta)) = sg \left(\phi(A_1(\theta)) \left(X - \frac{\phi((A_0 - \theta)R_l)}{\phi(A_1(\theta))} Y \right) \right) = sg \left(X - \frac{\phi((A_0 - \theta)R_l)}{\phi(A_1(\theta))} Y \right)$$

Since $\phi(\cdot)$ is the density function of a normal distribution with mean μ_e and variance σ_e^2 , we have that

$$\frac{\phi((A_0 - \theta)R_l)}{\phi(A_1(\theta))} = e^{\frac{(A_1(\theta) - \mu_e)^2 - ((A_0 - \theta)R_l - \mu_e)^2}{2\sigma_e^2}} \quad (2.5)$$

Let's Z be the term above the exponential. Z is a second order polynomial equation in θ . We are interested in the sign of the square term, which would give us the "orientation" of the equation's curve. We have

$$Z(\theta) = \frac{(A_1(\theta) - \mu_e)^2 - ((A_0 - \theta)R_l - \mu_e)^2}{2\sigma_e^2}$$

$$Z(\theta) = \frac{1}{2\sigma_e^2} \theta \ell \left(2(A_0 - \theta)R_l + \theta \ell - 2\mu_e \right)$$

Let's denote by Z_α the square term coefficient of the Z polynomial equation. We have

$$Z_\alpha = \frac{1}{2\sigma_e^2} (\ell^2 - 2\ell R_l) < 0$$

We then know that (2.5) has a shape similar to that of a density of a normal variable, with

$$\lim_{\theta \rightarrow -\infty} e^{Z(\theta)} = \lim_{\theta \rightarrow +\infty} e^{Z(\theta)} = 0 \quad \text{and} \quad e^{Z(0)} = 1$$

Now we are almost done because recall that

$$sg(G''(\theta)) = sg\left(X - e^{Z(\theta)}Y\right) = sg\left(\frac{X}{Y} - e^{Z(\theta)}\right)$$

And we proved earlier that

$$0 < \frac{X}{Y} < 1$$

So basically this means that for arbitrarily low or high value of θ , the second derivative $G''(\theta)$ is positive, and because the function is continuous, there has to be some values of θ for which the second derivative $G''(\theta)$ is negative.

This means that the objective function is first convex, then concave and then convex again. If we recall the limits of the objective function, equals to $-\infty$ for $\theta \rightarrow +\infty$ and to a constant for $\theta \rightarrow -\infty$, it has to be the case that the objective function is first increasing (first convex, then concave) and then decreasing (concave and then convex). Because the first derivative is continuous (sum of continuous functions), there is always a unique maximum.

2) Bounds on the optimal value

We now need to prove the bounds on the optimal value. The bounds of impor-

tance here are 0 and A_0 . To do so, we study the value of the first derivative in 0 and in A_0 . If the first derivative in those points is positive this means the maximum is higher and if the first derivative is negative this means the maximum is before. Recall the first derivative of the objective function

$$G'(\theta) = (R_i - R_l^2)\Phi((A_0 - \theta)R_l) + \left(\frac{\ell - R_l}{\ell}\right) R_i \left[\Phi(A_1(\theta)) - \Phi((A_0 - \theta)R_l) \right]$$

We have

$$G'(0) = (R_i - R_l^2)\Phi(A_0R_l) + \left(\frac{\ell - R_l}{\ell}\right) R_i \left[\Phi(A_0R_l) - \Phi(A_0R_l) \right]$$

$$G'(0) = (R_i - R_l^2)\Phi(A_0R_l) > 0$$

This proves the optimal quantity invested in the illiquid asset θ is always positive. Let's now study the first derivative in A_0 .

$$G'(A_0) = (R_i - R_l^2)\Phi(0) + \left(\frac{\ell - R_l}{\ell}\right) R_i \left[\Phi(A_0\ell) - \Phi(0) \right]$$

And we are interested in the case for which this quantity is negative, corresponding to a optimal $\theta < A_0$. This is the case when

$$G'(A_0) < 0$$

$$\iff (R_i - R_l^2)\Phi(0) - (R_l - \ell)R_i \frac{[\Phi(A_0\ell) - \Phi(0)]}{\ell} < 0$$

We have to keep in mind that ℓ could be equal to 0, but in this case we have previously shown that this is equivalent to taking the limit. So to simplify notation, let $\mathcal{H}(\ell)$ be such that

$$\mathcal{H}(\ell) = \frac{[\Phi(A_0\ell) - \Phi(0)]}{\ell}$$

with $\lim_{\ell \rightarrow 0} \mathcal{H}(\ell) = A_0\Phi'(0)$. We then have

$$G'(A_0) < 0 \iff \frac{\Phi(0)}{\mathcal{H}(\ell)} < \frac{R_i(R_l - \ell)}{R_i - R_l^2}$$

Q.E.D

Proof of Proposition 2.3 :

The first part of the proposition comes directly from Lemma 2.1. For the last part, protection sellers solve the following program.

$$\max_n n \Pi_u$$

By construction Π_u is independent of n . So the solution is trivial as the objective function is linear. Hence if $\Pi_u > 0$, the agent will choose the highest possible n , and if $\Pi_u < 0$ the agent will choose the lowest possible n . To make things interesting we need to require some exogenous constraints on the optimal number of contracts, $n^* \in [0, n_{max}]$.

Q.E.D

Proof of Proposition 2.4 :

By Lemma 2.2 we know that the derivative of the optimal quantity invested in the illiquid assets is always positive and lower than 1. The consequence is that if one removes some cash from the initial assets of a protection seller, he will invest relatively more in the liquid assets. Then since the assets initially removed are kept on the balance sheet of the CCP as cash, it cannot be that the global cash usage is lower with a CCP.

Q.E.D

Proof of Proposition 5 :

We want to compare the two following quantities.

$$EL^m = \int_{A_1(n\pi, \theta_m^*)}^{+\infty} m(e - A_1(n\pi, \theta_m^*)) \phi(e) de \quad EL^{ccp} = \int_Z^{+\infty} m(e - Z) \phi(e) de$$

where $Z = \mu_e + \frac{D - m(\mu_e - A_1(n^*\pi - d, \theta_{ccp}^*))}{m}$, $A_1(n^*\pi - d, \theta_{ccp}^*)$ is the available cash at the intermediary date in the case of central clearing, $A_1(n^*\pi, \theta_m^*)$ is the available cash at the intermediary date in the case without central clearing and $\phi(\cdot)$ the density of a normal random variable with mean μ_e and variance σ_e^2 .

Because of the symmetry in the two expressions, we can see that it is the same as comparing the lower bound on the integral in both cases.

$$A_1(n\pi, \theta_m^*) - \mu_e < \frac{D - m(\mu_e - A_1(n\pi - 2n\mu_e/m, \theta_{ccp}^*))}{m}$$

$$\begin{aligned}
 \Leftrightarrow \quad D &> m \left(A_1(n\pi, \theta_m^*) - A_1 \left(n\pi - \frac{2n\mu_e}{m}, \theta_{ccp}^* \right) \right) \\
 \Leftrightarrow \quad 2n\mu_e &> m \left(\frac{2n\mu_e}{m} R_l + (\theta_{ccp}^* - \theta_m^*)(R_l - \ell) \right) \\
 \Leftrightarrow \quad m &> \frac{2n\mu_e(R_l - 1)}{(R_l - \ell)(\theta_m^* - \theta_{ccp}^*)}
 \end{aligned}$$

If this condition is verified, then both the default probability and the expected loss are lower when there is a CCP in the market.

Q.E.D

Chapter 3

Policy Making with Hidden Information and Continuous-time Learning

joint work with Charles BERTUCCI

Abstract :

Policy uncertainty has a strong influence on the economic activity. We build a general framework allowing us to highlight the link between policy uncertainty and the institution's learning process. An institution must learn the agents' type through the agents' strategies to achieve optimal regulation. This endogenously creates some level of uncertainty. Overconfidence can have two effects. When the current policy forces the market to release a lot of information, overconfidence makes the policy converge smoothly toward its long term equilibrium. When the current policy makes it hard for the institution to learn, overconfidence creates inefficiencies. Regarding institution's design, we find that there is an optimal level of constraints for the institution that minimizes the uncertainty perceived by agents.

3.1 Introduction

Over the past three decades, central bank communication has become more and more important as a tool for central bankers and as a signal regarding future policy for economic participants¹. Nowadays, a central bank extensively relies on communication tools to conduct its policy. Especially in times of unusual policies like *quantitative easing*, when the interest rate reaches the zero-lower-bound, central bankers can still use communication as a way to influence the market². However, this can create uncertainty about the future of policy if different central bankers' interventions are not interpreted as going in the same direction. This can happen either because at the same time multiple senior directors issue contradictory statements or if there is a unique public disclosure that varies over time. The same mechanism applies to other types of state-level decisions like financial regulation. Indeed, there are a lot of communication and discussions around detail implementations of recommendations like the Dodd-Frank act or EMIR for instance³. This creates a significant amount of uncertainty among market participants⁴. The effects as well as the theoretical mechanisms behind this uncertainty are still mostly unknown.

This chapter aims at identifying the theoretical foundations behind policy uncertainty. We argue that uncertainty is inherent to optimal policy design and results solely from *i*) the lack of information (of both the policy designer and market participants) and *ii*) the lack of commitment. Indeed, to justify the first condition, we can say that there are no particular reasons why the decision maker would have perfect information about every aspects of the economy. It seems rather straightforward but against the traditional view of an omniscient social planner. The lack of commitment has been studied⁵ and is a feature of most regulatory actions. Indeed the regulator cannot commit to either do exactly what she said or commit to not change the current rule. The central idea of this chapter is that, under both conditions, the regulator will naturally use policy modifications, hence generating uncertainty, as a learning device to gain knowledge about the underlying model in

¹See Blinder, Ehrmann, Fratzscher, De Haan and Jansen (2008) for a survey about the central bank communication paradigm shift from a practical as well as an academic point of view.

²See Bernanke, Reinhart and Sack (2004) and Eggertsson and Woodford (2003) for monetary policies at the zero-lower-bound.

³Both Dodd-Frank and EMIR are recommendations issued as a response to the financial crisis, respectively in 2010 and 2011. As of the writing of this chapter there remains a number of actual recommendations that need to be confirm as actual regulations. Meanwhile economic agents can only interpret those texts and infer from the ongoing debates to what extent they will be subject to new business rules.

⁴See for instance the Financial CHOICE Act led by republicans in the US. This act passed the Congress on June 8th, 2017, and if enabled by higher level authorities would roll back many of the DF recommendations.

⁵See for instance Vartiainen (2014) for an application of mechanism design without commitment.

order to achieve optimal policy. For instance, how the market reacts to (minor) changes in the policy (or in the communication around future policies) will result in more information being captured by the regulatory institution. The literature on the effect of central bank communication justifies the impact of signaling through communication. It has been found that communication only, has a strong effect on economic variables⁶. We can see that a tradeoff starts to emerge as changing the policy might result in short term inefficiencies but will yield more information so most likely less inefficiencies be generated over the long run.

The goal of this chapter is to shed light on the inherent mechanisms behind the learning behavior of any central authorities like regulators or central banks. The main mechanism in our model could be summarized as follows. Because of imperfect information, the regulator cannot a priori assess the optimal policy, and need to learn the underlying market characteristics in order to achieve optimum. This process endogenously generates some level of uncertainty. The uncertainty is the consequence of the learning behavior of the regulator and is therefore necessary in order for the regulator to obtain information. Note that in the long run, once learning is achieved, the uncertainty should not remain. In our model, it is not a feature of anyone to like or dislike uncertainty, there is just some that will arise endogenously as long as the regulatory institution finds it useful as a mean to obtain more information. The cost of learning can be defined as the cost of *exploration* (to use the term of the Multi-Armed Bandit literature), and represents the potential loss incurred by the regulator when she changes the policy only to learn how market responds. Once this learning cost is too high compared to the expected gains from the current best estimate of the optimal policy, the level of uncertainty decreases.

We feature a model with a principal facing a continuum of myopic agents. The principal have control over a policy macro-variable and agents continuously react to this variable. The macro-variable could correspond to an interest rate or a regulation for instance, but as well to some signals being sent by the regulator. This lies on the assumption that agents respond in a similar way to a signal modification or an actual rule modification. Agents are myopic in the sense that they only respond to the current action of the principal without trying to anticipate the future evolution of the principal's policy. Non myopic agents might also bias their responses such that the learning ability of the regulator is reduced. For instance, in the case they have competing objectives, they would have to trade off the current gain associated with maximizing current utility with the loss incurred by the information released to the regulator that would reduce future gains. This is however out of the scope of this chapter.

The key aspect of our model is the imperfect information about the agents' type.

⁶See for instance Hansen and McMahon (2016).

Indeed, at the beginning of the model, the nature assigns the same value of private type for each agent. The principal does not know this value and this is what she needs to learn. Given a prior belief, the principal uses actions to obtain a signal from the agents in the form of their best response and then updates her belief according to Bayes' rule. Note that there are no exogenous signals received by anyone in the model. Only the agents' type and the prior belief of the institution are exogenous. The institution tries to learn the agents' type from their noisy response. If agents were not myopic they would try to learn the prior belief of the institution in order to deduce her internal state of mind, but again this is out of the scope of this chapter⁷. We provide a comprehensive way to account for the regulator learning process in a continuous fashion.

One important aspect of our model is that the tradeoff between short term inefficiencies due to learning and long term optimal strategy is driven only by the objective function of the central institution. Indeed, the principal is minimizing a quadratic loss function which increases when the agents' responses are far from some exogenous target response⁸. By opposition to Stein (2017) the central institution does not take into account any side effects resulting from its potentially volatile optimal strategy. Here, the central bank only targets some value and by extension does care about learning the agents' parameter but only to the extent that it reduces her loss. We could add more components to the objective loss function of the principal but so far this is out of the scope of this chapter.

The main dynamics of our model is the following. The institution receives a noisy version of the aggregate agents' response which is influenced by her policy. The institution then updates her belief accordingly. With the new belief, the institution can then set an optimal policy in order to obtain another noisy signal, and so on. All of this happens continuously. In the most basic version, we show that the volatility of the macro variable decreases with the uncertainty surrounding the value of the unknown parameter. The belief distribution has *full-support* so the Bayes rule indicates that learning is efficient and the distribution converges toward a Dirac measure on the true value. When it does so, the volatility of the macro variable is zero.

We show that overconfidence around the institution's belief can have mixed effects. Depending on the value of the macro variable, the institution can be in a *learning-* or a *non-learning-* state. First, when the institution is wrong but the state of the economy is such that the institution manages to extract a lot of information, the initial error will be corrected quickly. Second, when the state of the economy

⁷For an interesting way of modeling such two-sided interaction, although in a different setting, see Cisternas (2017).

⁸The effect is actually stronger as the institution's loss function also has a term that makes uncertainty costly. Therefore the resulting uncertainty comes directly from the learning process.

lowers the signal, this can dramatically harm learning. Overconfidence exacerbates this result. Recall that the institution is setting the state of the economy. However, the institution sets the state of the economy according to her belief. If associated with a poor-learning state yield a very poor learning process. Therefore the institution cannot update her belief too much, and the economy is stuck in a non-optimal state. On the other hand, if the institution is erroneously overconfident, but such that the signal is high and learning is strong, the economy converges quickly to the long term equilibrium.

Our results indicate that the noise in the signal received by the institution reduces the belief's influence on the optimal strategy. With a low quality signal, the institution's learning is poor. She knows that she is not learning much so she sees less value associated with moving the macro variable around and uncertainty is reduced. However, it is associated with long term inefficiencies as the principal needs a long time to learn. The signal quality can be interpreted as how good the institution is at extracting information from the economy. This might be due to information being hard to access, the institution needing time to assess it or if the consequences of the institution's policy are delayed. In an opaque market, policy changes are expected to be low and the *social loss* is expected to be high.

We provide a way to incorporate the usual tradeoff between exploitation and exploration in a continuous model. The principal's loss function is composed of two parts: the *exploitation cost* and the *exploration cost*. The *exploitation cost* represents the cost associated with being far from the expected best policy. The *exploration cost* is the cost associated with changing the policy. The latter cost is therefore also associated with uncertainty. With some *exploration cost*, the uncertainty is artificially reduced as it is costly for the institution to change the policy. Therefore, we should see any results as conditional on the given constraint on the institution's policy. By increasing the *exploration cost* we can simulate an economy in which the institution has less influence on the market she is meant to regulate⁹. We show that the control cost introduces some delays in the institution's policy with respect to changes in her belief. Another way to put it would be that with no *exploration cost* the institution would have chosen a policy different from the optimal one as her loss can be lowered by smoothing the policy.

Our results suggest that with no *exploration cost*, uncertainty is inefficiently high. As the *exploration cost* increases, the institution is more and more constrained and needs to find a way to keep learning without changing much the current policy.

⁹This can be due to higher order rules like constitutions or any kind of multilateral or bilateral agreement between several regulating institutions. It could also be linked to credibility. A market in which the institution's management has lost credibility will have less impact on the market, and therefore the economic variables will be less sensitive to policy changes. Increasing the *exploration cost* has the same effect.

Meanwhile the learning process can still be efficient. We conjecture that this happens until a cut off *exploration cost* above which the institution is too constrained and learning become inefficient. This suggests that some level of control should be applied to any regulating institution, and that uncertainty and learning goes together.

This chapter shed lights on the theoretical foundations of institution's design. We show that if the institution needs to learn something from the market, she will use uncertainty which she can do depending on the magnitude of the *exploration cost*. Therefore we can take one step back and study how an institution should be designed. Our results suggests that applying some level of constraints on what the institution can do will ultimately reduce the uncertainty perceived by agents, even with no moral hazard from the institution. This validates the use of *supra rules* that dictates policy changes limitations¹⁰. Note that this does not come from any deviations of the institution trying to use private information against the market. This is just a consequence of the learning that needs to happen.

This chapter contributes to the growing literature about central bank communication with partial information and coordination. Baeriswyl (2011) and Baeriswyl, Cornand and Ziliotto (2016) propose an extended version of the Keynesian beauty contest game that accounts for endogenous information. This framework has been introduced in a paper by Morris and Shin (2002) in which they feature a central bank that can only communicate with market participants. James and Lawler (2011) extended this framework to account for actions as well as communication. And most related to our model, Baeriswyl and Cornand (2010) subsequently extended this framework and introduces a central bank that can communicate through her actions. In a beauty contest game, agents exhibit some form of strategic coordination as they want to be close to the true value of the economic fundamental but they also do not like being far from each other in terms of strategy. We depart from this by making no particular assumptions on the utility function of agents. Our model features a continuum of agents and the institution, or central bank, observes a noisy average of agents' responses, so no agent can influence the signal observed by the institution. Therefore, they do not care about the other agents' strategies. Another important difference with Baeriswyl, Cornand and Ziliotto (2016), is that they use a static model, whereas we introduce a dynamic setting. This kind of framework is more suited to capture time variation in the learning process.

In a number of studies regarding central bank optimal policy and communication, it is often assumed that the central bank has some kind of private information with respect to the market. For instance, papers like Faust and Svensson (2001),

¹⁰By *supra-rules* we mean rules like those specified in the Treaty on European Union (1992) for the EU, or the Federal Reserve Act (1913) for the US Fed.

Orphanides and Williams (2006), Stein (2017), as well as the papers listed in the above paragraph, all assume that the central bank has either a target for its policy rate or an inflation target associated with a Taylor rule (linking inflation, unemployment and interest rate). Then the question becomes under what circumstances and how it would be optimal for the central bank to be transparent about her private knowledge.

In our model, however, the central bank has no exogenous private information about economic fundamentals or her internal response function. Communication under those assumptions refers to the central bank sending a signal (either a pure signal as in Morris and Shin (2002) or through signaling actions as in Baeriswyl, Cornand and Ziliotto (2016)). We rather take the point of view of the central banker who needs to learn about the economic driving forces in order to achieve her goal of optimal policy. Our notion of communication is different from the one used in this literature. Indeed, in our model the central bank does not have any private information valuable to agents. She does not communicate in order to reveal, partially or fully, any information. She rather uses communication as a learning device for the sole purpose of extracting more information out of the agents' responses.

Regarding the effect of communication on markets Eusepi and Preston (2010) provide a model-based evaluation strategy of central bank communication. Ehrmann and Fratzscher (2007) investigate the different communication strategies of the ECB, the Bank of England, and the Federal Reserve. Ranaldo and Rossi (2010) analyze the effect of the Swiss National Bank announcements on financial metrics. Hansen and McMahon (2016) analyze the communication of the Federal Reserve by studying the FOMC reports with a Natural Language Processing algorithm. Then they use a Factor-Augmented VAR to analyze the communication effect on market variables. They found that an important component of the change in economic output is in fact driven by communication. However most of the papers in this literature focus on communication as a way to stabilize markets or the optimal way a private information should be communicated to the market. This chapter rather argues that the central authority has no private information whatsoever. The only information that is not known by market participants is the prior belief of the central bank. Whether the central bank could or should communicate this prior belief in order to align all expectations and influence the learning process is an interesting question but out of the scope of this chapter.

Another important distinction is that in this literature, since the central bank has some private information, market participants are doing the learning of this private information from the central bank actions. In our model, the central bank is doing the learning on the underlying fundamentals. We feature myopic agents but even if they tried to strategically anticipate the future principal's actions, the principal

would still have to learn in the meantime, so this would be a kind of two sided learning game¹¹. In this chapter, communication of information (through actions) is more than just an information transmission channel, it is also an information acquisition mechanism because the central bank will learn from the way agents respond to her actions.

This chapter also contributes to the literature on general policy uncertainty. Bloom (2009) provides a structural framework to analyze uncertainty shocks. Baker, Bloom and Davis (2016) set up a Natural Language Processing algorithm that reads newspapers every day and compute a policy uncertainty index¹². By construction, this index is updated on a daily basis and is available in several countries and monetary areas. Among others, as contributions to this empirical literature we can cite Bailey, Zheng and Zhu (2014), Benati (2013), Brogaard and Detzel (2015), Pastor and Veronesi (2012) and Pastor and Veronesi (2013). While this literature generally accepts policy uncertainty shocks as exogenous in order to derive causal implications over the short run, we provide a framework for analyzing where this uncertainty comes from in the first place. While we do not have a reliable explanation for daily changes in uncertainty perceived by economic agents, we show that this kind of uncertainty naturally arises over the long run when policy makers have imperfect information about the underlying market. While, as most of this literature shows, uncertainty may be harmful in some situations, we argue that this also provides policy makers with more knowledge about the fundamentals. Our model does not take into account any direct cost of uncertainty for economic agents, further extensions could implement such changes and then be able to provide a real tradeoff between learning more information and having a negative impact on the economy while doing so.

Lastly, when analyzing decisions' timing in organization Grenadier, Malenko and Malenko (2016) also use a continuous time principal agent model. A biased agent sends a message to an uninformed principal who is responsible for taking a decision that also affects the agent's payoff. Our model shares some similarities regarding the way it is framed although some important differences remain. In their model, the agent has the ability to send a message to the principal, and they study incentive-compatible contracts that are also consistent with conditional probability updates from the principal. In our model, agents maximize a continuous reward and do not send any specific messages to the principal beside a noisy representation of their current aggregate response. Another main difference is that, in their model, the principal can only take a single decision regarding the final payoff. One of the main

¹¹See Cisternas (2017) for a similar approach with two-sided learning.

¹²Their index is constructed with basic techniques such as number of occurrence of sentences in which the terms *policy* and *uncertainty* are present. The detail methodology and data are available on their website: www.policyuncertainty.com.

contributions of this chapter is the signaling actions setup from the point of view of the principal. Indeed, contrary to Grenadier, Malenko and Malenko (2016) our model assumes that the principal sends a message, through her actions, to agents taking into account the impact of this signal onto their strategies. This allows the principal to control her learning rate and trade off the benefit and cost of a fast learning process.

The remaining of this chapter is organized as follows. Section 2 presents the model. Section 3 numerically solves for the optimal policy of the central institution. Section 4 presents a general interpretation of the solution. The analysis in terms of the *Multi-Armed Bandit* literature is conducted in section 5. Section 6 concludes.

3.2 The Model

Time is continuous and horizon is infinite¹³. There is a mass-1 continuum of agents. They all share the same “type”, $\Pi \in \mathbb{R}$, that is drawn by the nature at the beginning of the model. An agent’s strategy is defined by $\beta(\cdot)$. In this basic setup definition, we feature myopic agents. So at each time they simply maximize their current utility. The reward they get depends on their type as well as on a macro-variable X_t . They simply observe the current level of X_t and react to this without trying to anticipate future values of X_t , and their strategy can then be written as $\beta(\Pi, X_t)$.

There is a principal, also referred to as the *institution*, who has the ability to affect all agents’ payoffs using a single control. She discounts time at rate δ . More specifically, the institution can control the change in the macro-variable X_t . Her strategy profile is defined by $\{\alpha_t\}_t$ and X_t evolves according to the following process.

$$dX_t = \alpha_t dt$$

One possible interpretation would be to consider X_t to be the current interest rate set by the central bank. If we give the opportunity to the central bank to change the interest rate level at will, we would have to impose more constraints on how this interest rate could change. For instance, we would like to avoid situations in which the rate process X_t is discontinuous. We give the institution the possibility to control only the change in interest rate. This frees us from having to make extra continuity assumptions as now even if the institution control α_t changes a lot, the interest rate cannot jump. This is more realistic when dealing with the central bank’s optimal policy but also in the general case as there might be some external reasons why an institution would not want to change its policy back and forth too

¹³This setup is inspired by the work of Lions (2016).

often¹⁴.

Let's define the information structure. The agents' type is unknown to the principal. From her point of view, and at the beginning of the model, the agents' parameter Π is a random variable drawn out of a normal distribution with mean μ_0 and variance σ_0 . The distribution of the principal's prior belief does not really matter for learning¹⁵, and a normal distribution is more convenient. The principal continuously receives a noisy signal Y_t which is an aggregate of the agents' response functions plus a normal noise.

$$Y_t = \beta(\Pi, X_t) + W_t$$

with W_t a Wiener process, $dW_t \sim \mathcal{N}(0, \sigma_w)$. The institution may use this continuous signal to try to learn the true value of Π , the agents' parameter. More specifically, the institution will update the value of μ_0 and σ_0 according to the bayesian rule of conditional probabilities. We denote by μ_t and σ_t the best estimate that the institution can make of the normal distribution parameters given all history preceding time t . In order to simplify the interpretation of the learning process, we will use the inverse of the variance around μ_t , so that $\lambda_t \equiv \sigma_t^{-1}$ now represents the degree of confidence in the estimation of the agents' type Π . Then at each time t , we denote by f_t the institution's belief about the proportion Π .

$$f_t \sim \mathcal{N}(\mu_t, \lambda_t^{-1})$$

We do not explicitly characterize the maximization program of agents, as they are myopic. We simply assume a specific reaction function to the level of the macro-variable¹⁶ which depends on the agents' type Π .

$$\beta(\Pi, X_t) = \Pi X_t$$

We can plug this reaction function in the signal and differentiate with respect to time in order to obtain the signal generation process, dY_t .

$$dY_t = \Pi X_t dt + dW_t$$

The signal Y_t will be used by the principal to update her belief at each point in time.

¹⁴Such reasons could be concerns regarding market volatility in the case of central bank or lost of credibility in a more advanced setup with non-myopic agents for instance.

¹⁵With bayesian learning, as long as the true value is included in the support of the belief distribution, the belief process will converge toward a Dirac measure on the true value.

¹⁶Although, it would not be hard to set up a utility function that yields an optimal action such as the one assumed.

The continuous time bayesian updating rules for the belief parameters μ and λ are given by the following lemma.

Lemma 3.1 (*Bayesian updating*) *Given the information process $(Y_t)_{t \geq 0}$ and the evolution of the macro variable $(X_t)_{t \geq 0}$ the belief of the institution evolves according to*

$$\begin{aligned} d\lambda_t &= X_t^2 dt; \\ d\mu_t &= -\mu_t \frac{X_t^2}{\lambda_t} dt + \frac{X_t}{\lambda_t} dY_t; \end{aligned}$$

By Lemma 3.1, we see that the learning process is actually efficient because $d\lambda_t \geq 0$. As time goes, the institution becomes more and more confident on its estimation of the agent's parameter Π , as $(\lambda_t)_{t \geq 0}$ grows like X_t^2 . Intuitively this is because when observing the signal Y_t the principal is really extracting information from the quadratic variation of the signal and not from the absolute value of the signal. At each point in time, the institution then benefits a lot from keeping¹⁷ a large absolute value of the macro-variable X_t . This arises because of the assumption we made on the agent's response function, that is linear in X_t . We have that a high value of X_t will yield a fast learning process. The cost of having a large value of X_t will appear once we setup the optimization program of the principal.

Another interpretation is that, since the agents' response function is given by $\beta(\Pi, X_t) = \Pi X_t$, a low X_t will yield a low response, whatever the value of Π is, hence making the learning process hard for the institution. On the contrary a high value of X_t will help distinguish Π because the noisy response that the institution get will be better influenced by the value of Π rather than the value of X_t . This comes directly from the specification of the agents' response and is not a general result.

Let's now turn to the optimization program of the principal. For simplicity we assume that the institution is targeting a single value, $\bar{\beta}$, for the response function of agents, $\beta(\Pi, X_t)$. This assumption is simply a way to provide some kind of target to the principal. As an illustration, without such target, if the principal's aim was to learn the agents' parameter Π , she will simply play a very large absolute value of α which will make the macro variable very large, hence reducing the impact of noise (bounded by σ_w) in the agents' response, and he would be able to infer the value of the agents' parameter instantly.

Let's denote by \mathcal{L} the loss function that the principal is minimizing at each point

¹⁷Recall that the institution can only affect $\alpha_t = \frac{dX_t}{dt}$, the change of X_t .

in time.

$$\mathcal{L}(X_t, \Pi, \alpha) = (1 - \theta) (\Pi X_t - \bar{\beta})^2 + \theta \frac{1}{2} \alpha^2$$

where $\theta \in (0, 1)$. This loss function is composed of two terms. First there is the *economic part*, $(\Pi X_t - \bar{\beta})^2$. This is a strictly convex equation in X_t , so had the institution known the agents' parameter Π , she could easily have computed the optimal value of the macro variable and played an α_t so that the macro variable X_t converges to its optimal value. However, in the presence of uncertainty about the value of Π , the principal will minimize over all possible values of Π according to its current beliefs. Then it follows that if the confidence in the belief λ_t is small or if the belief's mean μ_t is far from the true Π , too much weight will be associated with value of Π that are far from the true one, thus resulting in a large loss. Second, there is a *control part*, $1/2\alpha^2$. This acts as a way to prevent the macro variable process X_t from jumping, by applying a cost on large variation of X_t . Intuitively, introducing a control cost in the loss function means that the institution would be better off with $\alpha_t = 0$, hence no changes in X_t , therefore no uncertainty. It also means that if some uncertainty remains in the long run, this is because the institution still needs to learn from the market, learning is not achieved. The parameter θ represents the relative importance of those two costs. A comprehensive analysis of the optimal strategy with respect to those parameters is provided in section 3.4.3.

We can now write the full optimization program. Recall that the institution discounts time at rate δ , so at any time t the minimization program faced by the institution is given by the following.

$$\alpha_t^* \equiv \arg \inf_{\alpha} \mathbb{E} \left[\int_0^{+\infty} e^{-\delta t} \int_{-\infty}^{+\infty} \mathcal{L}(X_t, \Pi, \alpha) f_t(\mu_t, \lambda_t^{-1})(\Pi) d\Pi dt \right] \quad (3.1)$$

The infimum is taken over all processes (α_t) progressively measurable with respect to the σ -algebra defined by the information process and the expectation is taken over the information process $(Y_t)_{t \geq 0}$. Note that a consequence of the bayesian updating hypothesis is that this information process is now seen as a mean-0 brownian motion by the principal. Indeed if we denote by $(\tilde{W}_t)_{t \geq 0}$ a standard brownian motion under this new measure, then the evolution equation satisfied by the variables in (3.1) are given by the following.

$$dX_t = \alpha_t dt \quad (3.2)$$

$$d\lambda_t = X_t^2 dt \quad (3.3)$$

$$d\mu_t = -\mu_t \frac{X_t^2}{\lambda_t} dt + \frac{X_t}{\lambda_t} dY_t \quad (3.4)$$

$$dY_t = \sigma_w d\tilde{W}_t \quad (3.5)$$

Note that because we have assumed that the principal is using the Bayes' rule to update its belief, the information process $(Y_t)_{t \geq 0}$ evolves (from the point of view of the principal) as a brownian motion. Therefore, the value of the process (Y_t) might not be considered a state variable. In order to solve for the optimal policy α_t^* we make full use of the dynamic programming principle. We denote by V the value function of the principal's problem. By definition V satisfies :

$$V(x, \lambda, \mu) = \inf_{\alpha} \mathbb{E} \left[\int_0^{+\infty} e^{-\delta t} \int_{-\infty}^{+\infty} \mathcal{L}(X_t, \Pi, \alpha) f_t(\mu_t, \lambda_t^{-1})(\Pi) d\Pi dt \right]$$

where $(X_t)_{t \geq 0}$, $(\lambda_t)_{t \geq 0}$ and $(\mu_t)_{t \geq 0}$ are driven by (3.2), (3.3) and (3.4) and satisfy :

$$\begin{cases} X_0 = x \\ \lambda_0 = \lambda \\ \mu_0 = \mu \end{cases}$$

The value function does not depend on the value of the information process $(Y_t)_{t \geq 0}$. This is a consequence of the fact that the cost does not depend on this variable either. For the principal, only the variations of the information process are important to compute the evolution of the other state variables. We recall that $V(x, \lambda, \mu)$ represents the expected loss of the principal if the initial conditions are (x, λ, μ) . By the dynamic programming principle, we expect this value function to be a solution of an Hamilton-Jacobi-Bellman equation (HJB). The HJB equation associated with the optimal control problem of the principal is the following.

$$\delta V - x^2 \frac{\partial V}{\partial \lambda} + \frac{\sigma x^2 \mu}{\lambda^2} \frac{\partial V}{\partial \mu} + \frac{\sigma^2 x^2}{\lambda^2} \frac{\partial^2 V}{\partial \mu^2} - \inf_{\alpha} \left\{ \alpha \frac{\partial V}{\partial x} + L(x, \lambda, \mu, \alpha) \right\} = 0;$$

with

$$L(x, \lambda, \mu, \alpha) = \int_{-\infty}^{+\infty} \left((\Pi x - \bar{\beta})^2 + \frac{1}{2} \alpha^2 \right) f_t(\mu_t, \lambda_t^{-1})(\Pi) d\Pi$$

Note that because the moments of a Gaussian distribution are known, a simplification of the equation is possible. The following result summarizes the links between the value function and the HJB equation.

Proposition 3.1 *If the value function V is smooth : globally \mathcal{C}^1 and \mathcal{C}^2 with respect to the variables μ , then it is the only smooth solution of the following HJB equation.*

$$\delta V - x^2 \frac{\partial V}{\partial \lambda} + \frac{\sigma x^2 \mu}{\lambda} \frac{\partial V}{\partial \mu} + \frac{\sigma^2 x^2}{\lambda^2} \frac{\partial^2 V}{\partial \mu^2} - \frac{1}{2\theta} \left(\frac{\partial V}{\partial x} \right)^2 - (1-\theta) \left((x\mu - \bar{\beta})^2 - \frac{x^2}{\lambda} \right) = 0 \quad (3.6)$$

where this equation is defined on $\mathbb{R} \times \mathbb{R}_+^* \times \mathbb{R}$. Moreover, in this situation, the optimal control of the institution is given by the feedback :

$$\alpha^*(x, \lambda, \mu) = -\frac{1}{\theta} \frac{\partial V}{\partial x}(x, \lambda, \mu). \quad (3.7)$$

Thus the optimal control problem of the institution reduces to the resolution of a stationary HJB equation in dimension 4. We solve this equation numerically in section 3.3.

3.3 Numerical Calculations

In this section, we present the numerical approach we use to solve equation (3.6). The main numerical issue is to tackle the problem of the boundary conditions. Indeed the domain in which equation (3.6) is posed is not bounded. Instead of solving (3.6) on the whole domain, we are going to solve it on the domain

$$(x_{\min}, x_{\max}) \times (0, \lambda_{\max}) \times (\mu_{\min}, \mu_{\max});$$

where x_{\min} , x_{\max} , λ_{\max} , μ_{\min} and μ_{\max} are five real numbers, together with the boundary conditions

$$\begin{aligned} V(x, \lambda_{\max}, \mu) &= g(x, \mu); \\ \frac{\partial V}{\partial \mu}(x, \lambda, \mu_{\min}) &= \frac{\partial V}{\partial \mu}(x, \lambda, \mu_{\max}) = 0; \end{aligned} \quad (3.8)$$

with state constraints on $\{x_{\min}; x_{\max}\} \times (0, \lambda_{\max}) \times (\mu_{\min}, \mu_{\max})$. In (3.8) g is a prescribed function. Let us recall that by imposing state constraints on the variable x , we are excluding the controls which allow the principal to exit the domain in x_{\min} or x_{\max} ; thus the value of V at this boundary is entirely determined by the value of V on the rest of the domain. We now justify, variables by variables, why such boundary conditions can lead to a suitable approximation of the solution on the whole domain.

- **The state constraints in x :** Intuitively, if the process $(\mu_t)_{t \geq 0}$ does not spend too much time near 0, then the principal has no interest in using high value of x . Hence, if x_{\min} and x_{\max} are big enough in absolute value, then the value function should not differ much from the addition of those state constraints.
- **The Dirichlet condition in λ :** In the asymptotic $\lambda \rightarrow \infty$, the value

function V converges toward the solution U of

$$\delta U - \frac{1}{2\theta} \left(\frac{\partial U}{\partial x} \right)^2 = (1 - \theta) ((x\mu - \beta)^2)$$

on the whole domain \mathbb{R}^2 . We use this function U as g in (3.8) to provide a good approximation of the value function V when λ_{\max} is big enough.

- **The Neumann conditions in μ :** We here state with the boundary condition on μ that the mean of the belief of the principal is "reflected" at μ_{\min} and μ_{\max} . Such a choice is arbitrary but if μ_{\min} and μ_{\max} are far enough from the true value of Π , then the actual process $(\mu_t)_{t \geq 0}$ stays far from those values with high probability and thus does not alter too much the behavior of the principal.

In order to solve (3.6) together with the boundary conditions (3.8) and state constraints in x , we use a standard Godunov's scheme to discretize the HJB equation and a Newton's method to solve it numerically.

3.4 Interpretation

In this section we provide a comprehensive analysis of the learning behaviour of the principal. The goal is to convince the reader that this specification gives interesting results in equilibrium that are close to the real life setting. We start by presenting the default case scenario, and we will then provide some interesting scenarii.

Recall that the loss function that the principal is minimizing is given by

$$\mathcal{L}(X_t, \Pi, \alpha) = (1 - \theta) \underbrace{(\Pi X_t - \bar{\beta})^2}_{\text{Economic Cost}} + \theta \underbrace{\frac{1}{2} \alpha^2}_{\text{Control Cost}}$$

with Π an unknown parameter. The value ΠX_t represents the response of the agents to the macro variable X_t set by the principal through the control α . The parameter $\bar{\beta}$ represents the principal's target for the agents' response. Had the principal known Π she would have played an α to drive the macro variable to $X_t = \bar{\beta}/\Pi$. The parameter θ represents the strength associated with moving the macro variable around. A higher θ means this is harder for the institution to set a non-zero control α , therefore harder to move the macro variable X_t .

3.4.1 General case

To wrap our head around the mechanisms underlying the principal's behaviour, let's first present the most basic version. Table 3.1 presents the selected values for those

parameters.

Parameters	Symbol	Value
Target response	$\bar{\beta}$	1
Agents' type	Π	1
Cost on the control	θ	0.5

Table 3.1: Parameters

With this setting, the principal should learn that $\Pi = 1$ so we anticipate the expected value of Π according to the principal beliefs, μ , to tend toward 1. We therefore also anticipate that X_t should tend toward 1. Indeed, had the principal known $\Pi = 1$ she would play $X_t = 1$ in the limit¹⁸.

Recall that the speed of learning is directly linked to the value of X_t . This comes from the specification of the agents' response function, $\beta(\Pi, X_t) = \Pi X_t$. A low value of X_t lowers the response whatever the value of Π is. The consequence is that the institution will try to get the macro variable X_t as far from zero as possible, but of course also close to the optimum value $\bar{\beta}/\Pi$. This relates to the traditional tradeoff between exploration and exploitation. We will study how this tradeoff shows up here in depth in section 3.5.

As a benchmark, we start by analyzing the simplest case in which the macro-variable, X_t , as well as the mean of the prior distribution, μ , starts off at 1. The confidence level on the prior distribution starts at 1. In this scenario, the mean of the prior distribution is already aligned to the true value of the unknown parameter, $\Pi = 1$, and the macro-variable, X_t , is already aligned to its target value. We expect to see some learning happening through the inverse standard deviation of the institution's belief, λ . The trajectories of the relevant variables over 100 periods of time are presented in figure 3.1. For comparison purposes we set the exact same noise for all the simulations in this chapter¹⁹.

The learning process is directly impacted by the value of X_t . By Lemma 3.1, we know that the dynamic of λ_t , the confidence level of the institution's belief around the mean μ_t , is proportional to the square of the macro variable, X_t . In this case, the first signals received by the institution makes her increase the mean of the belief distribution, μ_t , and as a consequence choose a negative α to decrease the macro variable X_t . Remember that the institution loss function is given the following.

$$\mathcal{L}(X_t, \Pi, \alpha) = (\Pi X_t - \bar{\beta})^2 + \frac{1}{2}\alpha^2$$

¹⁸Since there is a cost associated with moving X_t , $\frac{1}{2}\alpha^2$, the principal would tradeoff between having a response function close to the target and moving α to quickly.

¹⁹We shall make available a comprehensive version of the code so that the reader can use this model directly.




Figure 3.1: Base case scenario $[(\mu_0, X_0, \lambda_0) = (1, 1, 1)]$

The institution is trying to keep the product ΠX_t close to $\bar{\beta} = 1$. Since she doesn't know the value of Π , she will essentially be trying to keep the product $\mu_t X_t$ close to the target²⁰. Therefore, when the institution gets a signal that increases the mean of the belief, μ_t , she decreases the macro variable, X_t . In the meantime, the confidence level of the institution's belief, λ_t , increases at a slower rate. But after a while, the mean of the belief, μ_t , starts to decrease, and as a consequence the macro variable, X_t , goes up and the confidence level's rate of increase goes up. The institution is learning faster and the value of μ_t and X_t are converging toward 1.

The value of the macro variable reacts slowly compared to the changes in the mean of the belief because there is a cost associated with moving α around. Recall that the target of the product $X_t \mu_t$ is 1. The mean of the belief goes up to 2, so at that time the optimal value of X_t would be 0.5 and the actual value of X_t goes down only to around 0.8. The same happens when the value of μ_t goes down to 0.5, the value of X_t goes up to around 1.5 instead of 2. This is a feature of our setup. The strength of the cost associated with moving α_t around reduces the volatility of the macro variable. This shouldn't be interpreted as less uncertainty around X_t ,

²⁰In section 3.5, we discuss how the traditional tradeoff between exploration and exploitation actually plays a crucial role in the principal's objective.

this is just an artificial way to avoid uncertainty (by constraining the changes in the variable). Therefore, any result is relative to a degree of constraints around the way the institution can change the macro variable. In section 3.4.3 we study the learning behaviour as a function of the strength in the cost associated with moving α_t around.

As the institution becomes more confident about her belief, the volatility of X_t , and therefore the uncertainty surrounding its value, decreases. Figure 3.2 presents the main variables' dynamics for a situation in which the institution starts with a high confidence level around her belief, more precisely $\lambda_0 = 8$. We keep the initial value of the macro variable and the mean of the prior belief to 1. The institution is very confident about her belief but she is right about it because the true value is $\Pi = 1$. Being erroneously overconfident about the value of the unknown parameter will be studied in section 3.4.2. This is just to show what happens if the institution is already pretty confident about the value. One way to think about it is that it is a way to “move forward” in time. If the institution started with $\lambda_0 = 1$ and at a later time t , after some learning, end up in a situation in which $(X_t, \lambda_t, \mu_t) = (1, 8, 1)$, the dynamics will be identical to this case.




Figure 3.2: High confidence $[(X_0, \lambda_0, \mu_0) = (1, 8, 1)]$

When the confidence is high and the expected value is aligned to the true one²¹, the target value of X_t is well defined and doesn't change much. If, in addition, the

²¹More specifically, the condition is when the signal received by the institution doesn't make her update her belief too much.

macro variable X_t is already aligned to this target value, the volatility X_t is reduced. Therefore, the expected value is not updated too much and the macro variable doesn't change much either. In this example, the value of μ_t oscillates around roughly 0.8 and 1.2, while it was roughly around 0.5 and 2 in the previous example. The value of X_t oscillates between roughly 0.9 and 1.1, while it was oscillating between roughly 0.8 and 1.5 in the previous example. Because X_t doesn't change much, the learning process happens at a roughly constant pace. The uncertainty is reduced because the institution is more confident about her belief, which has an effect similar to reducing the asymmetry of information. In our model, uncertainty comes from the lack of information of the institution, so as the institution becomes more confident, the uncertainty is reduced.

Let's now turn to the more interesting case in which the macro variable, X_t , and the mean of the belief, μ_t , starts far from their (long term) equilibrium value but are still aligned in terms of the institution's loss function. More specifically, the macro variable starts at 3, while the mean of belief starts at $1/3$, so that the quantity $\mu_t X_t$ is already aligned to the target $\bar{\beta} = 1$. However, the prior belief, μ_0 , is far from the true value $\Pi = 1$. The dynamics of the relevant variables are presented in figure 3.3. For coherence, the noise on the signal that the institution receives is the same as in the previous example.




Figure 3.3: Mean of belief far from the true value $[(\mu_0, X_0, \lambda_0) = (\frac{1}{3}, 3, 1)]$

The value of the macro variable X_t starts high, so the signal received by the institution is quite high what makes her update her belief accordingly. Because the learning pace is directly linked with the value of X_t , the institution's confidence

grows fast at the beginning. The mean of the belief μ_t also increases fast, due to the high signal received. Because the quantity $\mu_t X_t$ is already aligned with the target 1, the institution finds it convenient to decrease X_t , to keep the product as close to 1 as possible. Once again, the cost associated with moving X_t around makes the adjustment on X_t to be delayed compared to how quickly μ_t is updated. As the macro variable decreases, the learning speed, represented by λ_t , decreases. However by comparing with figure 3.1, the total learning that happened after 100 periods of time is almost the same. The total amount learned is therefore not so dependent on the initial distance between the prior belief and the true value. This happens because overall the macro variable X_t stays relatively high and the learning rate depends directly on the value of X_t .

This first analysis introduces the most basic dynamics in our model. Let's define the *economic loss* by the loss associated with the true parameter, that is $(\Pi X_t - \bar{\beta})^2$. This represents the actual loss associated with a policy X_t . The speed of learning is directly impacted by the value of the macro variable, so when the economic loss is high, but in a state in which the value of the macro variable stays high, learning happens faster. Another way to put it would be that when the policy maker is wrong about a policy, everyone is better off being in a state in which the principal can learn from her mistakes. In some state, the principal is wrong but it will be harder to extract information from the market and learning will be poor. Moreover, as the institution becomes more correctly confident about her belief, the uncertainty perceived by the agents drops. Confidence makes the institution's belief less sensitive regarding new signals, so the macro variable does not need to change much.

3.4.2 Erroneously overconfident

Our setup allows us to study how overconfidence would impact the learning behaviour. In the last section we show how learning is impacted by overconfidence when the prior is aligned to the true value. In this section we study how the learning behaviour changes when the institution is erroneously overconfident. We always keep the loss function initialized at a minimum, namely the quantity $\mu_0 X_0 = 1$. This makes the first move by the institution to depend only on updates regarding the belief through μ_t . Because of the specific form of the response function of the agents we need to differentiate 2 cases. First, when the macro variable starts high and the mean of the belief starts low, and then the opposite.

When the value of the macro variable is high, the signal observed by the institution will be relatively high, and therefore the learning process will be fast, regardless of the prior. Even if the belief is not aligned to the correct value, the institution will extract a lot of information from changes in the signal. Overconfidence results

in putting a lot of weight on the “past”, and less weight on the new signal received when updating the belief. Figure 3.4 presents the dynamic of the relevant variables for a situation in which $(\mu_0, X_0, \lambda_0) = (\frac{1}{3}, 3, 5)$.




Figure 3.4: Mean of belief far from the true value $[(\mu_0, X_0, \lambda_0) = (\frac{1}{3}, 3, 5)]$

In this case the mean of the institution belief starts off relatively low at $\mu_0 = \frac{1}{3}$, while the macro variable starts off quite high, $X_0 = 3$. The situation is similar to the one in figure 3.3, but because the institution is more confident about the “past”, the adjustments on the mean of belief are less important. In figure 3.3, after some periods of time, the mean of belief peaks at around 2, what drives (with some delay) the macro variable to around 0.5. Here, because of the overconfidence, the mean of belief peaks at around 1.1, and is a lot more stable. In the meantime, the macro variable converges smoothly toward 1. Also, because the macro variable stays mostly above 1, the learning is very efficient and at the end of the 100 periods, the learning process reaches its maximum (numerically).

It seems that being overconfident reduces the overall uncertainty and helps reaching the long term equilibrium. Being overconfident reduces the impact of a new signal compared to the history of signal received. If in conjunction to that, the macro variable stays in a “high” range, this makes the learning process faster, and less uncertainty because less updates on the mean of the belief, so less changes in the macro variable. The convergence is smooth.

On the other hand, when the macro variable starts low, the learning process is slow, and the institution’s belief has hard time converging to the long term equilibrium point. Recall that the speed of the learning process is proportional to the

square of the macro variable. Therefore a low value of X_t , will make the process of updating the mean of belief μ_t slow, so X_t does not change much and the learning process stays slow. Figure 3.5 presents the dynamics of the relevant variables.




Figure 3.5: Mean of belief far from the true value $[(\mu_0, X_0, \lambda_0) = (3, \frac{1}{3}, 5)]$

The signal perceived is so low that the learning process is poor and even after a while the variables still haven't converged to their long term value. In this case, the mean of belief is updated slowly, and after 100 periods is still above 2. As a consequence, it is not optimal for the institution to change the macro variable a lot. However, the short term volatility seems to have increased. All the short term ups and downs are due to the institution trying to increase the macro variable in order to learn more, but the next period she realises the signal is still not high enough, and the mean of belief hasn't changed that much, therefore the macro variable goes back down. By comparing with figure 3.4, in which the confidence level on the belief reaches its maximum of 20, the confidence level hardly passes 8 in this case. Uncertainty is high due to the institution being overconfident but in region in which the macro variable (hence the learning process) is slow.

Overconfidence can yield very efficient, as well as inefficient, learning process. It all depends on how the institution is overconfident. It is better off being overconfident with a high macro variable, than the reverse. With a high macro variable, overconfidence will be "compensated" by the institution extracting a lot of information from the signal. With a low macro variable, the institution is erroneously overconfident and does not learn much at the same time, yielding very inefficient

learning.

The tradeoff between exploration and exploitation is at the heart of this remark. Overconfidence modifies the optimal strategy in such a way that it makes it optimal to explore more when the macro variable is high because it will yield a high signal and hence a high amount of information extracted from the agents. When the macro variable is low, exploration is too costly and the institution is better off exploiting what she knows because exploring with different X_t won't yield much information to be revealed.

3.4.3 Impact of the control cost and the signal's noise on the learning process

In our setup, an economy can be characterised with two parameters that are the amount of noise in the signal received by the institution, and the strength of the cost associated with moving the macro variable around. The noisiness of the signal represents the ease with which the institution can extract information from the market she is meant to regulate, for instance how good the data she is collecting are and how frequent data are available. The strength associated with the cost of moving the macro variable around represents how much influence the institution has on the market she is meant to regulate. Another way to interpret this is the freedom it has to change the current policy (or the signal associated with changing the policy). For instance, in the case the policy set by institution represents a signal, if the power of the signal depends a lot on how credible the institution is, changing the policy variable around will reduce the institution's commitment perceived by the market and therefore the impact on the market for future signals regarding the policy.

3.4.3.1 Control cost

First, we study the changes in the optimal policy with respect to the cost associated with the control. The institution's loss function is composed of two terms : the *economic cost*, $(\Pi X_t - \bar{\beta})^2$, and the *control cost*, $1/2 \alpha^2$. The parameter θ represents the relative strength of the control cost. In the base case example, it was set to $\theta = 0.5$ which means the economic cost and the control cost have the same importance. A higher (lower) θ means a higher (lower) control cost over the economic cost. The learning process, as well as the efficiency of the optimal policy, is expected to be impacted by the parameter θ .

When the control cost is high, it is more expensive for the institution to change the macro variable. Regarding the exploration/exploitation tradeoff, exploration should be more expensive. Figure 3.6 presents the dynamic of the relevant variables.

As it turns out, the institution, albeit more constrained on the control, manages to keep a policy that approaches the base case scenario, see figure 3.1. The control cost does not fundamentally change the learning process²². In this case, the institution managed to keep the macro variable in mostly the same range as in the base case scenario, but through less and smoother changes. Because the macro variable stays mostly the same, it hardly affects the learning process regarding the mean and the confidence level of the belief. The institution learns almost the same amount. The only difference is that the macro variable X_t is a lot smoother.




Figure 3.6: Low influence from the Institution ($\theta = 0.8$)

Constraining the institution over her control seems to have reduced the volatility of the macro variable, therefore the uncertainty. This result suggests that when the institution is sufficiently constrained, the policy is more efficient. What happens is that the principal has to target the economy with fewer changes in the control. The institution has to be more efficient at trading off exploration versus exploitation, which she manages to do in this example. However, it might not always be the case. As the control cost increases, the volatility of the macro variable mechanically decreases, so the uncertainty perceived by the agents also decreases. This shall not be taken as better efficiency because it says nothing on the institution's learning process. It seems that short run exploration is reduced because of the high control cost, while the long term objective of the economy is still achieved. At some point,

²²Recall that the learning process is directly impacted by the value of the macro variable, X_t . Constraining the control the institution has over the macro variable can therefore affect the learning process through the macro variable.

the control cost becomes so high that it is hard for the institution to align the macro variable to its belief, and the outcome is inefficient from a social point of view. We conjecture that there exists a cut off cost, above which the institution cannot maintain sufficient economic value²³, and under which too much *exploration* is done yielding an inefficient level of uncertainty. However, our setup does not allow us to derive this cut off cost, or even to show that it actually exists.




Figure 3.7: High influence from the Institution ($\theta = 0.2$)

3.4.3.2 Signal's noise

Second, let's study the changes in the optimal policy with respect to the noisiness of the signal received by the institution. This represents different market conditions regarding the ease at which the institution can extract information. As information becomes harder to extract, we expect a poorer learning process.

Figure 3.8 presents the dynamic of the relevant variables for a noisy signal, $\sigma = 4$. The increase in the noise makes it harder to learn, and the confidence level on the belief, λ , hardly gets to 8, while in the base case scenario (cf. figure 3.1) it reaches something around 12. The noise in the signal does not directly change the way the mean of the belief, μ_t , is updated. In the beginning, the macro variable drops to around $X_t = 0.6$, because the mean of belief increases (like in the base case scenario in figure 3.1). After that, the noisiness in the signal (as well as the poor confidence level on the belief) reduces the magnitude of the macro variable changes to changes

²³By economic value, we mean the ability to drive the economy in the target state.

in the belief's mean. It then struggles to go back up, so the learning increase rate stays low and the macro variable is a lot less volatile. Coincidentally, after about 100 periods, the macro variable is close to $X_t = 1$, which is its long term equilibrium value. To be clear, this is not a feature of the model, it just happened that in this case, reducing the macro variable volatility makes it stabilized around its long term value. However, the reduced volatility of the macro variable is a consequence of a signal with high noise.




Figure 3.8: High noise in the signal ($\sigma = 4$)

For completion purposes, figure 3.9 presents the dynamics of the relevant variables when the signal's noise is low, $\sigma = \frac{1}{3}$. With a low signal, the dynamics of the belief's mean, μ , is close to the base case scenario, as well as the macro variable. Therefore the learning rate is almost the same. There might be a slight difference in overall uncertainty but not meaningful enough to be discussed properly.

3.5 Exploration vs Exploitation

The traditional view of learning around the concepts of exploitation and exploration can be used here. Indeed, the *economic cost* could be interpreted as the *exploitation-part* of the cost function, and the *control cost* as its *exploration-part*. Recall that the institution is minimizing the expected loss function which is defined by the following.

$$L(x, y, z, \lambda, \alpha) = \int_{-\infty}^{+\infty} \left((1 - \theta)(\Pi x - \bar{\beta})^2 + \theta \frac{1}{2} \alpha^2 \right) f_t(\mu_t, \lambda_t^{-1})(\Pi) d\Pi$$




Figure 3.9: Low noise in the signal ($\sigma = \frac{1}{3}$)

with x the macro variable, α the control, y the signal received by the institution, λ the inverse of the belief's standard deviation, μ the belief's mean, and $z = \mu\lambda$.

The *Economic cost* represents the main objective of the principal, that is to drive the economy to a specific target, $\bar{\beta}$. It is therefore associated with the cost of exploiting what she knows regarding the unknown, namely the belief. If the institution plays a control α that makes the macro variable X far from its optimum value given μ , the belief's mean, the expected cost function will be higher.

The *Control cost* is the cost associated with moving around the macro variable. We have seen that if we increase the control cost there is a delay between updates in the mean of belief and the time the macro variable reaches the optimal value conditional on the mean of belief. But since it is also associated with any changes in X_t , it can also represent the cost of exploration.

If there were no learning happening she would have driven the macro variable to the *optimal instantaneous value*. The institution does not value any deviation from this value since it won't yield any changes in the belief. We denote by \tilde{X}_t^* this *optimal instantaneous value*.

$$\tilde{X}_t^* = \frac{\bar{\beta}}{\mu_t}$$

The dynamics of the macro variable will depend on the initial conditions. If the initial conditions are not already aligned to $X_0 = \bar{\beta}/\mu_0$, then whenever $\theta > 0$, it won't be possible for the institution to play its optimal \tilde{X}_t^* straightaway because of the quadratic control cost. But of course, in the long run the macro variable

will converge to the *optimal instantaneous value*, which won't change because there is no learning here, so no updates in μ_t . The control cost essentially introduces a delay between changes in μ_t , and changes in X_t . This delay is inherent to the cost structure of the institution whether some learning is happening or not.

With learning, the institution might find it optimal to play an $X_t^* \neq \tilde{X}_t^*$, even if $X_{t-1}^* = \tilde{X}_{t-1}^*$ and no new signal occurred. Indeed, by increasing the macro variable for instance, the institution knows that she will be able to learn more and therefore to increase λ which will reduce the expected loss function in later periods. This is the exploration part of the cost function. Figure 3.10 illustrates the delay between the *optimal instantaneous value* and the optimal value. The blue line corresponds to the optimal value given by our model, and the orange line corresponds the *instantaneous optimal value*, that is $\tilde{X}_t^* = \bar{\beta}/\mu_t$. The macro variable that would have been optimal if there were no control cost drops faster than the actual optimal value. The macro variable converges to \tilde{X}_t^* but slowly. Once again the uncertainty perceived by the agents, the macro variable's volatility, is reduced because of the constraint on the institution's control. This should not be taken as something positive, the real uncertainty obviously depends on the ability of the institution to change the macro variable. In a more general framework, uncertainty should be a function of the institution's control cost, not just the actual volatility which can be constrained exogenously by this cost. Emphasis should also be put on the fact that the agents will perceive changes in the macro variable more (less) if the institution is highly (poorly) constrained in her policy.




Figure 3.10: Optimal instantaneous value vs Optimal value for the Macro-variable

Therefore, if we abstract from the short run delay between each updates of the beliefs and changes in the macro variable, the control cost is associated with the exploration part of the cost function.

In this model, the relative strength of the constraint on exploration and exploitation is determined by the parameter θ . When $\theta = 1$, the institution cannot explore at all, and must play an $\alpha_t^* = 0$. As θ decreases, the constraint on exploration is relaxed and the institution can start to play around with α . When $\theta = 0$, there is no control cost and the institution is free to play whatever α minimizes the cost function.

3.6 Conclusion

This chapter develops a general framework for analyzing institutional policy making with uncertainty about the underlying economy. We build a model in which an institution facing an economy must drive the economy to an unknown state. The institution will influence a macro variable and continuously receive signals from the market which she uses to update her belief about the unknown agents' type. We numerically solve the equilibrium strategy and provide a deep analysis of the solution.

In our model learning is efficient. As a consequence, the long run uncertainty perceived by agents tends to decrease over time. However short term uncertainty arises in a more complex way. The institution uses policy changes to learn from the market so depending on the learning quality the institution will use short term changes of the policy to improve learning. The learning ability of the institution comes from the internal state of the economy²⁴ but can also be influenced by the magnitude of the control cost in the loss function. Therefore, if policy changes are not too costly and/or the economy is in a bad state regarding learning, the institution will find it optimal to use short term policy changes. We find that overconfidence²⁵ exacerbates this result.

We also find that noise in the signal the institution is receiving creates an economic loss. Indeed, the institution does not see much value in changing the policy because the signal is going to be noisy anyway. Therefore, uncertainty is reduced, but this is associated with an economic loss for agents because learning is very inefficient and therefore the economy can stay in a bad state for a while. The quality of the signal can be exogenously enhanced by increasing the frequency/quality of reporting from the market participants.

²⁴In our model, the institution learns from the value of the macro variable, but it controls the change in this macro variable. Therefore, if the institution does nothing, but the macro variable is in a "learning area", the institution will still learn a lot. The reverse is also true.

²⁵Overconfidence in our model is defined as a narrower initial belief distribution.

In our model, no one likes or dislikes uncertainty per se, this is just a consequence of the learning process. Uncertainty arises in equilibrium as a way to learn for the institution. An important issue regarding policy design is the credibility associated with policy uncertainty. If the regulator changes her policy too often it can be perceived as a lack of commitment, furthermore reducing the influence of the institution. With less influence, learning will be even poorer, and this can be persistent for a while²⁶. In our model the institution (as well as the agents) does not take into account any effect of credibility and the impact of a given policy is always the same. We just provide a way to think about uncertainty and analyze some of its causes.

Our model allows us to distinguish between the *exploitation* and the *exploration* part of the institution's behaviour. The *exploration* component represents the ease with which the institution can change the current policy. But recall that actual policy changes and signals about the future of policy are equivalent. So in terms of signals, the *exploration* component can be associated with the credibility of the institutions. With less credibility, the market won't react much to signals. A high exploration cost corresponds to the institution having lost a lot of credibility and lacking the ability to have a meaningful influence on the market. Therefore in our model, being constrained about the potential actions or being less credible from the market point of view have the exact same effect. We show that introducing some *exploration* cost artificially reduces uncertainty but we are able to distinguish between the decrease in uncertainty that comes from the *exploration* cost and the uncertainty that comes from the underlying learning. Our results suggest that as we increase the *exploration cost* the economic loss decreases until a cutoff cost above which the institution is too constrained and the economic loss increases.

This chapter provides an extended analysis of the learning process of any kind of institution. However, we do not provide any numerical *real-world* calibrations but the framework we develop should be general enough so that it applies to a lot of different applications. Even if the solution to the HJB equation is only numerical, and hard to characterise, we think that it could be possible to estimate the level of constraints of a given institution for instance, or even maybe make some forecast regarding future uncertainty.

Our work could also be applied to institutional design. Indeed, any institution is designed by an *Act* or a contract between one or several market participants and/or one or several governments. We show that the institution should be constrained regarding potential actions, and that there exists an optimal level of constraints. Constraining further the institution would increase the overall *expected* economic

²⁶Recall that in the long run, learning is efficient so uncertainty will be reduced eventually but can be very slow.

loss, as *exploration* and therefore policy changes become too costly. Without constraints, the exploration cost will be too low and the resulting uncertainty would be high.

3.7 Appendix

3.7.1 Proofs

Proof of Lemma 3.1 :

We denote by ν_t a un-normalized density of the belief f_t of the principal. We use the result of Zakai (1969) stating that this un-normalized density follows the so called Zakai equation which is a stochastic partial differential equation.

$$d\nu_t = \nu_t \Pi X_t dY_t.$$

Thus, ν_t satisfies :

$$\nu_t = \exp \left(\Pi \int_0^t X_s dY_s - \frac{1}{2} \Pi^2 \int_0^t X_s^2 ds \right) \nu_0.$$

By the definition of λ_0 , μ_0 , we have that $\nu_0 = \mathcal{N}(\mu_0, \lambda_0^{-1})$. Thus,

$$\nu_t = \exp \left(\Pi \int_0^t X_s dY_s - \frac{1}{2} \Pi^2 \int_0^t X_s^2 ds \right) \exp \left(-\frac{1}{2} (\Pi - \mu_0)^2 \lambda_0 \right); \quad (3.9)$$

$$\propto \exp \left[-\frac{1}{2} \left(\Pi - \frac{\int_0^t X_s dY_s + \mu_0 \lambda_0}{\int_0^t X_s^2 ds + \lambda_0} \right)^2 \left(\int_0^t X_s^2 ds + \lambda_0 \right) \right]. \quad (3.10)$$

This final expression leads to the introduction of the variable $z_t = \mu_t \lambda_t$. Because $\nu_t \propto \mathcal{N}(\mu_t, \lambda_t^{-1})$, we read on (3.10) that

$$\begin{cases} d\lambda_t = X_t^2 dt; \\ dz_t = X_t dY_t. \end{cases}$$

When then compute with Itô's lemma that

$$\begin{cases} d\lambda_t &= X_t^2 dt; \\ d\mu_t &= -\mu_t \frac{X_t^2}{\lambda_t} dt + \frac{X_t}{\lambda_t} dY_t. \end{cases}$$

Q.E.D

Proof of Proposition 3.1 :

Recall that the cost L is given by :

$$L(x, \lambda, \mu, \alpha) = \int_{-\infty}^{+\infty} \mathcal{L}(X_t, \Pi, \alpha_t) f(\mu, \lambda^{-1})(\Pi) d\Pi;$$

Hence rearranging and using the fact that $f_t(\mu_t, \lambda_t^{-1})$ is normalized we deduce :

$$\begin{aligned} L(x, \lambda, \mu, \alpha) &= (1 - \theta)\bar{\beta}^2 + \frac{\theta}{2}\alpha^2 + (1 - \theta)x^2 \int_{-\infty}^{+\infty} \Pi^2 f(\mu, \lambda^{-1})(\Pi) d\Pi \\ &\quad - 2(1 - \theta)\bar{\beta}x \int_{-\infty}^{+\infty} \Pi f(\mu, \lambda^{-1})(\Pi) d\Pi; \\ L(x, \lambda, \mu, \alpha) &= \frac{\theta}{2}\alpha^2 + (1 - \theta) \left((x\mu - \bar{\beta})^2 - \frac{x^2}{\lambda} \right) \end{aligned} \quad (3.11)$$

Recall that the HJB equation of the optimal control problem is given by

$$\delta V - x^2 \frac{\partial V}{\partial \lambda} + \frac{\sigma x^2 \mu}{\lambda} \frac{\partial V}{\partial \mu} + \frac{\sigma^2 x^2}{\lambda^2} \frac{\partial^2 V}{\partial \mu^2} + \frac{1}{2\theta} \left(\frac{\partial V}{\partial x} \right)^2 - \inf_{\alpha} \left\{ \alpha \frac{\partial V}{\partial x} + L(x, \lambda, \mu, \alpha) \right\} = 0$$

Solving for the minimization and plugging (3.11) in, yields the following HJB equation.

$$\delta V - x^2 \frac{\partial V}{\partial \lambda} + \frac{\sigma x^2 \mu}{\lambda} \frac{\partial V}{\partial \mu} + \frac{\sigma^2 x^2}{\lambda^2} \frac{\partial^2 V}{\partial \mu^2} + \frac{1}{2\theta} \left(\frac{\partial V}{\partial x} \right)^2 = (1 - \theta) \left((x\mu - \beta)^2 - \frac{x^2}{\lambda} \right)$$

The fact that V satisfies the HJB equation if it is smooth is a straightforward consequence of the dynamic programming principle. Because this HJB equation is elliptic, it is a direct application of the comparison principle for viscosity solutions that there is at most one solution to this equation. The optimal control $\alpha^*(x, \lambda, \mu)$ is given by (3.7) as a consequence of a verification argument.

$$\alpha^*(x, \lambda, \mu) = -\frac{1}{\theta} \frac{\partial V}{\partial x}(x, \lambda, \mu).$$

Q.E.D

Chapter 4

Incentives on the Lightning Network : A Blockchain-based Payment Network

Abstract :

The Lightning Network is a decentralized payment network built on top of a blockchain, in which intermediary nodes provide a trustless routing service for end users. We provide an overview of the current state of the network and show that it can be well approximated by a scale free generative model with a fitness parameter, which suggests that nodes behave strategically on the network. Those strategic interactions between nodes can be described by a Bertrand competition model with capacity constraints. We show that there is a unique equilibrium in which a centralized network is never optimal, and the routing fee is strictly greater than the marginal cost. When nodes are heterogeneous in their opportunity cost of capital only, the equilibrium network structure can match the current state of the network.

4.1 Introduction

During the past 10 years, blockchains have gone from a short post on a cryptography mailing list to a global industry with a total market capitalisation of roughly USD 275 billions¹. Introduced by Nakamoto (2008), a blockchain is a *communication protocol, a distributed database, and a consensus algorithm such that the communication protocol allows a group of untrusted agents to reach consensus over the content of the distributed database*. To each blockchain, there is an associated cryptocurrency, that is needed to provide the necessary incentives to the maintainers (i.e. miners). By extension, the database itself is also called the blockchain and it usually contains the list of transactions from the very first transaction in the denominated cryptocurrency. In the blockchain world a transaction can refer to *i*) a *pure* transaction (agent A sends some value to agent B), or *ii*) a *smart-contract* which is essentially a piece of software that can interact (send/receive value) with other agents (or other smart-contracts) the way it was meant at the contract creation². For security reasons, the blockchain is slow and costly to use³. To overcome this issue, Poon and Dryja (2016) show that with the right combination of smart contracts, it is possible to create a payment network, which they called the *Lightning Network*, in which participants can transact among each other at a cost supposed to be a lot lower. Intermediary nodes on the network will charge a fee to route a transaction and in this chapter we analyze the current state of the Lightning Network and investigate the incentives of nodes.

The term “blockchain” has been used to mean very different concepts. In this chapter, we use the term “blockchain” as first intended by Nakamoto (2008), that is a “proof-of-work chain”. *Proof-of-work* refers to the consensus algorithm used on the Bitcoin blockchain. Intuitively, *miners* need to brute force the solution of a hash based puzzle, that is designed to take a constant amount of time to be solved. If it is solved too quickly (slowly), the difficulty increases (decreases). The system does not rely on any trusted agent because miners prove their work by showing the solution of the puzzle, and anyone can validate this solution without the need to trust the miner who solved it. There are other consensus algorithms but, as of mid-2019, none of them is implemented at the scale of the bitcoin blockchain. There are other concepts named “blockchain” like “permissioned blockchains” that function like a proof-of-work chain but they rely on a group of trusted agents so there are very

¹Total market capitalisation of all cryptocurrencies as observed on coinmarketcap.com on May 28th, 2019.

²Otherwise stated, once a contract is added to the blockchain, it cannot be removed or changed. It *will* execute as long as the blockchain is still growing.

³The speed is artificially lowered to reduce the probability of forks. See Biais, Bisiere, Bouvard and Casamatta (2019) for an analysis of forks in the blockchain. The monetary cost of using a blockchain has to stay high in order to incentivize the miners.

different from a trustless service. The real technological advance is the ability to generate consensus among a group of untrusted agents, so any other kind of so-called blockchain is outside the scope of this chapter.

The Lightning Network is a network of bilateral *payment channels* allowing users to send/receive value (to/from anyone in the network) without all transactions being stored on the blockchain. In financial terms, this is a payment system. Users escrow some amount on the blockchain (that can be used only on the Lightning Network), then they send/receive value and when they need value out of the Lightning Network, the blockchain pays them the *net amount* given the history of bilateral transactions. At the blockchain level, the Lightning Network allows users to perform *net-settlement*, but there is an important difference with the Lightning Network and traditional payment systems. We sometimes refer to the blockchain as a *trustless* system, meaning users don't need to trust each other, as a consequence, there is no counterparty risk in the blockchain world⁴. The Lightning Network itself is very close to a *Real Time Gross Settlement* system. Because using the blockchain is costly, the Lightning Network allows users to move transactions to a less costly payment system, while keeping the absence of counterparty risk associated to a gross settlement system.

The contribution of this chapter is twofold. First, we provide an overview of the current state of the Lightning Network. As the time of writing this chapter, there are 2,865 nodes and 36,982 payment channels among them. Although the network is at its premises, we show that it already has good properties. From a static point of view, we show that the network is well interconnected in terms of clustering and centrality measures. A well interconnected network is necessary in order to be used as a global payment system. Indeed, if user A wants to pay user B, it needs to either *i*) have a direct payment channel with user B, or *ii*) find a route of payment channels within the network⁵. From a dynamic point of view, we estimate a generative model and we show that the node-specific component is an important determinant of the probability of getting a new edge (i.e. a payment channel with another node). This means that nodes are not connecting at random when they enter the network. The second contribution of this chapter is to build a model to analyze nodes' incentives on the Lightning Network. Since nodes do not connect at random, we are interested in understanding how they decide which node to connect to, and what are the implications regarding the network structure.

To be able to route transactions, a node needs to have enough funds on his payment channels. If Alice wants to pay 1 unit to Bob through Charlie, Alice needs

⁴More precisely, there is no counterparty risk as long as the blockchain is well functioning. Users will systematically ask for proof of ownership of the amount they are about to receive.

⁵Note that no node can steal any funds when they are routing transactions.

to have at least 1 unit on her payment channel with Charlie and Charlie needs to have at least 1 unit in her payment channel with Bob. Payment channels are independent from one another, so if Charlie does not have enough funds on her payment channel with Bob, she won't be able to route Alice's payment. The amount of funds on a payment channel is called the capacity. When routing a transaction, Charlie will be able to collect a fee, paid by Alice, the original sender of the payment.

Our model features a set of profit maximizing nodes that are competing in a two-stages model to route a fixed number of transactions. In the first stage, nodes enter the routing market, each opens at least two payment channels and assign them a capacity. During the second stage, given the capacity and the graph of payment channels, nodes compete on price to attract the next transaction. A buyer on the product market will pick the cheapest available route (with enough capacity) to pay the seller of a good or service. We assume the marginal cost of routing a transaction to be zero⁶.

In order to open a payment channel, nodes must send a transaction to the blockchain and therefore incur the associated cost. Also, when all the capacity of a payment channel is locked toward one side, say Charlie's side, Alice cannot pay Charlie anymore with this payment channel. There are then two solutions, Alice can either *i*) send a *refund* transaction to the blockchain (with the associated costs), or *ii*) if the network's graph allows it, she can pay herself within the Lightning Network and therefore rebalance her channels' capacity (for instance, Alice pays herself through David and Charlie, so that Charlie has to give back some capacity to Alice on the payment channel that was initially blocked). Note that in both cases, it is costly to maintain a well-funded channel, either because of the blockchain cost or the Lightning Network cost. It is supposed to be a lot cheaper on the Lightning Network though, so nodes will try to always have potential route to loop back to them in order to rebalance payment channels. We call this cost, *the cost of maintaining a given interconnectedness level*.

In our model, nodes are subjected to an opportunity cost of capital for the funds they lock inside payment channels. This is meant to represent the usual missed opportunities but also a cost linked to security. On the Lightning Network, nodes need to advertise their *ip-address*, as well as the capacity of each channel, in order for end users to be able to select them as part of the optimal route⁷. As a direct consequence, a node with a lot of funds locked inside payment channels will become very visible, most importantly to potential bad actors. Otherwise stated, it is likely to become the target of hackers, and will either have to incur additional security

⁶In practice, this is essentially zero, as it consists only of sending a few data packets over the internet.

⁷It is not mandatory for nodes to advertise those, but then they would be excluded from the public network, and no one could use them as intermediary node.

costs or be subject to the risk of loss. For this reason, we assumed the opportunity cost of capital to be quadratic.

We assume no barriers to entry in the routing market. This assumption is quite relevant for our use case, because the actual Lightning Network is completely open. This means that anyone can download the software, run it, and start to create payment channels with other nodes, and this does not require anything other than a computer and some funds. In particular, no barriers to entry means that if any outside agent sees a profitable opportunity, it will jump into the market and make a profit.

A key aspect of our model is that we identify 3 factors influencing the likelihood of a node to be on the optimal route. First, there has to exist an actual route (a set of payment channels) connecting the end-buyer with the end-seller through this node. Second, the route needs to have a sufficient capacity on all intermediary payment channels so that the transaction can take this route. And lastly, this route has to be the cheapest one. Among several routes that satisfy the first two conditions, the end-buyer would pick the cheapest one. In the first stage, nodes will individually invest in costly interconnectedness, and channels funding, to increase the probability that they route the next transaction. The probability of being selected as part of the optimal route is assumed to follow an exponential distribution. In the second stage, nodes compete on the third factor to attract transactions, the price. Note that this third factor is also costly because as a node decreases its price, to increase the probability of being selected, it lowers its revenue if it is selected.

We show that when the cost of maintaining a given interconnectedness level is linear, and the opportunity cost of capital is quadratic, there is a unique equilibrium for the two stages game. In this equilibrium, all nodes have the same optimal strategies when they are homogeneous, and the routing fee is strictly greater than zero, the marginal cost. Indeed, we show that there is an endogenous participation constraint for the nodes, that is if the routing fee is anticipated to be too low, nodes prefer not to participate in the routing market, and the Lightning Network collapses. Also, because of the price competition in the second stage, each node makes zero profit in equilibrium.

Our results indicate that when nodes are heterogeneous with respect to their opportunity cost of capital, the equilibrium is still unique but the heterogeneity is reflected in the level of interconnectedness and the amount of funds locked in payment channels. In particular, a high (low) opportunity cost of capital is associated with low (high) interconnection in the network, and low (high) funds in payment channels. Because we assumed there are no barriers to entry, no node makes a profit in equilibrium, even when they are heterogeneous.

Our model generates the equilibrium relationship between the opportunity cost

of capital distribution, the distribution of centrality/clustering in the network, and the distribution of channels' capacity. This has important implications regarding the structure of the network. In particular, we show that a lower bound on the distribution for the opportunity cost of capital will produce an upper bound on node's centrality as well as on the amount a node will lock in the Lightning Network. The upper bound on node's centrality predicts that a completely centralized network is never optimal in equilibrium. Moreover, our model is able to consistently generate network topology given a distribution for the opportunity cost of capital.

It should be pointed out that the Lightning Network is still in an early stage of development. A lot of nodes operate on the Lightning Network to test the technology and the overall demand is relatively low for a global payment network⁸. Also, the fees are close to zero and no one is really trying to profit from routing transactions. Therefore, we should not expect our model to be tested on real data before some strategic agents enter the network, and the demand for this payment system is higher.

This chapter contributes to the recent literature about blockchains. Although being an analysis of a so called second layer protocol, it relies on a blockchain to operate. The motivation behind Poon and Dryja (2016) is that the way Nakamoto (2008) designed the blockchain does not allow for a lot of scaling on the blockchain itself. Easley, O'Hara and Basu (2017) study how fees and delays are used as incentives for miners to operate the blockchain. Those *on-chain* costs are necessary for the security of the blockchain. As a consequence, small or micro payments are likely to be unfeasible if they have to be stored on a blockchain⁹. The Lightning Network, as presented by Poon and Dryja (2016) is called an *off-chain* scaling solution, and this chapter is the first attempt to analyze such a solution from an economic point of view.

The first part of this chapter provides a description of the current state of the Lightning Network. In line with Rohrer, Malliaris and Tschorsch (2019) and Seres, Gulyás, Nagy and Burcsi (2019) we find evidence of a well connected graph from a static point of view. From a dynamic point of view, our results also support a *scale-free* classification of the graph. In order to properly estimate the power law degree distribution resulting from a scale-free model, we use a dataset of channel openings and closings from mid-January 2018 (the early days of the Lightning Network) until mid-August 2018. Both paper cited above find a different coefficient for the power law exponent because they use a different, more recent period. However, contrary to those papers we also estimate the fitness distribution with a non parametric

⁸As we will see, the total capacity of the network is only around EUR 5 million and it is not completely used.

⁹The routing fee would be greater than the volume transacted.

estimation. The fitness distribution from Bianconi and Barabasi (2001) represents the idiosyncratic strength of each node. We find evidence of a significant non-constant fitness distribution which suggests that nodes behave in a strategic way. They do not simply connect at random to other nodes, they specifically choose who to connect to.

This chapter also contributes to the payment systems literature. Rochet and Tirole (1996) proposes an overview of the payment system mechanisms, Freixas and Parigi (1998) and Kahn and Roberds (1998) study the difference between a *net* payment system and a *gross* payment system. They show that a payment system with *gross* settlement is generally preferred as it lowers the risk of default and contagion. Leitner (2005) shows that the risk of contagion might provide enough incentives for participants to bail out one another, even if transactions are net-settled. Usually in payment network literature the issue of net settlement is associated with risk and losses for participants. For instance, if the network nets all transactions at the end of the day, some participant might be unable to make the required payment. If, on the other hand, all transactions are settled separately and at the time they are made, there is no risk of loss. The problem is that the total amount exchanged is higher and therefore a *gross* payment system requires more money to operate. The Lightning Network, coupled with a blockchain, is the first kind of network that is able to take advantage of both *net* and *gross* payment system. From the blockchain point of view, the Lightning Network is a netting system. Participants can exchange value on this network, and when they want their funds back, they can exit the network and the blockchain pays them the net amount given the history of Lightning transactions. The transaction volume can be dramatically decreased with this system. However, because everything sits on a blockchain, there are no counterparty risk. Indeed, on the Lightning Network itself, everything is *gross* settled because, once Alice sent Bob 1 unit on the Lightning Network, she cannot use this unit to pay someone else. The unit is effectively transferred right away but the blockchain does not need to know about it. As long as the blockchain is secure, all the funds on the Lightning Network are at no risk¹⁰. This chapter is the first attempt to analyze this particular network that separates the risk of the funds from the payment structure itself.

Our work also relates to the industrial organization literature. Indeed, nodes on the Lightning Network are competing on the routing market in Bertrand style, that is competition on price. However, as explained above, nodes are constrained on the quantity of transactions they can route by *i*) their interconnectedness and *ii*) the ca-

¹⁰There could actually be an issue regarding adversarial close of payment channels, but there are some ways around it. See “Unilateral closing of channel” in section 4.2 and a discussion on this point in section 4.7.

capacity of their payment channels. A number of papers have analyzed such games of Bertrand competition with capacity constraints like Kreps and Scheinkman (1983), Davidson and Deneckere (1986), or Moreno and Ubeda (2006). The main findings of this literature is that when firms are constrained on their capacity while they compete on price, the equilibrium is actually the Cournot outcome. Otherwise stated, the equilibrium is the same as if firms were competing on quantity, and the price is higher than the Bertrand price (i.e. the marginal cost). The Lightning Network is different from the usual assumptions made when analyzing traditional product market competition. Indeed, nodes do not, strictly speaking, provide a homogeneous service, because depending on their position in the graph and the number of nodes they are connected to, they will be more or less useful for a given transaction. Also, we do not assume that the demand is expressed in terms of the equilibrium quantity. We argue that the quantity demanded by end users does not depend on the routing price as long as it is cheaper than an alternative payment system¹¹. Our model features a two stages game very similar to Kreps and Scheinkman (1983), apart from the previous assumptions. Because we are less strict on the assumption we make, we do not have a Cournot outcome in equilibrium, but we are able to show that the equilibrium price is higher than the Bertrand price, which is zero (the marginal cost).

The remaining of this chapter is organized as follows. Section 4.2 provides some generalities about the Lightning Network. Section 4.3 describes the current state of the network. Section 4.4 presents the main model with a homogeneous cohort of agents while section 4.5 introduces heterogeneity in the nodes' opportunity cost of capital. Section 4.6 provides a comparison between our model and the current state of the network. Section 4.7 sheds light on future potential research. Section 4.8 concludes.

4.2 General Facts about the Lightning Network

This section presents an overview of the technical aspect of the Lightning Network and how it is implemented. We assume basic knowledge of blockchain and cryptocurrencies.

The Lightning Network is a communication protocol, built on top of a blockchain, that allow users to send and receive value without the need for the underlying blockchain to store every transaction. Because the Lightning Network needs an

¹¹A payment system has no value in itself. The value users get is from the physical product they are buying for instance. They only use the Lightning Network as a tool to pay for what they want. Therefore, it is important to note that everything is conditional on the outside option, that is to pay with another payment system. As long as the price on the Lightning Network is below the outside option, users don't care about the price itself, and the demand is constant.

underlying blockchain to operate, it is also called a second-layer protocol, with the blockchain being the first-layer protocol.

Value When we use the term value transfer and value storage for instance, it means value denominated in terms of a blockchain cryptocurrency. Whether a cryptocurrency constitutes money is outside the scope of this chapter and we assume it does.

Payment channel Payment channels are bilateral contracts in which both parties lock some amount $x \geq 0$ on the blockchain, of which the sum is called the *capacity*. To “open” the channel, they will have to broadcast a transaction (a contract) to the blockchain, hence incurring fees and delays inherent to blockchains¹², and they would broadcast another transaction to “close” the channel. Any transaction that changes the allocation of the channel capacity among both parties can be made bilaterally without the need to broadcast anything to the blockchain nodes. These transactions are fast (instantaneous) and a lot cheaper than transactions on the blockchain (the point of this chapter is actually to study the cost structure on the Lightning Network).

Payment channels network A network of payment channels can be used to route transactions using intermediary nodes. Despite them being real intermediaries, thanks to some cryptographic tools, no intermediaries have the ability to steal any funds. Moreover, the routing algorithm currently implemented is based on *Onion Routing*. Onion Routing is the same routing algorithm used on the TOR Network for instance. With this algorithm, minimum information for the intermediaries is achieved: any intermediate node just knows that he received a transaction from address `0x.....4e12a4` and that it has to forward it to address `0x.....5834bc4` for instance, it does not even know whether any of those two nodes are the actual *sender/recipient* of the original transaction¹³. The sender does the routing selection based on the information it has from the network nodes. The most known implementation of such a network is the Lightning Network¹⁴, and is currently (as of late 2018) being implemented on top of the bitcoin blockchain¹⁵. Otherwise stated, any mention of a payment channels network will refer to the Lightning Network on the Bitcoin blockchain.

¹²On the bitcoin blockchain for instance the block time interval is 10 minutes on average and one often expects around 5/6 confirmation blocks which then average to 1 hour. The fees are set by market forces using supply and demand and can be pretty high due to the limited space in each block. For instance on the bitcoin blockchain each block is limited to 1Mb.

¹³Also, for long payment routes, a given node does not know its position within the route.

¹⁴See Poon and Dryja (2016)

¹⁵There is another smaller implementation on the Litecoin blockchain.

Unilateral closing of channels If both parties want to close the channel, and if they agree on the history of bilateral transactions while the channel was open, they can each broadcast a closing transaction. The blockchain will “verify” that both transactions are correct and allocate the corresponding amount to each party. However, it is also possible to close the channel unilaterally. There has to be a way to do so because nothing prevents one party to disconnect from the network at any point in time, and the other party must still be able to retrieve his funds. However, unilateral close also means that one party might try to close the channel with a previous history of transactions. To prevent that, the current Lightning Network specifications work as follow. When a party tries to unilaterally close a channel, there is a 14-day window before the channel is actually closed by the blockchain itself. Within this window, if the other party broadcasts a more recent history to the blockchain¹⁶, then the “cheater” loses all his funds on this channel, and the other party gets access to the full channel capacity. This is meant as an incentive not to unilaterally close channels with previous history. But the consequence is that if an agent has an active channel, he must monitor the network to make sure the other party does not cheat, and he cannot disconnect for more than 14 days, the safety window. Although not part of the current protocol specifications, it is also possible to delegate this monitoring to someone else, a *watchtower*, that would be paid to monitor channels for other agents. Of course, this is also achieved in a trustless fashion. Section 4.7 develops this idea.

4.3 Graph Analysis

Several implementations of the Lightning Network have been released in late 2017, and have been since then fully functional and growing. This section gives an overview of the current network structure.

4.3.1 Data

Since information about the network is needed by a node for finding a route across multiple channels, the state of the (public¹⁷) Lightning Network is available to any node. This information consists in a list of public nodes (public key, network address) and a list of edges or payment channels (source, destination, capacity, fees). Table 4.1 presents an overview of this dataset as of May, 2nd 2019.

Nodes with less than 2 edges are not very useful for routing as they essentially

¹⁶Thanks to digital signatures, the blockchain will be able to verify that the more recent history is actually more recent.

¹⁷Anyone can open a channel and not all channels are made public. Some nodes may be willing to run a parallel private network. We focus on the public network, available for routing.

	Full network	Nodes with more than 2 edges
Number of nodes	4,298	2,865
Number of edges	38,929	36,982
Total capacity	1,063.23 BTC	1,045.83 BTC

Table 4.1: State of the Lightning Network as of early May 2019

connect two nodes. We distinguish between them and the full network. Most analyses below will be on the subset of nodes than can be used for routing. Note that as of early May 2019 $1BTC \simeq 5,000EUR$, so this corresponds to a total capacity of around EUR 5 million.

In addition to this dataset, we have been able to get the history of opening and closing of payment channels almost from the first release (January 16th, 2018) until late August 2018. This has been privately saved by an individual who gave us access to this dataset. All channels opening and closing transactions are broadcasted to the blockchain but, from an external point of view, they cannot be differentiated from other transactions. Therefore it is easy to check if a channel-related transaction is indeed included in a block, but we cannot make sure that all channel-related transactions are in this dataset. Table 4.2 presents some characteristics of this dataset.

	Value
Number of channel opened	35,809
Number of channel closed	24,368

Table 4.2: History of channels opening and closing from mid-January 2018 until late August 2018

This dataset will be used to analyze dynamic properties of the Lightning Network. In generative graph theory, the graph is obtained by starting with an initial set of nodes and adding successive node(s) and edges each round. In cases like *social network* or *citation network* this is generally not an issue but on the Lightning Network, edges (payment channels) can also be removed. Since the network is in a very early stage at the time of writing, some people are conducting some experiments by opening/closing channels. Table 4.3 separates channels by how long they stayed open.

	Value	Percentage
Number of channels closed	24,368	100%
Number of channels opened for less than 1 week	11,382	46,71%
Number of channels opened for less than 1 month	19,123	78,48%

Table 4.3: Dealing with closed channels

We see that around 78% of closed channels were opened for less than a month. At this stage it is likely that those channels were opened for testing purposes. For the remaining of this section we simply exclude the closed channels.

4.3.2 Static properties

A number of measures are used to describe graphs, but we can start by visualizing the current graph and qualitatively inspect its structure. Figure 4.1 and figure 4.2 show the state of the Lightning Network as of early May 2019. We can observe that a few nodes seem to be highly connected but most nodes seem to be well connected. On figure 4.2, we can clearly see all nodes with a single or a couple of payment channels surrounding the main network in the middle, the one that can be used for routing. For the remaining of this subsection, we remove nodes with 2 or less payment channels. Indeed, such nodes do not really bring any value regarding routing ability, therefore to measure and discuss some static properties of the network, we prefer to concentrate on the “useful” part of the network.



Figure 4.1: Vizualization of the Lightning Network (1) (Downloaded from <https://explorer.acinq.co> on May 4th, 2019)

4.3.2.1 Centrality

Centrality is a measure of a node’s importance and can be expressed by different indicators. In this section we focus on degree centrality and the betweenness-centrality as defined by Freeman (1977).

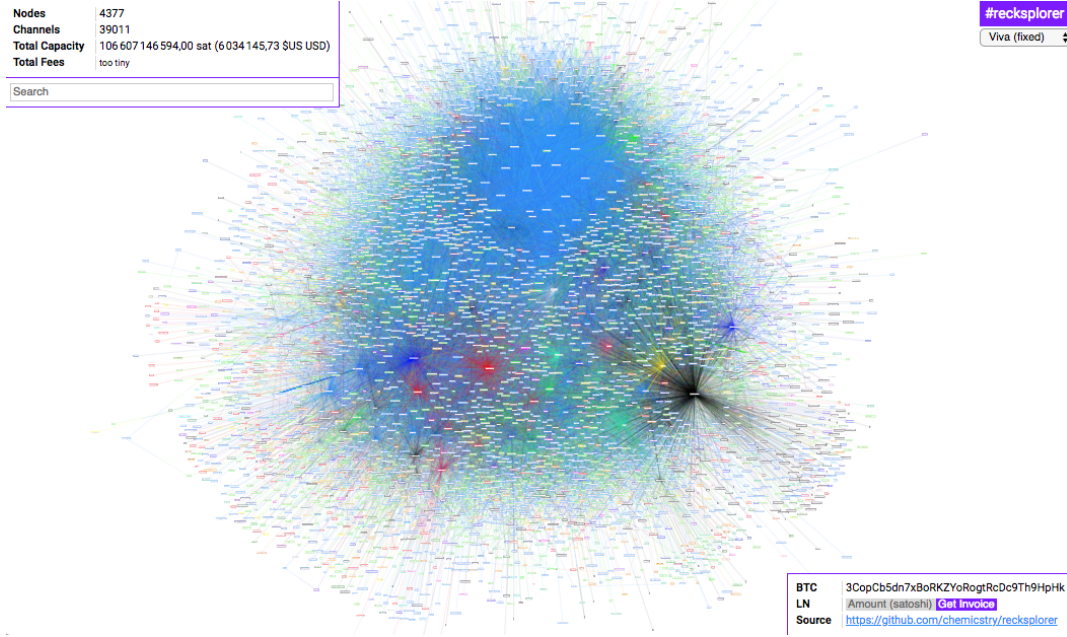


Figure 4.2: Vizualization of the Lightning Network (2) (Downloaded from <https://rompert.com/recksplorer/> on May 4th, 2019)

The degree centrality of a node n is defined as the proportion of nodes it is connected to. Figure 4.3 presents the distribution of degree centrality (on a log-scale) in the Lightning Network. The distribution looks like a power law, which we will formally test in a later section. The highest value is 0.4613 which indicates that a node is connected to 46.13% of the graph.

The betweenness-centrality of a node n measures how often the node n is part of the shortest path between two randomly selected nodes. This is a lot more relevant given the context of routing payment through the network. If a node were to be on the optimal route for too many transactions, it will start to have power over the network which obviously can be an issue. Here the shortest path is defined as the actual shortest path in terms of number of intermediaries but keep in mind that on the Lightning Network the “shortest path” will be computed regarding the fees associated with each intermediary. A rational user will select the cheapest route rather than the actual shortest one. However, the fees are quite low so far so the cheapest route will be very close to the shortest one. The betweenness-centrality, $c_B(v)$, is defined as

$$c_B(v) = \sum_{s,t \in V} \frac{\sigma(s, t|v)}{\sigma(s, t)}$$

where V is the set of nodes, $\sigma(s, t)$ is the number of shortest (s, t) -paths, and $\sigma(s, t|v)$ is the number of those paths passing through some node v other than s, t .

Figure 4.4 presents the distribution of node betweenness centrality (on a log-




Figure 4.3: Distribution of the degree centrality (log scale)

scale). It still looks like a power law even if the order of magnitude is much different for this measure. The highest value is 15.71% meaning that 15.71% of all of the shortest paths between nodes have the most important node on their path, this is also referred to as the Central Point Dominance. A low Central Point Dominance is preferable in this kind of network because it indicates that no node can suddenly break more than 15.71% of the optimal route on its own.




Figure 4.4: Distribution of the betweenness centrality (log scale)

4.3.2.2 Clustering

Another important measure for the Lightning Network is the level of interconnection within the graph. We will see that if the network is well interconnected, it will be easier for nodes to rebalance their channels and then keep routing payment across the network. It will be a key part of our model. In terms of measure, we can use the clustering coefficient measure as a proxy for interconnectedness. This measures the number of triangles through a specific node relative to the node's degree. Figure 4.5 presents the distribution of nodes clustering coefficients. We can see that the distribution of clustering coefficients is quite wide. Indeed, a lot of nodes have a quite high clustering coefficient. This is surprising because, as we previously saw in figure 4.4, the betweenness centrality is distributed like a power law. This suggests that nodes tend to prefer to be well interconnected within the graph even if they are not so often on the shortest path. The model will later show why nodes tend to favor such situations. This has to do with channel capacity rebalancing. It is cheaper for nodes to rebalance their channels if they are well interconnected (in the exact sense of the clustering coefficient).




Figure 4.5: Distribution of clustering coefficients

4.3.2.3 Articulation points

Resiliency against the failure of one node is also a relevant concept. In a sufficiently high connected graph, no single node failure would have an impact on the ability of nodes to find a route to any other node. To measure that we use the notion of

articulation point.

The Lightning Network is a connected graph, meaning that any node can reach any other node. A node n is an articulation point if the removal of n yields a disconnected graph, therefore preventing some node to reach the others. This is obviously very bad for the payment network. With the data we had we found that there are 0 articulation points.

From this point of view, the network seem well interconnected, and resilient to attack on one particular node.

4.3.2.4 Network capacity

A specific characteristic of the Lightning Network is the capacity. Each payment channel, edge in the graph, has a maximum capacity¹⁸. The distribution of total capacity per node gives an indication on how liquid the node is, or otherwise stated how likely it is to actually route transaction given a level of interconnectedness. Figure 4.6 presents the distribution of total capacity per node.




Figure 4.6: Distribution of average capacity per node

We can observe that a few nodes have a very high capacity and that the distribution looks like a power law. We previously saw that the distribution of centrality also looks like a power law. It turns out that the average capacity per channel is rather constant. Therefore the nodes with a high number of edges are also the ones

¹⁸There are actually capacities for both ways, *in* and *out*, but for privacy and security reasons there are hidden. In practice, a node select a route that it likely to have sufficient capacity, and if it fails, it tries another one, and so on.

that, on average, have a higher overall capacity. No nodes concentrate a lot of liquidity in a few payment channels. We see this graphically in figure 4.7 that presents a view of total capacity and number of payment channels on the same plot.




Figure 4.7: Distribution of total capacity with the associated number of edges

4.3.3 Dynamic properties

The static properties of the Lightning Network highlighted in the previous section tend to show that this network behaves like a scale-free network. This kind of complex network is best understood using generative models. Unlike random graphs, generative models feature the notion of node arrival as well as edges being added over time. We use the historical dataset of channel's openings and closings. Note that this dataset covers the period between January, 16th 2018 and August, 22nd 2018. This was the early stage of the Lightning Network, but we can still derive some interesting results as well as a method for later replications of this analysis.

4.3.3.1 Actual algorithm used in the Lightning Network

At this point, it is interesting to understand the actual behaviour of nodes on the network. The Lightning Network is actually no more than a standard of communication, a protocol, called *BOLT*. Therefore, anyone can, and is encourage to, develop his own application that will follow the *BOLT* standards¹⁹. There are currently

¹⁹As defined in <https://github.com/lightningnetwork/lightning-rfc>

3 major implementations: *Eclair* (made by ACINQ), *c-lightning* (made by Blockstream) and *lnd* (made by Lightning Labs). Anyone willing to join the network can download any of those 3 clients and be sure about compatibility among BOLT-compliant clients.

Although not being part of the BOLT standards, interestingly all of these clients feature an *autopilot* mode when first starting the node. If the node is fired with that option, it will try to connect with some nodes in the network using a preferential attachment algorithm similar to Barabasi and Albert (2001). In the following, we estimate this preferential attachment model as well as Bianconi and Barabasi (2001), a more general model that introduces a node-specific fitness parameter in the edge probability between two nodes.

4.3.3.2 Models and method

The model developed by Barabasi and Albert (2001) allows for a node centrality to depend on the degree of this node, the more a node has edges the more likely it is for a new node to connect with it. This is called *preferential attachment* (PA) and is often used to model social network, or other scale-free networks.

The network begins with a set m_0 of initial connected nodes. Time is infinite, and at each point in time a new node gets added to the network and connects to $m < m_0$ existing nodes with a probability proportional to the current degree (number of edges) of each node. The probability p_i that the new node is connected to node i is

$$p_i \propto \frac{k_i}{\sum_j k_j}$$

where k_i is the degree of node i . Barabasi and Albert (2001) shows that under such circumstances the degree distribution across the network, $P(k)$, follows

$$P(k) \propto k^{-\gamma}.$$

It already looks like it should be the case regarding figure 4.3 or figure 4.4, but we will check that this is the case and estimate the value of γ , the parameter of the distribution.

As a more general version of this model, we can include another parameter in the probability of connecting to a node. Bianconi and Barabasi (2001) add a node-specific parameter, called the *fitness*, meant to represent the strength of the node. In the previous model, there is a natural tendency of earlier nodes to be the largest, because the older a node is, the more connection it has, which will trigger even more connections. To represent the ability of some nodes to enter the network at a later stage and still be able to be among the largest ones, a fitness parameter, η_i ,

is assigned at a new node i drawn from a distribution. The probability p_i that the new node is connected to node i is then modified to be

$$p_i \propto \frac{\eta_i k_i}{\sum_j \eta_j k_j}$$

Bianconi and Barabasi (2001) show that the network degree distribution depends on the distribution of the fitness parameter.

To estimate both models, we use a method called *PAFit* introduced by Pham, Sheridan and Shimodaira (2015) which consists in a non-parametric maximum likelihood estimation of the preferential attachment in temporal complex networks. The method goes as follows. Define the generic probability that a new node will connect to node i by

$$p_i(t) \propto A_{k_i(t)} \times \eta_i$$

where A_k is a function of the degree of node i at time t , $k_i(t)$. We briefly describe this method which is a Bayesian estimation, and we encourage the reader to dive into the *PAFit* paper for a more detailed explanation. The method is essentially a maximization of the log-likelihood of the most general model with suitably added regularization terms to avoid overfitting. Those regularization parameters control the amount of regularization for both the preferential attachment and the fitness. The dataset is split into a *learning sample* and *testing sample*, and a grid of *regularization parameters* is built. Then we compute a *Maximum-a-Posteriori* likelihood for each couple of *regularization parameters*. The estimates of the regularization parameters and the distribution are then the one that maximizes the overall likelihood.

4.3.3.3 Results

We first estimate, also with the *PAFit* method, both the Barabasi and Albert (2001) model with only preferential attachment, and the Bianconi and Barabasi (2001) with fitness. Table 4.4 presents the estimation of the preference attachment function when *i*) it is estimated alone, and *ii*) when estimated jointly with the fitness.

$P(k)$		
	Pref. Attach. Alone (1)	Pref. Attach. & Fitness(2)
γ	0.973895	0.5730294
	[0.026]	[0.0068]

Table 4.4: Estimation of the attachment function (*standard deviation in brackets*)

For comparison purposes, table 4.5 presents different values for the power coefficient, γ , for several kinds of networks. Estimating the PA function alone yields a

power coefficient close to 1, but when it is jointly estimated with the fitness distribution, the coefficient drops to 0.57. This tends to show that the fitness is actually important in the Lightning Network. In each case, the coefficient γ is smaller than most *real-life* networks, so it seems that the current degree of a node is less important, on average, to predict future links with new nodes.

Network	γ_{in}	γ_{out}	Reference
World Wide Web	2.72	2.1	Broder et al. (2000)
Movie actors	2.3	2.3	Barabasi and Albert (1999)
Silwood Park (food web)	1.13	1.13	Montoya and Sole (2000)

Table 4.5: Estimation of γ_{in} and γ_{out} , which are respectively the exponent of the degree *in* and *out* distribution for a directed graph




Figure 4.8: Distribution of Fitness parameter

The PAFit method also returns the fitness distribution, as presented in figure 4.8. We can see that most nodes have an estimated fitness below 1, and the distribution is quite wide. The wideness of the distribution means that the fitness is an important determinant of the network's topology. If all nodes had the same fitness, it would have suggested that all nodes behave the same, however we see that some nodes have a tendency to connect more, and to gain more importance in the network.

Overall, this suggests that most nodes are not running in *autopilot* mode for link (payment channel) creations. Some nodes have a natural tendency to be more important in the network, and decide with who and when to open a payment channel. The point of this chapter is to study the behaviour of nodes on this market. We will

see in later sections, why some nodes may be willing to connect more or less than with only the preferential attachment algorithm.

4.4 The Model : Homogeneous Nodes

This section presents the general version of the model featuring a homogeneous cohort of nodes. Because all nodes are the same, it is also a way to analyze the incentive of one node on the network.

4.4.1 General setup - Exogenous fee

The model features a network of payment channels, $G = (V, E)$, taken as exogenous, with V being the set of nodes with cardinal N and E being the set of payment channels among them, with cardinal M . A payment channel is defined as a bilateral contract in which both parties have to self-manage the allocation of some value “locked” on the contract. We assume that this network is backed by a secure enough blockchain, which means that agents (nodes) do not perceive any counterparty risk regarding payment channels since any amount deposited in the Lightning Network is escrowed on the secure blockchain. Recall that even if they are bilateral channels, thanks to some cryptographic tricks²⁰ it is possible to select a route of payment channels from the sender to the receiver so that no one in between has the ability to block the payment and steal value. We assume that no one will try to steal any funds, hence whatever fund an agent locks on the chain, there is no risk of loss. As explained in section 4.2, the closing of a payment channel can be made bilaterally or unilaterally, and here we abstract from any issues regarding channels monitoring, that is when a node wants to close a channel, it is done at no other cost than the fees and delay inherent to the blockchain itself, and it does not incur monitoring costs.

We assume that there is an exogenous cohort of buyers and sellers of goods and services that are willing to exchange value and, to simplify, have to use the payment network G for that. We are interested in the behaviour of a profit-maximizer node $i \in [1, N]$ among the V nodes²¹. In order to route payments, node i needs to open bilateral payment channels with at least 2 others nodes and fund those channels.

We call d the number of payment channels opened by node i , and $x \geq 0$ the aggregate amount of funds locked inside those d payment channels. We assume

²⁰These are called *Hash TimeLock Contract*. If Alice has a payment channel with both Bob and Charlie, Bob can pay Charlie through Alice. For Charlie to claim the payment coming from Alice, she needs to make public the secret that Alice needs to claim the payment from Bob. Either both transactions are completed or nothing happens.

²¹Since we are only interested in this node, we might omit the subscript i .

there is a quadratic opportunity cost of capital for node i to leave x inside the d channels, and denote it by $c(x)$. Recall that the Lightning Network is a public network, and in particular it may contain adversarial nodes or potential thieves might just be monitoring the network. Therefore by having a lot of money locked inside those payment channels, a node will start to look like a potential target. Note that the security of the blockchain itself or the Lightning Network protocol is not an issue here, the risk we are referring to is an outside risk. Indeed, an attacker could use the network data to construct a list of public IP-address associated with the amount of funds locked in payment channels for each address, and therefore could use traditional hacking tools to try to steal the private keys associated with a given address in order to steal the funds. Therefore the node would have to invest more in cyber-security in order to protect its funds, hence the quadratic form. The following proposition summarizes this idea.

Assumption *The opportunity cost for node i to leave x in payments channels, called $c(x)$, is defined by*

$$c(x) = \frac{x^2}{2} \quad (4.1)$$

In order to route transactions, a node will have to open channels with existing nodes. Recall that the originator of the transaction will use the graph data to pick the optimal route. If a node wants to increase the volume of routed transactions, it will have to increase its interconnectedness in the sense of the probability to be on the optimal route for a random payment²². Instead of having to keep track of the whole graph structure, we define by $\lambda > 0$ the interconnectedness of node i . By assumptions, the case $\lambda = 0$ corresponds to node i not being connected to the graph, and the extreme case $\lambda \rightarrow \infty$ corresponds to node i being a “central point” on the graph. This notion is purposely loosely defined, but being a central point refers to situation in which all transactions would have to go through node i . The scale is arbitrary, and we will specify later how the expected revenue grows with λ but the idea is that node i will choose a level of interconnectedness λ and in order to maintain it through time, it will have to open and close channels depending on the evolution of the network structure. Indeed, if payments channels are closed in other parts of the graph, that would affect indirectly the interconnectedness of node i , hence the induced dynamic management of payment channels. We assume the cost of maintaining this interconnectedness level λ to be linear. The following

²²Recall that the optimal route is the cheapest one, but here we are referring to the interconnectedness everything else being equal, in particular the routing fee.

proposition summarizes this idea.

Assumption *The cost of maintaining λ interconnectedness level, $\beta(\lambda)$ is defined by*

$$\beta(\lambda) = \lambda \quad (4.2)$$

In practice, a node will need to fund all his channels independently. As transactions come through, channels balances will change. A channel balance is updated as soon as a transaction is routed. And because, from the blockchain point of view, the two legs of the routed transaction are 2 separate contracts, they cannot offset any funds from one another even if both channels are owned by the same agent²³. Remember that from the blockchain point of view, the Lightning Network is just a bunch of bilateral contracts that are opening and closing with different allocations. Imagine a node i always makes a payment to node j and never the opposite, then regardless of the original allocation the capacity of the channel will be driven toward node j . At this point node i cannot make a payment anymore (not enough funds on his side of the channel), but node j can, if he wants, make a payment back to node i . Nodes ability to keep channels well balanced will be an important determinant of their profits (from the individual perspective) and even overall liquidity of the network (from the network perspective). There are 2 main ways to get around this issue. First, the option to broadcast a “refill” transaction to the blockchain (with fees and delays) is always available, at a cost. Second, if the network is “well connected” it might be possible for a node to pay itself over the Lightning Network choosing a route that decreases his highly loaded channels and increases his poorly loaded channels, we call this *rebalancing*. We assume this cost of *rebalancing* to be included in the cost of maintaining a given interconnectedness level, $\beta(\lambda)$. More specifically, the cost of rebalancing would be driven by the number of loops that are available to node i so that it can rebalance its channels by paying itself on the Lightning Network. We assume that this number of loops is sufficiently correlated to the interconnectedness level λ so that the cost savings induced by rebalancing is included in the linear form of the cost of maintaining a given interconnectedness level λ .

On the revenue side, the node essentially collects routing fees for forwarding payment across the Lightning Network. When a node initiates a payment, it selects a route based on the public information of the Lightning Network, namely the list

²³Or at least it is costly. Recall that everything is always possible when you broadcast transactions to the blockchain, at a cost. We use the term “impossible” to mean impossible on the Lightning Network.

of public payment channels and the fees that each node set. We abstract from some minor complications regarding routing²⁴ and focus on pure optimal routing based on fees minimization.

We are interested in what happens with the next transaction to be routed through the Lightning Network. An end-buyer contracts with an end-seller for a good or service and decide to pay through the Lightning Network. We define by $p(\lambda)$ the probability that node i is selected as part of the optimal route for this next transaction. We assume this probability is directly linked to the interconnectedness level of node i , λ , and follows an exponential distribution. This assumption implies that the term “optimal route” should be understood as the optimal route based only on the interconnectedness level λ and is therefore not amount-dependent. The real optimal route would be such a route in which all intermediary nodes have enough funding to actually route the transaction.

$$p(\lambda) = 1 - e^{-\lambda}$$

Once the node has been selected as part of this optimal route, he will be able to route the payment if and only if he has enough funds in his channels. We define the amount of the next transaction by y and assume it follows an exponential distribution, $F_y(y)$, with a probability density function $f_y(y) = F'_y(y)$.

$$f_y(y) = e^{-y}$$

When the node is selected as part of the optimal route, and when it has enough funds to route the payment, it will earn a routing fee, $R > 0$. This routing fee does not depend on the amount being routed²⁵. We take the routing fee, R , as exogenous. Actual nodes on the Lightning Network are actually choosing fees but for now we abstract from that, in section 4.4.2 we will discuss implications of allowing the node to select the routing fee.

The node will maximize his expected revenue by solving for (x^*, λ^*) .

$$(x^*, \lambda^*) = \arg \max_{x \geq 0, \lambda \geq 0} p(\lambda) \int_0^x R f_y(y) dy - \beta(\lambda) - c(x)$$

If we solve the integral and plug the cost functions and the probability, we can

²⁴Routing is not perfect on the Lightning Network. Take for instance a sender that selects a node i on a route, by the time the payment reach the actual node i the state of the payment channel might have changed. This would result in a payment rejection and the sender would have to find another route.

²⁵On the Lightning Network, the fee does not depend on the amount being routed. From a theoretical point of view, the fee could be made amount-dependent as transferring large amount of money could unbalance channels. Here we assume this cost is included in the cost of maintaining a given interconnectedness level, $\beta(\lambda)$.

rewrite this maximization program as follows.

$$(x^*, \lambda^*) = \arg \max_{x \geq 0, \lambda \geq 0} (1 - e^{-\lambda}) (1 - e^{-x}) R - \frac{x^2}{2} - \lambda \quad (4.3)$$

The node solves this program and we have the following lemma.

Lemma 4.1 *(Unicity) On the Lightning Network, a profit-maximizer routing node chooses the optimal interconnectedness level, λ^* , and the amount of funding to lock in payment channels, x^* , and there is a unique solution (x^*, λ^*) that maximizes his expected payoff.*

By lemma 4.1, we know that the node's objective function behaves well. Intuitively this makes sense because the node's revenue is concave with respect to both control variables. The cost function for the interconnectedness level and the channel's funding however, are respectively linear and quadratic. The remaining of this section provides an analysis of the solution and its implications for the Lightning Network.

In this first general setup, the routing fee is taken as exogenous. This could represent a highly competitive market for which nodes are price-taker, or it could represent a situation of a new node entering the Lightning Network and using the network's current average routing fee. The next section will discuss implications of endogenous routing fees.

Here, both the cost of maintaining a given interconnectedness level and the opportunity cost of capital act like fixed costs, they essentially have to be incurred before the transaction is actually routed through the node. Therefore the node needs to make enough (expected revenue) otherwise it will be better off exiting the market and playing $(x^*, \lambda^*) = (0, 0)$. This idea is summarized in the following lemma.

Lemma 4.2 *(Exit condition) There is a minimum routing fees, $\bar{R} > 1$, under which the node exits the network, and chooses*

$$(x^*, \lambda^*) = (0, 0).$$

The very structure of the Lightning Network yields this exit condition in equilibrium. In equilibrium, and in order to operate, the nodes want a routing fee R greater than the minimum required routing fee \bar{R} . Compared to traditional production models, there is no variable cost and only fixed costs. Moreover, the revenue

generated by routing a transaction is independent of its amount and there is no quality differentiation. This particular setup can be thought of as a lottery. A transaction is going to come through the network and all nodes are trying to get their hands on it²⁶. They can do two things to increase the likelihood of being on the optimal route (*i.e. the selection likelihood*) : *i*) invest in interconnectedness, through β , and *ii*) invest in payment channels through x . They do that by trading off the increase in selection likelihood and the associated cost. If the lottery prize, the fee R , is lower than the minimum required fee, \bar{R} , even increasing a little bit the selection likelihood will be too costly and the node will simply prefer not to participate in this lottery and keep its selection likelihood at zero. This lottery story accurately describes the Lightning Network from the point of view of a single node and considering the marginal transaction.

Because the routing fee is known to the node when it makes investment decisions, it is also interesting to study how the optimal investment is affected by different routing fees. By lemma 4.2, we already know that for a low enough fee, the investment in interconnectedness and in payment channels will be null. The following proposition describes the node's investment with respect to the current routing fee, R .

Proposition 4.1 (*Complete solution*) *Let's define by $f : R \mapsto x_R^*$ the function that maps possible values of the routing fee R to equilibrium values of the amount funded in payment channels x_R^* . We have that :*

- *The function $f(R)$ is null for $R < \bar{R}$;*
- *For $R \geq \bar{R}$, the function $f(R)$ is strictly increasing and concave ;*
- *Figure 4.9 is the graph of the function $f(R)$;*
- *The optimal interconnectedness level λ^* is given by*

$$\lambda^* = \ln \left(1 + x^* (e^{x^*} - 1) \right). \quad (4.4)$$

There are no explicit solutions for the equilibrium quantities (x^*, λ^*) but fortunately we can describe the solution enough to analyze it. As previously stated, there is a minimum required fee under which $(x^*, \lambda^*) = (0, 0)$. Right above this threshold, the optimal funding amount is quite sensitive to changes in the routing fee. For larger values, because the function is concave, this sensitivity decreases. We

²⁶Note that this is not the usual type of lottery as there could be several “winners”, i.e. all nodes selected as part of the optimal route.




Figure 4.9: Graph of the function $f(R)$ representing the optimal amount invested in payment channels with respect to the current routing fee

will see in later sections that the routing fee is actually unlikely to be very high. This suggests that the network liquidity, as defined by the total amount locked inside all the network's payment channels, is likely to have a high sensitivity to short term (if any) variations of the average routing fee.

Note that the current state of the Lightning Network does not allow for any test for the funding-to-fee sensitivity. Indeed, as of early-2019, the network is essentially maintained by people trying to help the development of the network. As previously stated, the average routing fee is very low, and no one is really trying to benefit from transactions' routing at the moment.

To analyze the sensitivity of the interconnectedness level to changes in the routing fee we need proposition 4.1 which provides the equilibrium relationship between the interconnectedness level, λ , and the channels' funding amount, x , with equation 4.4. Figure 4.10 represents the graph of the function $\lambda^*(x^*)$. Loosely speaking we can see that this function behaves like a linear function with slope higher than 1, except for small values of x . This suggests that the interconnectedness-to-fee sensitivity will be even greater than funding-to-fee sensitivity. However, note that it might be complicated to interpret the interconnectedness in a dynamic context. This analysis should be taken as sensitivity to initial conditions, whether the interconnectedness level will adjust dynamically to changes in the average routing fee is out of the scope of this section. By looking at the market from a higher perspective, the section 4.4.2 provides some insights regarding this issue.




Figure 4.10: Equilibrium relationship between the interconnectedness level and the channels' funding

4.4.2 Two stages model - Endogenous routing fee

So far we have taken the routing fee, R , as exogenous. In reality, nodes set the fee R and are using it as a competition device. In this subsection we endogenize the price by analyzing a more general game in which nodes also compete on price.

The routing activity on the Lightning Network can be thought of as a *Bertrand competition* model with capacity constraints. In line with Kreps and Scheinkman (1983), hereafter KS83, such a game can be described with a 2-stages competition model. In the first stage, firms invest in capacity, then at the second stage all firms capacity become public and they engage in a price competition. On the Lightning Network, the interconnectedness λ , and the amount funded in channels, x , can be seen as capacity constraints, let's call them the *network capacity* and the *channel capacity* respectively. Indeed, once a node has set (λ, x) it will only be able to route so much transactions²⁷. The previous game can therefore be seen as the first stage of such a game. Nodes first compete on network and channel capacity, using the exact same setup as in the previous subsection, then they enter a price competition in *Bertrand* style.

There are some important differences with respect to a traditional product market. Like in KS83, the marginal production cost within the capacity is null²⁸, but the service sold is not homogeneous. Because we are analyzing a network, a specific node cannot be replaced by any other node, so the service they provide is not ho-

²⁷Again, nodes can always update dynamically their interconnectedness level and payment channel funding by going through the blockchain, at a cost. However, for the sake of the argument, we assume this cost is infinite.

²⁸Note that this assumption is a simplification in KS83 but it is actually true within the Lightning Network. More precisely, the marginal cost of routing a transaction is the cost a sending a few data packets on the internet, which is very close to zero.

homogeneous strictly speaking. However, before a transaction is initiated, all nodes with a given interconnectedness level look the same, therefore we say that the routing service is homogeneous given a particular value of interconnectedness level. As soon as an end-buyer has been assigned an end-seller, and wants to transact, each node becomes different from all others, depending on how likely it is to be on the optimal route, but nodes cannot anticipate that. Moreover, due to onion routing, nodes cannot know their position within the optimal route, nor whether there is an alternative route, so they cannot differentiate themselves with different service apart from the interconnectedness level. We might be tempted to think that the capacity constraint works here a little different than in KS83 because we assumed that the end-buyer would select another route if a node's capacity does not meet the transaction amount. In KS83, the consumer demands first to the lowest-pricing firm up to this firm's constraint, and the residual demand is going to the other firm. However, it is actually very similar thanks to *Atomic-Multi-Payments*, with which the end-buyer is able to select several routes and split the amount among them so that the end-seller receives all payments or nothing. This is out of the scope of this chapter but bear in mind that the channel capacity constraint behaves very much like a traditional capacity constraint like in KS83.

Another important difference is the interpretation of demand. Contrary to KS83 we do not make any assumptions regarding the demand for routing. We can focus only on the next transaction, because we only care about the initial decision of nodes, only the expected profit matters, which does not depend on the actual realisation of demand. On the Lightning Network, the profit of a node does not depend on the amount of the transactions routed, the fee is not amount-dependent. Also the volatility of the demand is already included in the channel and network capacity constraints, because with a given (x, λ) the node can only expect so much revenue.

Lastly, we assume the following for the price discrimination. For two nodes $i \neq j$, such that $x_i = x_j$ and $\lambda_i = \lambda_j$, then the node setting the lowest fee R is getting the next transaction, the other node expects no return at all. This represents the fact that nodes with similar capacity constraints are, *ex-ante*, the same. Because they are *ex-ante* homogeneous, Bertrand competition applies to them, and drives the price down in order to attract transactions. This idea is summarized in the following statement.

Assumption (*Price competition*) Take any two nodes $i \neq j$, such that $x_i = x_j$ and $\lambda_i = \lambda_j$. Assume, without loss of generality, that they announce routing fees $R_i \geq R_j$, then

- If $R_i > R_j$: node j attracts the next transaction, and node i expects a null profit

- If $R_i = R_j$: node i and node j share the volume among each other, and both can expect a positive profit as defined in equation (4.3).

Let's formally describe the 2-stages model. The first stage is exactly like the model in the last subsection, let's recall it briefly. There are N homogeneous nodes competing to route transactions. At the first stage, each of them sets an interconnectedness level and an amount to lock in payment channels. At the beginning of the second stage, all first stage choices are made public, and nodes announce a routing fee R_i . Then an end-buyer is matched with an end-seller and an optimal route is selected within the network. Each node still maximizes its profit at each stage.

At the second stage, the network topology is set and nodes are competing only on prices, taking what they know about the network as given. Let's define by Θ the information available at the end of the first stage.

$$\Theta = (\{x_i\}_{i \in I} ; \{\lambda_i\}_{i \in I})$$

To solve this game, we run backward by solving first the second stage subgame. For every outcome of the first stage game, we defined as the Θ -subgame, the price competition engaged by the nodes given the network state Θ , of which the resulting optimal routing fee is denoted by R^* . Since all nodes are homogeneous, they will take the same decision. Then we can restrict our attention to Θ -subgame in which all decisions are the same across nodes. Therefore, in this section, we can write $\Theta = (x, \lambda)$, and a (x, λ) -subgame.

In a (x, λ) -subgame, each node is maximizing its profit. Recall that nodes are solving this program under the *Price competition* assumption. In order for them to attract the next transaction (i.e. a non-zero volume), they need to take into account the fee sets by other nodes. The optimal fee set by node i , in the (x, λ) -subgame, denoted by $R_i^*(x, \lambda)$. We have the following lemma.

Lemma 4.3 *(Subgame) In the (x, λ) -subgame, price competition drives the equilibrium routing fee to low values and there is a continuum of Nash-equilibria regrouped into two categories :*

- All nodes make zero profit and $\forall i$:

$$R_i^*(x, \lambda) = \frac{x^2 + 2\lambda}{2(1 - e^{-\lambda})(1 - e^{-x})}$$

- All nodes make a loss and $\forall i$:

$$R_i^*(x, \lambda) \in \left[0; \frac{x^2 + 2\lambda}{2(1 - e^{-\lambda})(1 - e^{-x})} \right]$$

Note that this subgame can be viewed as a kind of prisoner's dilemma. The good outcome can be associated with zero profit for all nodes, and the bad outcome with a loss for all nodes. Cheating could be considered to be a price dumping strategy, to overtake the volume. However, a node does not benefit from "cheating", because its loss increases, so the "good" outcome is also a Nash-equilibrium. We will discuss the implication of having this continuum of equilibria once we solve the full game. To do this, we first restrict our attention to a subset of (x, λ) -subgame, and we establish the following proposition.

Proposition 4.2 *(Full game) There is only one Subgame Perfect Equilibrium (SPE) for the full game. The equilibrium is given by the following :*

- Define $\lambda(x) = \ln(1 + x(e^x - 1))$
- All nodes make the same decision (x, λ)
- The optimal amount locked in payment channel, x^* , is the unique solution to

$$x^* = \left\{ x \in \mathbb{R}_+^* ; \frac{1 + x(e^x - 1)}{1 - e^{-x}} = \frac{x^2 + 2\lambda(x)}{2(1 - e^{-x})(1 - e^{-\lambda(x)})} \right\}$$

- The optimal interconnectedness level, λ^* , is given

$$\lambda^* = \lambda(x^*)$$

- The optimal routing fee, R^* , is given by

$$R^* = \frac{x^{*2} + 2\lambda^*}{2(1 - e^{-x^*})(1 - e^{-\lambda^*})}$$

- No node makes a profit

Given this cost structure, there is a single equilibrium outcome that consists in a strictly positive routing fee and no profit for the nodes. All nodes are homogeneous so they make the same decision. Section 4.5 explains the consequences of heterogeneous agents.

In KS83, they show that Bertrand Competition with capacity constraints yields the Cournot outcome. In our model, surely the equilibrium price is higher than the one with the Bertrand price without constraints, which would be zero, the marginal cost. However, the way the model is specified makes our analysis different. In particular, the demand is perfectly inelastic, which means that the users would demand the same amount regardless of the price. Also, the quantity demanded by the users does not impact the profit of the nodes (the routing fee is amount-independent). Those assumptions are far from the usual ones so the analysis of a pure Cournot game is not very relevant to the case of the Lightning Network.

The second subgame equilibrium is not an equilibrium of the full game, and we explain intuitively why this is the case. In the (x, λ) -subgame, setting the fee to zero is an equilibrium because nodes cannot change the capacity limit. In the full game, however, they are free to select the capacity they want, and because having a zero fee produces loss, the nodes would be tempted to lower their capacity investments to reduce their loss. This would yield the following $(x^*, \lambda^*, R^*) = (0, 0, 0)$. However, this point is not an equilibrium because the demand is completely unsatisfied and any node could jump into the network and start routing at any fee. The proof of proposition 4.2 goes on to demonstrate that any other point is not an equilibrium.

We’ve described a model to represent a homogeneous Lightning Network. The unicity of equilibrium basically ensures that the network is well functioning and can be used to route payments. Note that we are studying a payment network and in real life there would be other options for agents to transact. So any equilibrium in this “narrow” situation is contingent on the outside option of using another payment network. If the outside option is too low, the demand would be null and the network wouldn’t exist. Whether this payment network would be stable and working in a more general setup, with other payment networks, is outside the scope of this chapter.

The parameters we chose to model the network should be taken cautiously. Indeed, we model the interconnectedness of nodes with λ , which represents the probability that the node is on the optimal route between any two other nodes. With homogeneous agents, we saw that they all take the same decision in equilibrium, especially the same interconnectedness level. Such a network can have a very bizarre structure and, as seen in section 4.3, is not at all representative of the Lightning Network which is very concentrated in terms of node centrality (i.e. the interconnectedness level). With this simpler version in mind, the next section will establish results for a more general setup with heterogeneous agents.

4.5 Heterogeneous Agents and Network Structure

So far we have assumed a homogeneous cohort of nodes. In this subsection we analyze agents that are heterogeneous with respect to their opportunity cost of capital, $c_i(x_i)$.

On the Lightning Network, nodes could be run by very different kinds of agents. As previously discussed, one of the main risks of running a node is the payment channels funding security. We assumed that the network is powered by a secure enough blockchain so that nodes do not perceive counterparty risk. Although, on the blockchain, ownership is defined by private key ownership and nothing else, so nodes are subject to *private key* thief risk (or *security risk*). Indeed, the network status, including payment channel balances, is made public²⁹ so an external attacker could locate the *ip-address* of a node and deploy resources to hack into it and steal the amount locked in payment channels³⁰. Having a large amount invested in payment channels would require to increase the security of the server itself, hence incurring another cost.

We assume that each node has a different ability to deal with security risk. Recall that we assumed that this risk is included in the node's opportunity cost of capital, hence its quadratic form. Therefore, we introduce some heterogeneity in the nodes' opportunity cost of capital to represent the difference in security risk. More specifically, let's assume that for each node $i \in [1, N]$, there is a value $\gamma_i > 0$ assigned by nature, and the node's opportunity cost of capital is now a function of γ (we shall omit the subscript i when we focus our attention only on node i).

$$c(x) = \frac{\gamma x^2}{2}$$

All the other assumptions remain. We describe how the equilibrium changes when γ changes like in the previous section, first the incentive of only one node with exogenous routing, and then the two stages game.

4.5.1 Exogenous fee

Let's first analyze the setup in which R is exogenously given. Nodes are maximizing their investment in channel and network capacity, (x, λ) . The objective program of

²⁹Only part of the Lightning Network is actually made public. However, only this subgraph can be used to route payments for an anonymous user. Some private network can be deployed but will be used only by agents aware of this *sub-network*.

³⁰The *ip-address* may not always be revealed. Some nodes are using TOR to basically hide their *ip-address* but still be able to route payments. However, the channel balance on the blockchain is assigned to a particular users. This is a way the node could be located, and it would require extra-caution for highly loaded payment channels.

the node, i.e. equation (4.3), becomes the following.

$$(x^*, \lambda^*) = \arg \max_{x \geq 0, \lambda \geq 0} (1 - e^{-\lambda}) (1 - e^{-x}) R - \frac{\gamma x^2}{2} - \lambda \quad (4.5)$$

Note that the changes are small with respect to the last setup. Therefore, most of the results of the previous section still hold. What is interesting is to see where this γ parameter shows up in equilibrium. Existence and unicity still hold.

Recall that under a fee, \bar{R} , the node just prefers to exit the market and play $(x, \lambda) = (0, 0)$. This is still true but we can establish a link between the parameter γ and this fee, denoted by $\bar{R}(\gamma)$.

Lemma 4.4 *When nodes are heterogeneous, the minimum required routing fee, \bar{R} , is an increasing function of the opportunity cost of capital, γ .*

$$\frac{d\bar{R}(\gamma)}{d\gamma} > 0$$

When a node is more constrained on its opportunity cost of capital, it makes sense that it requires a higher minimum fee to cover the increase in cost. The complete solution of proposition 4.1 is also still valid up to some parameter. Proposition 4.3 gives the main differences with respect to the previous section.

Proposition 4.3 *When nodes are heterogeneous, the function f that maps possible values of the routing fee R to equilibrium values of the amount funded in payment channel x^* , gets modified as follow.*

$$\forall \gamma_1 < \gamma_2 \ ; \ \forall R \geq \bar{R} \ ; \ \left. \frac{df(R)}{dR} \right|_{\gamma=\gamma_1} > \left. \frac{df(R)}{dR} \right|_{\gamma=\gamma_2} \quad (4.6)$$

If we denote by x_γ^* the optimal amount invested in payment channel for node i with opportunity cost parameter γ , we have

$$\forall \gamma_1 < \gamma_2 \ ; \ \forall R > \bar{R} \ ; \ x_{\gamma_1}^* > x_{\gamma_2}^* \quad (4.7)$$

Also, the equilibrium interconnectedness level, λ^* , is now given by

$$\lambda^* = \ln \left(1 + \gamma x^* (e^{x^*} - 1) \right) \quad (4.8)$$

The second part of proposition 4.3, (4.7), says that for a given routing fee, R , a more constrained node will invest less in equilibrium, both in channel and network capacity, (x, λ) . The first and last part of proposition 4.3, (4.6) and (4.8), tell us that when the nodes have a higher opportunity cost of capital, γ , the amount locked in payment channel is less sensitive to changes in the routing fee, and the interconnectedness level is less sensitive to changes in the funding amount. We can already see that the opportunity cost of capital will have a large impact on the structure of the network, through λ , and on its capacity, through x . We will discuss that once we solved the two stages game.

4.5.2 Two stages model - Endogenous routing fee

Let's now turn to the most interesting case, that is a two stages game with heterogeneous agents. In this section, we solve this two stages game and compare it with the homogeneous agents case. Again, most of the results still apply but we are interested in understanding where the opportunity cost parameter, γ , shows up.

We first have to clarify how the price competition works in a heterogeneous setup. We assumed in the previous section that for the same interconnectedness level, λ , a node displaying a higher price would not get any of the volume. It was enough for the homogeneous case because all nodes take the same decision in equilibrium. We assume that the service nodes are providing is differentiable with respect to the interconnectedness level. From the network point of view, nodes with a low interconnectedness level, λ , do not help the network as much as nodes with high λ . Therefore, they should be able to charge different price. And as a consequence, a node charging a higher routing fee could still attract some volume if it has a higher interconnectedness level. This idea is summarized in the assumption below.

Assumption *There is a constant exogenous number of transactions to be routed which, in equilibrium, is shared across all nodes with the lowest fee for a given interconnectedness level.*

We also need to define the notion of a *competitive network*. In this case, a *competitive network* is defined by a network in which barriers to entry are low, and a sufficiently big cohort of potential nodes is waiting outside the network in case there would be an opportunity. More precisely, a network with strictly more than one node for each interconnectedness level is also a *competitive network*.

Defintion 4.1 *A network is said to be competitive when either one of those 2 conditions are satisfied:*

- *There are always either zero or more than one nodes providing a given inter-connectedness level*
- *There is an exogenous cohort of potential nodes outside the Lightning Network, that could jump in the network and if there is a profitable opportunity for any interconnectedness level.*

A competitive network ensures that the zero profit condition will hold in equilibrium. For the remaining we assume that the network is competitive. We will discuss in section 4.6 the potential consequence of a non competitive network.

Recall that we defined the Θ -subgame to be subgame in which nodes compete only on price, given all nodes channel and network capacity. Like previously, we first solve the last stage and then the full game. The setup is exactly like in the previous section except for heterogeneity in nodes' opportunity cost of capital. The routing fee is now a function of the complete information Θ , as nodes could take different decision in equilibrium. The following lemma provides the Θ -subgame equilibrium.

Lemma 4.5 *(subgame) In the Θ -subgame, with a competitive network, price competition yields low value for the routing fee. In particular, the same mechanics applies as in lemma 4.3, but for each different opportunity cost of capital. All nodes make at most zero profit, and in some case, they have a loss in equilibrium.*

The equilibrium continuum is a little more complex than in the homogeneous case, because there are a lot of different possible combinations, given all the possible opportunity cost of capital γ . Similarly to the previous section, we derive the equilibrium of the two stages game in the following proposition.

Proposition 4.4 *On the Lightning Network, when nodes are heterogeneous in their opportunity cost and in equilibrium, nodes use the price to discriminate among them for the interconnectedness they provide, and for each γ we have a unique equilibrium similar to the one in proposition 4.2. In particular, we have the following.*

- *Define $\lambda(x) = \ln(1 + \gamma x(e^x - 1))$*
- *All nodes with the same opportunity cost of capital make the same decision*

$$\forall i \neq j \quad \text{if } \gamma_i = \gamma_j \quad \text{then} \quad (x_i^*, \lambda_i^*, R_i^*) = (x_j^*, \lambda_j^*, R_j^*)$$

- The optimal amount locked in payment channel, x_i^* , is the unique solution to

$$x_i^* = \left\{ x \in \mathbb{R}_+^* ; \frac{1 + \gamma_i x (e^x - 1)}{1 - e^{-x}} = \frac{\gamma_i x^2 + 2\lambda(x)}{2(1 - e^{-x})(1 - e^{-\lambda(x)})} \right\}$$

- The optimal interconnectedness level, λ_i^* , is given

$$\lambda^* = \lambda(x_i^*)$$

- The optimal routing fee, R_i^* , is given by

$$R_i^* = \frac{\gamma_i x_i^{*2} + 2\lambda_i^*}{2(1 - e^{-x_i^*})(1 - e^{-\lambda_i^*})}$$

- No node makes a profit

4.5.3 Implications for the network topology

Now that we have solved the full game, we interpret the equilibrium and derive some implications of this model for the network topology.

The main mechanics of the model can be expressed as follows. Nature assigns an opportunity cost parameter, γ_i , to each node i . Then each node sets capacity limits, in terms of interconnectedness level and funding capacity (x, λ) , and a fee, R , it will charge for routing transactions through the network. This produces a network topology that is then used by end users (buyers and sellers) to select optimal routes for their transactions. The network topology will, of course, be an important determinant of the nodes' ability to make payment in an efficient manner. End users benefit from a well funded and well connected network, as they would be able to route payments at a lower cost³¹.

In our model, the network topology is endogenous and fully depends on the nodes' optimal choices, (x^*, λ^*, R^*) , and the only exogenous parameter is the opportunity cost of capital. By proposition 4.4 we know that the equilibrium is unique under the model assumptions, therefore there exists a unique function that maps the opportunity cost of capital γ to an equilibrium value for each control variable. Let's call those equilibrium functions $x^*(\gamma)$, $\lambda^*(\gamma)$ and $R^*(\gamma)$. Given the structure

³¹Note that the cost of routing depends on the capacity and the interconnectedness. It depends on the interconnectedness as end users would, on average, have to pick a longer route in a poorly connected network. If the capacity on a given route is too low, they can still use the network but at a higher price as they would need to split the payment in a multitude of payments that fit the capacity.

of our model, we cannot obtain closed form solution for those functions. However, it can be shown that the functions $x^*(\gamma)$, $\lambda^*(\gamma)$ and $R^*(\gamma)$ are monotonic³². At the network level, we can impose a distribution for the opportunity cost parameter, call it f_γ , that would give the equilibrium distribution for the channel capacity, f_x , the network capacity, f_λ , and the routing fee, f_R .

$$f_\gamma \rightarrow \begin{cases} x^*(\gamma) \\ \lambda^*(\gamma) \\ R^*(\gamma) \end{cases} \rightarrow \begin{cases} f_x \\ f_\lambda \\ f_R \end{cases}$$

Even without closed formulation for those equilibrium functions, we can provide some insights on the network topology. Note that we impose $\forall i \ \gamma_i = 0$. If the opportunity cost were to be zero, the node would set an infinite channel capacity to maximizes its profit. Any $\gamma > 0$ will produce a finite equilibrium. Therefore, the support of the opportunity cost distribution, f_γ , will extend to the distribution of channel and network capacity, f_x and f_λ . The following proposition formalizes this idea.

Proposition 4.5 (*Upper/Lower bounds*) *Intrinsic characteristics of the Lightning Network create an upper and lower bound for nodes' the interconnectedness level, λ_i , and channels' funding, x_i , which corresponds to the minimum and maximum of γ_i .*

Because of the previous remark, we can show that there is an upper-bound on the channel and network capacity, (x, λ) . It is complex to derive straight topology consequences but the fact that there is an upper-bound for the interconnectedness level and channel's funding already rule out some kind of network. For instance, a completely centralized network, in which a node (or a group of nodes) gets to route all transactions, $\lambda \rightarrow \infty$, is never optimal given $\min_i \gamma_i > 0$. Similarly a completely fragmented network (i.e. low values of λ) would only come from a very high opportunity cost of capital and is therefore very unlikely.

It is complex to say much more given that we do not have any closed form for the equilibrium functions. However our model is able to consistently generate network topologies given an opportunity cost distribution. We leave a deeper analysis of the results from a network point of view to future research. Section 4.7 provides some other insights regarding future directions and extension of this model.

³²We don't have any analytical proof, but simulations show that the functions are indeed monotonic.

4.6 Bridge with the Real Lightning Network

In this section we present the main differences between our model and the actual Lightning Network, and discuss their implications on the equilibrium analysis.

First, the Lightning Network is still new. It was first introduced theoretically by Poon and Dryja (2016). After a couple of years, the first implementations were live from January 2018. Section 4.3 describes precisely the current state of the network, but note that it is overall still relatively small. Most active nodes are currently participating in the network for the sake of development and testing. The routing fees are overall close to zero, but the total volume is also low, no one is really trying to profit from the routing activity so far, as of early 2019. We have built a model in which heterogeneous nodes compete to route transactions. We shouldn't expect our model's results to hold with the current state of the market. Our model is more likely to predict nodes' behaviour in a highly competitive market in which the demand is also high. We leave for future research refinements of this model and empirical tests when the network would have grown sufficiently.

Second, our results are up to some proportional constant. In equilibrium, there is a direct relationship between the optimal routing fee and the investment in channel and network capacity, (x, λ) , but to simplify the equilibrium analysis we didn't put any proportional constant in the cost functions. Recall that we specified $\beta(\lambda) = \lambda$ for instance. It is not likely that the real cost of maintaining λ as interconnectedness level is as such. A more realistic formulation could have been $\beta(\lambda) = \alpha\lambda$ with $\alpha > 0$ a parameter. The same would apply for the channel capacity cost, x . Such a formulation would have been able to get a more realistic equilibrium in terms of the variable's scale. However, we argue that it won't change any results of existence and unicity in equilibrium, under normal parameter conditions³³. The global shape of the objective function won't change for instance, and therefore neither will the shape of *first order conditions*, etc. Equilibrium will prevail to a more general formulation. Note that this work of parametrizing the model would have to be done before any empirical work.

Lastly, our model relies on a fix exogenous demand, and we use traditional industrial organization literature to analyze this payment network. However, a payment network is different from a traditional product market. It would be more realistically represented by a needed service as it is a mean of payment allowing "real" transactions to happen on the physical market. Note that agents, end users, need to somehow pay the seller of any goods or services, but the Lightning Network is not the only option. Therefore, all our results are conditioned on the value of

³³Here, normal parameter conditions mean that the cost are not "too high". If the cost were too important, it will of course change the shape of the objective function, but those are unrealistic.

the outside option for end users. If the outside option is low enough, they would prefer to use it and the Lightning Network would probably die because of insufficient demand. Note that lemma 4.2 holds pretty much for the whole model. It says that if the routing fee R is too low, the nodes prefer not to participate in this market. The value of the routing fee could be generalized to the demand. If the demand is too low, nodes won't be able to extract enough value to pay for their expenses (channel and network capacity), therefore they would exit the network.

There can be two forms of outside options, the blockchain itself, or any other traditional payment networks (Visa, Mastercard, Paypal, etc.). Let's call the former *trustless* payment networks, while the latter would be the *trust based* payment networks as they require agents to trust a central party of any kind³⁴. The first tradeoff is whether agents value the absence of trust in the network. If no one cares about the absence of trust, it is likely that the Lightning Network will be an expensive solution that another network could do charging less. The difference between the cost in a *trust based* network and a *trustless* one has to be lower than the benefit of not having to trust anyone. Note that this analysis does not take into account the demand for other reasons than pure transactional ones. If agents have other interest in holding cryptocurrency, they might be more likely to use such a *trustless* solution.

If it turns out that agents value the absence of trust, they still have another option to transact, namely the blockchain itself. Recall that the blockchain is always available to anyone to send and/or receive payments. The issue is that the fees on the blockchain are also dynamic and depend mostly on the blockchain characteristics like the time per block, or the size of each block. If the cost of transacting on the blockchain is low enough and if there are some issues on the Lightning Network, it could be cheaper to use the blockchain for *i*) a brief period of time, in such a case the Lightning Network would probably recover, or *ii*) an extended period of time, and in such a case the Lightning Network would probably collapse. However, this is quite unlikely, because even if the fees on the blockchain drop a lot, there is still the issue of confirmation time (it usually takes several confirmations - blocks - to make sure the payment is definitive) and this will never go below a few *inter-block* time, which is for instance 10 minutes on the Bitcoin blockchain³⁵.

³⁴Some might argue that high transaction costs are also inherent to those traditional payment network but there exists some service like Paypal that are essentially fee-less. The real difference between the Lightning Network and other payment networks is the trust that is required to use traditional payment networks.

³⁵To fully confirm a transaction, it is considered good practice to wait around 5 to 6 blocks after the block containing this transaction. This significantly lowers the likelihood of *double-spending*.

4.7 Future Research

While computer scientists have been doing academic research on blockchain since the early days, it still remains a niche in other fields, especially economics and finance. This section provides a number of potential extensions of the model we presented and other future directions in general. Of course, there are also potential research on blockchain per se but we choose to focus on Lightning Network related issues. This is the first description and analysis of second layer protocols like the Lightning Network, so there are a large number of open questions, and we divided these remarks in two categories, the theoretical and the empirical side.

4.7.1 Theory

Let's begin with open questions on theory. This chapter is an attempt to model a network of payment channels. We characterize some aspects of the network topology but we do not have a proper endogenous network structure. Using graph and network theories should prove useful in filling this gap. However, it seems that graph theory needs to progress a little as, to the best of our knowledge, the Lightning Network properties do not seem to fit any usual graph frameworks. Indeed, mathematicians have worked extensively on graph with constraint on edges but the capacity constraint³⁶ work differently here. In traditional capacity constrained network (like the electricity network for which the constraint is the wire's size), constraints are fix and constant. In other words, the capacity between two nodes is constant and it is either full or partially used. In particular, the capacity does not depend on the actual quantity flowing through this edge. On the Lightning Network, this is not true. The capacity between two nodes changes dynamically depending on the transactions volume. At some point the capacity would have to be reset (either to the blockchain itself or with rebalancing), or the edge is essentially cut off. Taking those differences in consideration could yield interesting results on the network's structure like : "What conditions are needed for all channels to remain well balanced?", "Is there any possible attack scenario that would significantly harm the network liquidity?", or "Can such a model be used to analyze anything else than the Lightning Network?"

Another way to better approximate the equilibrium in such a payment network could be to use the Mean Field Game theory (MFG hereafter). Indeed, if we do not want to use complex network and graph models, we could extend the framework of this chapter to a MFG-like setup, as introduced by Lasry and Lions (2007). MFG is basically an extension of the Nash-equilibrium concept with an infinite number of

³⁶We mean the capacity constraint from the point of view of end users, that is the amount funded in payment channels.

players. MFG is proven to be very handy to model dynamic situations, or otherwise stated it could allow us to study the network properties dynamically. It could be a way to model attack scenarii and other dynamic issues related to the Lightning Network.

At the micro level, a transaction is very different from a bank transaction for instance. There is no central authority so the nodes have to, instead, do a lot of work to properly use the network. First, the sender needs to work its way through the network. There are a number of optimal routing algorithms but they all fail at some point. Indeed, the sender will request the network status and then pick the optimal route but by the time the payment is actually sent, the network status (and especially the capacity on the optimal route) might have changed. If the selected route is not able to handle the payment, the transaction is rejected and the sender tries to find another route. True optimal routing is therefore unlikely to be available but some algorithms can do *best effort* routing. However, depending on how those algorithms are implemented and the *best effort* route is chosen, there might be some ways for nodes to increase the probability of being on the *best effort* route without necessarily incurring the same cost as for the optimal route. Second, as explained in section 4.2, network nodes need to stay connected to prevent the payment channel counterparty from stealing the funds. In a payment channel, funds are secured to the extent that each node stays online to monitor the channel state³⁷. The incentive for a node to cheat could have important consequences on the network formation.

Both previous issues could be externalized. Channel monitoring can be trustlessly outsourced to some specific nodes called *watchtowers*. Such nodes would be in charge of monitoring payment channels for other nodes, and if they detect a cheating attempt, they would provide the blockchain with the cheating proof and be compensated with a percentage of the recovered funds. Externalization of optimal route computation is done for another reason though. In the future it might be impractical for a regular node to compute the optimal route as it would have to download the full state of the network each time he wants to send a payment, of which the size can be quite high³⁸. The state of the network could be downloaded on a regular basis by another node which will be paid by the sender to send back the optimal route. The incentives of each node and the optimal contract for both externalized services, channel monitoring and optimal route computation, remain largely unknown. Note that both services yield in somehow a loss in privacy for the

³⁷Recall that, a node can attempt to close unilaterally a payment channel with any previous history of transactions (one in which he has more money for instance). To prevent that the other node has to monitor those attempts. If a node detects a false closing he provides the most recent history of transactions to the blockchain and, if valid, the blockchain allocates the whole channel funding to the non-cheating node. This is meant as a punishment for the cheating node.

³⁸As of early 2019, it is around 25Mb.

original sender³⁹.

There are higher order protocols that allow to generalize the idea of the Lightning Network. The main idea of the Lightning Network is to lock some funds on the blockchain, tell all participants about this transaction, and then use those funds bilaterally without the need to tell anyone but the nodes on the payment route. The same idea can be taken further with *channel factories*. With *channel factories*, a group of nodes holds some amount on the blockchain that allows them to open and close payment channels among each other without telling anyone but the nodes in the same *channel factory*. Each of those payment channels can then be used as a regular payment channel on the Lightning Network. In principle, it is possible to push this idea further and create *factories of factory* and so on. Recall that those factories or even payment channel are simply a set of contracts between several agents. Those contracts and their implications for end users were never studied from an economic and financial perspective.

Lastly, let's shed light on *atomic swaps*. So far, we presented the Lightning Network as a payment channel running on a single blockchain. However, several blockchains are currently running and it is possible to use the Lightning Network as a way to communicate value across blockchains. Note that this also works on the blockchains themselves, without the need of the Lightning Network. On the blockchain, *atomic swaps* can be used as follows. On blockchain A, user A sends some amount of A-coin to user C, who then sends some amount of B-coin on blockchain B to user B. This *atomic swap* is done trustlessly⁴⁰. The same idea can be implemented on the Lightning Network so that users do not need to even know what type of coin someone else is accepting, the swap could be made automatically in the background. Therefore even with one payment network, there could be several options, in terms of possible coin, for a buyer to pay the seller. Whether those swap contracts have an impact on the structure of the Lightning Network is unknown.

4.7.2 Empirical

As explained in section 4.6, we should not expect any empirical results as long as the network is still mostly experimenting. However, note that most of the data is publicly available as anyone should be able to use the Lightning Network to route payments. Historical data might be hard to get but the current state of the network like the channels, their capacity, the fees, etc. can be downloaded by anyone.

In sub-section 4.5.3, we briefly discuss the relation between the equilibrium vari-

³⁹ *Watchtowers* would need to hold the current state of the payment channels, and nodes that compute optimal route need to know the recipient of the transaction.

⁴⁰It is called "atomic" because either both users receive their value or no one does. Otherwise stated, counterparty risk is fully eliminated.

ables. Recall that in this model, there are unique equilibrium relationships between the interconnectedness level, the channel total capacity and the fee rate. Once the network would have grown, it would be interesting to see whether those relations hold in the data.

On the liquidity side, there are also many questions that need to be answered empirically. Does the network produce a stable capacity? Can nodes rebalance channels in an efficient way? Are there liquidity traps in which all funds get locked in payment channels such that no rebalancing can occur? Liquidity will be an important determinant of the usefulness of the Lightning Network as end users should be able to send and receive payments without caring much about the underlying details. If those issues turn out to be unsolvable or too complex to solve, the network is likely to collapse.

Lastly, we want to emphasize the link between the payment channel network and the underlying blockchain. Since the Lightning Network can be reduced to a set of contracts on the blockchain, the cost of underwriting and managing those contracts will be directly impacted by the transaction cost on the blockchain⁴¹. It will therefore be interesting to know if and how shocks on the blockchain parameters (transaction cost, security) propagate to the Lightning Network.

4.8 Conclusion

The Lightning Network is a communication protocol allowing the creation of payment network on top of a blockchain. The payment network is a set of bilateral contracts, called payment channels, registered on the blockchain. Once the blockchain usual transaction costs (monetary cost and delay) have been incurred to open a channel, value can move instantly at a fee supposed to be smaller than the blockchain fees. Some nodes in this network will have for sole purpose the routing of transaction across the network. This chapter *i*) presents an overview of the current state of the Bitcoin Lightning Network, *ii*) builds a model for analyzing routing nodes' incentives and *iii*) derives some implications of the model regarding network characteristics.

There are currently 2,865 nodes and 36,982 payment channels in the network⁴². The total capacity of the network is 1,045.8298 BTC (around USD 6 millions) as early 2019. The network centrality, as measured by degree centrality and betweenness-centrality, has the shape of a power law. We show that the Lightning Network is already quite interconnected. We also estimate two dynamic models, Barabasi-Albert (2001) and Bianconi-Barabasi (2001), and both indicate that the network's

⁴¹A *contract* is simply a particular kind of transaction.

⁴²These figures do not take into account the nodes with only 1 or 2 open channels, as they are not very useful to route transactions.

dynamic structure is similar to the one of social networks. Essentially, this means that older node tends to have more connections which suggests that no one is really trying to profit from the routing activity so far. Note also that average routing fees are quite low (close to zero). Those results could have been expected as the network is too new and there is still a lot of testing happening.

We propose a theoretical framework to analyze nodes' incentives and the resulting structure of the network. We feature a model in which several nodes are competing on price to route transactions, however, they must beforehand set costly channel capacity and interconnectedness level. Indeed, due to the network mechanics a node can only route a transaction if it has enough funds available in both the *in* and *out* payment channel. The marginal cost of routing a transaction is assumed to be zero. This creates a game similar to a Bertrand competition game with capacity constraints.

We show that there is a single equilibrium in which all nodes make no profit and the routing is strictly positive. If the routing is expected to be too low, nodes prefer to not participate in the market as they cannot break even. This endogenous participation constraint provides a rationale for positive routing fees even with zero marginal cost. Otherwise stated, the capacity constraints prevent the equilibrium fee to reach marginal cost. Moreover, we show that, in equilibrium, the investment in both capacities (channel and network) are linked with a positive relationship. The more a node will be interconnected, the more it will have invested in its payment channels.

We show that heterogeneity in the nodes' opportunity cost of capital is enough to generate a wide variety of network topology and overall capacity. The only exogenous parameter is the opportunity cost of capital. We show that in equilibrium, there exists a mapping between the opportunity cost of capital and the optimal values for the control variables. In particular, we show that the bounds of the opportunity cost of capital will extend to the control variables including interconnectedness level and funding amount. For instance, it means that there is an upper-bound on the centrality of a node, which corresponds to the node with the lowest opportunity cost of capital. Another way to put it is that a completely centralized network is never optimal as long as the lowest opportunity cost is higher than zero.

The Lightning Network might very well become a widely used payment network. This chapter is a first attempt to formalize decisions in such a network. Our results suggest that there are stable equilibrium outputs in which the network is liquid, well connected and with low fees (in the sense that no nodes make a profit). However, there are a number of issues that should not be overlooked, like dynamics properties, network formation, etc. We believe that this model can be used as a reference for any further extensions and improvements. The specifics of the Lightning Network

(mainly capacity constraints that adjust dynamically) makes it also interesting from a pure mathematical point of view. The study of this particular network has the potential to significantly improve what we understand about networks in general.

The very nature of the Lightning Network and the blockchain itself could raise the question of the goal of this research. Indeed, those systems were built to be self regulated, i.e. regulated by the software itself. The blockchain builds consensus, therefore everything that is recorded on the blockchain is the ground truth, and the Lightning Network is essentially a set of contracts lying on the blockchain. The consequence is that any disagreement on the channel balances would ultimately be resolved by the blockchain and the rule written within its source code. No regulator is needed for the blockchain or the Lightning Network and everyone is free to use it however they want. Also, it can be very hard to change actual rules, because every user needs to agree on protocol changes. It is important to understand that the blockchain and the Lightning Network cannot be regulated in a traditional way. Ultimately the developers are the ones who decide the rules and write them inside the source code, after which they cannot be changed⁴³. What each government can do is to allow or forbid the use of those technologies. Therefore, the goal should be to understand the mechanics and see in what cases it produces a stable output that could eventually and safely be used by end users across the world.

⁴³We refer here as changes in consensus protocol and communication standards. Small changes like the algorithm used to optimally route transactions is not part of consensus and, therefore, it can be updated easily.

4.9 Appendix

4.9.1 Proofs of Lemmas

Proof of Lemma 4.1

The proof of this lemma is provided within the proof of proposition 4.1.

Proof of Lemma 4.2

The proof of this lemma is provided within the proof of proposition 4.1.

Proof of Lemma 4.3

In this subgame, nodes must find the routing fee that maximizes their profit, R^* , taking the channel and network constraints, (x, λ) , as given. With the assumption we made, nodes get to routes payment only if the routing fee they set is lower or equal to all other nodes' routing fee. Otherwise stated, nodes are competing on price to attract volume, and a price too high would yield zero volume as end-users would select nodes with lower price.

Let's define by $\hat{R}(x, \lambda)$ the routing fee that produces a zero profit given (x, λ) . Conditioned on attracting volume, if $R > \hat{R}$, the node will make a positive profit and a loss otherwise.

$$\hat{R}(x, \lambda) = \frac{x^2 + 2\lambda}{2(1 - e^{-x})(1 - e^{-\lambda})} \quad (4.9)$$

Let's consider node i , and check when he benefits from deviating in the candidate equilibrium. Assume node i expects all other nodes to play $R_{-i} > \hat{R}$. Then node i can increase its profit by attracting all the volume and playing R^* such that $\hat{R} < R^* < R_{-i}$. The profit would still be positive and the node gets to route more volume. The same mechanics can be applied to any node to show that the routing fee will be driven toward $R^* = \hat{R}$.

If node i expects all nodes to play $R_{-i} = \hat{R}$, he does not benefit from increasing or decreasing the fee. If he increases it, he would get no volume, and if he decreases it, he would take all the volume but at a loss.

If node i expects all nodes to play $R_{-i} < \hat{R}$, he cannot benefit from playing something other than $R^* = R_{-i}$. If it plays higher it won't attract any volume, and if it plays less, it would increase the loss. Therefore all prices $R^* \in [0; \hat{R}[$ are also Nash equilibrium of the subgame.

Q.E.D

Proof of Lemma 4.4

The proof of this lemma is provided within the proof of proposition 4.3.

Proof of Lemma 4.5

In this subgame, nodes must find the routing fee that maximizes their profit, R^* , taking the channel and network constraints, $(\{x_i\}_i; \{\lambda_i\}_i)$, as given.

The argument is exactly the same as for lemma 4.3, except that the mechanisms in lemma 4.3 applies for any interconnectedness level, λ_i , because of the assumption we made regarding the demand (infinite with any interconnectedness level).

Define $\hat{R}(x, \lambda)$ by

$$\hat{R}_i(x, \lambda) = \frac{\gamma_i x^2 + 2\lambda}{2(1 - e^{-x})(1 - e^{-\lambda})}$$

We can show that if $R_i^* \geq \hat{R}_i(x_i, \lambda_i)$, there will be an arbitrage opportunity. However, once again, for each value of γ_i , any price between 0 and $\hat{R}_i(x_i, \lambda_i)$ will be a Nash equilibrium, in which nodes make a loss.

Q.E.D

4.9.2 Proofs of Propositions

Proof of Proposition 4.1

Recall that we are in the one-shot version of the model, where R is taken as exogenous. The program the node needs to solve is the following.

$$(x^*, \lambda^*) = \arg \max_{x \geq 0, \lambda \geq 0} \Pi(x, \lambda).$$

with

$$\Pi(x, \lambda) = (1 - e^{-\lambda})(1 - e^{-x})R - \frac{x^2}{2} - \lambda$$

The first order condition gives

$$\begin{cases} \frac{\partial}{\partial x} \Pi(x, \lambda) = e^{-x}(1 - e^{-\lambda})R - x = 0 \\ \frac{\partial}{\partial \lambda} \Pi(x, \lambda) = e^{-\lambda}(1 - e^{-x})R - 1 = 0 \end{cases}$$

Combining the 2 equations and rearranging the terms yields the following FOC.

$$\lambda^* = \ln(1 + x^*(e^{x^*} - 1)) \quad (4.10)$$

We can now plug that relationship into the first derivative of the objective function with respect to λ .

$$\frac{1 + x^*(e^{-x^*} - 1)}{1 - e^{x^*}} = R \quad (4.11)$$

There is no closed form solution for this equation. However, we can numerically demonstrate some important properties. Let's define $g(x)$ by

$$g(x) := \frac{1 + x(e^{-x} - 1)}{1 - e^x}.$$

Figure 4.11 presents the graph of the function g . Given the definition of the function g , x^* is such that $g(x^*) = R$. The function g admits a minimum, denoted by $\bar{R} > 0$.

$$\bar{R} = \min_{x>0} g(x)$$

This tells us that for $R < \bar{R}$ there are no value of x that can satisfy equation 4.11. For values of R such that $R > \bar{R}$, there are two solutions to the first order conditions, x_1 and x_2 with $x_1 < x_2$. However it can be shown that x_1 is actually a saddle point. The second order condition tells us that $x^* = x_2$, when $R > \bar{R}$, is a maximum of the function $\Pi(x, \lambda)$ on the domain $x \geq 0$ and $\lambda \geq 0$, with λ^* given by (4.10).




Figure 4.11: Graph of the function $g(x)$

The full solution is therefore characterised by the following.

$$x^* = \begin{cases} 0 & \text{when } R < \bar{R} \\ x_2 & \text{when } R \geq \bar{R} \end{cases}$$

The function f defined in the proposition is the inverse of the g function for values of x greater than the $\arg \min$ of $\Pi(x, \lambda)$.

Q.E.D

Proof of Proposition 4.2

We start by focusing our attention on equilibrium such that $\forall i \neq j; x_i^* = x_j^*$. The argument bellow could be used to rule out those cases. So keep the notation (x, λ) for the subgames. Equilibria of the full game are defined as (x^*, λ^*, R^*) .

First, let's consider the following points

$$\forall R; (x, \lambda, R) = (0, 0, R).$$

For the same reason as in proposition 4.1, those points cannot be equilibrium because the demand is completely unsatisfied and any node could jump into the network with any arbitrary $(x, \lambda) \in \mathbb{R}_+^* \times \mathbb{R}_+^*$, charge a corresponding fee and make a positive profit.

With proposition 4.1 we know that x^* and λ^* are linked in equilibrium with the following relation.

$$\lambda^* = \lambda(x^*) = \ln(1 + x(e^x - 1))$$

therefore we can restrict our notation to

$$(x^*, \lambda^*, R^*) = (x^*, \lambda(x^*), R^*) = (x^*, R^*).$$

Let's now define, \bar{x} , a critical point for our study as follows.

$$\bar{x} := \left\{ x \in \mathbb{R}_+^* ; \frac{1 + x(e^x - 1)}{1 - e^{-x}} = \frac{x^2 + 2\lambda(x)}{2(1 - e^{-x})(1 - e^{-\lambda(x)})} \right\}$$

It is possible to show that \bar{x} has a single possible value. We cannot prove this analytically but, first we can rewrite the equation defining \bar{x} as $h(x) = 0$.

$$h(x) = x^2 + 2\ln(1 + x(e^x - 1)) - 2x(e^x - 1) = 0$$

Then, when we plot the graph of $h(x)$ in figure 4.12, we can see that it has a single solution. The point $x = 0$ is also a solution of this equation but we can safely discard it for the reason used to prove that $(0, 0, R)$ is not an equilibrium. Note that the model is up to some proportional constant, but changing the scale of variables

should not change the curvature of the function and therefore the existence of the solution should extend to any different scale.




Figure 4.12: Graph of the function $h(x)$

Note that \bar{x} and the corresponding fee $\bar{R}(\bar{x}, \lambda(\bar{x}))$, as defined by (4.9), define the point for which the profit is maximized but equal to zero. At this point the node cannot increase its profit by playing around with capacity constraints (x, λ) . Also the fee is such that the node makes zero profit. Decreasing this fee would increase the loss, and increasing the fee would not attract any volume, so resulting in a loss as well. Now consider possible values of x compared to \bar{x} .

Case 1 : When $x > \bar{x}$. By lemma 4.3 we know that either $R = \hat{R}$ as defined by (4.9), or $R \in [0; \hat{R}[$. Recall that \hat{R} yields zero profit. Values such that $R \neq \hat{R}$ cannot be equilibrium because all nodes would make a loss and therefore prefer a subgame such as $(x, \lambda, R) = (0, 0, R)$ which, as previously seen, cannot be an equilibrium. If $R = \hat{R}$, it cannot be considered an equilibrium either because node i could play an x' such that

$$\bar{x} < x' < x$$

which would increase its profit. Indeed, at $x > \bar{x}$ it is profitable for the nodes to play a smaller value of x , it will make the profit increase to its maximum as the node approaches $x = \bar{x}$.

Case 2 : When $x < \bar{x}$. For the same reason, the points $(x, \lambda(x), 0)$ cannot be considered equilibrium. If $R = \hat{R}$, it cannot be considered an equilibrium as well

because node i could play an x' such that

$$x < x' < \bar{x}$$

which would increase its profit. Indeed, at $x < \bar{x}$ it is profitable for the nodes to play a larger value of x , it will make the profit increases to its maximum as the node approaches $x = \bar{x}$.

Case 3 : When $x = \bar{x}$. For the same reason as above, $R < \hat{R}$ cannot be an equilibrium. When $R = \hat{R}$, no node can benefit from playing something else. If they move x , keeping the same R , they would start to make a loss, as \bar{x} is defined to be the objective function maximum, holding R constant. And moving R would either increase their loss or make them unable to attract any volume. Therefore this point is an equilibrium.

Only the last case is an equilibrium, so without loss of generalities we can say that the only equilibrium is the following

$$(x^*, \lambda^*, R^*) = (\bar{x}, \lambda(\bar{x}), \hat{R}(\bar{x}, \lambda(\bar{x}))) .$$

Q.E.D

Proof of Proposition 4.3

This proof is very similar to the proof of proposition 4.1. We don't detail everything, only the main difference compared to the homogeneous agents' case.

First of all, the first order conditions give

$$\begin{cases} \frac{\partial}{\partial x} \Pi(x, \lambda) = e^{-x}(1 - e^{-\lambda})R - \gamma x = 0 \\ \frac{\partial}{\partial \lambda} \Pi(x, \lambda) = e^{-\lambda}(1 - e^{-x})R - 1 = 0 \end{cases}$$

and it can be re written as

$$\lambda^* = \ln(1 + \gamma x^*(e^{x^*} - 1)) .$$

We can now plug that relationship into the first derivative of the objective function with respect to λ .

$$\frac{1 + \gamma x^*(e^{-x^*} - 1)}{1 - e^{x^*}} = R \tag{4.12}$$

Again, there is no closed form solution for this equation. However, we can

numerically demonstrate some important properties. Let's define $g(x)$ by

$$g(x) := \frac{1 + \gamma x(e^{-x} - 1)}{1 - e^x}.$$

We can graphically see that the minimum of this function for positive values of x , increases with γ . Given the definition of the function g , x^* is such that $g(x^*) = R$. The function g admits a minimum, denoted by $\bar{R} > 0$.

$$\bar{R} = \min_{x>0} g(x)$$

Relation (4.6) and (4.7) also come from graphical interpretation of the function $g(x)$.

Q.E.D

Proof of Proposition 4.4

The proof is very similar to the proof of proposition 4.2. The equilibrium mechanisms are exactly the same as for the homogeneous case except everything happens within each cohort that shares the same opportunity cost of capital, γ . Note that this is fine even if the opportunity cost is continuous because we assumed that a continuum of potential nodes can jump into the network with any γ if they see an arbitrage opportunity. This will drive the price further down until the no profit condition is satisfied, for each γ .

The rest of the proof is exactly the same as in proposition 4.2.

Q.E.D

Proof of Proposition 4.5

It can be shown that all the equilibrium applications are monotonic and continuous. Also, as long as $\gamma \in]0; \gamma_{max}]$ (and $\gamma_{max} < +\infty$), the equilibrium applications, $x(\gamma)$ and $\lambda(\gamma)$, will converge to finite values.

Therefore, if the support of the γ -distribution is finite, the equilibrium value of (x, λ) will also be finite, bounded by the γ -distribution's bound.

Q.E.D

Bibliography

- Acharya, Viral, and Alberto Bisin, 2014, Counterparty risk externality: Centralized versus over-the-counter markets, *Journal of Economic Theory* 149, 153–182.
- Akerlof, George A, 1978, The market for “lemons”: Quality uncertainty and the market mechanism, in *Uncertainty in economics* . pp. 235–251 (Elsevier).
- Albert, Réka, and Albert-László Barabási, 2002, Statistical mechanics of complex networks, *Reviews of modern physics* 74, 47.
- Algan, Yann, Sergei Guriev, Elias Papaioannou, and Evgenia Passari, 2017, The european trust crisis and the rise of populism, *Brookings Papers on Economic Activity* 2017, 309–400.
- Allen, Douglas W, 1999, Transaction costs, *Encyclopedia of law and economics*.
- Alvarez, Fernando, and Francesco Lippi, 2009, Financial innovation and the transactions demand for cash, *Econometrica* 77, 363–402.
- Amini, Hamed, Damir Filipović, and Andreea Minca, 2015, Systemic risk and central clearing counterparty design, *Swiss Finance Institute Research Paper*.
- Amini, Hamed, Damir Filipovic, and Andreea Minca, 2017, Systemic risk with central counterparty clearing, *Working Paper*.
- Anderson, Stacey, Jean-Philippe Dion, and Hector Perez Saiz, 2013, To link or not to link? netting and exposures between central counterparties, *Bank of Canada Working Paper*.
- Armakola, Angela, and Jean-Paul Laurent, 2017, Ccp resilience and clearing membership, *Working Paper*.
- Baeriswyl, Romain, 2011, Endogenous central bank information and the optimal degree of transparency, *International Journal of central banking* 7, 85–111.
- , and Camille Cornand, 2010, The signaling role of policy actions, *Journal of Monetary Economics* 57, 682–695.

- , and Bruno Ziliotto, 2016, Observing and shaping the market: the dilemma of central banks, *Working Papers*.
- Bailey, Warren, Lin Zheng, and Yinggang Zhou, 2014, What makes the vix tick?, *Working Papers*.
- Baker, Scott R, Nicholas Bloom, and Steven J Davis, 2016, Measuring economic policy uncertainty, *The quarterly journal of economics* 131, 1593–1636.
- Benati, Luca, 2016, Economic policy uncertainty, the great recession, and the great depression, *manuscript, University of Bern*.
- Bernanke, Ben, Vincent Reinhart, and Brian Sack, 2004, Monetary policy alternatives at the zero bound: An empirical assessment, *Brookings papers on economic activity* 2004, 1–100.
- Bernstein, Marver H, 1955, *Regulating business by independent commission* . vol. 2324 (Princeton University Press).
- Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, and Catherine Casamatta, 2019, The blockchain folk theorem, *The Review of Financial Studies* 32, 1662–1715.
- Biais, Bruno, Florian Heider, and Marie Hoerova, 2016, Risk-sharing or risk-taking? counterparty risk, incentives, and margins, *The Journal of Finance* 71, 1669–1698.
- Bianconi, Ginestra, and A-L Barabási, 2001, Competition and multiscaling in evolving networks, *EPL (Europhysics Letters)* 54, 436.
- Bignon, Vincent, and Guillaume Vuillemeys, 2017, The failure of a clearinghouse: Empirical evidence, .
- Blinder, Alan S, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen, 2008, Central bank communication and monetary policy: A survey of theory and evidence, *ECB Working Paper Series* 898.
- Bloom, Nicholas, 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.
- Brogaard, Jonathan, and Andrew Detzel, 2015, The asset-pricing implications of government economic policy uncertainty, *Management Science* 61, 3–18.
- Cisternas, Gonzalo, 2017, Two-sided learning and the ratchet principle, *The Review of Economic Studies* 85, 307–351.
- Coase, Ronald Harry, 1937, The nature of the firm, *economica* 4, 386–405.

- Coase, Ronald H, 1960, The problem of social cost, in *Classic papers in natural resource economics* . pp. 87–137 (Springer).
- Cont, Rama, and Thomas Kokholm, 2014, Central clearing of otc derivatives: bilateral vs multilateral netting, *Statistics & Risk Modeling* 31, 3–22.
- CPMI-IOSCO, 2012, Principles for financial market infrastructures., *Manuscript*.
- , 2014, Recovery of financial market infrastructures., *Manuscript*.
- , 2016, Resilience and recovery of central counterparties (ccps): Further guidance on the pfmi., *Manuscript*.
- Davidson, Carl, and Raymond Deneckere, 1986, Long-run competition in capacity, short-run competition in price, and the cournot model, *The Rand Journal of Economics* pp. 404–415.
- Duffie, Darrell, Martin Scheicher, and Guillaume Vuillemeys, 2015, Central clearing and collateral demand, *Journal of Financial Economics* 116, 237–256.
- Duffie, Darrell, and Haoxiang Zhu, 2011, Does a central clearing counterparty reduce counterparty risk?, *The Review of Asset Pricing Studies* 1, 74–95.
- Easley, David, Maureen O’Hara, and Soumya Basu, 2019, From mining to markets: The evolution of bitcoin transaction fees, *Journal of Financial Economics*.
- Eggertsson, Gauti B, and Michael Woodford, 2003, The zero bound on short-term interest rates and optimal monetary policy, *Brookings papers on economic activity* pp. 139–211.
- Ehrmann, Michael, and Marcel Fratzscher, 2007, Communication by central bank committee members: Different strategies, same effectiveness?, *Journal of Money, Credit and Banking* 39, 509–541.
- Ellickson, Robert C, and Charles DiA Thorland, 1995, Ancient land law: Mesopotamia, egypt, israel, *Chi.-Kent L. Rev.* 71, 321.
- Eusepi, Stefano, and Bruce Preston, 2010, Central bank communication and expectations stabilization, *American Economic Journal: Macroeconomics* 2, 235–271.
- Faust, Jon, and Lars EO Svensson, 2001, Transparency and credibility: Monetary policy with unobservable goals, *International Economic Review* 42, 369–397.
- Fischer, Michael J, Nancy A Lynch, and Michael S Paterson, 1985, Impossibility of distributed consensus with one faulty processor., *Journal of the ACM* 32(2), 374–382.

- Freeman, Linton C, 1977, A set of measures of centrality based on betweenness, *Sociometry* pp. 35–41.
- Freixas, Xavier, and Bruno Parigi, 1998, Contagion and efficiency in gross and net interbank payment systems, *Journal of Financial Intermediation* 7, 3–31.
- Gregory, Jon, 2014, *Central Counterparties: Mandatory Central Clearing and Initial Margin Requirements for OTC Derivatives* (John Wiley and Sons).
- Grenadier, Steven R, Andrey Malenko, and Nadya Malenko, 2016, Timing decisions in organizations: Communication and authority in a dynamic environment, *The American Economic Review* 106, 2552–2581.
- Grossman, Gene M, 2016, The purpose of trade agreements, in *Handbook of commercial policy* vol. 1 . pp. 379–434 (Elsevier).
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2008, Trusting the stock market, *the Journal of Finance* 63, 2557–2600.
- Hansen, Stephen, and Michael McMahon, 2016, Shocking language: Understanding the macroeconomic effects of central bank communication, *Journal of International Economics* 99, S114–S133.
- Heller, Daniel, and Nicholas Vause, 2012, Collateral requirements for mandatory central clearing of over-the-counter derivatives, *BIS Working Paper No 373*.
- Huang, Wenqian, 2016, Central counterparty capitalization and misaligned incentives, *Working Paper*.
- James, Jonathan G, and Phillip Lawler, 2011, Optimal policy intervention and the social value of public information, *The American Economic Review* 101, 1561–1574.
- Jensen, Michael C, and William H Meckling, 1976, Theory of the firm: Managerial behavior, agency costs and ownership structure, *Journal of financial economics* 3, 305–360.
- Joskow, Paul L, and Nancy L Rose, 1989, The effects of economic regulation, *Handbook of industrial organization* 2, 1449–1506.
- Kahn, Charles M, and William Roberds, 1998, Payment system settlement and bank incentives, *The review of financial studies* 11, 845–870.
- Keynes, John Maynard, 1937, The general theory of employment, *The quarterly journal of economics* 51, 209–223.

- Kreps, David M, and Jose A Scheinkman, 1983, Quantity precommitment and bertrand competition yield cournot outcomes, *The Bell Journal of Economics* pp. 326–337.
- Laffont, Jean-Jacques, and Jean Tirole, 1986, Using cost observation to regulate firms, *Journal of political Economy* 94, 614–641.
- , 1991, The politics of government decision-making: A theory of regulatory capture, *The quarterly journal of economics* 106, 1089–1127.
- , 1993, *A theory of incentives in procurement and regulation* (MIT press).
- Lamport, Leslie, Robert Shostak, and Marshall Pease, 1982, The byzantine generals problem, *ACM Transactions on Programming Languages and Systems (TOPLAS)* 4, 382–401.
- Lasry, Jean-Michel, and Pierre-Louis Lions, 2007, Mean field games, *Japanese journal of mathematics* 2, 229–260.
- Leitner, Yaron, 2005, Financial networks: Contagion, commitment, and private sector bailouts, *The Journal of Finance* 60, 2925–2953.
- Levine, Michael E, and Jennifer L Forrence, 1990, Regulatory capture, public interest, and the public agenda: Toward a synthesis, *JL Econ & Org.* 6, 167.
- Lions, Pierre-Louis, 2016, Equation de HJB et extensions de la théorie classique du controle stochastique, in *Public Lectures at College de France, Paris*.
- Loon, Yee Cheng, and Zhaodong Ken Zhong, 2014, The impact of central clearing on counterparty risk, liquidity, and trading: Evidence from the credit default swap market, *Journal of Financial Economics* 112, 91–115.
- Lucas, Robert E, et al., 1976, Econometric policy evaluation: A critique, in *Carnegie-Rochester conference series on public policy* vol. 1 pp. 19–46.
- Menkveld, Albert J, 2017, Crowded positions: An overlooked systemic risk for central clearing parties, *The Review of Asset Pricing Studies* 7, 209–242.
- Moore, Cristopher, Gourab Ghoshal, and Mark EJ Newman, 2006, Exact solutions for models of evolving networks with addition and deletion of nodes, *Physical Review E* 74, 036121.
- Moreno, Diego, and Luis Ubeda, 2006, Capacity precommitment and price competition yield the cournot outcome, *Games and Economic Behavior* 56, 323–332.

- Morris, Stephen, and Hyun Song Shin, 2002, Social value of public information, *The American Economic Review* 92, 1521–1521.
- Murphy, David, 2013, *OTC Derivatives: Bilateral Trading and Central Clearing: An Introduction to Regulatory Policy, Market Impact and Systemic Risk* (Springer).
- , 2017, I’ve got you under my skin: large central counterparty financial resources and the incentives they create, *Journal of Financial Market Infrastructures* 5, 57–74.
- Nakamoto, Satoshi, et al., 2008, Bitcoin: A peer-to-peer electronic cash system, .
- North, Douglass C, 1991, Institutions, *Journal of economic perspectives* 5, 97–112.
- North, Douglass Cecil, and Douglass Cecil North, 1992, *Transaction costs, institutions, and economic performance* (ICS Press San Francisco, CA).
- on Banking Supervision, Basel Committee, 2005, An explanatory note on the basel ii irb risk weight function, .
- Orphanides, Athanasios, and John C Williams, 2006, Monetary policy with imperfect knowledge, *Journal of the European Economic Association* 4, 366–375.
- Pastor, Lubos, and Pietro Veronesi, 2012, Uncertainty about government policy and stock prices, *The Journal of Finance* 67, 1219–1264.
- , 2013, Political uncertainty and risk premia, *Journal of Financial Economics* 110, 520–545.
- Pham, Thong, Paul Sheridan, and Hidetoshi Shimodaira, 2015, Pafit: A statistical method for measuring preferential attachment in temporal complex networks, *PloS one* 10, e0137796.
- Pirrong, Craig, 2009, The economics of clearing in derivatives markets: Netting, asymmetric information, and the sharing of default risks through a central counterparty, .
- Poon, Joseph, and Thaddeus Dryja, 2016, The bitcoin lightning network: Scalable off-chain instant payments, .
- Ranaldo, Angelo, and Enzo Rossi, 2010, The reaction of asset markets to swiss national bank communication, *Journal of International Money and Finance* 29, 486–503.
- Rochet, Jean-Charles, and Jean Tirole, 1996, Controlling risk in payment systems, *Journal of Money, Credit and Banking* 28, 832–862.

- Rohrer, Elias, Julian Malliaris, and Florian Tschorsch, 2019, Discharged payment channels: Quantifying the lightning network's resilience to topology-based attacks, *arXiv preprint arXiv:1904.10253*.
- Saleh, Fahad, 2018, Blockchain without waste: Proof-of-stake, *Working Paper*.
- Seres, István András, László Gulyás, Dániel A. Nagy, and Péter Burcsi, 2019, Topological analysis of bitcoin's lightning network, *CoRR* abs/1901.04972.
- Sidanius, Che, and Filip Zikes, 2012, Otc derivatives reform and collateral demand impact, *Financial Stability Paper* 18.
- Smith, Adam, 1822, *The theory of moral sentiments* . vol. 1 (J. Richardson).
- Stein, Jeremy C, and Adi Sunderam, (forthcoming), The fed, the bond market and gradualism in monetary policy, *Journal of Finance*.
- Stephens, Eric, and James R Thompson, 2013, Lemons and proud of it: Information asymmetry and risk transfer markets, *Working Paper*.
- Stigler, George J, 1971, The theory of economic regulation, *The Bell journal of economics and management science* pp. 3–21.
- Thompson, James R, 2010, Counterparty risk in financial contracts: Should the insured worry about the insurer?, *The Quarterly Journal of Economics* 125, 1195–1252.
- Vartiainen, Hannu, 2014, Mechanism design without commitment; general solution and application to bilateral bargaining, *Working Papers*.
- Williamson, Oliver E, 1975, Markets and hierarchies, *New York* 2630.
- , 1985, Reflections on the new institutional economics, *Zeitschrift für die gesamte Staatswissenschaft/Journal of Institutional and Theoretical Economics* pp. 187–195.
- , 1998, The institutions of governance, *The American Economic Review* 88, 75–79.
- Zakai, Moshe, 1969, On the optimal filtering of diffusion processes, *Probability Theory and Related Fields* 11, 230–243.

RÉSUMÉ

A travers les siècles, les institutions financières ont façonné le paysage financier et influencé l'activité économique. L'objectif de cette thèse est de mettre en évidence, d'un point de vue théorique, les limites fondamentales des institutions modernes et d'en déduire les implications concernant le futur rôle de ces institutions.

Le premier chapitre propose une analyse des Chambres de Compensation. A la suite de la crise financière de 2008, les autorités financières à travers le monde ont mis en place des réglementations imposant la compensation centrale sur la plupart des produits dérivés. Nous montrons que la compensation centrale nécessite un plus grand niveau de liquidité que la compensation bilatérale.

Le second chapitre présente un modèle d'apprentissage en temps continu censé représenter le processus d'apprentissage d'une institution par rapport à une information cachée détenue par le marché. A l'équilibre, le niveau d'incertitude perçue par les agents, constitue une limite fondamentale du rôle des institutions financières.

Le dernier chapitre introduit et analyse le Lightning Network qui est un réseau de paiement basé sur la Blockchain. Il permet aux utilisateurs de transférer de la valeur de façon instantanée sans avoir recours à un tiers de confiance. Nous discutons des implications à propos de la structure de ce réseau de paiement ainsi que de sa capacité à prendre une place importante dans le paysage financier.

MOTS CLÉS

Institutions financières ; Compensation centrale ; Risque de liquidité ; Apprentissage ; Blockchain ; Réseau de paiement

ABSTRACT

Throughout the centuries, financial institutions have shaped the financial landscape and influenced economic activity. The goal of this dissertation is to highlight, from a theoretical point of view, fundamental limitations of modern institutions and eventually derive implications regarding the future role of those institutions.

The first chapter provides an analysis of Central Clearing Platforms (CCP). In the aftermath of the financial crisis of 2008, financial authorities around the world implemented regulations imposing central clearing on most derivative products. It is shown that central clearing often requires a larger liquidity buffer than bilateral clearing.

The second chapter presents a continuous-time learning model meant to represent the learning process of an institution such as a central bank regarding a hidden information held by the market. The equilibrium level of uncertainty perceived by agents is shown to be an important limitation to the role of financial institutions.

The last chapter introduces a specific blockchain-based payment network called the Lightning Network. It allows users to transfer value instantly without relying on any trusted third party. We discuss the implications regarding the structure of this network as well as its ability to become an important part of the financial landscape.

KEYWORDS

Financial institutions ; Central clearing ; Liquidity risk ; Learning ; Blockchain ; Payment network