Money, Intermediation and Coordination in Decentralised Markets

by

Christian Hellwig

B.Sc. University of Lausanne (1998)
M.Sc. London School of Economics (1999)

Thesis submitted to the University of London
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

UNIVERSITY OF LONDON

LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

19 November 2002

Money, Intermediation, and Coordination in Decentralised Markets

by

Christian Hellwig

Submitted to the University of London
on 19 November 2002, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Overview:

This thesis studies the coordination of individuals' transactions in a large, decentralized market. The first half of the thesis examines the role of market institutions in an environment with frictions. In particular, it studies the interaction between "intermediaries" (banks, shops) and decentralized "equilibrium arrangements" such as money or credit. The second half of the thesis studies the role of public and private information in coordinating individual actions, as well as the macro-economic effects of such coordination.

First Half:

I study a search economy in which intermediaries are the driving force co-ordinating the economy on the use of a unique, common medium of exchange for transactions. If search frictions delay trade, intermediaries offering immediate exchange opportunities can make arbitrage gains from a price spread, but they have to solve the search market's allocation problem. Intermediaries solve this problem best by imposing a common medium of exchange to other agents, and a Cash-in-Advance constraint arises in equilibrium: Agents trade twice in order to consume, once to exchange their production against the medium of exchange, and once to purchase their consumption. I extend my analysis to the study of fiat currencies and, in particular, free banking regimes.

Second Half:

The second half consists of two essays studying the role of public and private information in coordination games. In the first, I relate the convergence of equilibria to the convergence of higher-order beliefs. The central result of this essay relates the convergence of players' higher-order beliefs (and hence equilibrium convergence) to the parameters of the signal structure. This provides a limit condition determining the uniqueness or multiplicity of equilibria in the coordination game.

Building on the previous chapter, the second essay studies the effects of information policies on output and inflation, when price-setters face higher-order uncertainty concerning the money supply. I show that this may lead to substantial delays in price-adjustment following a shock.

To the extent that public announcements coordinate expectations, they reduce this delay, and thereby reduce the effect and persistence of monetary shocks on output in the short-run. On the other hand, public announcements introduce a second component of noise, and may therefore increase short-run volatility.

Thesis Supervisor: Nobuhiro Kiyotaki
Title: Professor, London School of Economics

Contents

Acknowledgements

A thesis is never the work of a single person (even if the author at times would like to think so), and in my own case there are too many people that I am indebted to for their help, their encouragement, and for their support, to do justice to all of them in these few lines. First and foremost, I would like to thank my supervisor, Nobu Kiyotaki, for the many hours of discussion and plenty of very useful advice, the ton of little things that I learned from him over the years, as well as the one important thing: how to do research, and how to think about economics. One of my earliest acts at LSE was to visit his office to talk to him about monetary theory (I wasn't yet a graduate student then...), and at times, I wonder what course things would have taken, had not gone to talk to Nobu.

Besides Nobu, many faculty members, at LSE and elsewhere, were very supportive of my research and my studies. I am extremely grateful for the time, the teaching and the discussions with Margaret Bray, Godfrey Keller, Thomas Mariotti, John Moore and Hyun Shin at LSE, Bengt Holmstrom and Rob Townsend during my visit at MIT, as well as Martin Hellwig (Mannheim) and Stephen Morris (Yale). Over time, each of them significantly influenced the development of the ideas in this thesis, and of my thoughts about economics and research in general. I would further like to thank Bengt Holmstrom and Stephen Morris in particular for their support of my candidacy on the academic job market.

I would also like to thank my co-students and colleagues; many of them have become friends and a few co-authors over the course of these years. Through them, I have learned that research is a highly social activity that is most entertaining and enjoyable when done as a group. I have had valuable comments, new perspectives and insights from many people, but I would most like to thank Heski Bar-Isaac, George-Marios Angeletos, Guido Lorenzoni and Alessandro Pavan for the many shared hours bouncing ideas off each other and thinking about economics. Discussions with the last three have since led to the co-authoring of further papers which expand on ideas that are touched upon in this thesis, and I do hope that I will also get a chance to work together with the first in the group.

I am also grateful for the support that I received from the institutions I attended - the Economics department at LSE and the ESRC for financial support throughout my PhD. studies; the Financial Markets Group for providing a wonderful research environment, and the MIT for its hospitality and again a wonderful place to do research. Furthermore, I would like to thank the Economics department at the University of Lausanne, where I was able to work for a few weeks each summer. The calm and beauty of that place provided an excellent environment for intensive work, and the first and third chapter of this thesis in particular have been written and extensively revised there at various stages of their development.

Finally, I thank my wife Elisa and my family for their love and support throughout the last few years. Completing this thesis would have hardly been possible without, nor would it have been worthwhile.

0.1 Introduction and Overview

This thesis presents a theoretical investigation into the forces which determine how a large number of independent actors are able to *coordinate* their actions in environments, where decisions are taken independently. The coordination problems studied in this thesis go beyond the simple coordination implied by the notion of a Nash equilibrium, rather, the difficulties arise because of a multiplicity of Nash equilibria. I first study, how evolutionary forces may lead to efficient coordination on a medium of exchange in a market environment with intermediaries. In the second half of the thesis, I study how lack of information may prevent efficient coordination to take place. To summarize the results in this thesis in one sentence, "In a stable market environment, evolutionary forces eliminate suboptimal outcomes in the long-run, while in the short run, or in response to shocks, incomplete information may lead to substantial delays of adjustment, or a failure to coordinate on an efficient outcome."

Of course, this one-sentence summary is too abstract to be directly useful or applicable, and to the outside reader, the conclusions may not appear to be that novel either. The purpose of this introduction is therefore to discuss briefly in more detail the motivations for this study, discuss more specifically the environments that are used for establishing the main results, and comment on how, in my view, they fit into the large picture, i.e. the literature and the state of knowledge of economic theory.

In the first half of the thesis, I study the emergence of money and and intermediaries in a decentralized trading economy. In such an environment, the coordination problem lies in the selection of a medium of exchange, and complementarities arise between the trading strategies of different agents. The study was motivated by the historical observation that intermediaries play a central role in introducing and using media of exchange, or more generally are a constant force of innovation in the transactions process. The paper then proposes a reason for why this would be the case, in an environment where intermediaries respond to the existing market frictions. Since intermediaries are typically restricted in the number of goods that they deal with, i.e. they cannot handle all commodities at once, their success in alleviating frictions crucially depends on their ability to provide their customers ways of carrying purchasing power from one intermediary to another, which motivates the interest that the intermediaries have in promoting a common medium of exchange.

Furthermore, intermediaries have the ability to introduce such a medium of exchange, if they mutually coordinate on its use. Under certain conditions, I show that the unique stable equilibrium of the trading game that I study is one in which fiat money is used as a common medium of exchange, all producers and consumers trade through intermediaries, and fiat money is used in every transaction. There do exist other stationary equilibria, however, in all of them, the explicit coordination of strategies by even a small set of players will induce the remaining players to break away from the equilibrium strategies.

To understand the significance of these results, it may be useful to think of the existing prior frameworks for monetary economics. The classical approach consists in augmenting the atemporal Arrow-Debreu economy, or the dynamic general equilibrium model by some assumptions that justify the circulation of money. Typically, this involves the imposition of a Cash-in-Advance constraint, i.e. the assumption that cash must be used for the purchase of commodities, or in less explicit form, involves modelling the liquidity services of money directly by embedding them in the utility or production functions of households and firms.[1] This approach is not directly concerned with the justification of such a constraint, but rather serves as a workhorse which has been extremely influential in monetary macroeconomics.

Introducing the use of money into general equilibrium theory in this form leads to the challenge of providing a foundation for monetary theory, i.e. a foundation for the use of such constraints. But the general equilibrium model provides an even bigger challenge, namely that of thinking about what markets are, how they are established, how they evolve, and what the respective roles of different market participants are, i.e. who determines how commodities are exchanged, by whom they are exchanged, and at what prices. This question is shortcircuited entirely by the general equilibrium model, which assumes that prices are set to clear markets, i.e. equalize demand and supply for each commodity, and hence remains silent about the entire process of interaction.

[1] More elaborate versions of the same theme are the shopping time model, in which money saves on time spent purchasing goods, and hence leaves more time for production and leisure, and the cash-credit economy, in which only a subset of goods are bought using money.

The challenge of thinking about market evolution is theoretical more than it is empirical: it seems plausible that markets behave almost as if they were perfect, i.e. they clear, and the exact structure of interaction seems to have little effect on prices and allocation. Furthermore, once the perfect functioning of markets is assumed as empirically plausible, it is easy to agree that the role of money is precisely as a convenient way of exchanging goods, and that a Cash-in-Advance constraint exactly reflects the way in which market transactions are carried out.

The challenge of providing a foundation for money, as well as market interaction then implies going beyond the Walrasian model. One class of models that has taken up this challenge, and attempts to provide a formal foundation for the use of money, is the search model. This model assumes that transactions occur in bilateral meetings, and are temporarily delayed by a lack of suitable trading partners - market participants have to search (and find) them, before they can trade. It is thus shown that the use of a common medium of exchange reduces the delays in transactions that would arise from the double coincidence problem in a barter economy; it is then not only socially desirable, but also individually rational that market participants agree on the use of a common medium of exchange. The search model takes an explicitly evolutionary view of money, in the sense that it views the use of a medium of exchange as the steady-state equilibrium of an evolutionary process of individually rational transaction decisions. The model itself, however, runs up against some powerful criticisms: First, it inherently leads to multiple equilibria, and a selection problem: since the 'saleability' of goods (the relative ease with which they are traded) is endogenous to their use in transactions, in principle any good could become a medium of exchange, provided that it is sufficiently durable and not to costly to transact - as historical examples have shown, the minimum requirements of durability and ease to transact are not very stringent; how else could one reconcile the use of such goods as raw tobacco or cattle as media of exchange? Second, the model is based on a market environment that is entirely different from Walrasian markets, which raises the question of what is the role of money in close to perfect markets. Finally, current media of exchange are essentially fiat objects, which maintain its value only because they are used as media of exchange, but not because of any intrinsic use in consumption and production. Since the search model focuses on steady-states, it remains silent about how such objects would come into circulation, i.e. how one would evolve from a commodity money world to one with fiat money.

These questions mostly arise because the search model takes an evolutionary view on money, while taking the market environment as given. In contrast, the first half of the thesis studies the evolution of money and market jointly, and is thus able to provide answers. The evolutionary view of markets comes from the observation that markets are organized, and that so-called 'intermediaries'[2] perform an important role of organizing transactions and matching buyers and sellers. These intermediaries are often the ones who experiment with and innovate the way transactions are carried out, and historically were a driving force in the development and circulation of money. Formally, I capture this role of intermediaries by focusing on evolutionarily stable equilibria , i.e. steady-states, where no small set of agents by explicitly coordinating their actions is able to induce the rest of the population to change their strategies. Evolutionary stability has entirely different effects in an environment where intermediaries can coordinate, as compared to the pure search model without intermediation. As a central result, I find that under weak restrictions, the unique evolutionarily stable equilibrium of the decentralized trading economy with intermediaries has transaction patterns that effectively mimick the Walrasian Cash-in-Advance model. I also study, under what conditions this Cash-in-Advance equilibrium can be implemented by the coordination of intermediaries who start to circulate'bank notes' and accept them from each other, i.e. a free banking arrangement. The model derives some conditions under which such a free banking equilibrium is stable, and in the thesis, I briefly discuss those and compare them to the historical evidence on free banking regimes.[3]

I should note the evolutionary view of markets and money has a long-standing tradition in economics, most notably in the Austrian school. This view, for instance, underlies Menger's classical essay on the origin of money. Menger, however, did not recognize the multiplicity issue, and hence the need for explicit coordination to arise endogenously in some form. He also does not discuss the role of intermediaries in this context. In the classical literature on money,

[2]I use this term for any agent who performs a role as a middleman in a transaction, whether this is a store that buys goods from producers and sells them to consumers, a financial intermediary, or quite possibly takes on a different form. One of my personal favorites is the language school as an intermediary for language teaching. They essentially buy teaching hours from teachers, and sell them to customers, providing a some quality guarantee, but little else, other than a matching service.

[3]The discussion on free banking regimes has since led to subsequent work with Guido Lorenzoni, where the discussion is framed in the context of contractual frictions, liquidity provision and borrowing constraints, see [5].

the 'evolutionary theorists' are opposed to institutional economists who emphasize the role of governments in introducing a medium of exchange (see Goodhard [4] for an exposition of this alternative view); with the discussion of this thesis in mind, it should be noted that it remains undisputed that somehow the coordination problem of introducing a medium of exchange to the market needs to be internalized. This thesis takes a stand that evolutionary forces might be able to accomplish this if intermediaries are able to explicitly coordinate; however, this should not be interpreted as denying that the same role could be played (and has at times been played) by a government, but it does raise the question of what is, and what should be, the role of a government in the organization of markets and the provision of transaction services, in the presence of competitive intermediaries, or other private market institutions that provide such services.

Having studied forces that lead to efficient coordination in the long run, the second half of the thesis studies the role of information in coordination failures in the short run. Incomplete information has long been recognized as a reason why individuals may not be able to agree on an equilibrium, even though it Pareto-dominates other outcomes. The idea of coordination failures due to informational reasons has long been conceptualized by the notion of 'sunspots' i.e. informationally irrelevant events that lead market participants to coordinate on one of multiple equilibrium outcomes (depending on the environment, these equilibria may be Pareto-ranked, but they don't need to be).

The more recent literature on so-called global games has taken an alternative approach towards modelling incomplete information. The idea of the global games literature is to argue that individuals observe the payoff-relevant information or "fundamentals" with noise, more over they have different information about the state of the world. Breaking away from the assumption that the state is common knowledge, Carlsson and van Damme [3] introduce the idea of a "global" coordination game, where instead of observing the state of the world perfectly, players receive noisy signals about the state, moreover, conditional on the fundamental, signals are uncorrelated across players. In 2×2 coordination games, Carlsson and van Damme then show that there exists a unique equilibrium. The reason for uniqueness can be roughly sketched as follows: a single player may use his signal to draw some inference about the state of the world, but it does not enable him to learn anything about what the other players have observed *relative to him.* If a player is indifferent between two strategies in a coordination setting, his action will depend much on what he thinks the other players are likely to do; the fact that the information structure prevents him from gaining any such information implies that in the equilibrium of a 2×2 game, at the margin he must assign a probability of $\frac{1}{2}$ to the event that the other player plays either action, conditional on which the assumptions of strategic complementarities imply that there is a unique signal, at which the player is indifferent between playing either action, i.e. the equilibrium must be unique. Moreover, in the limit as the private signals become infinitely precise, the equilibrium is insensitive to the exact distribution of signals. At first hand, this unique limit equilibrium might therefore be used as a very convenient 'selection' for comparative static or dynamic purposes, when multiplicity in coordination games at first prevents any such analysis. Subsequent applications of global games, most prominently Morris and Shin [6], have done just that: use the global games approach as a selection mechanism in coordination games to study the effects of policy proposals, or perform other comparative statics exercises. There are plenty of potential applications which are essentially variations on similar static coordination games: Bank runs, currency crises, coordination among debt holders, to name a few, are just the most prominent ones; if the global games methods were to be properly extended to dynamic coordination games, the establishment of a medium of exchange that was discussed during the first half of the thesis might also prominently figure among the potential applications of this theory.

The second half of this thesis argues against this view. Instead of exploiting uniqueness, it takes the global games approach as what it is and takes an agnostic view on the issue of uniqueness vs. multiplicity. I rather argue that this modelling approach opens up the black box of expectations in macroeconomic models. By explicitly modelling different information structures, I study the ability of market participants to coordinate on one or the other outcome. In the second chapter of this thesis, I study the linkages between the information structure, equilibrium selection (uniqueness vs. multiplicity) and the structure of higher-order beliefs (beliefs that players have about each others' beliefs). In particular, I study the role of public information in this context, and show that if agents have access to sufficiently precise public information, then the multiplicity of equilibria will be maintained. In the limit, as both public and private signals become infinitely precise, I obtain a unique equilibrium if and only if the rate of convergence of the private signal variance is more than twice as fast as the rate of convergence of the public signal variance; otherwise, the set of equilibria converges to the one obtained under common knowledge.[4] Hence, even if the public signal contains almost no informational value concerning the fundamental, it has an enormous impact on the set of equilibria. The intuition for this result is that even if the public signal provides little information about the fundamental, it does provide information about the beliefs of other players, and that is what enables them to coordinate, i.e. leads to multiplicity. This convergence condition is then translated into its effect on higher-order beliefs, and I show that the same condition determines whether in the limit the belief structure will be "close" (in a well-defined sense) to common knowledge or not.

My results are first meant to strike a cautionary note with respect to the use of the private information limit as an equilibrium selection in coordination games; however, the translation of signal structures into belief structures has a second purpose: it illustrates the usefulness of the global games approach for thinking about how bayesian expectations are formed, not only with respect to the underlying states, but also with respect to the actions of other players; in macro-economic contexts, the latter problem of belief formation is typically 'solved' by the assumption of rational expectations and, possibly sunspots, assumptions which brush away the underlying coordination problems that may arise. While questioning the use of global games for equilibrium selection purposes, the paper stresses its importance for thinking about expectation formation, and the possibility for economic and information policies to influence expectations.

This approach towards modelling expectation formation is translated into a different context and pursued further in the last chapter of the thesis. There, I study the effects of higher-order uncertainty on the adjustment dynamics in a macro-economic game with complementarities; the discussion is embedded into a version of Lucas' incomplete information model of incomplete nominal adjustment, but the main points apply directly to other environments with higher-order uncertainty and strategic complementarities:[5] in a model of strategic complementarities

[4] The mathematical intuition for this condition is given in the text.

[5] The model that I study is a mix between Woodford [9] and Morris and Shin [8]. Woodford makes the point

in price-setting, I show that the speed of nominal adjustment following a monetary shock depends on the degree of higher-order uncertainty. Public information has two effects in this environment: by reducing higher-order uncertainty, better public information leads to faster price adjustment; on the other hand, public information is itself inherently noisy, and hence better information may expose an economy to additional sources of noise in the short run. To conclude, I embed the model into the Barro-Gordon framework to study the Welfare effects of information provision by the central bank. I show that a higher degree of transparency not only improves central bank monitoring, but has a second positive welfare effect by reducing the gains from "surprise monetary expansion". I further discuss the informational benefits of a formal framework for monetary policy, and argue that information provision by the central bank may have positive welfare effects through coordinating expectations.

The use of the global games framework to study expectation formation also appears in other papers, and in concluding this overview of my thesis, I shall briefly discuss those studies. In two closely related papers on static coordination problems, Morris and Shin [7,8] study the role of the information structure and the coordinating role of public information. These two papers are mostly motivated by observations of the coordinating and potentially destabilizing role of public information, in particular, the role of news media. Their analysis takes a setting with a unique equilibrium as given, and mostly focuses on identifying the relevant comparative statics effects. In contrast, in joint work with G.-M. Angeletos and A. Pavan [1,2], I further push this interpretation of global games as opening up the black box of expectations: In the first paper, we study the informational role of policies in environments where the policymaker has private information, and come to the conclusion that, since the policy publicly reveals private information, it may enable market participants to coordinate their actions, and opens up the door to multiple equilibria and 'policy traps', where the policy maker is forced to follow a certain equilibrium policy that is ex ante determined by the market's expectations. In the second paper, we study the dynamics of exogenous and endogenous information revelation in a currency crisis

that higher-order uncertainty leads to delays in price-adjustment, however his model does not allow to study the effects of different types of information, nro does it readily extend to other settings. Morris and Shin study the effects of public information in a static model similar to the one studied by Woodford. This paper translates Morris' and Shin's analysis into a dynamic environment, and discusses how the insights regarding higher-order uncertainty and persistence, as well as the effects public information, translate generally into macro-economic environments with strategic complementarities. I also substantially simplify the exposition from Woodford.

game that is played out over time. This analysis shows how the exogenous and endogenous flow of information naturally leads to a staggering of attacks: each unsuccessful speculative attacks reveals a large amount of information about the central bank, notably improving its reputation. It then takes time for new information to gradually offset this positive effect, until market participants are sufficiently uncertain about the central bank's defense policy to start a new attack. In both of these applications of global games, multiplicity is an integral feature of the model's outcomes; nevertheless, we are able to gain some insights into the forces that enable market participants to coordinate their actions in each of the environments.

Thinking about how expectations are formed in macro-economic environments or games is central for understanding the effects of information policies. In this sense, the second half of this thesis, as well as the other papers that I briefly discussed, propose a framework in which to think about these issues. Whether this framework stands the test of time depends on its usefulness in other applications - and to what extent they are able to guide the informed policymaker. Of course, this entire research program is still very much in the beginning, and the contributions of this thesis represent little more than a first step. It will be interesting to see how this discussion unfolds, and how its conclusions are altered by subsequent work - just as this thesis has tried to put Carlsson and van Damme's contribution into a different light. Similarly, it will be interesting to see to what extent the intermediation model of the first half of the thesis holds up against the evolutionary pressures that shape the market place for ideas in ways not too different from their effect on other markets.

0.1.1 Bibliography

[1] G.-M. Angeletos, C. Hellwig, and A. Pavan, "Coordination and Policy Traps in Currency Crises", MIT and Northwestern, 2002.

[2] G.-M. Angeletos, C. Hellwig and A. Pavan, "On the Dynamics of Speculative Attacks", MIT, 2002.

[3] H. Carlsson and E. van Damme, "Global Games and Equilibrium Selection", *Econometrica* **61** (1993), 989-1018.

[4] C. Goodhart, "The Two Concepts of Money: Implications for the Analysis of Optimal Currency Areas", *Europ. J. Polit. Economy* **14** (1998), 407-432

[5] C. Hellwig and G. Lorenzoni, "Bubbles and Private Liquidity", Princeton University, 2002.

[6] S. Morris and H. S. Shin, "Unique Equilibrium in a Model of Self-fulfilling Currency Attacks", *Amer. Econ. Rev.* **88** (1998), 587-597.

[7] S. Morris and H. S. Shin, "Coordination Risk and the Price of Debt", *FMG discussion paper* **373**, London School of Economics, 2001.

[8] S. Morris and H. S. Shin, "Social Value of Public Information", *Amer. Econ. Rev.*, forthcoming.

[9] M. Woodford, "Imperfect Common Knowledge and the Effects of Monetary Policy", *NBER working paper* **8673** (2001).

Part I

Money and Intermediaries

Chapter 1

Money, Intermediaries and Cash-in-Advance Constraints

Summary of Chapter 1 *I study a search economy in which intermediaries are the driving force coordinating the economy on the use of a unique, common medium of exchange for transactions. If search frictions delay trade, intermediaries offering immediate exchange opportunities can make arbitrage gains from a price spread, but they have to solve the search market's allocation problem. Intermediaries solve this problem best by imposing a common medium of exchange to other agents, and a Cash-in-Advance constraint arises in equilibrium: Agents trade twice in order to consume, once to exchange their production against the medium of exchange, and once to purchase their consumption. By studying the evolutionary stability of equilibria, I discuss which equilibria are likely to arise as long run outcomes. I extend my analysis to the study of fiat currencies and free banking systems.*

JEL classification numbers: D51, E40

Keywords: Monetary Exchange Economy, Intermediation, Search, Cash-in-Advance Constraints

[0]I am grateful for many discussions with my supervisor, Nobu Kiyotaki, as well as helpful comments from George-Marios Angeletos, Martin Hellwig, Godfrey Keller, Thomas Mariotti, Michele Piccione and Rob Townsend. Financial support through LSE and the Economic and Social Research Council (UK) is gratefully acknowledged. All remaining errors are mine.

1.1 Introduction

What accounts for the use of money in economic transactions in competitive markets? This simple, seemingly trivial question has been the cause of much debate and a rich tradition of research in economics. The answer that is typically given today starts from Jevons' [17] suggestion that the use of a medium of exchange eliminates the need for a "double coincidence of wants", if market participants trade bilaterally, and have to spend time and resources to find suitable trade partners. If all market participants instead agree on the use of a common medium of exchange, they will sell their production for the medium of exchange, and will use the medium of exchange to buy what they consume. In a large economy, this reduces the waiting time until an agent can receive his consumption good. In the purest statement of this idea, Kiyotaki and Wright [18] show that the use of a common medium of exchange may be the equilibrium outcome of individually rational transaction decisions in an economy where encounters between market participants are random.[1]

This paper attempts to provide an alternative explanation for the transactions role of money in an economy that appears close to the frictionless, competitive benchmark. The observation of transactions in markets will suggest that we rarely have to search randomly to find the goods that we want to consume, or to find a buyer for the products that we want to offer; indeed, in most instances, we don't even face a delay in the transactions we carry out. For most products, we know where we can buy or sell them, and we just go and buy them whenever we want to, and we expect to find them at that time and place. In other words, the same frictions that account for the use of money do not enter into the theoretical model for which we are trying to provide a foundation, nor do they seem to be relevant for the use of money in most transactions. The challenge of providing microfoundations thus lies not so much in providing an explanation for the transactions role of money, but in providing this explanation *in the context of an exchange economy that is perceived as frictionless and competitive.* This implies going outside the Walrasian framework and leads to a broader underlying question: How do individuals interact in a decentralized market, so that the market outcome appears to

[1] See Aiyagari and Wallace [1] for a general discussion of this statement in the context of the model of Kiyotaki and Wright [18].

be virtually without frictions; in other words, *how do markets evolve?*

A natural way of dealing with search frictions is the centralization of transactions through a system of intermediaries of known location and specialization. Of course, money and intermediaries are both essential features of transactions in markets. However, insofar as they are studied separately, they are implicitly regarded as substitutes in dealing with the frictions in the market. In other words, if intermediaries are capable of alleviating market frictions, what would be the role of money (or vice versa)? In order to fully account for the transactions role of money, one therefore has to ask how money interacts with intermediation in alleviating market frictions.[2]

Here, I discuss this interaction of money and intermediation in the context of a decentralized exchange economy where trade is bilateral and potentially subject to delay. Individuals can modify the trading environment by acting as intermediaries, thereby reducing the frictions to which other agents are subject. Intermediaries are immediately accessible, and delays in trade with them only depend on their ability to accommodate the transactions demanded. They centralize transactions more easily, if a common medium of exchange is used by the agents with whom they trade. On the other hand, they have the possibility to introduce it to all other agents, who, in turn, are willing to use it, if it allows them to buy from the intermediary whatever good they want to consume. The analysis thus points to *a complementarity* between the use of a medium of exchange and the centralization of transactions by intermediaries that roughly matches historical facts: Throughout history, intermediaries were often the ones who developed more efficient ways of exchanging goods, and they were particularly important in introducing and using money. On the other hand, they were also the primary beneficiaries of the introduction of a common medium of exchange.[3]

[2] This question could of course be raised in a variety of contexts, in which money helps to alleviate or even eliminate frictions. Consider the issue of record-keeping and contract enforcement in a dynamic environment: Kocherlakota [19] argues that "Money is Memory", i.e. in an environment where imperfect record-keeping limits the possibility of writing and enforcing contracts in the future, money may serve as a substitute of memory, and implements allocations that would only have been feasible under perfect memory. On the other hand, Dixit [8] also studies an enviroment with imperfect record-keeping, citing the mafia as an information intermediary who makes a profitable business out of centralizing record-keeping i.e. memory, and contract enforcement. To fully understand the role of money as a substitute for memory, one would have to ask how it interacts with the intermediary, who also may act as a possible substitute for memory. In the conclusion, I discuss why the mechanism at play in the search environment of this paper may also apply to other contexts.

[3] A particularly neat example of these effects, that also highlights the mechanisms in this paper, is Radford's

Intermediaries have been introduced into models with trade frictions in the past.[4] These models usually focus on the exchange of a single good with a given number of buyers and sellers who trade off the delay in the transaction against the price at which they trade. Intermediaries act as arbitrageurs who offer immediate transactions, but charge a mark-up for their services. With many commodities, the success of intermediaries depends on their ability to match buyers and sellers *for each good.* This becomes a problem, if the number of goods that an intermediary can trade is restricted, and consumers may not always be willing to consume what a given intermediary would be willing to offer them for their production. This transfers the double coincidence problem from the search market to the intermediaries. They solve it by introducing and promoting a common medium of exchange that enables consumers to transfer purchasing power from transactions with one intermediary to transactions with another. That an intermediary cannot trade with all commodities at once is an important part of the argument: Otherwise, one intermediary would be able to perfectly eliminate the frictions, and there would be no need for a common medium of exchange. Consumers could simply trade their excess demand in all goods at once with an intermediary, at the prices set by the latter. If the intermediary fixes market-clearing prices, then no medium of exchange is needed to buy some goods from a different intermediary. Similarly, if the medium of exchange solved the allocations problem perfectly, there would be no role for the intermediary. It is precisely the fact that each of them on its own is unable to perfectly alleviate frictions that makes them complementary.

The emergence of intermediaries alters the way in which trade decisions are made by other agents. Trade with intermediaries enables consumers and producers to direct their search towards a particular good, as opposed to the random search in economies without intermediation. By limiting their clients' choices to the use of a unique, common medium of exchange, intermediaries introduce its use to the entire economy. In an equilibrium of the economy considered

[24] description of exchanges in a Prisoner of War camp. He describes how economic institutions and markets developed within the completely unorganized environment of a PoW-camp, driven mainly by the scope for trade arising from differences in endowments (Red-Cross packages) and tastes. In the early days of the camp, some individuals who exploited the price margins between different parts of the camp ("intermediaries") promoted and established the cigarette as common money. This was fundamental for the later development of more sophisticated market institutions, such as a store, and even the introduction of a paper money, backed by the store's inventories of goods.

[4] Rubinstein and Wolinsky [26] explore intermediation in a search-theoretic model in which one good is traded. The present analysis is closer in spirit to Gehrig [9].

here, all agents trade twice to acquire what they want to consume: once to obtain the medium of exchange (sell their production), and once to buy their consumption good. Effectively, a *Cash-in-Advance constraint* for transactions with intermediaries is introduced, whereby market participants have to use the common medium of exchange. Since the medium of exchange enables intermediaries to match buyers and sellers, the latter face no waiting time to perform the desired transaction. As a result, the search market empties, since producers and consumers take advantage of the intermediaries' services. Equilibrium allocations bear the characteristics of Walrasian allocations, and the resulting transaction patterns resemble trade in frictionless Walrasian markets: At any time, almost all agents are able to trade, and consume in every other (unless they produce or consume the common medium of exchange, and only need one transaction).

Intermediation also provides a mechanism by which the economy can coordinate on the common use of an *efficient* medium of exchange. If a small set of agents coordinates their activities and offers some new organization of transactions, they may induce other agents, and eventually the entire economy, to adopt their innovation. This can be assimilated to the historical role of intermediaries in developing more efficient means of exchange. Formally, I study which of the resulting equilibria are evolutionarily stable. In contrast to the standard search model, evolutionary stability implies Pareto efficiency in an environment with intermediaries.[5] I also allow for the circulation of fiat money. Under a general set of conditions, *the unique evolutionarily stable steady-state is then a Cash-in-Advance equilibrium in which fiat money circulates as the common medium of exchange.*

Finally, I study how fiat money may come into circulation, and again illustrate the coordinating role of the intermediaries: assuming that these intermediaries can write out demandable debt certificates ("notes"), I discuss under what conditions they become perfect substitutes in a "free banking equilibrium". I illustrate how the clearing mechanism serves to monitor the note issue of banks. In practice, a free banking regime has to rely on (i) the clearing mechanism to monitor the competitive issue of notes, and (ii) reliable punishment mechanisms in

[5] In large population matching games, such as the search model of money, evolutionary stability considerations have little effect on equilibrium selection, since the "mutants" have no possibility to interact with each other to explicitly coordinate their actions. Intermediation provides such a channel.

case of default. It is important to note that with free entry into note issue, i.e. in a truly competitive environment, the loss of the banking licence is insufficient to prevent overissue and strategic defaults, since no rents are directly associated with being an intermediary. Historically, it seems that the most successful free banking regimes were the ones that effectively used the note clearing, and used harsh punishments in case of default. But free banking systems also faced difficulties, even when those conditions were met: the model illustrates a coordination problem arising in the clearing market, i.e. if notes are entirely safe, and costs are associated with clearing, banks may prefer to hold notes in reserve, or bring them back into circulation rather than return them to the issuer.

1.2 Related Literature

The results in this paper have various implications for existing equilibrium models of monetary exchange. Money is, of course, one of the essential elements of our understanding of macroeconomic fluctuations, and of the role of policy in response to them. Since the competitive Arrow-Debreu framework does not endogenously account for such a transactions demand for money, its existence is usually assumed into the model by way of a restriction on the transactions in which individuals engage: money must be used to buy consumption goods.[6] This leads to the use of money as a short-term store of value, and the demand for real balances will depend on the availability of other assets for short- or longer term savings, and on their liquidity when they are to be sold to satisfy consumption needs. Of course, the observation of how transactions take place in markets makes such a restriction appealing, and it has proven successful for the purposes of macroeconomic analysis. However, the same observations suggest that the main purpose of money is *not* its use as a store of value, but as a convenient medium

[6] Examples where such a Cash-in-Advance constraint is made explicit are Svensson [32] and Lucas and Stokey [20]. The constraint that money is used to buy goods also appears in Romer's [25] general equilibrium treatment of Baumol's [4] and Tobin's [33] inventory demand for Cash. An alternative approach assumes that the transactions services of money enter directly into the market participants' utility functions, following Sidrauski [28]. The overlapping generations model, introduced by Samuelson [27], provides one example where money is essential in improving allocations, (without such a constraint that it must be used in transactions, or a direct effect on the utility function) - however, it is used to transfer wealth between generations, and it looses its role once its rate of return is dominated by other assets.

Hellwig [13] provides a detailed, critical discussion of the recent and not-so-recent literature on monetary equilibrium theory, on which some of the ideas in this paper are based.

for transactions, and money is used as a short-term store of value *only because it has a primary purpose as a medium of exchange.*

Precisely such a Cash-in-Advance constraint is the result of equilibrium trading strategies in the present model, where market interaction is viewed as an ongoing evolutionary process.[7] This paper may therefore be viewed as providing a microfoundation for the structure of transactions in the macroeconomic applications that exogenously impose such a constraint. However, the efficiency of the constraint here is in stark contrast with its macroeconomic counterpart - in fact, viewing such a constraint as efficiency enhancing seems contradictory. Efficiency follows from the strategic interaction of intermediaries, as the consequence of an evolutionarily stable steady-state. It turns out that the Cash-in-Advance constraint enables intermediaries to achieve Pareto-efficiency in the trade process, if they can enforce it on all agents who want to trade with them. The constraint is observed in all intermediated exchange, but is not binding for exchange outside intermediation. Formally, I do not assume away the possibility that two agents exchange "goods" for "goods" outside intermediated transactions, but in equilibrium, they never incur a situation in which they agree to exchange "goods" for "goods".

The paper also responds to some of the weaknesses of existing search models of money that follow Kiyotaki and Wright [18]. While they succeed in explaining why it may be individually rational and socially efficient that all agents use a common medium of exchange, they cannot account for the fact that the vast majority of transactions involves the exchange of goods for money - indeed, one of the conclusions from the literature following Kiyotaki and Wright is that such a Cash-in-Advance constraint where "goods" are only traded for "money" fails to materialize (Aiyagari and Wallace, [2]), since the delays in trade provide a sufficient incentive to accept "goods" for further exchange, instead of immediate consumption. The same search frictions which motivate the use of a medium of exchange render the existence of a Cash-in-Advance constraint impossible.[8]

[7] Such an evolutionary view of markets has a long tradition in the Austrian school. For example, in his classical article on the origin of money, Menger [22] views money as the determinate outcome of an evolutionary process; however his analysis does not recognize the potential for multiplicity inherent in coordination problems, nor does he explicitly refer to intermediaries as a coordinating force in the market.

[8] While this result obviously clashes with the observation of Cash-in-Advance constraints in quasi-perfect markets, it has some intuitive appeal with respect to the importance of barter trade in environments, in which markets are far from frictionless.

A second weakness of search models is the multiplicity of equilibria. The strategic complementarity that exists between players for using a single good as a common medium of exchange also implies that players may coordinate on *any* good as the common medium of exchange in equilibrium; in other words, the model remains silent about the choice of a medium of exchange. Similarly, while the search model can be used to show that a fiat object, which no one consumes and no one produces, may be valued and used as money, the very same set-up always implies that this need not be the case in equilibrium. Hence, the search model is unable to say anything about how a fiat money comes into circulation in a decentralized exchange economy. In contrast, intermediation arguably provides a natural framework for studying these selection issues, as well as the emergence of fiat money.

It should be noted that the general equilibrium as well as the search models of money have multiple equilibria. As discussed in Hahn [10], this multiplicity of equilibria is the manifestation of an intertemporal coordination problem: the acceptance of money today is based on the acceptance of money tomorrow. The evolutionary approach taken here resolves the multiplicity. It should be noted, however, that the solution relies on the assumption that individuals are able to coordinate their strategies explicitly not only within a single period, but also *across time,* at least on a small scale.

Formally, this paper is most closely related to, and shares much of its motivation with, a series of contemporaneous papers that discuss monetary trade in a "trading-post" environment. In such an environment, markets are in separate locations, and typically each location represents a different pair of goods that can be traded at that location. Iwai [16] studies such an environment with search frictions. In Starr [31], as well as Howitt [14] and Howitt and Clower [15], these trading posts are run by intermediaries similar to the ones encountered here.[9] In all these papers, a combination of increasing returns to scale in the intermediary's transaction technology and the double coincidence problem lead to a concentration on the smallest possible number of trading posts and the thereby the use of a common medium of exchange. In another paper that uses a trading-post environment, Matsui and Shimizu [21] discuss the emergence of money in a market place environment, where the location, rather than an intermediary defines

[9] See also Corbae, Temzelides and Wright [7] for an endogenous matching environment with similar outcomes.

the trading post. As in this paper, they study the evolutionary stability of equilibria, and show that a unique "single-price equilibrium" survives, in which the supply of fiat money is equal in value to its demand for market transactions. All these papers take the trading post structure as given, and are more concerned with the properties of the resulting equilibria, i.e. when a monetary equilibrium exists and what its properties are, as well as the relation of money and prices.

In contrast, this paper abstracts from the problems of price formation, and instead concentrates on the evolutionary aspects of the emergence of money and markets.[10] In order to embed intermediation into the search framework of Kiyotaki and Wright [18], I restrict transactions to one-for-one swaps, emphasizing the role of the double coincidence problem in the exchange process. The trading posts generate exchange opportunities only insofar as intermediaries become active, the choice of becoming an intermediary is itself endogenous in this model. This environment provides a few simplifying insights: In equilibrium, the specialization between production and intermediation takes place in such a way that intermediaries optimize the number of transactions they carry out, and no transactions take place as a result of pure search. From a much less structured trading environment, we therefore obtain the same transaction patterns, but using the search-theoretical framework as a background, we give a strategic account as to how intermediation develops and induces improvements in the transaction process until at some point, transaction patterns and allocations closely resemble Walrasian equilibrium allocations.

An alternative account of the role of money in Walrasian markets is Banerjee and Maskin [3]. In contrast to our paper, they focus on asymmetric information (a lemons problem between buyers and sellers) as the source of frictions. While their paper shares many of the motivations and results with this one, including a foundation for Cash-in-Advance constraints, the assumption of Walrasian markets leaves open the question of how money interacts with other institutions that may serve to alleviate the lemons problem, namely intermediaries.

The rest of this paper is organized as follows: Section 3 describes the basic economic environment and introduces the notion of steady-state equilibrium. I then derive some preliminary

[10]In this sense, I view the afore-mentioned papers very much as complementary to this one.

results, to provide conditions that a steady-state has to satisfy. Section 4 considers one type of equilibrium, in which a particular good is used as a common medium of exchange. I contrast the findings of the economy with intermediation with the monetary equilibria resulting from pure search. Section 5 introduces evolutionary stability, and shows that any evolutionarily stable equilibrium must be Pareto efficient. Section 6 extends the initial set-up to allow for the circulation of fiat money. Under general conditions, it is then shown that the unique evolutionarily stable equilibrium has a Cash-in-Advance constraint for fiat money. I also discuss the implementation of this equilibrium in a free banking environment. I conclude the paper with some remarks on how the mechanism described here may be extended, or apply to other contexts.

1.3 The Model

1.3.1 The physical environment

I consider a continuum of measure 1 of infinitely-lived agents. There are $N \geq 4$ different goods and N types of agents in the economy. Type i agents always consume good i. There is a measure of $\frac{1}{N}$ of each type.

In this economy, time is discrete and infinite, and all goods are perfectly durable. In order to consume, an agent chooses to act either as a *producer* or as an *intermediary*. A producer always holds one unit of a good, and tries to obtain, after a sequence of one-for-one exchanges, a unit of his own consumption good, labelled i. He then consumes and immediately thereafter produces a unit of good $i+1$ (good N consumers produce good 1). An intermediary does not produce, but instead holds one unit of her own consumption good, and has a shop, where she trades.[11][12] In each period, the intermediary uses her inventory to offer a one-for-one exchange between her own consumption good i and some other good j and vice versa. She can immediately be located by all other producers and intermediaries. Whenever she gives out her consumption

[11] I use male pronouns to refer to producers, and female pronouns to refer to intermediaries.

[12] An earlier version of this paper allowed intermediaries to accumulate inventories for transactions. The main results of this paper go through almost identically, but the restriction to one unit substantially simplifies the analysis.

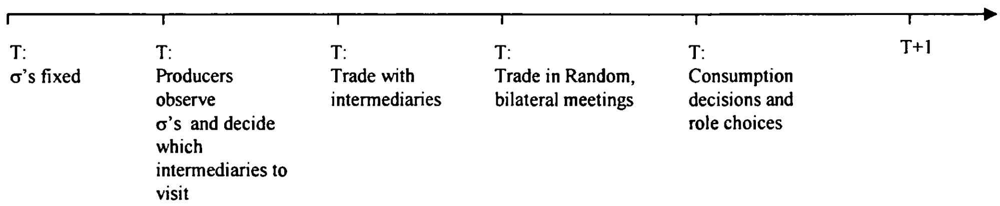

Figure 1-1: Timing of actions within each period

good and acquires the other good, the unit she gives out is reduced by a fraction σ_{ij}. Whenever she acquires her consumption good and gives out the other good, she sells a full unit for a full unit. The intermediary lives off this mark-up.

Within a period, the sequence of actions is described in figure 1. At the end of each period, every intermediary decides on the size of the unit of her consumption good that she offers during the next period, and consumes the residual. She thus starts with a unit of her own consumption good that is reduced by a fraction σ_{ij}. She then lets all producers know that during the period, she is willing to exchange one unit of size $1-\sigma_{ij}$ of good i against an integer unit of good j, and an integer unit of good j against an integer unit of good i. Producers observe these, and then decide if they want to trade with an intermediary, and which one they want to trade with. Then, transactions with intermediaries take place. To visualize this transactions process, it is useful to order all intermediaries for a given ij-transaction according to the size of their mark-up. All producers who intend to carry out a transaction of i for j then form a queue (the position of each individual within the queue being random), and each one chooses the intermediary with the lowest mark-up available to him to carry out the transaction. Furthermore, producers cannot use the intermediary to coordinate on a location for carrying out their transactions without going through the intermediary (i.e. producers cannot coordinate amongst each other to "meet in front of the shop"). Once producers are allocated to intermediaries, each intermediary first sells her unit of good i for an integer unit of good j, and then sells the unit of good j to acquire an integer unit of good i. I assume that an intermediary is willing to trade only if she is able to carry out a two-way transaction, that is, if she is allocated one producer for each side of the transaction. If there is an insufficient number of intermediaries to carry out all desired two-way transactions, or if there is an excess of producers on one or the other side of the market, then some producers are not allocated to an intermediary and are unable to carry out the desired transaction. On the other hand, if there are too many intermediaries compared to the total number of two-way transactions demanded by producers, then some intermediaries are unable to trade. Within each period, the total measure of two-way transactions between any pair of goods i and j is thus bounded by the total measure of ij-intermediaries, as well as by the number of agents who want to complete the transaction in each direction.

If the intermediaries' total inventory is insufficient to accommodate all transactions, or if there is a difference between the total demand for exchanging good i for good j and the demand for the opposite exchange, some producers are unable to acquire their desired good from an intermediary. In a second stage of the period, after all transactions with intermediaries are completed, every producer, who was either unable or unwilling to trade with an intermediary, randomly meets someone else who also has not traded with an intermediary, and they have a second opportunity for a transaction. In such a random match, each producer observes which good his trading partner holds, and decides whether or not to swap his inventory good for the good held by the other agent. An exchange takes place if and only if both agree to it. Since the matches are random, the probability of encountering a type i producer who holds good j (henceforth called ij-agent) in such a meeting is given by the measure of ij-agents who enter the bilateral matching stage.

After trading in the random meetings has taken place, all agents decide on their consumption, and on their role during the following period. A producer can consume only if he has acquired a full or reduced unit of his own consumption good. In this case, he can also choose to become an intermediary, simply by using his unit as inventory for intermediation. An intermediary decides what proportion of her own consumption good unit to consume, and thereby, with what size of a unit she wants to enter the following period. An intermediary has the option to become a producer in any period, simply by consuming her entire inventory and producing one unit of the production good. There is no cost of setting up or abandoning intermediation.

To complete the description of trading, I assume that no agent (intermediary or producer) ever accepts a reduced unit of a good in a transaction. Reduced units are therefore acceptable only for immediate consumption, and are never held in inventory.

Preferences are assumed to be symmetric across types. Consumption utility is linear: An intermediary obtains an instantaneous utility cU from consuming a fraction c of her inventory unit. A producer obtains utility $U(1-\sigma)$ from consuming a unit of his consumption good of size $1-\sigma$. Consuming any other good yields 0 utility. All agents discount time by a constant rate δ smaller than but close to 1. Whenever an agent trades, he incurs a direct transaction cost. Producers incur a cost of β_i, whenever they accept good i in a one-for-one exchange, or in a transaction with an intermediary. Goods are strictly ranked by transaction costs, and for further reference, good 1 is defined as the good which has the lowest cost of acceptance. Intermediaries incur a cost of $\beta_i+\beta_j$ from carrying out a two-way transaction between good i for good j. In addition, there is no fixed cost involved in setting up or maintaining intermediation.

The main innovation with respect to the original search-theoretic model of money by Kiyotaki and Wright [18] is the formal introduction of intermediation. As in their framework, this paper's aim is to analyze transaction patterns and the emergence of a common medium of exchange within a decentralized economy. This requires an environment in which all goods are durable, and no commodity is predestined by its storability qualities to become a medium of exchange. Many of the seemingly ad hoc modelling choices are motivated by the aim to allow for trade in random matches as well as through intermediaries within the same environment, at the expense of a formal modelling of price setting by intermediaries. It should be noted that the steady-state results in this paper are not affected by the way in which producers are assigned to intermediaries. Related papers that abstract from search trade (Howitt [14], Starr [31], Matsui and Shimizu [21]) suggest that the conclusions presented here for intermediated transactions apply to settings, where intermediaries set prices and act as "market-makers", and producers determine the quantities that they want to trade.

1.3.2 Strategies and equilibrium

I now introduce the notation for strategic variables to describe individual behavior, as well as the distribution of individual inventories to describe the evolution of the entire economy, given individual behavior. Throughout the paper, I restrict attention to symmetric strategy profiles, in which (i) all producers of the same type choose the same mixed-strategy profile for transaction strategies, and (ii) in each period, all intermediaries who offer the same transaction set the same markup. This is a weak restriction, but simplifies our notation considerably: In a stationary environment, i.e. one where the distribution of individual inventories across the population remains identical over time, identical and stationary markups will naturally come as a consequence of Bertrand competition among intermediaries.

Individual behavior and the aggregate state of this economy are described as follows: Trading strategies are (probabilistic) decision rules that indicate to a producer, which intermediary to visit, given his inventory good and the mark-ups. Under the given assumptions concerning the matching between intermediaries and producers, these decision rules only need to indicate which transaction a producer intends to carry out, but not the intermediary's identity. I denote these decision rules by functions ϕ_{ij}, where $\phi_{ij}(k,l)$ denotes the probability that an ij-agent (type i producer who holds good j) visits a kl-intermediary, and τ_{ij} for trade in random meetings, where $\tau_{ij}(k)$ indicates the probability that an ij-agent accepts good k for good j in a bilateral meeting. Intermediaries determine the mark-up σ_{ij}. Finally, the aggregate state of the economy is given by the distribution of inventories and role choices, i.e. (i) the measures of type i intermediaries who trade good i for good j (henceforth: ij-intermediaries), denoted by ν_{ij}, and (ii) the distribution of inventories across producers, where we let μ_{ij} denote the measure of ij-agents. This notation leaves aside a formalization of the decision problem for role choices, which is given by an indifference condition between the two activities: any type who holds one unit of his own consumption good has to be indifferent between either becoming an intermediary or consuming and becoming a producer.

Introducing the strategies in this way implicitly assumes symmetric and stationary behavior, but this will also be the outcome of optimizing behavior in a symmetric, stationary environment. Since each agent has no direct effect on the aggregate state, we can consider the optimization problem for each type of intermediaries and producers separately, taking the behavior of others and the aggregate states as given. Furthermore, our set-up implies that mark-ups are determined by Bertrand competition, and an open-entry condition is at work. In equilibrium, mark-ups will be at a level where (i) no producer has an incentive to become an intermediary, and no intermediary has an incentive to become a producer, and (ii) no intermediary has an incentive to slightly undercut all other intermediaries to offer the same transaction, in order to increase his trading volume. Finally, a stationarity condition determines the equilibrium inventory distribution of intermediaries. Summing up, this leads to the following informal definition of a symmetric stationary equilibrium:

Definition 1 *A symmetric, stationary equilibrium is a vector* $\{\mu_{ij}, \nu_{ij}, \phi_{ij}, \sigma_{ij}, \tau_{ij}\}_{i,j=1}^{N}$*, such that*

(i) *for all* ij*-intermediaries, it is optimal to set their mark-up equal to* σ_{ij} *in each period,*

(ii) ϕ_{ij} *and* τ_{ij} *are optimal trading strategies for producers,*

(iii) *no intermediary wants to become a producer, and no producer wants to become an intermediary, and*

(iv) *the distribution of inventories and role choices* $\{\nu_{ij}, \mu_{ij}\}_{i,j=1}^{N}$ *remains constant over time, given these strategy choices.*

1.3.3 Preliminary Results

In this section, I derive some preliminary results that formalize the definition of a symmetric steady-state equilibrium given above. Formally, a trading strategy for an ij-agent consists of two decision rules, one that relates the current inventory to the choice of visiting an intermediary, and one that indicates the probability with which a good is accepted in exchange for another one in a bilateral meeting. Formally, for a type i producer, a strategy for trade with intermediaries is a vector of functions $\{\phi_{ij}\}_{j\neq i} : \{1, 2, ..., N\}^2 \to [0, 1]$, where $\phi_{ij}(k, l)$ indicates the probability that a type i producer with good j (an ij-agent) visits a kl-intermediary. The residual probability $1 - \sum_{k,l=0}^{N} \phi_{ij}(k, l)$ is assigned to the event that he chooses not to visit an intermediary. Since a kl-intermediary does not trade good j, we observe immediately that $\phi_{ij}(k, l) = 0$, whenever $j \neq k$ and $j \neq l$. Also, $\phi_{ij}(k, j) = 0$, whenever $k \neq i$, since no producer is willing to acquire a reduced unit of a good other than his consumption good. This leaves as possible choices the visit of any jk-intermediary, to trade j for a full unit of k (either for further exchange, or for consumption of a full unit of i), and the visit of an ij-intermediary to acquire a reduced unit of good i for consumption.

Similarly, trading rules for bilateral meetings are described by a collection of functions $\{\tau_{ij}\}_{j\neq i} : \{1, 2, ..., N\} \rightarrow [0, 1]$, where $\tau_{ij}(k)$ indicates the probability that an ij-agent accepts good k for good j. When an ij-agent meets an lk-agent, trade occurs with probability $\tau_{ij}(k)\,\tau_{lk}(j)$.

For given values of $\{\mu_{ij}, \nu_{ij}, \phi_{ij}, \sigma_{ij}, \tau_{ij}\}$, it is straight-forward to express a producer's optimization problem by a set of Bellman equations. The aggregate state and the strategy profile determine conditional trading probabilities for trade with intermediaries: for this purpose, define $\pi_{ij}(k)$ as the probability that an ij-intermediary is able to deliver a unit of good $k \in \{i, j\}$ to a producer who wants to trade with her. Define $V_i(j)$ as the life-time discounted utility of a producer of type i who holds good j at the end of a period. Then, $V_i(j)$ satisfies: [13]

$$
\begin{aligned}
(1-\delta)\,V_i(j) \;=\; & \delta \max_{\phi_{ij},\tau_{ij}} \left\{ \sum_{l=1}^{N} \left(V_i(l) - \beta_l - V_i(j)\right)\phi_{ij}(j,l)\,\pi_{jl}(l) \right. \\
& + \left(V_i(i) - \sigma_{ij}U - \beta_i - V_i(j)\right)\phi_{ij}(i,j)\,\pi_{ij}(i) \\
& + \left(1 - \sum_{l=1}^{N} \phi_{ij}(j,l)\,\pi_{jl}(l) - \phi_{ij}(i,j)\,\pi_{ij}(i)\right) \\
& \left. \sum_{k,l} \frac{\mu'_{kl}\left(V_i(l) - \beta_l - V_i(j)\right)\tau_{ij}(l)\,\tau_{kl}(j)}{\sum_{k,l}\mu'_{kl}} \right\}
\end{aligned}
\tag{1.1}
$$

and

$$V_i(i) = U + V_i(i+1)\,;\; V_N(N) = U + V_N(1)$$

where

$$\mu'_{ij} = \mu_{ij}\left(1 - \sum_{l=1}^{N} \phi_{ij}(j,l)\,\pi_{jl}(l) - \phi_{ij}(i,j)\,\pi_{ij}(i)\right)$$

[13]Standard results imply that under stationarity, the solution to this set of Bellman equations is equivalent to the corresponding sequential optimization problem.

denotes the measure of ij-agents who did not visit an intermediary, and as a result enter a bilateral match. Before discussing the implications of (1.1) in more detail, it is worth noting that for any agent, trading away his own consumption good against some other good is strictly dominated by immediate consumption.

Several results follow from (1.1). One observes that for every producer, not visiting an intermediary is a weakly dominated strategy. The analysis of optimal strategies for trade with intermediaries and bilateral trade can be separated. The following proposition summarizes the findings and gives simple rules which optimal trading strategies have to satisfy:

Lemma 1 *If $\{\tau_i, \phi_i\}_{j \neq i}$ is an optimal trade strategy for a producer of type i, then the following must be true:*

(i) If $\phi_{ij}(j,k) > 0$, then $k \in \arg\max_l (V_i(l) - \beta_l - V_i(j))\, \pi_{jl}(l)$, and

$$\max_l (V_i(l) - \beta_l - V_i(j))\, \pi_{jl}(l) \geq (V_i(i) - \sigma_{ij}U - \beta_i - V_i(j))\, \pi_{ij}(i)$$

If k is a unique maximizer, then $\phi_{ij}(j,k) = 1$

(ii) If $\phi_{ij}(i,j) > 0$, then

$$(V_i(i) - \sigma_{ij}U - \beta_i - V_i(j))\, \pi_{ij}(i) \geq \max_l (V_i(l) - \beta_l - V_i(j))\, \pi_{jl}(l)$$

where $\phi_{ij}(i,j) = 1$, if the inequality is strict

(iii) If $\tau_{ij}(k) > 0$, then $V_i(k) - \beta_k - V_i(j) \geq 0$; and $V_i(k) - \beta_k - V_i(j) > 0$ implies $\tau_{ij}(k) = 1$

While evident in its content, lemma 1 highlights the main difference between trade with intermediaries and random bilateral trade: Strategies for the latter amount to simple decision rules that indicate whether one good is accepted in exchange for another, and agents might be willing to accept more than one good in exchange for their current inventory. As a result, trading patterns remain indeterminate, and there may be many possible sequences of exchanges which lead a producer from his current inventory to his consumption good. Trading with an intermediary enables a producer to follow a different strategy and direct himself towards the one trade where his expected surplus is maximized. The producer can follow a predetermined sequence of intermediated exchanges in order to eventually receive his consumption good, and generically, only one such sequence is optimal. Thus, intermediation replicates the structure of models with deterministic trading zones, where agents need to visit an "ij-island" in order to trade good i for good j. However, in contrast to those models, the structure here arises *endogenously* from the activity of intermediaries. Consequently, any delay in trade results from the inability of intermediaries to accommodate all the transactions demanded by producers.

Next consider the requirements for stationarity of the distribution of inventories. Taking as given $\{\nu_{ij}, \phi_{ij}, \tau_{ij}\}$, $\{\mu_{ij}\}$ has to satisfy:

$$\begin{aligned}\mu_{ij} &= \sum_{l=1}^{N} \phi_{il}(l,j)\,\pi_{lj}(j)\,\mu_{il} + \frac{1}{\sum_{k,l}\mu'_{kl}} \sum_{l} \mu'_{il} \sum_{k} \mu'_{kj}\tau_{il}(j)\,\tau_{kj}(l) \\ &\quad + \mu'_{ij}\left(1 - \frac{1}{\sum_{k,l}\mu'_{kl}} \sum_{k,l} \mu'_{kl}\tau_{ij}(l)\,\tau_{kl}(j)\right) \end{aligned} \tag{1.2}$$

whenever $j \neq i, i+1$, and

$$\begin{aligned}\mu_{i,i+1} &= \sum_{l=1}^{N} \left[\phi_{il}(l,i)\,\pi_{li}(i) + \phi_{il}(i,l)\,\pi_{il}(i)\right]\mu_{il} + \frac{1}{\sum_{k,l}\mu'_{kl}} \sum_{l} \mu'_{il} \sum_{k} \mu'_{ki}\tau_{il}(i)\,\tau_{ki}(l) \\ &\quad + \mu'_{i,i+1}\left(1 - \frac{1}{\sum_{k,l}\mu'_{kl}} \sum_{k,l} \mu'_{kl}\tau_{ij}(l)\,\tau_{kl}(j)\right) \end{aligned} \tag{1.3}$$

An i-agent's production good is treated separately from all other goods he may hold as an inventory. Condition (1.2) can be explained as follows: $\sum_{l=1}^{N} \phi_{ij}(l,j)\,\pi_{lj}(j)\,\mu_{il}$ is the measure of i-agents who acquire good j from an intermediary. μ'_{ij} is the set of ij-agents who are unsuccessful in trading with an intermediary, and was already defined above. A fraction $1 - \frac{1}{\sum_{k,l} \mu'_{kl}} \sum_{k,l} \mu'_{kl} \tau_{ij}(l) \tau_{kl}(j)$ of μ'_{ij} is unsuccessful in bilateral exchange as well, which gives the third term. Finally, a measure of $\frac{1}{\sum_{k,l} \mu'_{kl}} \sum_{l} \mu'_{il} \sum_{k} \mu'_{kj} \tau_{il}(j) \tau_{kj}(l)$ acquires good j through a bilateral match. Similarly, $\mu_{i,i+1}$ can be decomposed into those agents who were able to consume after visiting an intermediary, or after a successful bilateral meeting, and those who held good $i+1$ at the start of the period, but were unable to trade. Since holding one's own production good stands at the beginning of any sequence of trades, no agent will trade in his inventory for good $i+1$.

Trading probabilities for trade with intermediaries can be derived from the distribution of inventories and role choices as

$$\pi_{ij}(i) = \frac{\min\left\{\nu_{ij}, \mu_{ij}\phi_{ij}(i,j), \sum_{l} \mu_{li}\phi_{li}(i,j)\right\}}{\mu_{ij}\phi_{ij}(i,j)}$$

$$\text{and } \pi_{ij}(j) = \frac{\min\left\{\nu_{ij}, \mu_{ij}\phi_{ij}(i,j), \sum_{l} \mu_{li}\phi_{li}(i,j)\right\}}{\sum_{l} \mu_{li}\phi_{li}(i,j)}$$

respectively, simply the maximum possible measure of two-way transactions divided by the measure of agents wishing to perform the same transaction.

Finally, I look at the equilibrium implications of competition among intermediaries, which are summarized in the following lemma:

Lemma 2 *Bertrand competition and open entry into and exit from intermediation imply:*

(i)

$$V_i(i) = \frac{1}{1-\delta}\left[\sigma_{ij}U - \delta\left(\beta_i + \beta_j\right)\right] \tag{1.4}$$

(ii)

$$\nu_{ij} = \min\left\{\mu_{ij}\phi_{ij}(i,j), \sum_{l} \mu_{li}\phi_{li}(i,j)\right\}. \tag{1.5}$$

(i) follows from the fact that in equilibrium, any agent who holds a unit of size 1 or $1-\sigma_{ij}$ of his consumption good i has to be just indifferent between consuming everything and remaining a producer, and retaining $1-\sigma_{ij}$ to become an intermediary.[14] (ii) follows as a consequence of Bertrand competition: It states that ν_{ij} has to equal the measure of two-way transactions demanded between goods i and j. If $\nu_{ij} > \min\left\{\mu_{ij}\phi_{ij}(i,j), \sum_l \mu_{li}\phi_{li}(i,j)\right\}$, then some intermediaries would not be able to trade within the period, and would be better off either lowering their mark-up, or leaving intermediation altogether. If $\nu_{ij} < \min\left\{\mu_{ij}\phi_{ij}(i,j), \sum_l \mu_{li}\phi_{li}(i,j)\right\}$, then the demand for transactions exceeds the intermediaries' capacity, and price competition has no bite, i.e. every intermediary would be free to raise his price. Since

$$\nu_{ij} = \min\left\{\mu_{ij}\phi_{ij}(i,j), \sum_l \mu_{li}\phi_{li}(i,j)\right\},$$

the trading probabilities can be rewritten as

$$\pi_{ij}(i) = \min\left\{1, \frac{\sum_l \mu_{li}\phi_{li}(i,j)}{\mu_{ij}\phi_{ij}(i,j)}\right\} \text{ and } \pi_{ij}(j) = \min\left\{1, \frac{\mu_{ij}\phi_{ij}(i,j)}{\sum_l \mu_{li}\phi_{li}(i,j)}\right\}.$$

Summing up, in any steady-state equilibrium, lemma 1 provides necessary and sufficient conditions for the optimality of trading strategies ϕ_{ij} and τ_{ij}, (1.2) and (1.3) have to hold for stationarity of $\{\mu_{ij}\}$, and lemma 2 states that (1.4) and (1.5) determine σ_{ij} and ν_{ij}.

[14] A problem possibly arises in a non-stationary environment, if σ_{ij} decreases from one period to the next. A producer would then be unable to use his unit of consumption to start as an intermediary. Since, in a steady-state, no intermediary has an incentive to lower her mark-up, however, this concern is irrelevant for steady-state analysis.

1.4 Commodity Money

1.4.1 Characterization

In this section, I discuss the development of a common medium of exchange as an equilibrium property of the economy outlined above. The concept of money referred to is commodity money, i.e. a good which is used by all producers for indirect exchange. For an economy with intermediaries, it is straight-forward to conjecture the existence of a type of equilibrium, where an arbitrary good m circulates as money. Any producer chooses to first trade his production good for a full unit of good m, and, once he has acquired m, exchanges m for a reduced unit of his own consumption good. Type m producers trade their production good $m+1$ directly for good m, and producers of good m trade directly for their consumption good $m-1$. In such an equilibrium, $\nu_{im} > 0$, for all $i \neq m$, and $\nu_{ij} = 0$ otherwise, i.e. for each good different from m, there exists an active set of im-intermediaries. I call this equilibrium a Cash-in-Advance equilibrium for good m, because transactions with intermediaries necessarily involve the use of good m as a medium of exchange.

Definition 2 *A stationary equilibrium exhibits a Cash-in-Advance constraint for some good m, if and only if $\{\phi_{ij}\}_{i=1}^{N}$ satisfies $\phi_{ij}(j,m) = 1$, and $\phi_{im}(i,m) = 1$, for all i,j.*

Proposition 1 *If transaction costs are sufficiently small, then for any good m, there exists a stationary equilibrium in which $\{\phi_i\}_{i=1}^{N}$ exhibits a Cash-in-Advance constraint for m.*

Proof. A simple way to prove this proposition is to proceed by guessing and verifying. Given the sets of intermediaries active in equilibrium, any producer has only one possible choice for transactions with intermediaries. If this transaction sequence leaves him with positive lifetime utility, it will therefore be optimal for him to use these strategies. Now, let μ_{im} be the measure of type i producers who hold good m, and $\mu_{i,i+1}$ the measure of type i producers who hold their production good $i+1$. Conjecture further that intermediaries are able to carry out all the transactions demanded by producers with probability 1, for all producer types different from $m-1$ or m. This requires $\nu_{im} = \mu_{im} = \mu_{i,i+1} = \frac{1}{2}\left(\frac{1}{N} - \nu_{im}\right)$, or $\nu_{im} = \mu_{im} = \mu_{i,i+1} = \frac{1}{3N}$, and one easily confirms the conjecture that trading probabilities for types $i \neq m-1, m$ are indeed equal to 1. Type $m-1$ acquires good $m-1$ with probability $\frac{1}{2}$ in each period, and type m acquires his consumption good with probability $\frac{1}{3}$. These types are the only ones that enter random matches, but they will never agree to swap their inventories in a random meeting, since type $m-1$ would never be willing to give up good m for a good that is not his own consumption good - we therefore conclude that no trade will ever take place as a result of a random meeting.

I next determine the mark-up for each intermediary. Whenever a consumer of type $i \neq m$ holds a full unit of his consumption good, he has the choice of becoming an intermediary in the following period, in which case he consumes σ_{ij} now, and uses the remainder to trade in the following period, or consuming the entire unit to become a producer in the following period. In equilibrium, he must be indifferent between the two. For types $i \neq m-1, m$, the life-time utility of an im-intermediary is $\frac{1}{1-\delta}\left(\sigma_{im}U - \delta\left(\beta_i + \beta_m\right)\right)$, where $\sigma_{im}U$ is the utility from consuming the proceeds of one two-way transaction that the intermediary consumes at the end of each period, and $\delta\left(\beta_i + \beta_m\right)$ is the discounted cost of the next period's two-way transaction. The life-time utility of a type i producer before consuming his unit is $U + \frac{\delta^2}{1-\delta^2}\left(U\left(1-\sigma_{im}\right) - \beta_i - \frac{\beta_m}{\delta}\right)$: The type i producer consumes a unit of size $1-\sigma_{im}$ in every other period, and during the intermediate periods he first incurs a transaction cost β_m for acquiring the medium of exchange, and then a cost β_i of acquiring his own consumption good. Equating the two and solving for σ_{im} yields

$$\sigma_{im}U = \frac{1}{1+\delta+\delta^2}\left[U + \delta\beta_i + \delta^2\beta_m\right]. \tag{1.6}$$

The life-time utility of a type $m-1$ intermediary and a type $m-1$ producer are

$$\frac{1}{1-\delta}\left(\sigma_{m-1,m}U - \delta\left(\beta_{m-1} + \beta_m\right)\right)$$

and

$$U + \frac{\frac{1}{2}\delta}{1-\delta}\left(U\left(1-\sigma_{m-1,m}\right) - \beta_{m-1}\right),$$

respectively. In this case,

$$\sigma_{m-1,m}U = U\frac{2-\delta}{2+\delta} + \delta\beta_{m-1}\frac{1}{2+\delta} + \delta\beta_m\frac{2}{2+\delta}. \tag{1.7}$$

To complete the proof, I derive the equilibrium welfare levels denoted by W_i, which must be positive at the point right after a producer has consumed:

$$\begin{aligned}
(1-\delta)\,W_i &= \frac{\delta^2}{1+\delta+\delta^2}\left[\delta\,(U-\beta_i)-\beta_i-\beta_m\frac{1+\delta^2}{\delta}\right], \text{ if } i\neq m-1,m \\
(1-\delta)\,W_{m-1} &= \frac{\delta}{2+\delta}\left[\delta\left(U-\beta_{m-1}\right)-\beta_{m-1}-\delta\beta_m\right], \\
(1-\delta)\,W_m &= \frac{1}{3}\delta\,(U-\beta_m)\,.
\end{aligned} \tag{1.8}$$

As long as transaction costs are sufficiently small, these life-time utility levels are strictly positive, and hence an equilibrium with good m as medium of exchange exists. ■

In this equilibrium, the medium of exchange is the result of the intermediaries' strategies: Their implicit coordination favors one good for the use as a medium of exchange. Intermediaries can deliver this good much more quickly than the search market. If transaction costs are small enough, Bertrand competition among intermediaries guarantees that the benefits of intermediation exceed its costs, so that producers have no incentive to deviate from the proposed trading sequence.

The characterization of Cash-in-Advance equilibria in the previous proposition leads to several immediate observations. First, (1.8) gives an upper bound on transaction costs that must be satisfied so that a Cash-in-Advance equilibrium for good m exists. If all transaction costs are sufficiently small and the discount rate is close to 1, $(1-\delta)\,W_i$ is close to $\frac{1}{3}U$ for all types. We also observe that the producer and consumer of good m enjoy a kind of "rent" in equilibrium: If δ is close to 1, these two types will always prefer to be in a Cash-in-Advance equilibrium for good m, rather than in a Cash-in-Advance equilibrium in which they have to trade twice.

The trading patterns in a Cash-in-Advance equilibrium exhibit a form of *"market-clearing"*: for all types except m and $m-1$, trading probabilities for transactions with intermediaries are equal to 1, i.e. apart from the money producers and money consumers, no one faces delays in carrying out the desired transactions. Thus, in the Cash-in-Advance equilibria of proposition 5, almost all desired transactions are carried out at the prices at which a Walrasian market would clear. *Complete market-clearing* is impossible due to a disequilibrium in transaction sequences: The producers and consumers of the commodity money only trade once in order to consume, while all other types trade at least twice between the time they produce and the time they consume. Thus, the commodity money equilibrium distorts demand and supply of goods m and $m+1$ away from equality at the prices at which Walrasian markets for these goods would clear. Obviously, this result would not be robust, if the environment were altered in such a way that intermediaries could change prices to equate the aggregate quantities demanded and supplied for each transaction. Nevertheless, this discussion leads to an important insight: the liquidity demand for the commodity money distorts such market-clearing prices away from underlying Walrasian prices.

Finally, note the influence of intermediation on random transactions: In particular, observe that in the Cash-in-Advance equilibrium, transactions cease to occur in bilateral meetings. Here, this observation follows from the fact that the search market empties, but it is important to note that the conclusion continues to hold if intermediated transactions were only approximately able to clear the search market, i.e. trading probabilities for intermediated transactions are close to (but smaller than) 1 and some agents enter the matching stage:[15] Each producer who enters the random matching phase either holds the medium of exchange, or some other good. If he holds the medium of exchange, he is only willing to exchange it against his own consumption good. Any agent who holds a good other than the medium of exchange will only exchange it against the medium of exchange or his own consumption good in a random meeting.[16] The only possible trade is then between a pair of agents where one acquires the medium of exchange for his production good, while the other one acquires his own consumption good against the medium of exchange, but those two types would have visited the same intermediary earlier in the period,

[15] In an earlier version of this model, that allows for inventory accumulation by intermediaries, markets clear only approximately. The discussion here is based on the formal results there, which can be found in Hellwig [11].

[16] Trading for a different good has to be dominated, since this can only be motivated by a reduction in future search costs, but with approximate market-clearing, expected future trading probabilities are already close to 1, and hence any reduction in search cost is more than outweighed by the direct cost of an additional transaction.

and hence cannot both enter into a random match. The assumptions about who enters into the random matching phase imply that anyone who has successfully traded with an intermediary gets opportunities to trade in "chance" encounters. A more plausible, yet technically intractable assumption about the random matching process would have been to let all agents participate in a random meeting (independently of whether they were successful in trading with an intermediary or not). We would then come to the conclusion that *if intermediation leads to approximate market-clearing, then transactions occurring in random meetings have to enable the trading parties to save on the costs of intermediation*, i.e. had the two trade partners not met by chance, they would have chosen to carry out the same transaction indirectly through an intermediary. The existence of market institutions that successfully eliminate search frictions therefore has a deep impact on the transactions that arise out of random meetings, and the Cash-in-Advance property of intermediated transactions also extends to random transactions.

1.4.2 Intermediated vs. Pure Search Economies

These observations about the Cash-in-Advance equilibria with intermediaries are in contrast with the characteristics of commodity money in a pure search economy.[17] Since these properties have been studied extensively by Kiyotaki and Wright [18] and Aiyagari and Wallace [1,2], their main results will only briefly be reviewed here. In a pure search economy, a commodity money is defined as a good that is accepted by all producers, whenever it is offered in an exchange. A strategy profile entails a Cash-in-Advance constraint, if the commodity money is part of every transaction that takes place in equilibrium.

The aforementioned papers on search economies show that while commodity money equilibria do exist in pure search economies for at least some goods, these equilibria typically do not entail a Cash-in-Advance constraint: A Cash-in-Advance constraint implies that any agent

[17]Note that intermediation widens the possible set of equilibria of this economy from the one originally studied in Kiyotaki and Wright. If no agent acts as an intermediary, it is weakly optimal for producers not to visit an intermediary. But then, no agent has an incentive for becoming intermediary and trade will only take place in bilateral meetings. Thus, any steady-state equilibrium of the original Kiyotaki-Wright economy can be supported as an equilibrium of this economy with intermediation, setting $\nu_{ij} = \sigma_{ij} = 0$ and $\phi_{ij}(k,l) = 0$, for all i, j, k, l. This reduces the equilibrium definition to the distribution of inventories and to the search strategy profile $\{\mu_{ij}, \tau_{ij}\}_{i,j=1}^{N}$.

must first acquire the medium of exchange, before he can acquire his own consumption good. However, due to asymmetries across goods in the steady-state inventory distribution, goods are endogenously characterized by different qualities for indirect exchange, and there is an advantage in terms of expected waiting time for holding one good rather than another. Since agents are unable to direct their search towards a predetermined sequence of transactions, there is a non-negligible probability that the most preferred transaction cannot be carried out rapidly, which gives an incentive to swap inventories, even if neither of the goods is a universal medium of exchange. The Cash-in-Advance constraint breaks down because the existing search frictions give producers an incentive to trade "goods" for "goods" in an attempt to reduce expected waiting time, before they acquire the medium of exchange.

This incentive is missing when intermediaries successfully eliminate waiting times. In the intermediated economy, the medium of exchange results from the strategic interaction of intermediaries. Producers can direct their transaction strategy towards a predetermined sequence of trades, in this case the one imposed by intermediaries. If the intermediaries are efficient in carrying out transactions, producers are able to carry out the exchange proposed by the trade sequence (almost) immediately. Holding a particular good at time t becomes equivalent in value to exchanging it against the next good of the trading sequence at time $t+1$. In the Cash-in-Advance equilibrium, any good can almost directly be exchanged against the commodity money, so that there is no incentive to reduce search frictions by goods-for-goods trade, as in the pure search model.

1.4.3 Other Equilibria

We can give a simple representation of the Cash-in-Advance equilibrium, that will also be useful to characterize other equilibria: In figure 2, we represent the activity of ij-intermediaries by an arrow leading from i to j. A feasible trading strategy for some producer type is represented by a sequence of arrows that lead from the producer's production good to his consumption good, and only for the last arrow in the sequence, when he acquires his consumption good, the producer can move against the direction of the arrow (that is: accept a good in reduced units from the intermediary).

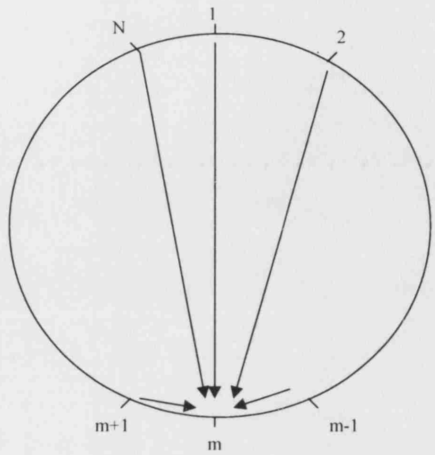

Figure 1-2: Cash-in-Advance equilibrium

In addition to the Cash-in-Advance equilibrium, other equilibria with intermediation exist. Any network of intermediaries that gives every producer type a positive welfare level and exactly one trading sequence by which he can acquire his consumption good can be supported as an equilibrium. In figures 3 through 6, we consider just a few alternative examples of equilibrium intermediation networks, without attempting to provide an exhaustive analysis of all stationary Nash equilibria.

Among these alternative examples, the two-money equilibria are the most interesting ones, and it will be useful for further analysis to provide a characterization. Consider an arbitrary two-money equilibrium with goods l and m as media of exchange, and suppose that type m never becomes an intermediary. In this equilibrium, every producer type trades at most twice between the time when he produces, and the time when he consumes. Again, a simple guess-and-verify procedure shows that each set of intermediaries is of measure $\frac{1}{3N}$, probabilities of trade with intermediaries are equal to 1, with the exception of types $l-1$ and $m-1$, who only trade once, and type m, who does not enter into intermediation. Types $m-1$ and $l-1$ trade their production good directly for their consumption good with probability $\frac{1}{2}$ in each period. For type m, the equilibrium inventory distribution (and hence the trading probabilities) are indeterminate, with possible solutions $(\mu_{m,m+1}, \mu_{m,l}) = (\frac{1}{3N} + \zeta, \frac{2}{3N} - \zeta)$, for any $\zeta \in [0, \frac{1}{3N}]$.

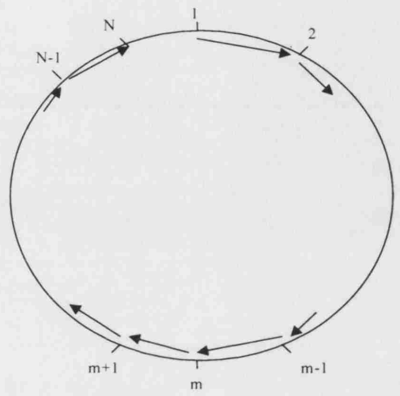

Figure 1-3: "Trade-one-up" equilibrium: for $i = 1, ..., N-1$, there exist $i, i+1$-intermediaries. All agents trade their production good directly against their consumption good, except for type N, who trade good 1 for good 2, then 3, etc. until they receive good N.

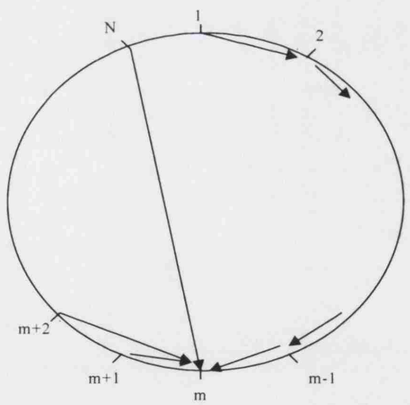

Figure 1-4: This combines a Cash-in-Advance constraint with case (i): Types 1 to $m-1$ trade their production good directly against their consumption good, type N trades good 1 for good 2, then good 3 and so on, until he receives the medium of exchange m, which is used as a medium of exchange by types m to $N-1$.

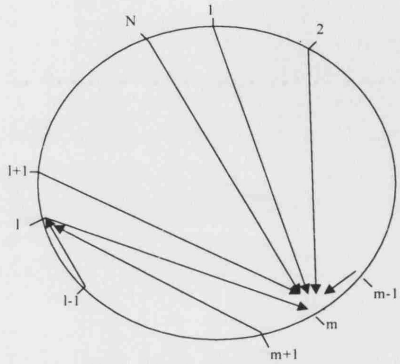

Figure 1-5: Two-money equilibrium: for $i = l+1, ..., m-1$, there exist im-intermediaries, and for $i = m, ..., l-1$, there exist il-intermediaries. In this case, goods l and m are both locally used as medium of exchange, l by types m to $l-1$, and m by types l to $m-1$.

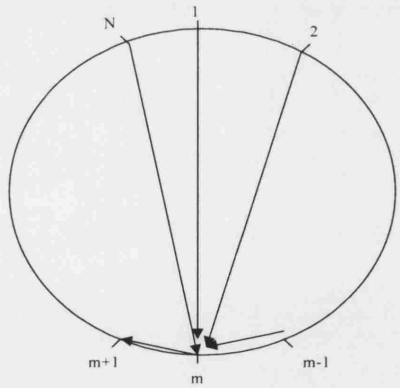

Figure 1-6: Alternative Cash-in-Advance equilibrium, in which $m, m+1$-intermediaries replace the $m+1, m$-intermediaries. Note that this also represents a special case of the two-money equilibrium, where goods m and $m+1$ are used as media of exchange.

Thus, on average, type m consumes every third period. For all types different from m, the indifference condition for each type is equivalent to those obtained in the Cash-in-Advance equilibrium, (when adjusting the indices for the good which each type uses as a medium of exchange). Thus, for $i \neq l-1, m-1, m$ the equilibrium mark-ups are

$$\sigma_{ij} U = \frac{1}{1+\delta+\delta^2}\left[U + \delta\beta_i + \delta^2\beta_j\right]. \tag{1.9}$$

where $j \in \{l, m\}$ represents the medium of exchange for which i is traded. For $i \in \{l-1, m-1\}$, the mark-ups are

$$\sigma_{i,i+1} U = U\frac{2-\delta}{2+\delta} + \delta\beta_i\frac{1}{2+\delta} + \delta\beta_{i+1}\frac{2}{2+\delta}. \tag{1.10}$$

and welfare levels are

$$\begin{aligned}
(1-\delta) W_i &= \frac{\delta^2}{1+\delta+\delta^2}\left[\delta\left(U-\beta_i\right) - \beta_i - \beta_j\frac{1+\delta^2}{\delta}\right], \text{ if } i \neq m-1, l-1, m \\
(1-\delta) W_i &= \frac{\delta}{2+\delta}\left[\delta\left(U-\beta_i\right) - \beta_i - \delta\beta_{i+1}\right], \text{ if } i = m-1, l-1
\end{aligned} \tag{1.11}$$

The equilibrium welfare level for type m is indeterminate, since it depends on the equilibrium inventory distribution, i.e. on ζ. In the simplest case where $\zeta = \frac{1}{3N}$ (i.e. type m trades good $m+1$ for l with probability $\frac{1}{2}$, and then l for m with probability 1), his welfare level is[18]

$$(1-\delta) W_m = \frac{\delta^2}{2+\delta}\left[\delta\left(U-\beta_m\right) - \frac{1}{\delta}\beta_l\right]$$

The same observations that applied to the Cash-in-Advance equilibria also apply to any two-money equilibrium. In particular, there are equilibrium "rents" accruing to three types: those who produce the two media of exchange, and the one type who is able to consume in integer units. I conclude this section by showing that any equilibrium, in which producers trade at most twice in order to consume has to be either a Cash-in-Advance equilibrium or a two-money equilibrium. Furthermore, I show that any mixed strategy equilibrium is Pareto-dominated.

[18] For other values of ζ, the welfare level is slightly lower, but for any ζ, this discrepancy vanishes, if δ is close to 1.

For δ close to 1, and a small level of transaction costs, it then follows that these two classes Pareto-dominate any other stationary equilibrium profile. On the other hand, the existence of equilibrium "rents" prevents a Pareto ranking between equilibria within these two classes.

Lemma 3 *(i) Any pure-strategy equilibrium, in which producers trade at most twice in order to consume is either a Cash-in-Advance equilibrium or a two-money equilibrium.*

(ii) Any mixed-strategy equilibrium is Pareto-dominated by some Cash-in-Advance equilibrium or two-money equilibrium.

Proof. (i) In any pure-strategy equilibrium, an intermediation network consists of exactly $N-1$ sets of intermediaries.[19] $N-2$ types trade twice, while the remaining two types trade once. If type i and $i+1$ both trade twice, they use the same good as a medium of exchange, and in equilibrium, at most two goods are used as commodity money. (ii) Consider an arbitrary mixed-strategy equilibrium, in which all types trade at most twice between the moment when they produce, and the moment, when they consume. The equilibrium transactions network has to connect all N types so as to enable them to trade their production good for their consumption good in some sequence of transactions. This implies that the transaction structure of some Cash-in-Advance or two-money equilibrium (which are the minimal transactions networks) has to be included in any equilibrium transactions network; moreover, if the equilibrium is mixed, the inclusion is strict, and since some types follow two different transaction patterns in equilibrium, there must be at least two different Cash-in-Advance or two-money transactions networks that are embedded in the mixed strategy network. But since for each type, the welfare attained in the mixed strategy equilibrium has to be lower than either of the two pure-strategy networks, it follows that the mixed strategy equilibrium has to be Pareto-dominated. ■

[19] $N-1$ is actually the minimum to sustain a complete intermediation network that enables everyone to carry out all transactions through intermediaries. If there are more sets of intermediaries active, at least one type must have at least two alternatives to trade from his production good to his consumption good, and hence must be indifferent and mix in equilibrium.

1.5 Efficiency and Evolutionary Stability

The previous discussion of Cash-in-Advance equilibria with and without intermediaries raises the problem of multiple equilibria, selection of a medium of exchange, and coordination of strategies. Related to the selection issue are the welfare properties of the various equilibria. In this section, I argue that intermediation may enable the market participants to coordinate equilibrium strategies so as to achieve steady-state welfare levels that are Pareto-efficient. Loosely speaking, I study whether an arbitrarily small set of agents (containing both intermediaries and producers), by deviating from an equilibrium and coordinating their actions off the equilibrium path, can improve their welfare level and consequently induce other agents to join their deviation. This idea is formalized by the definition of evolutionary stability, which requires that no small set of "mutants" can induce a large population of players to alter their behavior. However, note that in contrast to the traditional application of evolutionary stability to large population matching games, it here acts through the coordination of strategies among *various types of players*, as considered by definitions of *coalition proofness* concepts. Below, I briefly comment on the relation between the two concepts in the present environment.

Definition 3 *A Stationary Equilibrium is evolutionarily stable, if for every $\varepsilon > 0$, there exists $\eta^* > 0$, such that whenever the strategies of a set of players of measure $\eta \leq \eta^*$ is exogenously fixed (i.e. a set of measure η of players "explicitly coordinates their strategies" or "deviates"), there exists a stationary equilibrium in this modified game, whose Euclidean distance from the original equilibrium does not exceed ε.*

Using a standard continuity notion, this definition thus formulates the requirement that the coordination of a small set of players should not have a large (discontinuous) effect on the equilibrium. From a historical perspective, the idea of small deviating coalitions is meant to capture the innovating role of intermediaries, whether it comes through explicit innovation and coordination, or through arbitrary "experimenting", i.e. evolutionary forces: Someone proposes a new system for organizing transactions. If others find that this arrangement is efficient, they will also start using it. Since media of exchange, and more generally trading strategies are complementary across agents, everyone will start using the new system, if it leads to a Pareto-improvement. Clearly, intermediation is essential in promoting an innovation in the transactions system.

One should note that in any kind of decentralized trading economy, explicit coordination of several *types* of players is necessary for any successful deviation from an equilibrium. Trading environments, however, differ in *how many* agents need to coordinate to break out of a given equilibrium. In a pure search economy a la Kiyotaki-Wright, equilibrium payoffs are continuous in the strategy profile and the equilibrium inventory distribution, and since the "mutants" have no way to directly trade with each other (even if they can agree on how to deviate), they have only a marginal impact on the payoffs of the non-mutant population. Hence any equilibrium where payoffs are strictly higher than the next-best alternative is immune to the invasion by a small set of mutants, and, in order to break out of an inefficient equilibrium, the explicit coordination of a *large* coalition of players is needed.

This conclusion is fundamentally altered in the intermediated economy that we study here: To be specific, suppose that some equilibrium network of intermediaries is dominated by another one. In that case, a small set of agents may become intermediaries and coordinate their actions on the new intermediation network with some small set of consumer-producers. If, by doing so, both the intermediaries and their clients realize a higher life-time utility than the equilibrium strategy profile, other agents have an incentive to deviate from their initial strategy profile to take advantage of the higher life-time utility offered by the new intermediaries. If eventually, a non-zero measure of agents decides to abandon their equilibrium strategies, the old equilibrium is no longer stable and will be replaced by the new one. This type of coordination is more explicit than the one resulting from Nash equilibrium strategies, yet it only requires coordination of an arbitrarily small, positive measure of agents. The key insight here is that intermediation enables the mutants not only to coordinate their trading strategies, but also to coordinate on *trading with each other* so as to make the deviation profitable for the mutants, and to provide others an incentive to join the deviation.

In a trading economy with intermediaries, it turns out that evolutionary stability is actually equivalent to a more general form of coalition proofness, where we assume that players can coordinate only their strategies with respect to intermediation. To see this, suppose that a large coalition could deviate from an existing equilibrium and leave all its participants better off by proposing a new network of intermediation. Then, this change could also be implemented by an arbitrarily small coalition that starts to form the new intermediation network. Over time, other agents will be induced to switch their strategies, until the entire large coalition has deviated from the existing network to a new one.

It follows immediately from the above discussion that a steady-state equilibrium is unstable, if there is an intermediation network, which Pareto-dominates it. To formalize this idea, I use a definition of Pareto efficiency that is constrained to comparisons of intermediation networks (i.e. changes in the strategy profile for transactions with intermediaries, $\{\phi_{ij}\}_{i,j=1}^{N}$. This excludes inefficiencies resulting from the transactions in bilateral search meetings, however, in an equilibrium, in which all agents trade with intermediaries with very high probability, the resulting strategies τ_{ij} are prescribed by the network of intermediaries, and have only minor welfare implications.

Definition 4 *A stationary equilibrium* $\{\mu_{ij}, \nu_{ij}, \phi_{ij}, \sigma_{ij}, \tau_{ij}\}$ *is constrained Pareto-efficient, if there does not exist* $\left\{\mu_{ij}^{\circ}, \nu_{ij}^{\circ}, \phi_{ij}^{\circ}, \sigma_{ij}^{\circ}, \tau_{ij}^{\circ}\right\}$, *such that*

(i) $\phi_{ij} \neq \phi_{ij}^{\circ}$ *for at least one pair* i, j.

(ii) For all i, τ_{ij}° *is optimal given* $\left\{\mu_{ij}^{\circ}, \nu_{ij}^{\circ}, \phi_{ij}^{\circ}, \sigma_{ij}^{\circ}, \tau_{ij}^{\circ}\right\}$.

(iii) $\left\{\mu_{ij}^{\circ}, \nu_{ij}^{\circ}, \phi_{ij}^{\circ}, \sigma_{ij}^{\circ}, \tau_{ij}^{\circ}\right\}$ *is a Pareto improvement over* $\{\mu_{ij}, \nu_{ij}, \phi_{ij}, \sigma_{ij}, \tau_{ij}\}$.

Proposition 2 *A stationary equilibrium is evolutionarily stable, if and only if it is constrained Pareto-efficient.*

Proof. Clearly, if some Pareto improvement can be implemented by a change in the intermediation network, then groups of intermediaries and producers can implement this change on a small scale, and increase their personal welfare. Everyone else now individually has an interest in changing to the new strategies.

To show the converse, note that by virtue of lemma 3, in any constrained Pareto-efficient steady-state, at most two different goods are used as media of exchange, and if type i uses good

m as a medium of exchange, then type $i+1$ either consumes good m or also uses m as a medium of exchange. Similarly, in a successful deviation, types i and $i+1$ use the same medium of exchange (different from m). It follows that for a deviation to be successful, all N types have to participate in the deviation, i.e. be made no worse of than initially, which contradicts the definition of constrained Pareto efficiency. ∎

This result diverges from the main conclusions about search economies without intermediaries, where the continuity of objective functions with respect to strategies implies that small deviations change overall utility only marginally. Changes in the intermediation network may lead to discontinuous changes in payoff, and thus to strategy changes by a large fraction of the population. The second half of proposition 2 critically relies on the assumption of full specialization of production, i.e. good $i+1$ is produced *only* by type i. In the next section, we relax this assumption. In that case, a deviating coalition does not have to rely on the participation of all producer and consumer types, and hence an equilibrium may be Pareto-efficient, but not evolutionarily stable, if the implemented changes lead to welfare losses for agents who do not participate in the change. Some agents may strictly prefer the old equilibrium over the innovation, but once the innovation is introduced, they will change, because their trade partners also start using the new medium of exchange. Loosely speaking, different media of exchange are substitutes, but there are complementarities in using a medium of exchange.

What are the implications for the steady-state equilibria considered in the previous section? As an immediate consequence of lemma 3 and proposition 2, we have the following result:

Proposition 3 *If transaction costs are sufficiently low, the set of evolutionarily stable equilibria is a subset of the Cash-in-Advance and two-money equilibria.*

We thus come to the conclusion that the monetary structure of transactions appears in any evolutionarily stable equilibrium of the previous section. In either case, trading probabilities equal 1, i.e. markets clear. What importance can be attached to the evolutionarily stable equilibria where good 1 is not a universally accepted medium of exchange? In all minimally coalition-proof equilibria, one type i of agents does not offer any intermediation, and as a consequence, consumes his consumption good in integer units. Due to the assumption of full specialization of consumption and production in this model, a deviating coalition can impose a good as the universal medium of exchange, only if this type participates in the deviation. Among this remaining set of equilibria, an equilibrium is unstable, only if it is dominated for all types, *including the types who benefit from a rent as producers or consumers of a medium of exchange.* Precisely the existence of such rents makes it impossible to break away from some of the two-money and Cash-in-Advance equilibria.

1.6 General Results

1.6.1 Less than full Specialization

In this section, I discuss how the previous results are affected by a generalization of the assumptions concerning consumption and production. The point of departure for this discussion is the observation that full specialization of production and consumption choices, i.e. the assumption that type i is the only type to produce good $i+1$, protects equilibrium rents to money producers and consumers and thereby induces multiple evolutionarily stable equilibria. The following set of assumptions departs from full specialization of production and consumption patterns:

($A0$) *There are N goods and $M \geq N$ types of measure $\frac{1}{M}$ of agents. Each type i is characterized by a production good $p(i)$ and a consumption good $c(i)$.*

($A1$) *For each good, the total number of types consuming the good equals the total number of types who produce it.*

($A2$) *For every pair of types i, j, if $c(i) = p(j)$, then $p(i) \neq c(j)$.*

($A3$) *For every triple of types i, j, k, if $c(i) = p(j)$ and $c(j) = p(k)$, then $c(k) \neq p(i)$.*

($A4$) *Each good is produced by at least two types.*

Under ($A1$), this market would clear at relative prices of one for one, if the market environment was Walrasian. ($A2$) introduces the well-known "double coincidence problem", that there are no two types of agents who could just produce for each other. ($A3$) excludes the possibility of "three-way coincidences", i.e. situations, where a single type could successfully act as a middleman between two other types; effectively, ($A3$) implies that, to get exchange off the ground, at least two types must coordinate their transaction activities and agree on one good as a medium of exchange. ($A4$) rules out full specialization.

Under these assumptions, one notes that the only candidate for an evolutionarily stable pure strategy equilibrium is the Cash-in-Advance equilibrium, in which good 1 is used as a common medium of exchange. All other pure strategy equilibria are destabilized by a small group of players comprising a strict subset of types, who coordinate on using good 1 as a medium of exchange, but don't have to rely on the participation of a type who enjoys an equilibrium rent (formally, the converse of proposition 2 no longer applies). If transaction costs are small, these rents are small, and eventually all types will prefer the more efficient equilibrium trading network.

It remains to be shown that the Cash-in-Advance equilibrium for good 1 exists, and is evolutionarily stable. While existence is immediate, the properties of Cash-in-Advance equilibria do not automatically carry over: As can be shown by example, the transaction probabilities for transactions with intermediaries need not equal 1, and hence the equilibrium may fail to exhibit market-clearing. Intuitively, the imbalance in transaction sequences induced by the use of one good as a common medium of exchange now affects trading probabilities throughout all markets, and the overall frequency of consumption then remains suboptimal. But since some small deviation that uses a more costly good as a medium of exchange can offer its members a higher frequency of consumption, the Cash-in-Advance equilibrium for good 1 will not be evolutionarily stable, if it fails to lead to market-clearing.

1.6.2 Fiat Money

I now extend the model to allow for the circulation of a fiat money, labelled good 0, and traded with a transaction cost β_0. Following the search literature, I assume that a fraction S of the population each holds one indivisible unit of fiat money at any point in time. Following along the same lines as before, one can now discuss the existence of a Cash-in-Advance equilibrium for fiat money:

Proposition 4 *Suppose that (i) $S = \frac{1}{3}$ and (ii) $\beta_0 < \beta_1$. Then, under assumptions $(A0)-(A4)$, the Cash-in-Advance equilibrium for fiat money clears markets and is the unique evolutionarily stable steady state equilibrium.*

Proof. Proceeding along the lines of proposition 1, it is straight-forward to show the existence of a Cash-in-Advance equilibrium for fiat money. If $S = \frac{1}{3}$, markets clear exactly, and mark-ups and welfare levels are given by

$$\sigma_{i0}U = \frac{1}{1+\delta+\delta^2}\left[U + \delta\beta_i + \delta^2\beta_0\right]$$

and

$$(1-\delta)\,W_i = \frac{\delta^2}{1+\delta+\delta^2}\left[\delta\,(U-\beta_i) - \beta_i - \beta_0\frac{1+\delta^2}{\delta}\right].$$

Since no type produces or consumes fiat money, there are no rents associated with its production or consumption. Note that this equilibrium is evolutionarily stable, if and only if $\beta_0 < \beta_1$: In that case, $(A2)$ and $(A3)$ together imply that any coalition that tries to deviate from the fiat money Cash-in-Advance equilibrium has to include one type who is willing to accept a higher-cost good as a medium of exchange, without enjoying a rent as a money producer or consumer. But then he must be made worse of, and no one will be willing to follow his strategy - on the other hand, if $\beta_0 \geq \beta_1$, such a coalition may exist, and successfully deviate from the proposed equilibrium.

Finally, note that in any other equilibrium, $(A0) - (A4)$ imply that there has to exist a subset of types $i_1, i_2, ..., i_n$ who form a circle, i.e. $c(i_1) = p(i_2)$, $c(i_2) = p(i_3)$,..., $c(i_n) = p(i_1)$, where neither of these types produces or consumes a medium of exchange in equilibrium. This subset of types can successfully mutate to start using fiat money in a steady-state. ∎

This proposition states the central theoretical result of this paper and provides a foundation for a Cash-in-Advance equilibrium with fiat money as the unique evolutionarily stable equilibrium in a decentralized trading economy with intermediaries. The result is tied to a series of conditions, which are arguably of a technical nature. The assumptions $(A0)$ - $(A4)$ rule out possibilities for double or three-way coincidences, and assume less than full specialization.

Their role is to rule out coordination among a small number of types[20], or the protection of rents, when a deviation has to rely on the participant of a money producer or consumer. The condition $S = \frac{1}{3}$ states that the supply of real balances has to equal the demand for transactions purposes. Matsui and Shimizu [21] show that such a condition arises as the unique evolutionarily stable steady-state in a related model that makes stronger assumptions about the structure in which transactions take place and notably allows for nominal price adjustments, . Finally, the condition that $\beta_0 < \beta_1$ states that it has to be desirable from an individual perspective to use fiat money as a medium of exchange - alternatively, individuals always have an incentive to use commodities for indirect exchange to save on transaction costs.

Proposition 4 should not be read as a statement that rules out the observation of anything but a fiat money equilibrium in a steady-state. Rather, it states that as the long-run outcome of an evolutionary process, in which intermediaries play a central role in coordinating transaction strategies, one should expect the most efficient transactions arrangement (in this case a fiat money equilibrium) to prevail. Since the proposition emphasizes the uniqueness of the long-run outcome, it also contrasts with the multiplicity of equilibria within Walrasian and search models of monetary exchange, a 'problem' that was first discussed by Hahn [10]. Of course, the result remains silent about the question of how the long-run equilibrium is obtained, or what is observed in the interim stages. In this respect, the multiplicity of equilibria retains its relevance, as many of the observations made within the context of the search literature, or earlier in this paper with respect to commodity money may remain relevant as descriptions of intermediate stages of the evolutionary process, or as the consequence of aggregate shocks leading to a temporary break-down of market institutions and intermediation.

Nor should this proposition be viewed as stating that the final stage of the evolutionary process of market transactions will be a fiat money equilibrium as the one described here. Indeed, one of the constants of the history of market organization and transactions is innovation and change, and virtually every innovation is promoted or coordinated by some kind of inter-

[20] With three-way coincidences three types could coordinate a deviation towards a "local" commodity money, in such a way that one type trades twice, but consumes in integer units, while the other two types trade only once. Each type then gets to enjoy a small rent, which might be enough to offset the loss of using a higher-transaction cost medium of exchange. An even simpler argument applies, of course, for double coincidences.

mediary. The introduction of credit cards and other cashless means of transactions for example can be viewed as a move away from cash towards more efficient alternatives.

1.6.3 Free Banking

While the previous section discussed the emergence of a fiat money Cash-in-Advance equilibrium as an evolutionarily stable steady state, it does not answer the question of how fiat objects get into circulation in the first place. This is also a question that existing search-theoretic models do not answer appropriately: by focusing on steady-state equilibria, they focus on environments, where fiat objects have been around forever in the past, and are valued, because they are expected to be valued forever into the future. The purpose of this section is to illustrate how fiat money may *come into circulation* in a "free banking" equilibrium. For this purpose, I adapt the model by enabling intermediaries to issue debt certificates on which they promise to pay a unit of physical goods in the future.

To be specific, suppose that every intermediary has the ability to write out demandable debt claims, i.e. notes that are backed by her inventory of goods, and that are sold in transactions. Under what conditions do these notes start to circulate as media of exchange, and become perfect substitutes from the producers' perspective? To make such a system viable, it is necessary that in steady-state, intermediaries have an incentive to discipline their note issue, and not overissue notes to default in the future. This incentive compatibility requires that notes eventually return to their issuer. Again for the sake of concreteness, I start by assuming that this occurs at the end of each period, when all intermediaries participate in a clearing market, where they return any notes to the initial issuer. Below, I will also take into consideration other clearing mechanisms, as well as different assumptions concerning note-issuing privileges.

Given a transactions cost of β_0 for accepting notes (and assuming that there are no costs involved in the clearing process), proposition 4 characterizes the Cash-in-Advance equilibrium, *provided that intermediaries have an incentive to refrain from overissuing.* The equilibrium behavior of intermediaries and the circulation of notes is characterized as follows: Each period begins with half of the producers holding their production good (those who previously consumed) and half of the producers holding fiat object, i.e. a note issued by the intermediary to whom they sold their production in the previous period. Intermediaries begin each period with one note outstanding and a reduced unit of their consumption good. They then sell this unit in exchange for the bank notes held by visiting producers who purchase their consumption good, and withdraw this note from circulation. Afterwards, they purchase an integer unit of their consumption good from some producer who wishes to sell his production, and pay for it by issuing a new note. After all transactions with producers are completed, intermediaries meet in a clearing market and exchange the notes they withdrew from circulation. Since every intermediary had one note outstanding, the market clears, and the following period begins with each intermediary having one note outstanding.

Under the conditions of proposition 4, such a free banking equilibrium is evolutionarily stable and may account for the emergence of a fiat object, if it provides intermediaries with the right incentives to participate in the clearing mechanism and not overissue. These incentives depend on the comparison between the short-term gains from overissuing, and the potential long-run punishment in case of default. In the present case, if an intermediary decides to overissue and default, she can issue one note during one period, and not accept someone else's note in return for her consumption good; she is then found out at the end of the period, when she fails to clear her note in the clearing market. Her short-term benefit is then equal to $\beta_0 + (1 - \sigma_{i0})\,U$. The cost of default depends on the punishment structure. In the most lenient case, this punishment might simply involve the loss of her note-issuing privileges; given the open entry condition, the intermediary could become a producer, and continue without any welfare losses.[21] Alternatively, the most severe punishment might involve total exclusion from the market (as a producer or an intermediary) for all future periods - which would amount to the loss of all future consumption. Whatever the punishment mechanism, the cost of punishment must exceed the short-term gains of over-issuing.

Under alternative note-issuing and clearing arrangements, notes may not be redeemed im-

[21] This statement relies critically on the existence of open entry into intermediation, i.e. note-issuing. A related discussion by Cavalcanti, Erosa and Temzelides [5] of private note-issuing in a search model discusses a stable private money equilibrium relying solely on the withdrawal of note-issuing priviledges as an incentive mechanism, but in their environment there is no open entry into the banking sector, i.e. note-issuing priviledges are exogenously fixed, and default leads to the loss of seignorage (i.e. essentially scarcity) rents. See also Hellwig and Lorenzoni [12], who discuss reputational mechanisms for sustaining note circulation in a Walrasian equilibrium with borrowing constraints.

mediately, but circulate for several periods. This may happen, for instance, if the clearing market opens only infrequently, and instead notes are returned to the producers within the same period. Alternatively, only a limited number of types may issue and clear notes, and hence notes stay with the public or non-issuing intermediaries for several periods until they are returned to the issuer. To provide a specific example, return to the environment studied initial, where type i produces good $i + 1$. If notes are issued only by type 1 intermediaries (and only type 1 agents participate in the clearing), these notes circulate for N periods before they return to the initial issuer (From type 2 to type 3 to type 4 to... until they reach type N producers, and then type 1). In that case, a single producer might be willing to become a "rogue" intermediary, and start issuing notes, without ever exchanging back the consumption units. Since the notes take N periods until they return to the issuer, the intermediary does not have to redeem any notes immediately. Hence, it takes N periods to detect someone who overissues notes, and thus the short-run benefit of over-issuing increases. Similarly, the time of circulation of a note increases, when note clearing takes place less often, or not at all. Since an overissuing intermediary is detected only once the note is redeemed, the short-run benefits of over-issuing are proportional to the expected time of circulation of the notes.

While far from complete as a theory of free banking, this discussion points to some of the features that determine the stability of a free banking system. Clearing arrangements and the length of time a note circulates determine the short-run gains from over-issuing. These short-run benefits are contrasted with the long-run costs of default, determined by the harshness of punishment, as well as the potential loss of seigniorage or monopoly rents. Note that the clearing arrangements have no direct allocational purpose, but simply serve to decentralize the monitoring of the intermediaries' note-issuing activities. Finally, the model points to a coordination problem that arises in the clearing of notes: the equilibrium described above relied on the participation of intermediaries in the clearing market, and this participation was individually rational, since in order to clear her note, the intermediary had to return a note she collected to the initial issuer. There is, however, an alternative equilibrium, in which all intermediaries, instead of returning the notes they collected to the issuer, decide to return them to the producers within the same period, or withdraw them without clearing them, keeping them as reserves. Since no intermediary is clearing any notes, there is no reason for any intermediary to participate in the clearing, and if clearing notes is associated even with a small cost, intermediaries collectively prefer the no-clearing equilibrium. But then, the clearing market ceases to perform its monitoring role, and some intermediaries may find it optimal to default. Summing up, the model suggests that free banking regimes are stable, when:

(*i*) Notes circulate only for short periods, and quickly return to the issuer through a well-functioning clearing system, and

(*ii*) The loss of seigniorage or monopoly rents after a default provides incentives not to overissue.

I conclude this discussion with a brief review of some historical free banking episodes. Proponents of free banking typically point to Scotland as a country where free banking was extremely stable throughout several centuries. As is extensively discussed by Smith [30] in her classic analysis of free banking, the Scottish banking system indeed fulfills the conditions laid out here, providing the most clear-cut example: Although labelled as "free" banking, the banking sector really had an oligopolistic, almost cartell-like structure with a small number of large players. These bankers met on a very regular (weekly) base to clear notes, and notes stayed in circulation for short time periods only. In addition, they were subject to unlimited liability in case of a default. Over a stretch of approximately three centuries, Scotland had virtually no banking panics or defaults. Another example of free-banking success was the Suffolk bank system in nineteenth century New England, described for example by Smith and Weber [29]. The Suffolk Bank, one of the biggest note-issuing banks in the area, internalized the cost of running a clearing market by accepting the notes of other banks at parity, if these banks made a large up-front deposit. Other regions in America did not have as sophisticated clearing mechanisms during the free banking era in the nineteenth century, and thus had longer times of note circulation, and coupled with a legal system that made default more acceptable than in most European countries and free entry into banking, this lead to a higher degree of instability, banking panics, and defaults.

An intriguing final example for the ultimate failure of a free banking system despite initial success is Switzerland during the nineteenth century, as described by Neldner [23]. Although the system performed reasonably well by all conventional accounts (even though highly competitive, it was unusually safe, didn't lead to bank panics or failures throughout almost the entire century, and no noteholder ever incurred a loss due to default), it ultimately failed and was replaced by the Swiss National Bank in 1907. According to Neldner, while the system initially was very successful, during the last third of the nineteenth century, it suffered from overissuing of notes and a malfunctioning of the clearing market, even though very sophisticated clearing arrangements did exist. During this time period, the position of note-issuing banks was weakened by the arrival of non-issuing commercial banks, who held a competitive advantage in the market for loans. This gradually led note-issuing banks to reduce the clearing of notes, in fear of receiving their own notes in return, and having to pay in species. Within a very competitive environment, this led to an overissue of notes, and ultimately the deterioration of the exchange rate towards the French franc and an outflow of species from the country. The banks, however, did not return the notes which they received in exchange for the species to the issuer, preferring to return them directly to the market, and thus further slowing down the clearing process. While the exact causal link between these events is not entirely clear from Neldner's analysis, one possible interpretation might be that a weakening of the note-issuing banks led them to gradually reduce the clearing activities, i.e. switch from one equilibrium in the clearing market to another, in a collective attempt to maintain their economic viability.[22]

1.7 Conclusion

This paper studied a decentralized trading economy in which intermediaries induce the use of a common medium of exchange. As such, intermediation and money are complementary phenomena. Strategic interaction of intermediaries leads to a Cash-in-Advance constraint, such that trade sequences with intermediaries follow the well-known pattern that "goods buy money and money buys goods, but goods don't buy goods" (Clower [7]). As opposed to many other models of monetary exchange, this pattern comes as an equilibrium outcome and not an assumption of the model. The second central result is that the characteristics of a monetary equilibrium with intermediaries differ fundamentally from those of equilibrium models without

[22]Remarkably, the note-issuing banks were willing to collectively restrain from clearing notes, but they were unable to coordinate to limit the amount of note-issuing, even though this would have improved their competitive position relative to the commercial banks.

intermediaries. By coordinating its deviations, a small coalition of intermediaries can induce producers to use transaction strategies that ultimately lead to an efficient equilibrium. Under some additional conditions, the unique evolutionarily stable equilibrium leads to a Cash-in-Advance constraint for a fiat currency. I further study how this equilibrium can be implemented in a free banking equilibrium, in which fiat money is brought into circulation as debt certificates issued by intermediaries.

A series of questions cannot be properly addressed within the framework of this model. Most importantly, production and consumption choices remain outside the model. As in many related models, I have simply assumed the existence of an underlying Walrasian equilibrium, which in the absence of search frictions also represents an optimum. Production and Consumption choices are exogenously given in such a way that in a frictionless economy, all markets would clear at the relative price of 1. In an environment where trade is subject to frictions, assuming that consumption and production decisions do not depend on decisions about trade is problematic, since decision-makers would take into account their opportunities for trade when they decide what goods to produce or to consume. They may decide to produce one good because it is easy to trade, even though they are more efficient at producing a different, less marketable good. This problem does not appear, however, in discussions on exchange in decentralized economies with frictions, neither here, nor in the literature on which this paper builds. Furthermore, I have effectively abstracted from the problem of embedding a theory of prices into this decentralized trading economy, assuming that transactions take place at the implicit Walrasian prices. Under what conditions these prices prevail in a search or otherwise non-walrasian economy is an open question, since within each transaction, price formation would have to be modelled the result of some bilateral bargaining process, and hence also depend on the trading partners' outside options, which in turn depend on the trading environment. However, as discussed in the context of proposition 1, the liquidity demand for the medium of exchange creates some inherent price distortions away from the Walrasian equilibrium.

Finally, it should be noted that the model presented here relies on some ad hoc assumptions about intermediation. Most importantly, the idea that a limit to intermediation generates a need for intermediaries to introduce a common medium of exchange requires some reflection.

Precisely which technical restrictions affect the behavior of intermediaries, and how do they alter a given trading environment? While such constraints are taken as given in this context, further thought is needed in order to assess the validity of the way intermediation is introduced into the trading environment here, and the robustness of the results that follow from it. Also, the evolutionary stability arguments rely on the assumption that intermediaries are immediately accessible to producers; this effectively enables them to coordinate on an arbitrarily small scale. In an environment, where intermediaries are not immediately accessible to their customers, say, when intermediaries are simply more efficient at searching for matches, the evolutionary stability argument breaks down.[23] While this paper takes the stand that such coordination will eventually occur, whenever it is feasible, questions arise as to how intermediaries coordinate, in short, what happens along the transition towards the long-run steady-state.

Despite these technical shortcomings and open questions, the results presented here provide some more general insight into the role of intermediation. The complementarity of the medium of exchange and intermediation and the non-stability of inefficient transaction patterns both follow from three basic assumptions about the environment:

(*i*) A Pareto-optimal, market-clearing allocation, which would result from a competitive equilibrium in perfect markets, cannot be attained because of a form of market imperfection,

(*ii*) some agents have a technology to alleviate the imperfection by offering intermediation, and by offering this technology to the economy, they can make arbitrage profits from a price spread, and

(*iii*) the success of intermediaries depends on how they can deal with their own constraints.

In general, we know many reasons for frictions in a competitive economy, and the many facets and different forms of intermediation all respond to these imperfections. Here, I have considered search frictions as the underlying imperfection, however, one might try to apply the same logic

[23] Consider for example a random matching environment, where intermediaries find matches at a faster rate than producers. One would conjecture that the measure of agents who need to coordinate to break an inefficient equilibrium is decreasing in the frequency with which they encounter an intermediary, i.e. the more efficient intermediaries are in the matching process, the easier it is to break out of an equilibrium. This conjecture is open to future research.

to study how intermediation interacts with decentralized market instruments to alleviate other frictions, such as the lack of public memory mentioned initially, informational asymmetries, contracting constraints or other forms of credit market frictions. When these forms of market imperfections arise, intermediation performs a screening activity between both sides of the market, for which a price spread is charged. The success of intermediaries depends mostly on appropriating a large volume of transactions, and on establishing a repeated, credible interaction with their customers. This transfers the problems of price-setting and market allocation to the intermediation sector. Many features traditionally attributed to competitive markets, such as market clearing, the use of money and Cash-in-Advance constraints, can thus be explained as being in the interest of intermediaries who organize market exchange to alleviate an imperfection and take arbitrage gains from it.

Beyond these implications for the theory of intermediation, the results developed here also have some implications for existing Walrasian macroeconomic and monetary theory. The intermediation model combines frictionless market transactions a la Walras with an explicit, bilateral structure of exchanges. It thereby provides a channel, by which price-setting and information transmission can plausibly be discussed (although this is beyond the scope of this paper). The model further provides an evolutionary approach towards the development and structure of competitive markets. Extensions and simplifications of the intermediation model may thus prove useful to analyze questions in monetary and macroeconomic theory for which the existing theory has come to its limits due to the ad hoc structure of markets and monetary exchange.

1.8 Bibliography

[1] S. R. Aiyagari and N. Wallace, "Existence of Steady-States with Positive Consumption in the Kiyotaki-Wright Model", *Rev. Econ. Stud.* **58** (1991), 901-916.

[2] S. R. Aiyagari and N. Wallace, "Fiat Money in the Kiyotaki-Wright Model", *Econ. Theory* **2** (1992), 447-464.

[3] A. Banerjee and E. Maskin "A Walrasian Theory of Money and Barter", *Quart. J. Econ.* **111** (1996), 955-1005.

[4] W. Baumol, "The Transactions Demand for Cash", *Quart. J. Econ.* **67** (1952), 545-556.

[5] R. Cavalcanti, A. Erosa, and T. Temzelides, "Private Money and Reserve Management in a Random Matching Model", *J. Polit. Economy* **107** (1999), 929-945.

[6] R. W. Clower, "A Reconsideration of the Microfoundations of Monetary Theory", *Western Econ. J.* **6** (1965), 1-9.

[7] D. Corbae, T. Temzelides, and R. Wright, "Matching and Money", University of Pennsylvania, 2000.

[8] A. Dixit, "On Modes of Economic Governance", Presidential address given at the Econometric Society European Meeting, Lausanne, August 2001.

[9] T. Gehrig, "Intermediation in Search Markets", *J. Econ. Manage. Strategy* **2** (1993), 97-120.

[10] F. Hahn, "On Some Problems of Proving the Existence of an Equilibrium in a Monetary Economy", *in* Hahn and Brechling, eds., "Theory of Interest Rates", London: MacMillan, 1965.

[11] C. Hellwig, "Money, Intermediaries, and Cash-in-Advance constraints, *FMG discussion paper* **349**, London School of Economics, 2000.

[12] C. Hellwig and G. Lorenzoni, "Bubbles and Private Liquidity", Princeton University, 2002.

[13] M. Hellwig, "The Challenge of Monetary Theory", *Europ. Econ. Rev.* **37** (1993), 215-242.

[14] P. Howitt, "Beyond Search: Fiat Money in Organized Exchange", Brown University, 2000.

[15] P. Howitt and R. W. Clower, "The Emergence of Economic Organization", *J. Econ. Behav. Organ.* **41** (2000), 55-84.

[16] K. Iwai, "The Evolution of Money: A Search-Theoretic Foundation of Monetary Economics", *CARESS working paper* **88-03**, University of Pennsylvania, 1988.

[17] W. Jevons, "Money and the Mechanism of Exchange", London: Appleton, 1875.

[18] N. Kiyotaki and R. Wright, "On Money as a Medium of Exchange", *J. Polit. Economy* **97** (1989), 927-954.

[19] N. Kocherlakota, "Money is Memory", *J. Econ. Theory* **81** (1998), 231-251.

[20] R. Lucas, Jr., and N. Stokey, "Money and Interest in a Cash-in-Advance Economy", *Econometrica* **53** (1987), 491-514.

[21] A. Matsui and T. Shimizu, "A Theory of Money with Market Places", University of Tokyo, 2001.

[22] K. Menger, "On the Origin of Money", *Econ. J.* **2** (1892), 239-255.

[23] M. Neldner, "Lessons from the Free Banking Era in Switzerland: The Law of Adverse Clearings and the Role of the Non-issuing Banks", *Europ. Rev. Econ. Hist.* **2** (1998), 289-308.

[24] R. Radford, "On the Economic Organisation of a P.O.W. Camp", *Economica* **12**, 189-201.

[25] D. Romer, "A Simple General Equilibrium Version of the Baumol-Tobin Model", *Quart. J. Econ.* **101** (1986), 663-686.

[26] A. Rubinstein and A. Wolinsky, "Middlemen", *Quart. J. Econ.* **102** (1987), 581-593.

[27] P. Samuelson, "An Exact Consumption Loan Model of Interest with or without the Social Contrivance of Money", *J. Polit. Economy* **66** (1958), 1002-1011.

[28] M. Sidrauski, "Rational Choice and Patterns of Growth in a Monetary Economy", *Amer. Econ. Rev.* **57** (1967), 534-544.

[29] B. Smith and W. Weber, "Private Money Creation and the Suffolk Banking System", *J. Money, Credit, Banking* **31** (1999), 624-659.

[30] V. Smith, "The Rationale of Central Banking and the Free Banking Alternative", Ph.D. Dissertation, London School of Economics, 1936, reprinted by Liberty Press, 1990.

[31] R. Starr, "Why is there Money? Convergence to a Monetary Equilibrium in a General Equilibrium Model with Transaction Costs", *discussion paper* **99-23**, Dept. of Economics, U.C. San Diego, 1999.

[32] L. Svensson, "Money and Asset Prices in a Cash-in-Advance Economy", *J. Polit. Economy* **93** (1985), 919-944.

[33] J. Tobin, "The Interest Elasticity of the Transactions Demand for Cash", *Rev. Econ. Statist.* **38** (1956), 241-247.

Part II

Incomplete Information and Coordination Failures

Chapter 2

Public Information, Private Information, and the Multiplicity of Equilibria in Coordination Games

Summary of Chapter 2 *I study coordination games with incomplete public and private information and relate equilibrium convergence to convergence of higher-order beliefs. As the players' signals become more and more precise, the equilibrium manifold converges to the correspondence of common knowledge equilibria, whenever the variance of the public signal converges to 0 at a rate faster than one half the rate of convergence of the variance of private signals. The same condition also determines the convergence of common p-belief to common knowledge, which leads to a simple intuition for its origin, and an immediate generalization of the former results about equilibrium convergence.*[1]

JEL classification numbers: C72, D82

Keywords: Global Games, Equilibrium Convergence, Common Knowledge, Higher-Order Beliefs

[1]This paper has been accepted for publication in the Journal of Economic Theory, and copyright has been transferred to the Journal's publisher. I wish to thank George-Marios Angeletos, Tilman Boergers, Elena Carletti, Vicente Cunat, Martin Hellwig, Thomas Philippon, Hyun Shin and, especially, Stephen Morris for helpful discussions and comments. The remarks of one anonymous referee substantially improved the discussion in this paper. All remaining errors are my responsibility.

2.1 Introduction

Over the past couple of decades, models with multiple equilibria have been prominent in thinking about coordination failures. Arguably the best known example are the bank runs in Diamond and Dybvig [4]. In the related context of speculative attacks against a currency peg, multiple equilibria are at the heart of so-called second generation models, following Obstfeld [19]. In the simplest version of the multiple equilibrium model of speculative attacks, if all investors run on the currency, the central bank will abandon the peg. If investors anticipate this, they will all run. On the other hand, if investors anticipate that the peg will be maintained, they will not run on the currency, and the central bank has no reason to devalue. Typically, the existence of multiple equilibria depends on a state parameter, such as the amount of reserves that a Central Bank is willing to commit in case of an attack. Equilibrium multiplicity makes it impossible to draw determinate economic predictions or policy implications from dynamic analysis or comparative statics, and qualitative conclusions are often restricted to informal predictions as to how policies may alter expectations, or act as a coordination device towards one equilibrium.

In recent years, the multiplicity of equilibria in coordination games has been criticized as the artefact of an implicit assumption that the state is common knowledge. Carlsson and van Damme [2] argue that the coordination game without uncertainty should be viewed as the limit of a sequence of incomplete information games. They posit a so-called "global game" structure, in which players receive noisy, private signals about the game's payoffs, and show that this incomplete information game has a unique equilibrium. Moreover, equilibrium actions at each state converge to one of the equilibria of the common knowledge game, and this limit is independent of the distribution of private signals. Therefore, there exists a unique equilibrium in the coordination game with common knowledge that is also consistent with Carlsson and van Damme's limit continuity postulate with respect to noisy private signals. Uniqueness follows from the higher-order uncertainty inherent in the information structure: Whereas the private signal may provide very precise information about the fundamental, it provides less precise information about what the other player has observed, and it provides no information at all about one player's private signal relative to the other players' private signals. Even if uncertainty about the fundamental is small, at the margin, players remain highly uncertain about each other's actions. This precludes the possibility of multiple equilibria.

This paper reconsiders the equilibrium selection procedure proposed by Carlsson and van Damme, and further explores the link between higher-order uncertainty and equilibrium selection in the context of global games with public and private information. Consider the case where the no-uncertainty limit is approached by a sequence of incomplete information structures, in which players have access only to a more and more precise public signal. The structure of the game is then essentially identical to the common knowledge game: Players face uncertainty only about payoffs, but not about the information of other players. The game has multiple equilibria, which converge exactly to the ones of the game under common knowledge, as the public signal gets more and more precise. The convergence of equilibria in the game with incomplete public information thus differs substantially from the private information game. In general, the no-uncertainty limit of a sequence of incomplete information games will depend on the limiting path. This raises the first question underlying the analysis in this paper: Under what conditions on the information structure will the coordination game have multiple equilibria, and when will the equilibrium be unique? The arguments for uniqueness of the equilibrium in the private information game, as well as for multiplicity under public information relied on higher-order uncertainty, i.e. what players know about each others' information. This leads to a second question: When does a sequence of information structures converge to common knowledge? This paper studies these questions separately, first, by explicitly considering equilibrium convergence in the context of a well-understood example, then by deriving the level of common belief (cf. Monderer and Samet [12]) in the information structure. The results on common belief convergence then lead to an immediate generalization of the equilibrium convergence result beyond the initial example, without making explicit payoff assumptions.[2]

It turns out that the ratio of the *standard deviation of private signals* to the *variance of public signals* determines the answer to both questions.[3] Whenever this ratio becomes infinite

[2]Throughout this paper, I study information structures with normally distributed signals. Extensions to other distributions are possible, but beyond the scope of this paper.

[3]This condition comes as a consequence of the statistical properties of the information structure. The discussion leading up to Theorem 3 below provides some intuition for its origin, as well as its necessity and

in the limit, i.e. whenever the variance of the *public* signal converges to zero at a rate faster than $\frac{1}{2}$ the rate of convergence of the variance of *private* signals, the incomplete information game converges to the common knowledge game, and higher-order uncertainty vanishes. Whenever the above ratio converges to zero, the incomplete information game converges to the private information game studied by Carlsson and van Damme, and higher-order uncertainty remains at a maximal level. Whenever the ratio converges to a finite constant, this limit determines the number of equilibria, as well as equilibrium strategies. Once again, these results are best understood by considering the effect of public information on higher-order uncertainty: In contrast to the private information environment, the public signal enables a player to make some inference about other players' beliefs relative to his own, and therefore about the other players' actions. If this public information is sufficiently precise, relative to private signals, its coordinating effect becomes so large that it induces multiple equilibria. If the existing information is very precise, only a small amount of public information is needed to create a large coordinating effect, and thereby multiple equilibria.

2.2 Related Literature

To understand the implications of these results, it will be useful to place them within the context of the theoretical and applied global games literature. A series of authors have used Carlsson and van Damme's private information equilibrium to perform a comparative or dynamic analysis in a coordination game of economic interest. Among others, Morris and Shin [15] apply it to speculative attacks in a game with a continuum of players. Goldstein and Pauzner [5] analyze how the immunity of a bank against runs depends on interest rates and debt contracts. Rochet and Vives [20] discuss the effects of a lender of last resort policy. Morris and Shin [17] provide an excellent survey that does better justice to this rapidly growing literature than this short list. The results in this paper cast doubt on the validity of the use of the private information limit as an equilibrium selection in a coordination game with multiple Nash equilibria under common knowledge. Whenever the public information retains some informative value relative to private

sufficiency for common belief convergence.

information, the incomplete information game converges to the common knowledge game, as information becomes more and more precise. Even if the equilibrium is unique, its limit is still dependent on the convergence of the afore mentioned variance to standard deviation ratio, and the public signal has considerable effect on equilibrium actions. The private information equilibrium is thus in stark contrast with the public information limit, and any comparative statics conclusions that are derived from the private information equilibrium do not appear to be robust to the equilibrium effects of public information.

It should be noted that this paper is not the first to study the effects of public information in a global games context, however, it is the first to explicitly consider equilibrium convergence in a public information setting. Its predecessors are mostly interested in comparative statics with respect to the information structure in environments with a unique equilibrium away from the limit. The empirical relevance of public disclosures in the context of most coordination games of economic interest provides a natural motivation for such a study, both to understand its equilibrium effects (Morris and Shin [16]), and to draw conclusions about the design of disclosure policies, eg. Morris and Shin [18], Metz [10]. In an analysis of coordination effects in the pricing of debt, Morris and Shin [16] are the first to point out that with normally distributed signals, it is the ratio of the standard deviation of private information to the variance of public information that matters for determining the number of equilibria and equilibrium strategies. In their survey paper, they extend this result to a more general class of symmetric infinite-player, 2-action games. Our result that the variance to standard deviation ratio also determines whether higher-order uncertainty vanishes leads to an interpretation and generalization of their results that does not rely on payoff assumptions.

This leads us to relate the results in this paper to the literature on higher-order uncertainty and robustness of equilibria to incomplete information. This literature starts outs from a number of examples, where small payoff uncertainty has large strategic effects because of a lack of common knowledge (eg. Rubinstein [21], Carlsson and van Damme [2]). This apparent disparity between the equilibria of the common knowledge game and its incomplete information perturbation raises two questions: First, what kind of common knowledge equilibria are robust to incomplete information perturbations, and second, when are two information structures similar, in the sense that they lead to similar outcomes in the same game? In particular, when does an incomplete information structure with small pay-off uncertainty lead to equilibria similar to common knowledge? The answers to both of these underlying questions are typically phrased in terms of the players' belief structures. In response to the first question, Morris, Rob and Shin [14] define the notion of belief potential[4] for 2-player information structures, and show how it can be used to iteratively eliminate actions over a part of or the entire state space. Kajii and Morris [8] study the robustness of equilibria to arbitrary incomplete information structures, and also use belief potential to give a characterization of equilibrium robustness in terms of p-dominance.[5] These papers both set out to derive uniqueness conditions for games with a large amount of higher-order uncertainty. While they contribute to our understanding of how higher-order uncertainty affects equilibrium selection, uniqueness and robustness conditions are usually given in terms of the belief structure of players, and thereby fail to fully make the connection between the signal structure and the role of higher-order uncertainty in equilibrium selection.

Likewise, definitions of strategic proximity of information structures (Kajii and Morris [9]) or convergence to common knowledge are typically spelled out as conditions on the structure of higher-order beliefs, rather than the signal structure. As a useful way of formalizing higher-order uncertainty, Monderer and Samet [12] introduce the notion of belief operators and common p-belief in analogy with knowledge operators and common knowledge. An event E is common p-belief among a set of players S at a state ω, if at ω, all players in S believe with probability at least p (henceforth: p-believe) that E has occurred, all players in S p-believe that all others p-believe that E has occurred... and so on. For any information structure, we can thus define a common belief function $p(E, F, S)$ by assigning to each couple of events E, F and set of players S the highest $p \in [0, 1]$ such that E is common p-belief among players in S at every state $\omega \in F$. In the limit of a sequence of information structures, an event E is then said to be common knowledge in F, if $p(E, F, S) \to 1$. Despite its seemingly complicated definition

[4] The belief potential of an information structure is the highest probability p, such that, for any information set of either player, some statement of the form "player i believes with probability at least p that player $-i$ believes with probability at least p that player i believes with probability at least p ... that the true state is in the original information set" is true at every state.

[5] A strategy profile a is a p-dominant equilibrium at state ω, if, for every player i, it is optimal to play a_i, whenever a_{-i} is played with probability at least p.

on the entire hierarchy of beliefs, common p-belief has a number of qualities that make it appealing for the use in the context of global games. A simple fixed point characterization dispenses us from taking into account the entire hierarchy in determining common beliefs - we will make much use of this property to determine common beliefs in the signal structures of global games.[6] Common p-belief also preserves lower hemi-continuity with respect to the set of equilibria. To be precise, if an action profile a is a p-dominant equilibrium over a set of states E, and $p(E, F, S) \geq p$, then there exists an equilibrium where a is played (at least) in F. It follows that, if $p(E, F, S) \rightarrow 1$, i.e. E becomes limit common knowledge at F, every strict equilibrium in E can be locally supported at F. The common belief function thereby provides a summary statistic of higher-order uncertainty that preserves limit continuity with respect to the equilibrium manifold.[7]

Focusing on players' beliefs rather than signals leaves open the question how a signal structure translates into a belief structure. Our discussion of convergence of common belief resolves this issue for information structures with public and private signals, such as they appear in most global game applications, and thereby links equilibrium selection issues to the level of higher-order uncertainty. We find that, as long as the variance of public information converges to 0 at a rate faster than $\frac{1}{2}$ the rate of convergence of the variance of private information, the common belief converges to 1 in probability. Furthermore, the logic underlying the proof of our common belief convergence result provides a simple intuition for the origin of this condition on convergence of the signal variances. *Common belief convergence thus provides a reinterpretation of the uniqueness condition that is more general and (arguably) easier to interpret in the context of the models to which it is applied than the condition on the variances of public and private information.* If, as Morris and Shin argue, common information effects contribute to explain puzzles in debt pricing, price data may actually lead to quantitative estimates of common p-belief among market participants, and it may be possible to assess the applicability of the global games equilibrium selection by estimating the level of common belief that markets

[6]In their analysis, Monderer and Samet focus on a countable state space. Kajii and Morris [7] extend their results to uncountable state spaces that are relevant here.

[7]Upper hemi-continuity is already implied by the disappearance of payoff uncertainty. Monderer and Samet's original result stated that every equilibrium is the limit of a sequence of ε-equilibria. Rephrasing this in terms of p-dominance or strict equilibria is immediate.

are capable of generating.

That public announcements are important in creating (approximate) common knowledge has been recognized in a variety of examples. Furthermore, following Monderer and Samet's equilibrium approximation as common belief approximates common knowledge, it will come as no surprise that the same condition on the information structure should determine equilibrium selection and common belief convergence. However, beyond merely confirming this intuition, our independent proof of common belief convergence, coupled with lower hemi-continuity of the equilibrium manifold, provides an immediate extension of the previous conclusions to a much larger class of games, without having to rely on payoff assumptions. It thereby provides a new methodology for characterizing equilibria in incomplete information games. Separating the effects of assumptions about the information structure from the effects of payoff assumptions also highlights the effects of each on equilibrium selection. In particular, we can relate the derivation of common p-beliefs in an information structure with only private signals to equilibrium selection in the private information game: In this case, common p-belief is independent of the overall level of noise, and bounded away from 1. In particular, we can relate our derivation of common belief in an information structure with private signals only to equilibrium selection in a private signal information structure to equilibrium selection in private information games: In this case, the level of common belief is independent of the overall level of noise, and bounded away from 1. The uniqueness result of Carlsson and van Damme is therefore not at odds with Monderer and Samet's approximation of strict equilibria under incomplete information with common belief close to 1. However, the view that the sequence of private information games approximates the game under common knowledge is inaccurate.

The remainder of this paper is divided into four parts, followed by a conclusion, which discusses some of the practical implications for global game applications. In section 3, we introduce the general set-up for incomplete information games. In section 4, we study as an example the game considered in Morris and Shin [16]. However, while their analysis focuses on comparative statics in an environment with a unique equilibrium, we explicitly discuss the possibility of multiple equilibria and convergence to common knowledge in anticipation of the more general results that follow. Section 5 defines the common belief function and examines its convergence properties. Section 6 introduces the notion of p-dominance, which is then used to generalize the equilibrium convergence result encountered in the initial example.

2.3 The General Set-up

Consider a normal form game $G = (S, A, \pi)$ where S denotes the (possibly infinite) set of all players, $A = \times_{i \in S} A_i$, and $A_i = \{1, ..., N_i\}$ denotes the finite set of actions for each player $i \in S$, and $\pi_i(\theta, a_i, a_{-i})$ denotes payoff to player i of playing action a_i when all other players follow action profile a_{-i}, for all (a_i, a_{-i}). Payoffs depend on a fundamental variable $\theta \in \mathbb{R}$.

In a finite player game, we let the number of players be I. In the infinite-player version of this game, we impose a little bit more structure: Consider a game with I different classes of players, each represented by a continuum of players (normalized to have a measure 1 for each class, so that the total population size is I). Players in each class are identical. Each player in class i has access to actions $A_i = \{1, ..., N_i\}$,and payoffs for players in class i are $\pi_i(\theta, a_i, \mu)$, where $\mu = (\mu_1, ..., \mu_I)$ and $\mu_i = \{\mu_i(1), ..., \mu_i(N_i)\}$ represents the distribution of actions across types, i.e. $\mu_i(a)$ indicates the measure of players in class i that play action a.

The timing of this finite or infinite player game is as follows: Initially, nature draws a fundamental $\theta \in \mathbb{R}$. To simplify the exposition, assume that the prior distribution of θ is uniform over the entire real line.[8] Each player then receives noisy signals about θ. After forming their beliefs about θ, players simultaneously decide on their action.

We make the following assumptions concerning the information structure: Generically, each player observes a private signal (observed only by him) and a common signal (observed by all players) about θ. In the limiting case, where the variance of the common signal is infinite, we return to the information structure with only private information. Similarly, as the variance of the private signal becomes infinite, the information structure converges to a situation in which there exists only public information.

[8]Such an "improper prior" is not essential for our results. One could easily adapt the model to include a proper prior distribution of θ.

Explicitly, assume that all private signals are independent and identically normally distributed. Given a state θ, player j observes $x_j = \theta + v_j$, where $v_j \sim N(0, \varepsilon^2)$. In addition, all players observe $y = \theta + v$, where $v \sim N(0, \eta^2)$. To complete the notation, we let $\Phi(\cdot)$ and $\varphi(\cdot)$ denote the cdf. and pdf. of a standard normal distribution, respectively. Finally, $\alpha^2 = \frac{\eta^2}{\varepsilon^2+\eta^2}$ denotes the relative importance of the private information, and $\alpha\varepsilon$ the ex post uncertainty about the state.

Throughout this paper, information structures are parametrized by α^2 and $\alpha\varepsilon$. Define $x_j^1 = \alpha^2 x_j + (1-\alpha^2)\, y$ as the posterior expectation of player j about θ. Standard statistical results imply that ex post, player j believes that θ, conditional on y and x_j is normally distributed with mean x_j^1 and variance $\alpha^2\varepsilon^2$.

$$\theta \mid x_j, y \sim N\left(x_j^1, \alpha^2\varepsilon^2\right). \tag{2.1}$$

This information structure also appears in Morris and Shin [16], where a common prior takes the place of the public signal. The various cases mentioned above can be interpreted as special cases of this formulation. The private information equilibrium corresponds to the case, where $\alpha^2 = 1$. In that case, $x_j^1 = x_j$, and agents draw no information from the public signal. If $\alpha^2 = 0$, then $x_j^1 = y$ and all players have identical posterior beliefs. For $\alpha \in (0,1)$, players draw some information from both signals.

We should note that the framework with an infinite number of players has some analytical advantages. Assuming that the law of large numbers holds,[9] θ fully characterizes the distribution of signals over the population, and the proportion of agents whose signal falls below a certain threshold is given by the value of a normal cdf. For a given realization of θ and y, posterior expectations are normally distributed over the population and given by

$$x_j^1 \mid \theta \sim N\left(\alpha^2\theta + \left(1-\alpha^2\right) y, \alpha^4\varepsilon^2\right). \tag{2.2}$$

[9] See Judd [6] for a discussion of the role of the law of large numbers for a continuum of random variables. One might be concerned whether the order of limits that are taken (with respect to the number of players and with respect to the information structure) affects the results in the infinite player game. Throughout the paper, I will show that the same results apply for games with a finite or infinite number of players.

We impose the following regularity conditions on expected payoffs in the finite or infinite player incomplete information game:

C 1 *For an arbitrary player i, posterior expectation x_i, and strategy profile (a_i, a_{-i}), as $\alpha\varepsilon \to 0$, $E_\theta(\pi_i(\theta, a_i, a_{-i}) \mid x_i) \to \pi_i(x_i, a_i, a_{-i})$.*

$(C1)$ requires that for all players, expected payoffs converge to realized payoffs whenever the level of noise in the information structure vanishes. Roughly speaking, it rules out the possibility that unboundedly large, but highly unlikely negative payoffs influence a player's choice of action. A sufficient (but not necessary) condition for $(C1)$ is that payoffs are bounded over the state space. It is easy to check that $(C1)$ is necessary and sufficient to guarantee that, if (i) the state is θ^*, and (ii) a is played with probability close to 1, the expected payoffs will be close to $\pi_i(\theta^*, a)$. We impose the following additional restriction:

C 2 *There exist $\underline{\theta}$ and $\overline{\theta}$, and action profiles $\underline{a}$ and $\overline{a}$, such that for all i, under complete information, $\underline{a}_i$ is a strictly dominant action for player i, whenever $\theta < \underline{\theta}$, and $\overline{a}_i$ is a strictly dominant action for player i, whenever $\theta > \overline{\theta}$.*

Under $(C1)$ and $(C2)$, theorem 1 in Milgrom and Weber [11] establishes the existence of a perfect bayesian equilibrium in the incomplete information game, provided that $\alpha\varepsilon$ is sufficiently small. The following section studies an example of this general incomplete information game, and discusses equilibrium convergence directly from the convergence of conditions on equilibrium payoffs. The example thereby anticipates the general convergence result that we set out to prove in the remainder of the paper.

2.4 A Simple Example

Consider a game played by a $[0, 1]$-continuum of risk-neutral players, who all choose between two actions a and b. For any player, the payoff to playing a is always $r > 0$. The payoff to playing b depends on the proportion of players playing b, as well as a parameter θ ("economic fundamental"): it is $f > r$, if the proportion of players who choose to play b exceeds $c(\theta)$, and 0, if less than $c(\theta)$ players play b. If exactly $c(\theta)$ players play b, they receive a payoff of r.[10] $c(\theta)$ is strictly increasing and continuously boundedly differentiable, and there exist $\underline{\theta}$ and $\overline{\theta}$, such that $c(\underline{\theta}) = 0$ and $c(\overline{\theta}) = 1$. We assume a timing of this game and an information structure as introduced above for the general case, with nature drawing a fundamental $\theta \in \mathbb{R}$, and players observing normally distributed private signals and a normally distributed public signal. Players then simultaneously decide whether to play a or b.

This framework replicates the reduced-form model of speculative attacks against a currency,[11] and appears, with slight variations, in most global games applications. We observe that this game has multiple equilibria, if θ is common knowledge and falls inside a "critical region" $[\underline{\theta}, \overline{\theta}]$. In one equilibrium, all players play a, in a second equilibrium, all players play b, and finally, there exists a mixed strategy equilibrium, in which players randomize with probability $c(\theta)$. If $\theta > \overline{\theta}$, all players play a in the unique equilibrium, while if $\theta < \underline{\theta}$, all agents play b in the unique equilibrium.

Equilibria of this game under incomplete information, as well as their convergence properties for different paths of convergence for $\alpha\varepsilon$ and α^2 are easily derived by directly considering the payoff structure. For simplicity, we restrict attention to symmetric "threshold" equilibria, i.e. pairs of thresholds (x^*, θ^*) such that a player plays b, whenever his posterior expectation falls below x^*, and receives a payoff of f, whenever $\theta < \theta^*$. In other words, in a symmetric threshold equilibrium, whenever $\theta < \theta^*$, at least $c(\theta^*)$ players play b.

For the moment, take y as given. If the threshold state, below which the payoff to b is f, is θ^*, a player will choose to play b, as long as the expected payoff, given his posterior belief, satisfies $r \leq f \cdot \Pr\left(\theta < \theta^* \mid x_j^1\right)$. One observes that the right-hand side is strictly decreasing in

[10] Under uncertainty, the payoff for exactly $c(\theta)$ players playing b will not affect strategies and payoffs, but under common knowledge, this assumption preserves the third, mixed-strategy equilibrium and guarantees that the equilibrium manifold of this game is connected.

[11] For instance, in the speculative attacks interpretation of this model, f corresponds to the discrete gain to speculators after a successful devaluation, r to the interest rate premium on a domestic bond, and $c(\theta)$ to the level of reserves that a central bank would be willing to commit to defend the currency peg. In the speculative attacks paper of Morris and Shin [15], the devaluation premium $f(\theta)$ is a function of the fundamental. Assuming $f(\theta)$ constant simplifies the exposition without changing the results.

x_j^1. It follows that there exists a unique threshold expectation x^*, given by the *payoff indifference condition* (PI)

$$r = f \cdot \Pr(\theta < \theta^* \mid x^*), \tag{2.3}$$

at which a player is indifferent between a and b. By monotonicity, all agents with expectations below x^* strictly prefer b, and all agents with expectations higher than x^* strictly prefer a. (2.3) can be rewritten as

$$r = f \cdot \Phi\left(\frac{\theta^* - x^*}{\alpha\varepsilon}\right). \tag{2.4}$$

Rewriting (2.4) defines $x_{PI}(\theta^*)$:

$$x_{PI}(\theta^*) = \theta^* - \alpha\varepsilon\Phi^{-1}\left(\frac{r}{f}\right). \tag{2.5}$$

Next, suppose all agents follow a threshold strategy x^*. From (2.2), the fraction of players playing b given θ and x^* is $\Phi\left(\frac{x^*-\alpha^2\theta-(1-\alpha^2)y}{\alpha^2\varepsilon}\right)$. The return to b is f, iff $c(\theta) < \Phi\left(\frac{x^*-\alpha^2\theta-(1-\alpha^2)y}{\alpha^2\varepsilon}\right)$. From the monotonicity in θ and x^*, it follows that the return to b is f, if $\theta < \theta^*$, where θ^* is given by the *critical mass condition* (CM)

$$c(\theta^*) = \Phi\left(\frac{x^* - \alpha^2\theta^* - (1-\alpha^2)\,y}{\alpha^2\varepsilon}\right). \tag{2.6}$$

This can be rewritten as

$$x_{CM}(\theta^*, y) = (1-\alpha^2)\,y + \alpha^2\theta^* + \alpha^2\varepsilon\Phi^{-1}(c(\theta^*)). \tag{2.7}$$

Since $\Phi^{-1}(\cdot)$ is defined on $(0,1)$, x_{CM} is well-defined and strictly increasing for values of $\theta^* \in (\underline{\theta}, \overline{\theta})$. Differentiating (CM) yields

$$\frac{\partial x_{CM}(\theta^*, y)}{\partial\theta^*} = \alpha^2\left(1 + \frac{\varepsilon c'(\theta^*)}{\varphi(\Phi^{-1}(c(\theta^*)))}\right) \geq \alpha^2. \tag{2.8}$$

For given y, actions in any threshold equilibrium are implicitly defined by the intersection of the PI-condition with the CM-condition. Since $x_{CM}(\theta^*, y)$ converges to $-\infty$ or $+\infty$ as θ

approaches $\underline{\theta}$ and $\overline{\theta}$, respectively, and $x_{PI}(\theta^*)$ is linear, continuity implies that $x_{PI}(\theta^*) = x_{CM}(\theta^*, y)$ for some $\theta^* \in (\underline{\theta}, \overline{\theta})$. Therefore, there always exists a symmetric threshold equilibrium. By examining (PI) and (CM) more in detail, we can give conditions for uniqueness or multiplicity. We have the following lemma (Morris and Shin [16]):

Lemma 4 *There exist multiple equilibria, if and only if*

$$\frac{1-\alpha^2}{\alpha^2\varepsilon} \geq \min_{\theta \in (\underline{\theta},\overline{\theta})} \frac{c'(\theta)}{\varphi\left(\Phi^{-1}\left(c\left(\theta\right)\right)\right)}. \tag{2.9}$$

Proof. There exist multiple equilibria, if and only if for some $\theta^* \in (\underline{\theta}, \overline{\theta})$, $x_{PI}(\theta^*) = x_{CM}(\theta^*, y)$ and $\frac{dx_{PI}(\theta^*)}{d\theta^*} \geq \frac{\partial x_{CM}(\theta^*,y)}{\partial\theta^*}$. Moreover, if $\frac{dx_{PI}(\theta^*)}{d\theta^*} \geq \frac{\partial x_{CM}(\theta^*,y)}{\partial\theta^*}$ for some $\theta^* \in (\underline{\theta}, \overline{\theta})$, then there exists y such that $x_{PI}(\theta^*) = x_{CM}(\theta^*, y)$. Thus $\frac{dx_{PI}(\theta^*)}{d\theta^*} \geq \frac{\partial x_{CM}(\theta^*,y)}{\partial\theta^*}$ for some θ^* is necessary and sufficient for the existence of multiple equilibria. Using (2.8) and the fact that (2.5) implies $\frac{dx_{PI}(\theta^*)}{d\theta^*} = 1$, one rewrites this condition as (2.9). ∎

We determine equilibrium convergence as $\alpha\varepsilon \to 0$ by considering the joint convergence of the (PI) and (CM) curves. Substituting (PI) into (CM) and rearranging, one obtains

$$\theta^* - \frac{\alpha^2\varepsilon}{1-\alpha^2}\Phi^{-1}\left(c\left(\theta^*\right)\right) = y + \frac{\alpha^2\varepsilon}{1-\alpha^2}\frac{1}{\alpha}\Phi^{-1}\left(\frac{r}{f}\right), \tag{2.10}$$

and

$$\frac{d\theta^*}{dy} = -\frac{1-\alpha^2}{\alpha^2\left(1+\frac{\varepsilon c'(\theta^*)}{\varphi(\Phi^{-1}(c(\theta^*)))}\right)-1}. \tag{2.11}$$

Observe that the r.h.s. of (2.10) approaches $+\infty$ and $-\infty$, as θ^* approaches $\underline{\theta}$ and $\overline{\theta}$, respectively. For intermediate values of θ^*, the r.h.s of (2.10) is increasing, if $\frac{1-\alpha^2}{\alpha^2\varepsilon}$ is sufficiently large. The following three lemmas discuss the limit behavior of equilibria of this coordination game.

Lemma 5 *Suppose $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to 0$. Then the unique limit equilibrium θ_{pr} is the private information equilibrium, implicitly defined by $f\left(1-c\left(\theta_{pr}\right)\right) = r$.*

Proof. $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to 0$ implies $\alpha \to 1$, and it follows from (2.10) that $\Phi^{-1}\left(c\left(\theta_{pr}\right)\right) \to \Phi^{-1}\left(\frac{r}{f}\right)$, or $1 - c\left(\theta_{pr}\right) \to \frac{r}{f}$. ■

Lemma 5 revisits the uniqueness result in an environment with private information only. Morris and Shin [15] proceed to show by interim elimination of strictly dominated strategies that this threshold equilibrium is the unique equilibrium of the game.

Lemma 6 *Suppose $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to \infty$, as $\alpha\varepsilon \to 0$. Then, there are three equilibria in the limit, if $y \in \left(\underline{\theta}, \overline{\theta}\right)$. The set of limit equilibrium thresholds is $\left\{\underline{\theta}, \overline{\theta}, y\right\}$.*

Proof. Fix $k = \frac{1-\alpha^2}{\alpha^2\varepsilon}$. For sufficiently large k, $k = \frac{c'(\theta)}{\varphi(\Phi^{-1}(c(\theta)))}$ has two solutions, which we denote by $\overline{\theta}\left(k\right)$ and $\underline{\theta}\left(k\right)$, i.e. the local maximum and minimum of the r.h.s of (2.10). Clearly there exist values of y, such that some $\theta_1^* \in \left(\overline{\theta}\left(k\right), \overline{\theta}\right)$ and $\theta_2^* \in \left(\underline{\theta}, \underline{\theta}\left(k\right)\right)$ are simultaneously supported as equilibria. Since $\overline{\theta}\left(k\right) \to \overline{\theta}$ and $\underline{\theta}\left(k\right) \to \underline{\theta}$, as $k \to \infty$, it follows that there are limit equilibria at $\underline{\theta}$ and $\overline{\theta}$. Since at both of these, $k < \frac{c'(\theta)}{\varphi(\Phi^{-1}(c(\theta)))}$, there also exists an intermediate equilibrium, for which $k > \frac{c'(\theta)}{\varphi(\Phi^{-1}(c(\theta)))}$. As $k \to \infty$, (2.10) implies that this equilibrium converges to $\theta^* = y$. Now, let $\overline{y}\left(k\right)$ and $\underline{y}\left(k\right)$ denote the upper and lower bounds of the critical region of y. We have:

$$
\begin{aligned}
\overline{y}\left(k\right) &= \overline{\theta}\left(k\right) - \frac{1}{k}\Phi^{-1}\left(c\left(\overline{\theta}\left(k\right)\right)\right) - \frac{1}{k}\Phi^{-1}\left(\frac{r}{f}\right), \\
\underline{y}\left(k\right) &= \underline{\theta}\left(k\right) - \frac{1}{k}\Phi^{-1}\left(c\left(\underline{\theta}\left(k\right)\right)\right) - \frac{1}{k}\Phi^{-1}\left(\frac{r}{f}\right).
\end{aligned}
$$

As $k \to \infty$, $\frac{1}{k}\Phi^{-1}\left(c\left(\overline{\theta}\left(k\right)\right)\right) \to 0$ and $\frac{1}{k}\Phi^{-1}\left(\frac{r}{f}\right) \to 0$, and hence $\overline{y}\left(k\right) \to \overline{\theta}$. Similarly, $\underline{y}\left(k\right) \to \underline{\theta}$, since $\frac{1}{k}\Phi^{-1}\left(c\left(\underline{\theta}\left(k\right)\right)\right) \to 0$. ■

Among other cases, lemma 6 considers the case in which there exists some valuable public information in the limit ($\alpha < 1$), and shows that under this structure, there exist multiple equilibria, if the common signal falls inside a critical region. These three equilibria converge in probability to the equilibria of the common knowledge game: Whenever $\theta \in \left(\underline{\theta}, \overline{\theta}\right)$, then, whenever y falls inside the critical region, there exist three equilibria. As $\alpha\varepsilon \to 0$, in one of these, the probability of devaluation goes to 0, in a second the probability of devaluation goes to 1. The third equilibrium converges to the mixed strategy equilibrium. As $\alpha\varepsilon$ converges to 0, θ^* converges to y in such a way that the probability of devaluation converges to $c(\theta^*)$.

Using the same approach as above, lemma 4 discusses the case where $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to k < \infty$. If $k = \frac{c'(\theta)}{\varphi(\Phi^{-1}(c(\theta)))}$ has more than two solutions, there may be more than three equilibria. We rule this out by assuming that $\frac{c'(\theta)}{\varphi(\Phi^{-1}(c(\theta)))}$ has a unique inflection point. Let $k^* = \min\limits_{\theta\in[\underline{\theta},\overline{\theta}]} \frac{c'(\theta)}{\varphi(\Phi^{-1}(c(\theta)))}$. Then:

Lemma 7 *(i) Suppose* $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to k < k^*$, *as* $\alpha \to 1$ *and* $\alpha\varepsilon \to 0$. *Then, there is a unique limit equilibrium, whose threshold varies continuously with* y *over* $(\underline{\theta},\overline{\theta})$.

(ii) Suppose $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to k > k^*$, *as* $\alpha \to 1$ *and* $\alpha\varepsilon \to 0$. *Then, there are multiple limit equilibria, all with thresholds continuous in* y. *One equilibrium threshold* $\theta_1^* \in (\overline{\theta}(k), \overline{\theta})$ *is decreasing in* y *for* $y \in (-\infty, \overline{y}(k))$, *the second equilibrium,* $\theta_2^* \in (\underline{\theta}, \underline{\theta}(k))$ *is also decreasing in* y *for* $y \in (\underline{y}(k), \infty)$. *A third equilibrium* $\theta_3^* \in (\underline{\theta}(k), \overline{\theta}(k))$ *is increasing in* y, *for* $y \in (\underline{y}(k), \overline{y}(k))$. *For sufficiently large* k, $\underline{y}(k) > \underline{\theta}(k)$ *and* $\overline{y}(k) < \overline{\theta}(k)$.

Proof. Except for the last statement, all of lemma 7 follows directly from (2.10), (2.11) and the proof of lemma 3. $\underline{y}(k) > \underline{\theta}(k)$ and $\overline{y}(k) < \overline{\theta}(k)$ in the limit, iff $\Phi^{-1}\left(c\left(\overline{\theta}(k)\right)\right) > -\Phi^{-1}\left(\frac{r}{f}\right) > \Phi^{-1}\left(c\left(\underline{\theta}(k)\right)\right)$, or $\overline{\theta}(k) > \theta_{pr} > \underline{\theta}(k)$, where θ_{pr} denotes the private information equilibrium defined above. This will necessarily be the case for sufficiently large k. ■

If $\frac{1-\alpha^2}{\alpha^2\varepsilon}$ is sufficiently large, the critical region of the private signal is strictly included in the critical region of the fundamental.[12] Since the limit equilibrium depends both on the convergence of $\frac{1-\alpha^2}{\alpha^2\varepsilon}$ and the realization of y, the coordination game with public information no longer satisfies noise independence. It follows from (2.11) that the effect of y on θ^* is potentially much stronger than the relative importance of the public signal. To be precise, the relative importance of the public signal $1 - \alpha^2$ is multiplied by a publicity multiplier, which is strictly larger than 1, if ε is sufficiently small. This publicity multiplier locally goes to infinity, as the game approaches the zone of multiple equilibria, and y approaches the boundaries

[12] This will not generally be true: note that $\overline{\theta}(k) \to \underline{\theta}(k)$, as $k \to k^*$, but their limit need not be equal to the private information equilibrium, so $\overline{\theta}(k) > \theta_{*r} > \underline{\theta}(k)$ will no longer be satisfied.

$\underline{y}(k)$ and $\overline{y}(k)$. *Thus, even if there is a unique equilibrium, public information may have a disproportionately large impact on equilibrium strategies, both in and away from the limit.* Morris and Shin [16] use precisely this effect of public information in order to account for coordination effects in the context of debt pricing. We summarize our results in Theorem 1:

Theorem 1 *(i) If $\frac{1-\alpha^2}{\alpha^2\varepsilon}$ is sufficiently high and y falls in a critical region, this game has three equilibria. As $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to \infty$, the critical region for y converges to $(\underline{\theta}, \overline{\theta})$, and the set of limit equilibrium thresholds is $\{\underline{\theta}, \overline{\theta}, y\}$.*

(ii) If $\frac{1-\alpha^2}{\alpha^2\varepsilon}$ is sufficiently low, there is a unique equilibrium. As $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to 0$, the equilibrium threshold θ^* converges to the private information equilibrium.

(iii) If $\frac{1-\alpha^2}{\alpha^2\varepsilon}$ converges to a finite constant, there may be one or multiple equilibria in the limit. The critical region for the fundamental and for the public signal, for which multiple equilibria can be supported are both increasing in the limit of $\frac{1-\alpha^2}{\alpha^2\varepsilon}$.

Theorem 1 contains the first central result in this paper. (i) states that the multiple equilibria obtained in the benchmark model with common knowledge continue to exist under incomplete information, as long as players have access to a valuable common signal, and the overall level of noise is sufficiently low. If $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to \infty$, equilibrium actions converge to the ones obtained under common knowledge, and the equilibrium manifold converges to the equilibrium manifold under common knowledge.[13] In the limit, multiple equilibria exist, if y falls inside the same critical region as θ under common knowledge. Thus, any sequence of incomplete information games for which $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to \infty$, as $\alpha\varepsilon \to 0$, approximates the common knowledge game. (ii) states the analogous result for the private information equilibrium, which is approximated as the limit of any sequence of information structures, for which $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to 0$, and (iii) discusses equilibrium convergence for finite limits of $\frac{1-\alpha^2}{\alpha^2\varepsilon}$. Contrary to the limits to both the common knowledge game and the private information game, equilibrium actions in all limit equilibria vary with the realization of y in this case.

[13] It is important to note that y cannot be interpreted as a sunspot. If y has no informative value at all, then the Morris and Shin's result continues to hold. Rather, in the case where the existence of sufficiently precise common information leads to multiple equilibria, traditional sunspots might be viewed as one way of selecting among them.

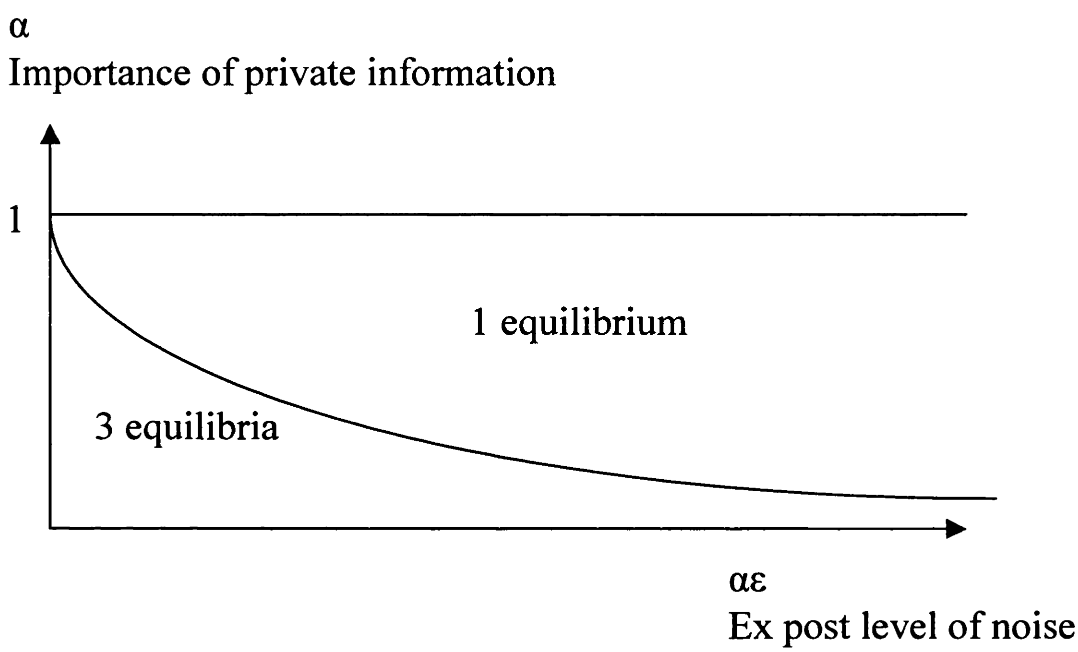

Figure 2-1: Multiplicity vs. Uniqueness

Returning to the initial parametrization, $\frac{1-\alpha^2}{\alpha^2\varepsilon} = \frac{\varepsilon}{\eta^2}$. We find that the *variance* of the public signal must go to 0 no faster than the *standard deviation* of the private signals, in order to maintain uniqueness. This condition first appears in Morris and Shin [16]. In their survey paper, Morris and Shin [17] show that the same condition applies to any infinite-player, 2-action games, provided that payoffs are boundedly differentiable. As long as $\frac{\varepsilon}{\eta^2} \to \infty$, the limit equilibria also converge to the equilibria under common knowledge. Theorem 1 is summarized in figure 1.

Overall, multiple equilibria are more likely to exist in this game, when the ex post level of noise is low, and when the relative importance of the public signal is high. As the level of noise goes to zero, the economy always exhibits multiple equilibria, unless the public signal becomes perfectly uninformative. Both statements can be explained by the role of information in explicit coordination on one equilibrium. The public signal gives a player information not only about the state, but, much more importantly, about what other players have observed relative to one's own private signal (and what they are therefore likely to play). As signals becomes more and more precise, the players ultimately have better and better information about other players' beliefs possibly even if the public signal becomes infinitely less informative than the private signal. In the next section, we return to the general set-up to formalize this connection between the equilibrium structure and higher-order uncertainty for a very large class of finite and infinite player games.

2.5 Higher-Order Beliefs

2.5.1 Definitions

Let us now return to the game outlined initially. The literature on higher-order beliefs typically defines a "state space" Ω sufficiently rich to include all uncertainty within the game, and the information that each player has is defined by a partition on that state space. In a finite player game, we can easily replicate this formulation by defining a state as a realization $(\theta, y, x_1, ..., x_I)$, the state space as $\Omega = \mathbb{R}^{I+2}$, and player i's information partition by realizations of y and x_i.

Following Monderer and Samet [12], at a state ω, we say that an event E is common p-belief among a set of players S, if (i) each player in S p-believes that E has occurred, (ii) each player in S also p-believes that all players in S p-believe that E has occurred, (iii) each player in S also p-believes that all agents in S p-believe that all players in S p-believe that E has occurred... and so on. *We thus define the common belief function $p(E, F, S)$ by assigning to events E and F and a set of players S the highest probability p, such that E is common p-belief among players in S, for all $\omega \in F$.* This definition provides a continuous transition from incomplete information to common knowledge, where the latter corresponds to the case where $p(E, F, S) \to 1$, for all E, F and S.

In most environments, characterizing the common belief function from the entire hierarchy of beliefs may be a very difficult task. Using a simple fixed point argument, we can provide an alternative approach to derive the common belief function. Monderer and Samet [12] define an event E to be p-evident for a set of players S, if, whenever E occurs, it also p-believed by all players in S. They show that an event E is common p-belief among players in a set S at a state ω, if and only if there exists a p-evident event F (for players in S) such that $\omega \in F$, and whenever F occurs, E is also p-believed by all players in S. In analogy with the common belief function, we can define an evidence function $\widehat{p}(E,S)$ that assigns to each set E and group of players S the highest p, such that E is p-evident in S. In fact, we can easily prove that $p(E,E,S) = \widehat{p}(E,S)$, for all E and S. For simplicity, we will write $\widehat{p}(E,S) = \widehat{p}(E)$ and $p(E,F,S) = p(E,F)$ whenever S refers to the whole set of players.

Defining a state space as above is rather inconvenient for the infinite-player game. Instead, we make use of the law of large numbers to summarize the distribution of signals across the population by the fundamental θ, and we can define a state as a realization of θ and y, and the state space as $\mathbb{R}^2$. Also, since payoffs depend on the measures of agents of each type that play each action, and not on their exact identities, it is more natural to define the common belief function and the evidence function with respect to the measure of agents that have a common belief, i.e. we let $p(E,F,m)$ be the highest p, such that, whenever F occurs, E is common p-belief among some measure m of players, and we let $\widehat{p}(F,m)$ be the highest p, such that, whenever F occurs, a measure m of players also p-believes F. One of the appealing consequences of the infinite player game is that for any fundamental event $\theta \in E_0$ all events of the form "a measure m of players p-believes E_0" can be reduced to events of the form $\theta \in E_1$, for an appropriately defined E_1. We will make use of this reduction of higher-order events into fundamental events to derive the common belief function for threshold events[14] directly from the entire sequence of higher-order events, and thereby provide an illustration of Monderer and Samet's fixed-point characterization of common p-belief by p-evidence.

To unify the notation between the two cases, we will use the notation for the infinite player game, $p(E,F,m)$, and include the finite player case in this notation by noting that $p(E,F)$ in the finite player case is equivalent to $p(E,F,1)$ in the infinite player case.

[14]events where the state falls below or above a given threshold in the set of fundamentals

2.5.2 The infinite player case

In a game with a continuum of players, two observations considerably simplify the derivation of the common belief function $p(E, F, m)$. First, once again by virtue of the law of large numbers, θ is a sufficient statistic for the posterior distribution of expectations in the population of players, and therefore for their first-order beliefs. Second, the belief operator preserves the relative order of players within the distribution of beliefs. For any two players i and j, and events $E : \theta \leq \theta^*$ and $E' : \theta \leq \theta'$, if i attaches higher probability to E_1 than j, he will also attach a higher probability to E_2. The first observation implies that a set of measure m of players p-believes that $\theta \leq \theta_0^*$, whenever θ falls below some θ_1^*. Thus, for any first-order event $E_0 : \theta \leq \theta_0^*$, the second-order event "a set of measure m of players p-believes E_0" is equivalent to some suitably defined other first-order event $E_1 : \theta \leq \theta_1^*$. By iteration, we can therefore represent the entire sequence of higher-order events by first-order events. It follows from the second observation that at their intersection in the state space, E_0 is *common* p-belief among a set of players of measure m.

We now use this procedure to derive the $p(E, F, m)$ for threshold events of the types $\theta \leq \theta_0^*$ and $\theta \geq \theta_0^*$. Consider an arbitrary event $E_0 : \theta \leq \theta_0^*$. A player p-believes that $\theta \leq \theta_0^*$, whenever his posterior belief x_j^1 satisfies

$$\Phi\left(\frac{\theta_0^* - x_j^1}{\alpha\varepsilon}\right) \geq p,$$

which defines the first-order threshold expectation x_0 as

$$x_0 = \theta_0^* - \alpha\varepsilon\Phi^{-1}(p).$$

A mass m of players p-believes that $\theta \leq \theta_0^*$, whenever

$$\Phi\left(\frac{x_0 - \alpha^2\theta - (1-\alpha^2)\,y}{\alpha^2\varepsilon}\right) \geq m,$$

which defines a new first-order event $E_1 : \theta \leq \theta_1^*$, where θ_1^* is given by

$$\theta_1^* = \frac{1}{\alpha^2}\left(x_0 - (1-\alpha^2)\,y\right) - \varepsilon\Phi^{-1}(m).$$

A player p-believes that a mass m of players p-believe that $\theta \leq \theta_0^*$, whenever

$$\Phi\left(\frac{\theta_1^* - x_j^1}{\alpha\varepsilon}\right) \geq p.$$

Proceeding in this way, we can recursively define the sequence of threshold expectations and fundamentals by the following two equations:

$$x_k = \theta_k^* - \alpha\varepsilon\Phi^{-1}(p) \tag{2.12}$$

$$\theta_k^* = \frac{1}{\alpha^2}\left(x_{k-1} - \left(1-\alpha^2\right)y\right) - \varepsilon\Phi^{-1}(m). \tag{2.13}$$

Substituting (2.12) into (2.13), and rearranging, we find

$$\theta_{k+1}^* - y = \frac{1}{\alpha^2}\left(\theta_k^* - y\right) - \frac{\alpha\varepsilon}{\alpha^2}\left(\Phi^{-1}(p) + \alpha\Phi^{-1}(m)\right).$$

This difference equation has as a solution

$$\theta_k^* - y - \frac{\alpha\varepsilon}{1-\alpha^2}\left(\Phi^{-1}(p) + \alpha\Phi^{-1}(m)\right) = \frac{1}{\alpha^{2k}}\left(\theta_0^* - y - \frac{\alpha\varepsilon}{1-\alpha^2}\left(\Phi^{-1}(p) + \alpha\Phi^{-1}(m)\right)\right). \tag{2.14}$$

This sequence of thresholds diverges monotonically, as long as

$$\theta_0^* - y - \frac{\alpha\varepsilon}{1-\alpha^2}\Phi^{-1}(p) - \frac{\alpha^2\varepsilon}{1-\alpha^2}\Phi^{-1}(m) \neq 0.$$

In order to sustain a common belief p with respect to $\theta \leq \theta_0^*$, it is necessary that the sequence of thresholds is monotonically increasing or stationary — otherwise the intersection of all higher-order events will be empty. If $\{\theta_k^*\}_{k=1}^{\infty}$ is monotonically increasing, $E_0 : \theta \leq \theta_0^*$ is common p-belief among a set of players of measure m, iff $\theta \leq \theta_1^*$, where θ_1^* is derived from (2.14) for $k = 1$. θ_1^* is decreasing in p and m. Therefore, for any event E_0, increasing p or m reduces the set of states at which E_0 becomes common p-belief. Common p-belief about E_0 of a proportion m of players about an event $\theta \leq \theta_0^*$ is maximized when $\theta_1^* = \theta_0^*$, or

$$\theta_0^* - y - \frac{\alpha\varepsilon}{1-\alpha^2}\Phi^{-1}(p) - \frac{\alpha^2\varepsilon}{1-\alpha^2}\Phi^{-1}(m) = 0, \tag{2.15}$$

which yields

$$p(E_0, E_0, m) = \Phi\left((\theta_0^* - y)\frac{1-\alpha^2}{\alpha\varepsilon} - \alpha\Phi^{-1}(m)\right). \tag{2.16}$$

Let $p(\theta \leq \theta_0^*, m) = p(E_0, E_0, m)$ be the common belief of a mass m of players about event E_0, whenever E_0 has occurred, as defined by (2.16). Note that this is equivalent to the evidence function that we have defined above - a manifestation of the fixed point characterization of Monderer and Samet. One observes that the common belief function is decreasing in m and increases, as the prior probability that $\theta \leq \theta_0^*$ increases (y decreases). Theorem 2 now follows immediately:

Theorem 2 *Assume that* $\frac{1-\alpha^2}{\alpha^2\varepsilon} \to \infty$ *and* $\alpha\varepsilon \to 0$. *Then,*

(i) For any $\theta_0^* > y$, $p(\theta \leq \theta_0^*, m) \to 1$. *For finite* $\alpha\varepsilon$, $p(\theta \leq \theta_0^*, m)$ *is decreasing in* $\alpha\varepsilon$.

(ii) For any $\theta_0^* < y$, $p(\theta \leq \theta_0^*, m) \to 0$. *For finite* $\alpha\varepsilon$, $p(\theta \leq \theta_0^*, m)$ *is increasing in* $\alpha\varepsilon$.

(iii) $p(\theta \leq \theta_0^*, m) \to \Phi\left(\frac{\theta_0^* - y}{\alpha\varepsilon}\right)$, *as* $\alpha \to 0$, *holding* $\alpha\varepsilon$ *constant.* $p(\theta \leq \theta_0^*, m)$ *is increasing in* α *(holding* $\alpha\varepsilon$ *constant).*

(iv) If $\alpha = 1$, $p(\theta \leq \theta_0^*, m) = 1 - m$.

This theorem has various implications. First, it states that the common signal generates convergence of $p(\theta \leq \theta_0^*, m)$ to 1, as long as its variance converges to 0 at a rate at least as fast as the rate of convergence of the standard deviation of the private signals, provided the public signal falls inside E_0. The theorem also highlights the very different nature of common p-belief in an environment with only private signals, and shows that in this case, higher-order beliefs do not depend on the variance of private noise. We also observe that a lower level of noise reduces higher-order uncertainty (common beliefs are closer to 1), and that a higher relative informativeness of the private signal results in more higher-order uncertainty. Whether the common belief about an event converges to 0 or 1 is determined by the location of the common signal.

We can apply a simple symmetry argument to derive common beliefs about events of the type $\theta > \theta_0^*$: At θ_0^*, the mass of players $1 - m$, who do not $p(\theta \leq \theta_0^*, m)$-believe that $\theta \leq \theta_0^*$ must $1 - p(E_0, m)$-believe that $\theta > \theta_0^*$, or $p(\theta > \theta_0^*, 1 - m) = 1 - p(\theta \leq \theta_0^*, m)$.

Let us briefly return to the previous example to link higher-order beliefs to equilibrium strategies. In the context of the previous example, consider an event $E : \theta \leq \theta^*$. If $p(\theta \leq \theta^*, m) \geq \frac{r}{f}$ for $m \geq c(\theta^*)$, then there exists an equilibrium, in which players play b (at least) whenever they p-believe that $\theta \leq \theta^*$. To see this, it suffices to recognize that if players play this strategy profile, the proportion of players playing b will exceed $c(\theta^*)$, whenever $\theta \leq \theta^*$. Since $p(\theta \leq \theta^*, m) \geq \frac{r}{f}$, it is then optimal to play b, whenever E is $p(\theta \leq \theta^*, m)$-believed. Moreover, since E is common $p(\theta \leq \theta^*, m)$-belief, it becomes common knowledge that all players optimally play b, whenever they $p(\theta \leq \theta^*, m)$-believe $\theta \leq \theta^*$, and that the return to b is f, whenever $\theta \leq \theta^*$. Similarly, whenever $1 - p(\theta > \theta_0^*, m) \leq \frac{r}{f}$ for some $m \geq 1 - c(\theta^*)$, there exists an equilibrium in which players do not play b (at least) whenever they $p(\theta > \theta_0^*, m)$-believe that $\theta > \theta_0^*$. We can then give a reinterpretation of the Nash equilibrium conditions in terms of common beliefs. Comparing the condition for a Nash equilibrium (2.10) with the derivation of common beliefs in (2.15), we find immediately that any Nash equilibrium solves $p(\theta \leq \theta_0^*, c(\theta_0^*)) = \frac{r}{f}$.

To complete the derivation of the common belief function in the infinite player case, consider events of the type $E : \theta \in [\theta_1, \theta_2]$. $1 > m_2 > m_1 > 0$. Whenever $\theta \in [\theta_1, \theta_2]$, there is a measure $m_2 - m_1$ of agents who have common $p(\theta \leq \theta_2, m_2)$-belief that $\theta \leq \theta_2$, but they all attach probability less than $p(\theta \leq \theta_1, m_1)$ to the event $\theta \leq \theta_1$. Therefore,

$$p(\theta_1 < \theta \leq \theta_2, m_2 - m_1) \geq p(\theta \leq \theta_2, m_2) - p(\theta \leq \theta_1, m_1).$$

We also have $p(\theta_1 < \theta \leq \theta_2, m) \leq p(\theta \leq \theta_2, m)$ and $p(\theta_1 < \theta \leq \theta_2, m) \leq p(\theta > \theta_1, m) = 1 - p(\theta \leq \theta_1, 1 - m)$. The following corollary follows immediately from theorem 2 and these observations:

Corollary 1 *Consider an arbitrary sequence of incomplete information structures, such that* $\frac{1-\alpha^2}{\alpha^2 \varepsilon} \to \infty$ *and* $\alpha\varepsilon \to 0$. *Then the following are true for all* $m < 1$:

(i) $p(\theta_1 \leq \theta \leq \theta_2, m) \to 1$ *for all* $m < 1$, *iff* $y \in (\theta_1, \theta_2)$.

(ii) $p(\theta_1 \leq \theta \leq \theta_2, m) \to 0$, *if* $y \notin [\theta_1, \theta_2]$.

2.5.3 The finite player case

In an information structure with a finite number of players, we can no longer appeal to the law of large numbers to summarize the distribution of private information in the state variable, and considering the entire hierarchy of beliefs becomes rather difficult. Instead, using Moderer and Samet's fixed-point characterization of common p-belief through p-evidence, we can provide a simple proof that extends the common belief convergence result in theorem 2 and corollary 1 to finite player games. This proof has the further advantage of highlighting the origin of the convergence condition that we have encountered throughout this paper.

Return to the finite player information structure introduced in section 3 and consider a "fundamental event" $E : \theta \in [\theta_1, \theta_2]$. Given a realization of y, any higher-order event of the form "all players believe with probability at least p that E has occurred" is defined on the space of private signals, and we can easily compute its evidence from the posterior beliefs. Note that p-evidence is not a useful concept for fundamental events E, since there always exist signal realizations such that ex post a player attaches very low probability to the event E, even though E has occurred. For the information structures under consideration, the evidence function for any interval in the fundamental space will therefore typically be 0. However, the evidence function is very easy to determine for "signal events", i.e. events of the form $X : x_i \in X_i, \forall i$, $X_i \subseteq \mathbb{R}$. In this case, $\widehat{p}(X)$ is simply given by

$$\widehat{p}(X) = \min_{i; x_i \in X_i} \Pr\left(x_j \in X_j, \forall j \neq i \mid x_i\right),$$

i.e. the lowest posterior probability that any player i attaches to the event that all other players have received signals within X_{-i}, given that his signal lies in X_i,.and X is common $\widehat{p}(X)$-belief, whenever it occurs. Now, consider a sequence of information structures $\{I_n\}_{n=1}^{\infty} = \left\{\alpha_n^2, \alpha_n \varepsilon_n\right\}_{n=1}^{\infty}$, and let $\widehat{p}_n(\cdot)$ denote the evidence function defined by information structure I_n. Using Monderer and Samet's fixed point characterization, we have the following definition of convergence of common p-belief:

Definition 5 *As $n \to \infty$, E becomes common p-belief with probability* 1, *if there exists a sequence of sets X^n, such that (i) $\min_{\theta \in E} \Pr\left(X^n \mid \theta\right) \to 1$, (ii) $\lim_{n\to\infty} \widehat{p}_n\left(X^n\right) \geq p$, and (iii)*

$\liminf_{n\to\infty, x_i\in X_i^n} \Pr(\theta \in E \mid x_i, y) \geq p$ *for all* i.

We are now able to prove the following useful lemma:

Lemma 8 *As* $n \to \infty$, E *becomes common p-belief with probability* 1, *if and only if there exists a sequence of signal events* X^n, *such that* $\liminf_{n\to\infty, x_i\in X_i^n} \Pr(\theta \in E \mid x_i, y) \geq p$ *for all* i *and* $\min_{\theta\in E} \Pr(X^n \mid \theta) \to 1$.

Proof. We need to show that in the above definition, (i) and (ii) imply (iii). But note that for any x_i,

$$\begin{aligned}
\Pr(X^n \mid x_i) &= \int \prod_{j\neq i} \Pr\left(x_j \in X_j^n \mid \theta\right) \frac{1}{\alpha_n \varepsilon_n} \varphi\left(\frac{\theta - \alpha_n^2 x_i - \left(1-\alpha_n^2\right) y}{\alpha_n \varepsilon_n}\right) d\theta \\
&\geq \int_{\theta\in E} \prod_{j\neq i} \Pr\left(x_j \in X_j^n \mid \theta\right) \frac{1}{\alpha_n \varepsilon_n} \varphi\left(\frac{\theta - \alpha_n^2 x_i - \left(1-\alpha_n^2\right) y}{\alpha_n \varepsilon_n}\right) d\theta \\
&\geq \min_{\theta\in E} \prod_{j\neq i} \Pr\left(x_j \in X_j^n \mid \theta\right) \cdot \Pr(\theta \in E \mid x_i, y)
\end{aligned}$$

and lemma 8 follows immediately. ■

As a direct consequence of lemma 8, we find that we can construct the sequence of signal events separately for each player. For each i, we need to construct a sequence $X_i^n = [\underline{x}_n, \overline{x}_n]$ in such a way that $\min_{\theta\in E} \Pr(X_i^n \mid \theta) \to 1$ and $\min_{x_i\in X_i^n} \Pr(\theta \in E \mid x_i, y) \to 1$. A necessary and sufficient condition that $\min_{\theta\in E} \Pr(X_i^n \mid \theta) \to 1$ is that the upper and lower bounds of the signal event converge to the end points of E from above and below at a rate strictly slower than the rate at which $\varepsilon_n \to 0$; formally, this condition is equivalent to $\frac{\overline{x}_n - \theta_2}{\varepsilon_n} \to \infty$, and $\frac{\underline{x}_n - \theta_1}{\varepsilon_n} \to -\infty$. $\min_{i; x_i\in X_i^n} \Pr(\theta \in E \mid x_i, y) \to 1$ is satisfied, whenever $\min_{x_i\in X_i^n} \frac{\theta_2 - E(\theta|x_i,y)}{\alpha_n \varepsilon_n} \to \infty$ and $\max_{x_i\in X_i^n} \frac{\theta_1 - E(\theta|x_i,y)}{\alpha_n \varepsilon_n} \to -\infty$. Since some $x_i \in X_i^n$ lie outside E, this clearly necessitates that $y \in (\theta_1, \theta_2)$, and $y \in (\theta_1, \theta_2)$ is also sufficient, when the variances of the public and private signals converge at the same rates, i.e. $\lim_{n\to\infty} \alpha_n < 1$. If the variance of the public signal converges at a rate slower than the variance of the private signal, we also need that $\theta_1 - E(\theta \mid \underline{x}_n, y) \to 0$ and $\theta_2 - E(\theta \mid \overline{x}_n, y) \to 0$ at a rate strictly slower than the rate at which

$\alpha_n\varepsilon_n \to 0$. Normality implies that $E(\theta \mid \underline{x}_n, y) = \alpha_n^2 \underline{x}_n + (1-\alpha_n^2)\, y$, which converges to θ_1 at the same rate as $(1-\alpha_n^2) \to 0$, provided that $(1-\alpha_n^2) \to 0$ at a rate slower than $\varepsilon_n \to 0$. For convergence to common knowledge, we therefore need that $(1-\alpha_n^2) \to 0$ at a rate slower than the rate at which $\alpha_n\varepsilon_n \to 0$, or $\frac{1-\alpha_n^2}{\alpha_n^2\varepsilon_n} \to \infty$. This finding is fully spelled out and proved in the following theorem whose content parallels theorem 2 and its corollary:

Theorem 3 *Consider an arbitrary sequence of information structures $\{I_n\}_{n=1}^{\infty} = \{\alpha_n^2, \alpha_n\varepsilon_n\}_{n=1}^{\infty}$, such that $\frac{1-\alpha_n^2}{\alpha_n^2\varepsilon_n} \to \infty$ and $\alpha_n\varepsilon_n \to 0$, as $n \to \infty$, and event $\theta \in [\theta_1, \theta_2]$:*

(i) If $\theta_2 > y > \theta_1$, then, with probability 1, $[\theta_1, \theta_2]$ becomes common 1-belief among all players in the limit, as $n \to \infty$.

(ii) If $y < \theta_1$ or $y > \theta_2$, then there exists no $p > 0$, such that in the limit, $[\theta_1, \theta_2]$ is common p-belief with probability 1 among all players.

(iii) If $\frac{1-\alpha_n^2}{\alpha_n^2\varepsilon_n} \to k < \infty$, then there exists no $p > 0$, such that, in the limit, $[\theta_1, \theta_2]$ is common p-belief with probability 1 among all players.

Proof. Let $\widehat{p}_n(\cdot)$ denote the evidence function generated by I_n. For an arbitrary fundamental event $E : \theta \in [\theta_1, \theta_2]$ construct a sequence of signal events $X^n : x_i \in [\underline{x}_n, \overline{x}_n], \forall i$ by setting $\underline{x}_n = \theta_1 - \chi_n$ and $\overline{x}_n = \theta_2 + \psi_n$, for a pair of sequences of positive real numbers χ_n and ψ_n. $\min_{\theta \in [\theta_n', \theta_n'']} \Pr(X^n \mid \theta) \to 1$ if and only if $\frac{\psi_n}{\varepsilon_n} \to \infty$ and $\frac{\chi_n}{\varepsilon_n} \to \infty$. In addition,

$$
\begin{aligned}
\Pr(E \mid x_i, y) &= \Phi\left(\frac{\theta_2 - \alpha_n^2 x_i - (1-\alpha_n^2)\, y}{\alpha_n\varepsilon_n}\right) - \Phi\left(\frac{\theta_1 - \alpha_n^2 x_i - (1-\alpha_n^2)\, y}{\alpha_n\varepsilon_n}\right) \\
&= \Phi\left(\frac{-\alpha_n^2\psi_n + (1-\alpha_n^2)(\theta_2 - y) + \alpha_n^2(\overline{x}_n - x_i)}{\alpha_n\varepsilon_n}\right) \\
&\quad -\Phi\left(\frac{\alpha_n^2\chi_n + (1-\alpha_n^2)(\theta_1 - y) + \alpha_n^2(\underline{x}_n - x_i)}{\alpha_n\varepsilon_n}\right)
\end{aligned}
$$

for arbitrary ψ_n and χ_n. Setting

$$\psi_n = \frac{1}{2}(1-\alpha_n^2)(\overline{x} - y)$$

and

$$\chi_n = -\frac{1}{2}\left(1-\alpha_n^2\right)\left(\underline{x}-y\right)$$

we find that $\Pr(E \mid x_i, y) \to 1$, as $\frac{1-\alpha_n^2}{\alpha_n^2 \varepsilon_n} \to \infty$ and $\alpha_n \varepsilon_n \to 0$, whenever $y \in (\theta_1, \theta_2)$ for any sequence $x_i \in [\underline{x}_n, \overline{x}_n]$, so that that $\widehat{p}_n(X^n) \to 1$. Moreover, for arbitrary ψ_n and χ_n, such that $\frac{\psi_n}{\varepsilon_n} \to \infty$ and $\frac{\chi_n}{\varepsilon_n} \to \infty$, $\Pr(E \mid x_i, y) \to 0$, as $\frac{1-\alpha_n^2}{\alpha_n^2 \varepsilon_n} \to \infty$ and $\alpha_n \varepsilon_n \to 0$, whenever $y < \theta_1$ or $y > \theta_2$, or when $\frac{1-\alpha_n^2}{\alpha_n \varepsilon_n} \to k < \infty$. ■

In addition to providing an intuition for the necessity and sufficiency of the condition that $\frac{1-\alpha_n^2}{\alpha_n^2 \varepsilon_n} \to \infty$ and $\alpha_n \varepsilon_n \to 0$, this argument also suggests that the same condition applies beyond the case of normally distributed signals that we study here. Note that up to lemma 8, no step of our argument hinged on the use of normal distributions. In theorem 3, we simply compared the rates of convergence of $\underline{x}_n$ ($\overline{x}_n$) to θ_1 (θ_2) and of $E(\theta \mid x_i, y)$ to x_i. To satisfy $\min_{\theta \in [\theta_1, \theta_2]} \Pr(X^n \mid \theta) \to 1$ in general, the sequence X^n will still have to satisfy that $\underline{x}_n$ ($\overline{x}_n$) converges to θ_1 (θ_2) at a rate slower than $\varepsilon_n \to 0$. We can also form a "best linear estimator" $\widehat{\theta} = \alpha_n^2 x_i + \left(1-\alpha_n^2\right) y$ and note that $\widehat{\theta} \to x_i$ at the same rate as $\left(1-\alpha_n^2\right) \to 0$. Whenever the fully bayesian $E(\theta \mid x_i, y)$ converges to x_i at the same rate as $\widehat{\theta}$, therefore, the same condition will apply, regardless of distributional assumptions.

We conclude this section with a brief discussion of common belief convergence when $\frac{1-\alpha_n^2}{\alpha_n^2 \varepsilon_n} \to k < \infty$. The last part of theorem 3 shows that sustaining any level of common belief (other than 0) about a fundamental event E with probability 1, when E has occurred, is impossible. However, for signal events, we have already observed that the evidence function and hence common beliefs are well-defined and easy to determine, and generically take on strictly positive values between 0 and 1. Without going into the details of the derivations, in the case where $\frac{1-\alpha_n^2}{\alpha_n^2 \varepsilon_n} = 0$, i.e. an environment with private signals only,the p-evidence of a symmetric signal event $X : x_i \leq \widehat{x}, \forall i$, is given by $\widehat{p}(X) = \frac{1}{I}$, where I denotes the number of players, i.e. the common belief about a threshold event for player signals is $\frac{1}{I}$. Moreover, whenever a player observes a signal $x_i \leq \widehat{x}$, he attaches a probability at least $1 - \frac{k}{I}$ to the event that at least $k-1$ other players have also observed signals below $\widehat{x}$ - note that these formulations correspond exactly to those obtained for the p-evidence function in the infinite player case. For bounded intervals in the signal space, these values also serve as upper bounds for the common-belief function (and in the limit as $\varepsilon_n \to 0$, the common belief function converges exactly to these values for any symmetric signal event).

2.6 Equilibrium Convergence

2.6.1 p-Dominance

In a finite-player game, we use the definition of p-dominance provided by Morris, Rob and Shin [14]: In a complete information normal form game $G = (S, A, \pi)$ with a finite set of players $S = \{1, 2, ..., I\}$, action profile $a = (a_1, ..., a_I)$ is a *strictly p-dominant equilibrium*[15] at state θ, if, for all $q \geq p$,

$$\left[\pi_i\left(\theta, a_i, a_{-i}\right) - \pi_i\left(\theta, a_i', a_{-i}\right)\right] \cdot q + \min_{a_{-i}'}\left[\pi_i\left(\theta, a_i, a_{-i}'\right) - \pi_i\left(\theta, a_i', a_{-i}'\right)\right] \cdot (1-q) > 0,$$

i.e. if θ is common knowledge, and if all other players play a_{-i} with probability at least p, then action a_i is optimal for player i, for all i: Any strict Nash equilibrium is p-dominant for some $p < 1$.

For a game with an infinite number of players, we need to use a slightly different definition of p-dominance. In an infinite-player normal form game $G = (S, A, \pi)$, let $a = (a_1, ..., a_I)$ be a pure action profile, and let $m \leq 1$. With a slight abuse of notation, a is said to be a *strictly p, m-dominant equilibrium* of this game at θ, if, for all $q \geq p$ and all μ, such that $\mu_i\left(a_i\right) > m$ for all i, and arbitrary μ',

$$\left[\pi_i\left(\theta, a_i, \mu\right) - \pi_i\left(\theta, a_i', \mu\right)\right] \cdot q + \left[\pi_i\left(\theta, a_i, \mu'\right) - \pi_i\left(\theta, a_i', \mu'\right)\right] \cdot (1-q) > 0,$$

i.e. if θ is common knowledge, agents in class i have a strictly dominant action a_i, whenever with probability p or higher, the measure of players playing a_{-i} exceeds m, for all i. Any strict Nash-equilibrium at θ^* is p, m-dominant for some values of $p < 1, m \leq 1$, in a neighborhood of θ^*.

[15] *We will be using the strict version of p-dominance. Weak p-dominance would be defined by replacing the strict with a weak inequality in the definition.*

To unify the notation between the finite and infinite player game, note that *p-dominance* in the finite player game is equivalent to $p,1$-dominance in the infinite player game. We will therefore simply refer to p,m-dominance from now on.

2.6.2 Equilibrium Convergence

Consider now a sequence of information structures $\{I_n\}_{n=1}^{\infty} = \left\{\alpha_n^2, \alpha_n\varepsilon_n\right\}_{n=1}^{\infty}$, such that $\alpha_n\varepsilon_n \rightarrow 0$, and let $p_n(\cdot,\cdot,m)$ be the common belief function derived by information structure I_n. We are now able to state and prove the following equilibrium convergence proposition for the sequence of incomplete information games $G \times \{I_n\}_{n=1}^{\infty}$:

Proposition 5 *Consider a game $G = (S, A, \pi)$ satisfying $(C1)$ and $(C2)$ and suppose a is a strictly p,m-dominant Nash equilibrium of G for all $\theta \in E = [\theta_1, \theta_2]$. If the sequence of common belief functions satisfies $\lim_{n\to\infty} p_n(E,F,m) = p(E,F,m) \geq p$ for some $F \subseteq E$, then $G \times \{I_n\}_{n=1}^{\infty}$ has a sequence of Nash Equilibria, in which, for sufficiently high n, all players play according to a, (at least) whenever they p-believe that F has occurred.*

Proof. Note first that for a given posterior expectation $x \notin [\theta_1, \theta_2]$, as $\alpha_n\varepsilon_n \rightarrow 0$, $\Pr(E \mid x) \rightarrow 0$. Thus, in the limit, if E is a common $p(E,F,m)$-belief, and $p(E,F,m) > 0$ necessarily, in the limit, a player $p(E,F,m)$-believes E, only if his posterior expectation x satisfies $x \in [\theta_1, \theta_2]$. Now, suppose player i after signal x believes with probability at least $p_n(E,F,m)$ that at least a measure m of players of each type play according to a_{-i}. The expected payoff difference $\Delta_n(a_i, a_i')$ between actions a_i and an arbitrary alternative a_i' satisfies

$$\begin{aligned}
&\Delta_n(a_i, a_i') \\
\geqslant\; & \min_{\mu:\mu_i(a_i)\geq m} \left[E_\theta\left(\pi_i(\theta, a_i, \mu) \mid x, \theta \in E\right) - E_\theta\left(\pi_i\left(\theta, a_i', \mu\right) \mid x, \theta \in E\right)\right] \cdot p_n(E,F,m) \\
& + \min_{\mu'} \left[E_\theta\left(\pi_i\left(\theta, a_i, \mu'\right) \mid x\right) - E_\theta\left(\pi_i\left(\theta, a_i', \mu'\right) \mid x\right)\right]\left(1 - p_n(E,F,m)\right).
\end{aligned}$$

Given x, as $\alpha\varepsilon \rightarrow 0$, $E_\theta(\pi_i(\theta, a_i, \mu) \mid x, \theta \in E) \rightarrow \pi_i(x, a_i, \mu)$, $E_\theta(\pi_i(\theta, a_i', \mu) \mid x, \theta \in E) \rightarrow \pi_i(x, a_i', \mu)$, $E_\theta(\pi_i(\theta, a_i, \mu') \mid x) \rightarrow \pi_i(x, a_i, \mu')$ and $E_\theta(\pi_i(\theta, a_i', \mu') \mid x) \rightarrow \pi_i(\theta, a_i', \mu')$. As

$n \to \infty$, therefore, for all $\theta \in [\theta_1, \theta_2]$

$$\begin{aligned}
\Delta_n\left(a_i, a_i'\right) &\to \Delta\left(a_i, a_i'\right) \\
&\geqslant \min_{\mu:\mu_i(a_i)\geq m}\left[\pi_i\left(\theta, a_i, \mu\right) - \pi_i\left(\theta, a_i', \mu\right)\right] \cdot p\left(E, F, m\right) \\
&\quad + \min_{\mu'}\left[\pi_i\left(\theta, a_i, \mu'\right) - \pi_i\left(\theta, a_i', \mu'\right)\right] \cdot \left(1 - p\left(E, F, m\right)\right) \\
&> 0
\end{aligned}$$

by the definition of strict p-dominance, and since $p\left(E, F, m\right) \geq p$. ■

As an immediate consequence of proposition 5, we observe that we can separate the effects of payoff assumptions from the effects of assumptions concerning the information structure. We can now combine the common belief convergence theorems with proposition 5 characterizing equilibrium convergence to generalize the characterization of equilibria in theorem 1.

Corollary 2 *Consider a sequence of incomplete information games $G \times \{I_n\}_{n=1}^{\infty}$ satisfying $(C1)$ and $(C2)$. Suppose a is a p, m-dominant equilibrium in some interval $E = [\theta_1, \theta_2]$, and $p < 1$. Suppose also that there exists a sequence F_n, such that $p_n\left(E, F_n, m\right) \to 1$ and $\Pr\left(F_n \mid E\right) \to 1$, as $\alpha\varepsilon \to 0$. Then $G \times \{I_n\}_{n=1}^{\infty}$ has a sequence of Nash Equilibria, where, in the limit, whenever $\theta \in E$, with probability 1, a is played by a measure of players at least m .*

The following theorem now follows from this corollary and common belief convergence, and discusses equilibrium convergence for a game with payoff structure G and a sequence of information structures $\{I_n\}_{n=1}^{\infty} = \left\{\alpha_n^2, \alpha_n\varepsilon_n\right\}_{n=1}^{\infty}$:

Theorem 4 *Consider a sequence of finite (infinite) player, incomplete information games $G \times \{I_n\}_{n=1}^{\infty}$ satisfying $(C1)$ and $(C2)$. Suppose a is a p, m-dominant equilibrium in some interval $E = [\theta_1, \theta_2]$, and $p < 1$. Then, $G \times \{I_n\}_{n=1}^{\infty}$ has a sequence of Nash equilibria, in which a is played with probability 1, if and only if $\{I_n\}_{n=1}^{\infty}$ satisfies $\alpha_n\varepsilon_n \to 0$ and $\frac{1-\alpha_n^2}{\alpha_n^2\varepsilon_n} \to \infty$, as $n \to \infty$, and $\theta_1 < y < \theta_2$.*

Theorem 4 provides, in various ways, a generalization of Theorem 1 (i). For symmetric (finite or infinite player) 2-action global games with the unique pure strategy equilibrium a for $\theta < \overline{\theta}$, pure strategy equilibrium b for $\theta > \underline{\theta}$, the theorem implies that as $\alpha_n\varepsilon_n \to 0$ *and* $\frac{1-\alpha_n^2}{\alpha_n^2\varepsilon_n} \to \infty$, there are multiple threshold equilibria, with two of the thresholds converging to $\underline{\theta}$, and $\overline{\theta}$, as seen in the initial example. An immediate extension of theorem 4 also implies that we can sustain non-monotonic equilibria. As an example, let $\underline{\theta} < \underline{\theta}_1 < \overline{\theta}_1 < \underline{\theta}_2 < \overline{\theta}_2 < \overline{\theta}$ in a 2-action global game. It then follows that, as $\alpha_n\varepsilon_n \to 0$ *and* $\frac{1-\alpha_n^2}{\alpha_n^2\varepsilon_n} \to \infty$, there also exists a non-monotonic limit equilibrium, in which a is played with probability 1 in the limit, (at least) whenever $\theta \leq \underline{\theta}$ or $\theta \in \left[\underline{\theta}_2, \overline{\theta}_2\right]$, and b is played (at least) whenever $\theta \geq \overline{\theta}$ or $\theta \in \left[\underline{\theta}_1, \overline{\theta}_1\right]$ - in fact any such collection of disjoint compact intervals coupled with a locally p-dominant equilibrium for each interval can be sustained as an equilibrium in the limit. In other words, convergence of common belief to common knowledge guarantees not only the existence of three threshold equilibria with action profiles locally converging to the common knowledge equilibria, but also the approximation of non-monotonic equilibria of the common knowledge global game.

2.7 Conclusion

This paper studied incomplete information games and examined equilibrium convergence, as the level of noise in the information structure vanishes. In particular, we analyze the convergence of equilibria, as common beliefs converge to common knowledge, and provide an intuition into the determinant underlying convergence condition that, if the variance of the public signal converges to 0 at a faster rate than the standard deviation of the private signals, common beliefs of players converge to common knowledge. With respect to global game applications, the following conclusions emerge from the analysis: First, a unique equilibrium is less likely to occur if public information is informative relative to private information, and/or if the overall level of noise is small. Second, the equilibrium convergence results cast doubt on recent papers, in which the private information limit is used as an equilibrium selection to perform a comparative static or dynamic analysis. Such a selection presumes that common beliefs are low and bounded away from 1. However, it is not a priori clear that in the environments under consideration, public information would not enable players to generate high common belief, and thereby coordinate their actions, nor, if the equilibrium is unique, that it will be independent of existing public information.

One could in principle extend the analysis in the present paper to information structures that move away from the separation into fully public and fully private information, or add uncertainty about the extent to which certain signals are observed within the population. It appears that any such extension will validate the principal conclusions that emerge from this analysis: *As long as the information structure generates common p-belief sufficiently close to* 1, *one will be able to sustain any equilibrium that also exists in a common knowledge game.* Moreover, for small degrees of payoff uncertainty, any small change in the information structure will affect equilibrium strategies mostly through its effect on common belief. Finally, as long as at some level, players have a common prior about the information structure, presumably, the formal analysis will be similar to this paper.

There are economic arguments in favor and against the existence of truly "common information." One view is that, even if public information exists, this information may be interpreted in different ways by different players, and therefore leads to different beliefs ex post. To be specific, suppose in the currency crisis example used in section 2, that the fundamental θ represents the "toughness" of the central bank in defending against a speculative attack, and $\theta = R + \mu\pi$ depends on the central bank reserves R and the rate of inflation π. R and π may be publicly announced, but the true weighting parameter μ is not known. Instead, each player has his own "model" μ_i, which is a noisy signal of μ. This model uncertainty translates into private signals as in the environment studied by Morris and Shin.[16]

But this example, as much of the global games literature, including this paper, relies on the game being static and one-shot. In reality, many of the episodes to which static global games are applied are dynamic. It is not clear that model biases of this kind would survive, if the game was played repeatedly over time, and players either commonly observed outcomes, or had the possibility to communicate with each other. Monderer and Samet [13] provide an example, in which the common observation of a stochastic process over time leads to convergence of

[16] I am indebted to one referee for suggesting this example.

common p-belief to common knowledge, or, in the language of the previous example, there is very high common belief about which is the right model to apply. Angeletos, Hellwig and Pavan [1] study a simple dynamic model of speculative attacks, in which common knowledge about equilibrium strategies (i.e. common knowledge about what each player will do, conditional on his information) is sufficient to break equilibrium uniqueness in each period. However, the incidence of devaluation in the long run still very much depends on the convergence condition for the variances of public and private signals encountered here.[17]

Unfortunately, we know relatively little about how markets aggregate information over time, and how common beliefs are formed. From a theoretical perspective, it is easy to construct environments, in which public information may or may not be sufficiently precise to allow common beliefs to converge to common knowledge over time. To be explicit, consider the following scenario: Let time be discrete and infinite, and the number of players be infinite. Suppose agents observe a common signal in each period with a variance constant over time, and in addition start also with a private signal. The variance of the public information declines at a rate t^{-1}. If each player now has the possibility to communicate his private information to one other randomly drawn agent in the population, and in turn, learns his private information, the variance of private information will decline at an exponential rate, much faster than the variance of the public signal (This relies on the fact that private communication becomes much more efficient at revealing information, but never leads to public revelation). On the other hand, if some player has the possibility to communicate his information to a non-zero measure of agents in each period, then in time, all players would share the information this player has, and the variance of public information would also decline at an exponential rate, thus making convergence to common knowledge a possibility.

[17] In defense of the argument against truly common information, one might argue that players not only disagree about the right model to apply, but also, disagree about how to interpret observations drawn over time, so as to update their model bias and achieve common belief about the underlying parameter. In principle, this simply places the same opposition of arguments at a higher level in the belief hierarchy. The more general point is that, whenever Bayesian players have a common prior about the process ex ante, and they commonly observe public outcomes over time, their posterior beliefs will converge so as to generate common belief. If a common prior does not exist, or if the arrival of new private information more than offsets the convergence of posterior beliefs, the flow of public information may not be able to sustain high common belief over time. However, convergence of common p-belief over time occurs in the long run, whereas, for most applications, we are rather interested in short run outcomes with a possibly limited degree of common belief.

How markets aggregate information to form common beliefs is mostly an applied question. In reality, information revelation in markets will depend on the exogenously given environment, as well as on strategic choices by players. As conjectured in the previous paragraph, the ability of some players to act "publicly", or to communicate with a large number of other players may be sufficient to generate common beliefs close to 1. Corsetti, Dasgupta, Morris and Shin [3] study an environment, where a single investor has substantial financial wealth ("Soros"), so that his decision influences other investors and ultimately the likelihood of devaluation (even when his action is not directly observable, the existence of Soros makes small investors more aggressive). It would be interesting to compare this strategic effect to an environment, in which, in addition, Soros may influence other players by making public announcements. Formally one would have to examine whether public information can be generated by strategic revelation of private information.

Alternatively, one may argue that the central bank can publicize information to the markets. Listed companies and banks are legally obliged to publish data on cash flows and profitability on a regular basis. If the present model is anything to judge on, one would conclude that, if increased transparency of either the markets or the central bank (in the form of better public information being released to market participants) leads to an increase in common beliefs, it may have a destabilizing impact on the market - a conclusion rather opposite to standard wisdom, and which follows from the reduction of higher-order uncertainty with more precise public information.[18]

Summing up these comments, we conclude that, in order to fully evaluate the global games approach as an equilibrium selection in applications of coordination games, we will need a better understanding of how private information is revealed to the public over time, either strategically, or through exogenous announcements (by legal requirements, for instance), and whether, within the environment of interest, the revelation of public information is sufficient to generate common beliefs close to common knowledge.

[18]see Morris and Shin [18] for a formalization of disclosures in a principal-agent environment similar to this. The notion of transparency employed here, as well as in their papers, requires that players agree on how to translate a series of data into information about θ.

2.8 Bibliography

[1] G.-M. Angeletos, C. Hellwig and A. Pavan, "On the Dynamics of Speculative Attacks", MIT, 2002.

[2] H. Carlsson and E. van Damme, "Global Games and Equilibrium Selection", *Econometrica* **61** (1993), 989-1018.

[3] G. Corsetti, A. Dasgupta, S. Morris, and H.S. Shin, "Does one Soros make a Difference? A Theory of Currency Crises with Large and Small Traders", *Rev. Econ. Stud.,* forthcoming.

[4] D. Diamond and P. Dybvig, "Bank Runs, Deposit Insurance and Liquidity", *J. Polit. Econ.* **91** (1983), 401-419.

[5] I. Goldstein and A. Pauzner, "Endogenous Probability of Bank Runs in a Rational Expectations Model", University of Tel Aviv, 2000.

[6] K.L. Judd, "The Law of Large Numbers with a Continuum of IID Random Variables", *J. Econ. Theory* **35** (2000), 19-25.

[7] A. Kajii and S. Morris, "Common p-Belief: the General Case", *Games Econ. Behav.* **18** (1997), 73-82.

[8] A. Kajii and S. Morris, "The Robustness of Equilibria to Incomplete Information", *Econometrica* **65** (1997), 1283-1309.

[9] A. Kajii and S. Morris, "Payoff Continuity in Incomplete Information Games", *J. Econ. Theory* **82** (1998), 267-276.

[10] C. Metz, "Public and Private Information in Self-fulfilling Currency Crises", *J. Econ. (Z. Nationalökon.)* **76** (2002), 65-85.

[11] P. Milgrom and R. Weber, "Distributional Strategies for Games with Incomplete Information", *Math. Oper. Res.* **10** (1985), 619-632.

[12] D. Monderer and D. Samet, "Approximating Common Knowledge with Common Beliefs", *Games Econ. Behav.* **1** (1989), 170-190.

[13] D. Monderer and D. Samet, "Stochastic Common Learning", *Games Econ. Behav.* **9** (1995), 161-171.

[14] S. Morris, R. Rob, and H. S. Shin, "p-Dominance and Belief Potential", *Econometrica* **63** (1995), 145-157.

[15] S. Morris and H. S. Shin, "Unique Equilibrium in a Model of Self-fulfilling Currency Attacks", *Amer. Econ. Rev.* **88** (1998), 587-597.

[16] S. Morris and H. S. Shin, "Coordination Risk and the Price of Debt", *FMG discussion paper* **373**, London School of Economics, 2001.

[17] S. Morris and H. S. Shin, "Global games: Theory and Applications", *in* "Advances in Economics and Econometrics, the Eighth World Congress" (M. Dewatripont, L. Hansen, and S. Turnovsky, Eds.), Cambridge University Press, forthcoming.

[18] S. Morris and H. S. Shin, "Social Value of Public Information", *Amer. Econ. Rev.*, forthcoming.

[19] M. Obstfeld, "Rational and Self-fulfilling Balance-of-payments Crises", *Amer. Econ. Rev.* **76** (1986), 72-81.

[20] J.-C. Rochet and X. Vives, "Coordination Failures and the Lender of Last Resort: was Bagehot right after all?", *CEPR discussion paper* **3233** (2002).

[21] A. Rubinstein, "The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge", *Amer. Econ. Rev.* **79** (1989), 385-391.

Chapter 3

Public Announcements, Adjustment Delays, and the Business Cycle

Summary of Chapter 3 *Building on Woodford (2001), I study the effects of a lack of common knowledge on nominal adjustment. In particular, I show how the speed of price adjustment following a shock depends on the information structure among price-setters. The provision of public information leads to a reduction of higher-order uncertainty, and hence to more rapid price adjustments, but it potentially comes at the cost of an increased exposure to informational noise. I extend my analysis to allow for other disturbances, showing that higher-order uncertainty may account for the persistence of any kind of shock; and I briefly discuss some implications for the design of monetary policy, most importantly showing that an increased degree of monetary transparency reduces the short-run output gains of unanticipated inflation, and hence may serve as an implicit commitment device for a discretionary central bank.*[1]

JEL classification numbers: D82, D84, E31, E32

Keywords: Higher-Order Uncertainty, Incomplete Nominal Adjustment, Transparency

[1] I am grateful to Olivier Blanchard and Andre Kurmann for helpful discussions. The usual disclaimer applies.

3.1 Introduction

What is the relation between the supply of money, prices and real output in the short run? In particular, why do prices not adjust immediately after a money supply shock? The theoretical literature around these two questions has emphasized two potential causes of incomplete nominal adjustment, each of which has led to important subsequent insights on the dynamics of price adjustments and expectations about the conduct of monetary policy: lack of information or "misperceptions" originally developed by Phelps [12] and Lucas [9], and adjustment costs or real rigidities that prevent an immediate adjustment of pricing decisions.

Despite its theoretical success, the incomplete information model runs up against a powerful criticism, when used as a descriptive model of business cycles: The theoretical model predicts that prices should fully adjust once the information about aggregate shocks becomes available; however most macro-economic data is available after only short delays, and incomplete information can therefore not account on its own for the observed delays of price adjustment.

In the original version of Lucas, uncertainty about monetary shocks is revealed with a delay of one period, and prices are determined by market-clearing in a competitive dynamic general equilibrium model with incomplete information. In contrast, the more recent new Keynesian models emphasize the role of prices as strategic variables in an environment of imperfect competition. This change of emphasis has profound implications for the underlying strategic interaction, since prices are strategic complements.

A recent article by Woodford [16] introduces strategic pricing into an incomplete information model similar to Lucas. Woodford also alters the information structure, assuming that information comes in the form of private signals to the price-setting decision makers, but the true state never becomes common knowledge. His analysis develops a simple intuition why monetary shocks have persistent real effects, even when they are accurately observed by price-setters: although firms may have precise information about the policy shock, they lack information about each other's beliefs; in fact, they have no information at all about what their beliefs are relative to the population average. In an environment of strategic complementarity, however, precisely such higher-order beliefs are necessary to forecast the behavior of other agents, and Woodford shows that the existence of higher-order uncertainty can lead to substantial nominal adjustment delays.

The purpose of this paper is to expand and develop Woodford's idea. My first objective is to explore the effects of higher-order uncertainty in as simple and accessible a model as possible, that is also sufficiently flexible to extend or adapt to other contexts. The second objective is to study the role of the information structure in detail: Since the composition of the information structure determines the degree of higher-order uncertainty, as quantified by the departure from common knowledge, Woodford's results suggest that it should also have an influence on nominal adjustment. Using some recent insights from the theory of global games, which emphasizes the coordinating effect of public information (cf. Morris and Shin, [10]; Hellwig [7]), this paper then explores this link between the parameters of the information structure, in particular the precision of public and private information, and the process of nominal adjustment. For these purposes, the paper develops a version of the Lucas-Woodford model that is sufficiently simple and flexible to study the dynamic implications of a whole range of information structures.

Before highlighting the paper's main results, it will be useful to motivate my approach towards modelling the information structure, in particular the separation of information into public and private signals. Woodford studies an environment, in which individuals have access only to private information. He bases his information structure on the famous "island" paradigm, which is meant to represent the informational differences between agents. Moreover, he appeals to limits in individual information processing capacities. As argued by Sims [13], this can account for a "private signal" information structure like Woodford's, even when the relevant economic data is publicly observed. Here, I take a more literal view of the information structure. While it is important to emphasize the role of differential information, the island paradigm has the drawback that it allows for no informational interaction among decisionmakers; in other words, in Woodford's economy a price-setter has no clue about how his information compares to the population average. More realistically, information processing within a market environment relies to a large extent on interaction and communication, and in the process, decisionmakers do learn about each other. In this respect, public disclosures and the processing of information by the media play an important role, and Morris and Shin [11] emphasize the importance of such publicly available information as focal points for beliefs, even if the information does not filter through to everyone. The hypothesis that a decisionmaker has access to idiosyncratic and public signals is therefore motivated not only as a more accurate description of reality, but also as an attempt to capture the informational differences, whether they are the result of decisionmakers using different sources of information, or the result of limited information processing a la Sims, at the same time as taking into account the fact that various channels of communication serve to coordinate expectations. The degree to which the population is capable of processing information is captured by the parameters of the information structure, in particular the relative importance of public information, and the overall degree of noise. As discussed in Hellwig [7], these informational parameters are related to the degree of common p-belief, which quantifies the departure from common knowledge in the information structure, and hence the degree to which individual decisionmakers are capable of efficiently coordinating their decisions.

On theoretical grounds, the analysis discusses, how the inflation-output trade-off depends on the importance of higher-order uncertainty, measured as a function of (i) the degree to which pricing decisions are strategic complements, and (ii) the parameters of the information structure. In particular, the provision of public information reduces higher-order uncertainty and therefore leads to a faster adjustment of prices and smaller, less persistent effects of monetary shocks on output; on the other hand, a higher precision of public information may increase the macroeconomic exposure to informational noise. This second effect becomes important in particular when public information is relatively noisy. The informational noise effect is at the heart of the static model by Morris and Shin [11]; indeed the formal analysis in this paper extends some of their results into a context that is of interest to dynamic macroeconomic theory. I also explore some of the welfare implications of changes in the market's information structure. Augmenting the model by an objective function for the central bank along the lines of Kydland and Prescott [8] and Barro and Gordon [2], I show that the provision of precise public information may serve as an implicit commitment device against inflationary biases: By committing to disclose public information, the central bank reduces the effects of monetary shocks on output, thereby reducing the temptation to use monetary policy to stimulate output. However, the provision of public information may come at the cost of increasing informational noise, if there is a lower bound on the precision of public information.

The model provides closed-form solutions for prices and output in response to the underlying aggregate demand and supply disturbances, as well as informational noise; the underlying informational parameters can potentially be inferred from the data, and the model itself leads to some interesting testable implications. Moreover, the modelling approach appears to be sufficiently flexible to be applied in other macroeconomic contexts, in which strategic complementarities play a role, for instance investment or demand spill-overs. The paper thus makes the additional methodological contribution of proposing a solution technique for embedding higher-order uncertainty into dynamic macroeconomic models, and the arguments suggest that higher-order uncertainty coupled with strategic complementarities may be the cause of persistent effects not only of monetary shocks, but of other aggregate disturbances as well.

The remainder of the paper is organized as follows: Section 2 introduces the main model, the informational assumptions, and discusses the main analytical building blocks. Section 3 presents the paper's main theoretical results regarding the link between the information structure about monetary shocks and the inflation-output trade-off. Section 4 extends the analysis to allow for higher-order uncertainty regarding other disturbances; in particular, it is argued that the Woodford's insight regarding the persistence of monetary shocks applies also to supply shocks. Section 5 augments the initial model to discuss the welfare implications of information provision by the central bank, and informally discusses the role that the monetary policy regime, and in particular explicit monetary targets, have in reducing higher-order uncertainty. Section 6 concludes by discussing the paper's main implications and potential other applications.

3.2 The Model

3.2.1 Set-up

There is a large number of price-setters in monopolistic competition a la Dixit-Stiglitz.[2] When solving the price-setter's optimization problem, the first-order condition implies that each price-

[2] see Woodford [17] for a detailed discussion of the micro-foundations of household and firm behavior, on which this model is based.

setter sets the log of his own price p_t^i according to

$$p_t^i = E_t^i(p_t) + (1-r) E_t^i(y_t) \tag{3.1}$$

y_t denotes the log of real output relative to its steady-state value (which here is normalized to 0), p_t denotes the population average of log-price, and $E_t^i(\cdot) = E\left(\cdot \mid \Im_t^i\right)$ denotes the expectations operator conditional on i's information set as of date t, $\Im_t^i$, and $r \in (0,1)$. The monetary authority targets the log of nominal output, denoted θ_t, which is assumed to be generated as an (exogenous) linear process from a sequence of monetary policy shocks $\{\varepsilon_t\}_{t=-\infty}^{\infty}$. Allowing for some finite degree k of integration,

$$\Delta^k \theta_t = \sigma \left[\varepsilon_t + \sum_{s=1}^{\infty} b_s \varepsilon_{t-s}\right], \tag{3.2}$$

where $\varepsilon_t \sim N(0,1)$, and ε_t is iid over time. Substituting $\theta_t = y_t + p_t$ into (3.1) yields

$$p_t^i = r E_t^i(p_t) + (1-r) E_t^i(\theta_t). \tag{3.3}$$

r thus measures the degree to which individual pricing decisions are strategic complements. As a consequence of the Dixit-Stiglitz model, r is increasing in the elasticity of substitution between different goods (i.e. in the degree of competition), and decreasing in the degree of convexity of the cost function. As the economy becomes perfectly competitive (or as the cost function becomes linear), r converges to 1.

In period t, a price-setter has access to noisy information about θ_t, to be precise, one process of private information $\left\{x_{t-s}^i\right\}_{s=0}^{\infty}$ and a process of public information $\{z_{t-s}\}_{s=0}^{\infty}$:

$$x_t^i = \theta_t + \sigma_u u_t^i;\ u_t^i \sim N(0,1)$$

and

$$z_t = \theta_t + \sigma_v v_t;\ v_t \sim N(0,1),$$

where $\left\{u_{t-s}^i\right\}_{s=0}^{\infty}$ and $\{v_{t-s}\}_{s=0}^{\infty}$ are iid processes, independent of each other, as well as of

$\{\varepsilon_{t-s}\}_{s=0}^{\infty}$. Finally, θ_t becomes commonly observable with a delay of T periods, that is, at time t, $\{\varepsilon_{t-T-s}\}_{s=0}^{\infty}$ (or equivalently, $\{\theta_{t-T-s}\}_{s=0}^{\infty}$) is common knowledge among all price-setters. I make this last assumption for pure convenience, and for computational reasons. Nothing prevents T from being very large, in which case we approach an environment, in which θ never becomes fully observable.

It will be convenient to introduce a vector notation for both the fundamental process and the processes of public and private information: define Θ_t as the column vector of realizations of θ from period $t-T+1$ up to period t; let $m_{t-\tau}$ denote the expectation of $\theta_{t-\tau}$, based on the common knowledge of $\{\varepsilon_{t-T-s}\}_{s=0}^{\infty}$, and let M_t denote the vector of $m_{t-\tau}$, for $\tau = 0, ..., T-1$. Finally, let E_t be the vector of monetary policy shocks from period $t-T+1$ up to period t. Then,

$$\Theta_t = \begin{pmatrix} \theta_t \\ \theta_{t-1} \\ . \\ \theta_{t-T+1} \end{pmatrix}, \; M_t = \begin{pmatrix} m_t \\ m_{t-1} \\ . \\ m_{t-T+1} \end{pmatrix} \text{ and } E_t = \begin{pmatrix} \varepsilon_t \\ \varepsilon_{t-1} \\ . \\ \varepsilon_{t-T+1} \end{pmatrix};$$

and the monetary policy process can be expressed as

$$\Theta_t = M_t + \sigma B E_t,$$

where B is some $T \times T$ upper-triangular matrix whose entries are derived from (3.2), with $b_{ii} = 1$ for $i = 1, ..., T$. Similarly, it will be convenient to express the signal process in vector form. Let

$$X_t^i = \begin{pmatrix} x_t^i \\ x_{t-1}^i \\ . \\ x_{t-T+1}^i \end{pmatrix}, \; Z_t = \begin{pmatrix} z_t \\ z_{t-1} \\ . \\ z_{t-T+1} \end{pmatrix} \text{ and } V_t = \begin{pmatrix} v_t \\ v_{t-1} \\ . \\ v_{t-T+1} \end{pmatrix},$$

i.e. X_t^i and Z_t denote the vectors of public and private signals available to price-setter i. The vector of public signals is then written as

$$Z_t = \Theta_t + \sigma_v V_t.$$

At time t, the state of the economy is summarized by $\left(\{\theta_{t-T-s}\}_{s=0}^{\infty}, \Theta_t, Z_t\right)$, while i's information set $\Im_t^i = \left(\{\theta_{t-T-s}\}_{s=0}^{\infty}, X_t^i, Z_t\right)$.

I now return to the pricing equation. Averaging (3.3) over i, and substituting forward yields

$$p_t^i = (1-r)\sum_{s=0}^{\infty} r^s E_t^i\left[\overline{E_t}^{(s)}(\theta_t)\right], \tag{3.4}$$

where $\overline{E_t}^{(s)}(\theta_t)$ denotes the s-th order *average expectation*, i.e. $\overline{E_t}^{(0)}(\theta_t) = \theta_t$, $\overline{E_t}^{(1)}(\theta_t) = \overline{E_t}(\theta_t)$ denotes the population average expectation over θ_t, and $\overline{E_t}^{(s+1)}(\theta_t) = \overline{E_t}\left[\overline{E_t}^{(s)}(\theta_t)\right]$. In words, $\overline{E_t}^{(s)}(\theta_t)$ is the population average expectation of the population average expectation of the ... (repeat s times) ... of the population average expectation of θ_t. Note that the average expectations operator in general does not satisfy the law of iterated expectations; in fact it satisfies it if and only if all available information is public. The average price is given by

$$p_t = (1-r)\sum_{s=0}^{\infty} r^s \overline{E_t}^{(s+1)}(\theta_t) \tag{3.5}$$

and the log of real output y_t is given by

$$y_t = (1-r)\sum_{s=0}^{\infty} r^s\left[\theta_t - \overline{E_t}^{(s+1)}(\theta_t)\right]. \tag{3.6}$$

The deviation of real GDP from its trend level is therefore a weighted average of the deviation of all average higher-order expectations. In order to fully derive the dynamics of price and output adjustments following a monetary policy shock, we need to work out the dynamics of higher-order average expectations. This is done in two steps: I first derive a linear filtering equation for $E_t^i(\Theta_t)$ as a function of the signal processes X_t^i and Z_t. By averaging over the filtering equation, I then find a linear relation between Θ_t and $\overline{E_t}(\Theta_t)$, which is iterated to solve for (3.5) and (3.6). These steps are carried out in the two subsequent lemmas.

3.2.2 Optimal Filtering

We begin the analysis by deriving a linear filtering equation for $E_t^i(\Theta_t)$. Standard results imply that the signal process $\{X_t^i, Z_t\}$ can be aggregated into

$$\Xi_t^i = \frac{\sigma_v^2}{\sigma_v^2 + \sigma_u^2} X_t^i + \frac{\sigma_u^2}{\sigma_v^2 + \sigma_u^2} Z_t$$

For further reference, define $\Sigma \equiv \frac{\sigma_v^2 \sigma_u^2}{\sigma_v^2 + \sigma_u^2}$ as the ex post noise in the signal process Ξ_t^i, and $\alpha \equiv \frac{\sigma_v^2}{\sigma_v^2 + \sigma_u^2}$ as the relative importance of private information. A simple way to express $E_t^i(\Theta_t)$ in the required matrix form is to proceed by maximum likelihood estimation. The log likelihood function $L\left(\Theta_t; \Xi_t^i . M_t\right)$ for the inference problem is given by

$$\begin{aligned} L\left(\Theta_t; \Xi_t^i . M_t\right) &= -\frac{1}{2\sigma^2}\left(\Theta_t - M_t\right)' \left[B^{-1}\right]' B^{-1}\left(\Theta_t - M_t\right) \\ &\quad -\frac{1}{2\Sigma}\left(\Theta_t - \Xi_t^i\right)'\left(\Theta_t - \Xi_t^i\right). \end{aligned} \tag{3.7}$$

Maximizing L with respect to Θ_t to solve for $E_t^i(\Theta_t)$ yields as a first-order condition

$$\left(BB'\right)^{-1}\left(E_t^i(\Theta_t) - M_t\right) + \frac{\sigma^2}{\Sigma}\left(E_t^i(\Theta_t) - \Xi_t^i\right) = 0$$

which has as a solution

$$E_t^i(\Theta_t) - M_t = \left[I_T + \frac{\Sigma}{\sigma^2}\left(BB'\right)^{-1}\right]^{-1}\left(\Xi_t^i - M_t\right)$$

We have shown

Lemma 9 *For an information structure satisfying the assumptions of the previous section, the posterior expectation of individual i about Θ_t satisfies*

$$E_t^i(\Theta_t) - M_t = \alpha\Delta\left(X_t^i - M_t\right) + (1-\alpha)\Delta\left(Z_t - M_t\right) \tag{3.8}$$

where

$$\Delta = \left[I_T + \frac{\Sigma}{\sigma^2}\left(BB'\right)^{-1}\right]^{-1}. \tag{3.9}$$

It should be noted that maximum likelihood methods can be used to obtain a similar linear filtering equation far more generally, for instance to account for "learning", i.e. a gradual increase of the public and private signal precisions over time. Also, note that, if v is an eigenvector of BB', with corresponding eigenvalue λ, then v is also an eigenvector of Δ, corresponding to an eigenvalue of $\widehat{\lambda} = \lambda\left(\lambda + \frac{\Sigma}{\sigma^2}\right)^{-1}$. Since BB' is positive definite, it follows that all eigenvalues of Δ are positive and strictly between 0 and 1. Again, this property can be shown to hold generally.

The matrix Δ determines the weights that a Bayesian estimate of Θ_t attributes to past observations. The coefficients in Δ only depend on the ratio between Σ and σ^2, i.e. the importance of signal noise *relative to fundamental noise.* We can thus separate the effects resulting from the *composition of the information structure* (parametrized by α) from the effects coming from signal noise in the inference problem, parametrized by σ^2/Σ, and the effects of fundamental shocks, i.e. σ^2.

3.2.3 Higher-order expectations

In the next lemma, I use the linear filtering equation (3.8) to recursively derive an expression for $(1-r)\sum_{s=0}^{\infty} r^s E_t^i\left[\overline{E_t}^{(s)}(\Theta_t)\right]$:

Lemma 10 *Suppose that i's Bayesian posterior of Θ_t satisfies*

$$E_t^i(\Theta_t) - M_t = N_1\left(X_t^i - M_t\right) + N_2\left(Z_t - M_t\right)$$

where N_1 is positive definite and all its eigenvalues are strictly smaller than 1. *Then,*

$$(1-r)\sum_{s=0}^{\infty} r^s\left[E_t^i\left[\overline{E_t}^{(s)}(\Theta_t)\right] - M_t\right] = \left[I_T - rN_1\right]^{-1}\left[(1-r)N_1\left(X_t^i - M_t\right) + N_2\left(Z_t - M_t\right)\right]. \tag{3.10}$$

Proof. Averaging (3.8), we determine the average expectation:

$$\overline{E_t}(\Theta_t) - M_t = N_1(\Theta_t - M_t) + N_2(Z_t - M_t) \tag{3.11}$$

and i's expectation of the average expectation equals

$$\begin{aligned} E_t^i\left[\overline{E_t}(\Theta_t)\right] - M_t &= N_1\left[N_1\left(X_t^i - M_t\right) + N_2(Z_t - M_t)\right] + N_2(Z_t - M_t) \\ &= N_1^2\left(X_t^i - M_t\right) + [I_T + N_1] N_2(Z_t - M_t). \end{aligned}$$

Iterating the procedure provides expressions for all higher-order expectations,[3]

$$E_t^i\left[\overline{E_t}^{(s)}(\Theta_t)\right] - M_t = N_1^{s+1}\left(X_t^i - M_t\right) + [I_T - N_1]^{-1}\left[I_T - N_1^{s+1}\right] N_2(Z_t - M_t)$$

and

$$\overline{E_t}^{(s+1)}(\Theta_t) - M_t = N_1^{s+1}(\Theta_t - M_t) + [I_T - N_1]^{-1}\left[I_T - N_1^{s+1}\right] N_2(Z_t - M_t). \tag{3.12}$$

Substituting into (3.4) and (3.5) yields

$$\begin{aligned} & (1-r)\sum_{s=0}^{\infty} r^s\left[E_t^i\left[\overline{E_t}^{(s)}(\Theta_t)\right] - M_t\right] \\ = \; & (1-r)\sum_{s=0}^{\infty} r^s N_1^{s+1}\left(X_t^i - M_t\right) \\ & + (1-r)\sum_{s=0}^{\infty} r^s [I_T - N_1]^{-1}\left[I_T - N_1^{s+1}\right] N_2(Z_t - M_t) \\ = \; & (1-r)[I_T - rN_1]^{-1} N_1\left(X_t^i - M_t\right) + [I_T - N_1]^{-1} N_2(Z_t - M_t) \\ & - (1-r)[I_T - N_1]^{-1} N_1 [I_T - rN_1]^{-1} N_2(Z_t - M_t) \\ = \; & (1-r)[I_T - rN_1]^{-1} N_1\left(X_t^i - M_t\right) + [I_T - rN_1]^{-1} N_2(Z_t - M_t) \\ = \; & [I_T - rN_1]^{-1}\left[(1-r) N_1\left(X_t^i - M_t\right) + N_2(Z_t - M_t)\right]. \end{aligned}$$

■

[3]Since $I_T - \lambda N_1$ is invertible, for $|\lambda| \leq 1$, all the matrix operations below are well-defined.

p_t^i is then given by the first entry of (3.10). As a first result, one observes that this dynamic model delivers (qualitatively and formally) the same implication as the static model of Morris and Shin [11]: *when taking his pricing decision, the price-setter discounts his private information by a factor* $(1-r)$*, and he thus over-reacts to public information* (i.e. reacts more than if the coordination motive were absent).

3.2.4 Impulse Responses

Applying the previous lemma to the present information structure and averaging over i, we then find an expression for p_t as a function of the processes of the fundamental and the public information process. The dynamics of price and output adjustment then depend on the dynamic processes of Θ_t, Z_t, and X_t^i, which can now be substituted to compute the impulse response functions of y_t, p_t^i, and p_t with respect to the shocks ε_t, u_t^i, and v_t. Using the above results, the average price is given by the first entry of

$$\begin{aligned}(1-r)\sum_{s=0}^{\infty} r^s \left[\overline{E_t}^{(s+1)}(\Theta_t) - M_t\right] &= \left[I_T - r\alpha\Delta\right]^{-1}\Delta\left[(1-r)\,\alpha\,(\Theta_t - M_t) + (1-\alpha)\,(Z_t - M_t)\right] \\ &= \sigma\,(1-r\alpha)\left[I_T - r\alpha\Delta\right]^{-1}\Delta B E_t \\ &\quad + (1-\alpha)\,\sigma_v\left[I_T - r\alpha\Delta\right]^{-1}\Delta V_t \qquad (3.13)\end{aligned}$$

and output is the first entry of

$$\begin{aligned}(1-r)\sum_{s=0}^{\infty} r^s \left[\Theta_t - \overline{E_t}^{(s+1)}(\Theta_t)\right] &= \sigma\left[I_T - (1-r\alpha)\left[I_T - r\alpha\Delta\right]^{-1}\Delta\right] B E_t \\ &\quad - (1-\alpha)\,\sigma_v\left[I_T - r\alpha\Delta\right]^{-1}\Delta V_t \qquad (3.14)\end{aligned}$$

The first row of the matrix $(1-r\alpha)\left[I_T - r\alpha\Delta\right]^{-1}\Delta B$ therefore measures the response of prices to a current or past monetary shock. The impulse response of output and prices to informational shocks is given by the first row of $(1-\alpha)\,\sigma_v\left[I_T - r\alpha\Delta\right]^{-1}\Delta$. Using (3.9) to solve for $(1-r\alpha)\left[I_T - r\alpha\Delta\right]^{-1}\Delta$ gives

$$(1-r\alpha)\left[I_T - r\alpha\Delta\right]^{-1}\Delta = \left[I_T + \gamma\left(BB'\right)^{-1}\right]^{-1}, \qquad (3.15)$$

where $\gamma \equiv \frac{\Sigma}{\sigma^2}\frac{1}{1-r\alpha}$. Thus, up to a scaling effect of σ, the impulse response of prices and output to monetary shocks only depends on γ, which I interpret as an index measuring the importance of higher-order uncertainty: If $\alpha r = 0$, then higher-order uncertainty is either irrelevant ($r = 0$ implies that there is no coordination motive) or inexistent ($\alpha = 0$ implies that all information is common, hence there is no higher-order uncertainty). In that case, $\gamma = \frac{\Sigma}{\sigma^2}$, and the impulse responses correspond to the ones obtained if price-setters set their prices equal to their bayesian estimate of θ_t; formally, impulse responses are given by the filtering matrix Δ. Alternatively, if $\alpha = 1$, we find ourselves in the environment of maximal higher-order uncertainty, studied by Woodford [16], in which $\gamma = \frac{\Sigma}{\sigma^2}\frac{1}{1-r}$. γ is (i) increasing in $\frac{\Sigma}{\sigma^2}$, i.e. the relative importance of signal noise, (ii) increasing in α, i.e. the relative importance of private information, and (iii) increasing in r, the importance of strategic complementarities. In the extreme case, where r is close to 1 and $\alpha = 1$, i.e. in a highly competitive market with no public information, γ can become arbitrarily large, even for low values of Σ. Since r is a function of the degree of competition, we thus conclude that more competition lead to more higher-order uncertainty. Furthermore, γ is increasing in σ_v.(since an increase in σ_v raises both α and Σ), and in σ_u: Taking the derivative of γ with respect to σ_u, while holding σ_v and r fixed, we find

$$\begin{aligned}
\frac{\partial\gamma}{\partial\sigma_u} &= \frac{\partial}{\partial\sigma_u}\left[\frac{\sigma_v^2}{\sigma^2}\frac{1-\alpha}{1-\alpha r}\right] = \frac{\sigma_v^2}{\sigma^2}\frac{\partial}{\partial\sigma_u}\left[\frac{1-\alpha}{1-\alpha r}\right]\cdot\frac{\partial\alpha}{\partial\sigma_u} \\
&= -\frac{\sigma_v^2}{\sigma^2}\frac{1-r}{(1-\alpha r)^2}\cdot\frac{\partial\alpha}{\partial\sigma_u} > 0, \text{ since } \frac{\partial\alpha}{\partial\sigma_u} < 0.
\end{aligned}$$

When changing σ_u, the effect of improving information always dominates the compositional effect, due to an increase in the private information component. A reduction of σ_u therefore leads to an overall reduction in higher-order uncertainty, and thus reduces adjustment delays and the exposure to informational noise.

The impulse responses of output to informational shocks, on the other hand, depend on the above matrix, as well as a scaling factor $(1-\alpha)\sigma_v$. Solving as a function of γ and σ_v, the impulse response to informational shocks is given by the first row of

$$(1-\alpha)\sigma_v\left[I_T - r\alpha\Delta\right]^{-1}\Delta = \frac{\sigma^2}{\sigma_v}\gamma\left[I_T + \gamma\left(BB'\right)^{-1}\right]^{-1}. \tag{3.16}$$

Apart from a scaling factor σ, the impulse response function thus depends on the one hand on higher-order uncertainty through γ, on the other hand directly on $\frac{\sigma}{\sigma_v}$, i.e. the informativeness of public information relative to the fundamental process. Changing the composition of the information structure thus has both a direct effect and an indirect effect on the impulse response of prices and output to informational shocks.

3.3 Main Results

General solutions are now easily computable for any given matrix B. Using the previous computations, the processes for prices and output are written as

$$\begin{aligned}(1-r)\sum_{s=0}^{\infty} r^s \left[\overline{E}_t^{(s+1)}(\Theta_t) - M_t\right] &= \sigma \left[I_T + \gamma (BB')^{-1}\right]^{-1} BE_t \\ &\quad + \frac{\sigma^2}{\sigma_v}\gamma \left[I_T + \gamma (BB')^{-1}\right]^{-1} V_t \end{aligned} \tag{3.17}$$

$$\begin{aligned}(1-r)\sum_{s=0}^{\infty} r^s \left[\Theta_t - \overline{E}_t^{(s+1)}(\Theta_t)\right] &= \sigma\gamma (BB')^{-1} \left[I_T + \gamma (BB')^{-1}\right]^{-1} BE_t \\ &\quad - \frac{\sigma^2}{\sigma_v}\gamma \left[I_T + \gamma (BB')^{-1}\right]^{-1} V_t \end{aligned} \tag{3.18}$$

(3.17) and (3.18) illustrate the effect of higher-order uncertainty on the delays in price adjustment: If $\gamma = 0$, i.e. if information is disseminated infinitely quickly, and the fundamental immediately becomes common knowledge among the price-setters, then prices adjust to the full information level without delay, and monetary shocks have no effect on output. If $\gamma > 0$, then there is uncertainty and a lack of common knowledge of fundamentals. Only then output is affected by monetary shocks, and higher values of γ lead to longer adjustment delays and more important output effects. How important output effects are depends on the ex post noise in the information structure and on the importance of private information: The more important private information is, the longer the adjustment delays are, due to higher-order uncertainty. In this respect, the benchmark case where all information is common and prices are set equal to the Bayesian posterior in each period, provides an upper bound for the speed of adjustment. *Since* $\gamma \geq \frac{\Sigma}{\sigma^2}$, *prices adjust less than they would, if there was no higher-order uncertainty. Higher-order uncertainty thus amplifies price stickyness.*

In addition to the monetary shocks, informational shocks affect prices and output. These shocks have effects similar to "supply shocks", insofar as any increase of prices also leads to a corresponding decrease in output. The impact of informational shocks depends positively on γ: The more important higher-order uncertainty is, the more important is the influence of noise in public signals. Moreover, holding γ fixed, a decrease in σ_v also leads to an increase of exposure to informational noise. As emphasized by Morris and Shin [11], changes in σ_v therefore have an ambiguous effect, for fixed values of r and σ_u: On the one hand, a decrease in σ_v decreases higher-order uncertainty, and therefore reduces delays of price adjustment and informational noise indirectly; on the other hand, improved public information may lead to an over-exposure to informational noise, since price-setters over-react to public information. Note that the overall exposure to informational noise is non-monotonic with respect to σ_v: If σ_v is very large, price-setters pay little attention to public signals, and hence, informational noise has little effect on output. When σ_v is very small, there is little first and higher-order uncertainty about Θ_t, which means that price-setters are able to react almost immediately to monetary shocks, and coordinate their price adjustments. [4]

To complete the discussion, I briefly comment on the effects of r close to 1, i.e. a high degree of strategic complementarities, or market competition. We have already observed that this leads to large values of γ, and hence to more delays in price adjustment, however note that in the special case where $\alpha = 1$, i.e. in a highly competitive market with no common information, prices can take arbitrarily long to adjust, even when the private information is very precise.

In the remainder of this section, I illustrate these points explicitly, by setting $T = 1$ and $T = 2$, and solving in closed form. I then show numerical solutions for impulse responses, when

[4] The same discussion applies, if we consider a reduction of σ_u and σ_v in equal proportions, i.e. reduce Σ while holding α fixed, or if we hold Σ fixed, while reducing α, i.e. increase the public information component of the information structure.

T is large.

3.3.1 $T = 1$

The case $T = 1$ provides a simple extension of the static model of Morris and Shin [11] into a dynamic contexts. In this case, $B = [1]$, and it can easily be checked that $\left[I_T + \gamma \left(BB'\right)^{-1}\right]^{-1} = (1+\gamma)^{-1}$. Writing the current average price and output as functions of the aggregate shocks, we find

$$p_t - m_t = \frac{1}{1+\gamma}\sigma\varepsilon_t + \frac{\sigma^2}{\sigma_v}\frac{\gamma}{1+\gamma}v_t \tag{3.19}$$

$$y_t = \frac{\gamma}{1+\gamma}\sigma\varepsilon_t - \frac{\sigma^2}{\sigma_v}\frac{\gamma}{1+\gamma}v_t \tag{3.20}$$

Substituting (3.19) into (3.20), one obtains

$$y_t = \gamma\left(p_t - m_t\right) - \frac{\sigma^2}{\sigma_v}\gamma v_t \tag{3.21}$$

Since $p_t - m_t$ is unexpected inflation, (3.21) is an expectations-augmented Phillips curve.[5] The slope of the Phillips curve depends on higher-order uncertainty. Thus, as higher-order uncertainty increases (γ increases), output becomes more sensitive to unexpected inflation.

The short-run volatilities of unexpected inflation and output are

$$E\left(p_t - m_t\right)^2 = \frac{\sigma^2}{(1+\gamma)^2}\left[1 + \gamma^2\left(\frac{\sigma}{\sigma_v}\right)^2\right]$$

$$Ey_t^2 = \frac{\gamma^2\sigma^2}{(1+\gamma)^2}\left[1 + \left(\frac{\sigma}{\sigma_v}\right)^2\right]$$

whereas the correlation between output and unexpected inflation is

$$E\left[y_t\left(p_t - m_t\right)\right] = \frac{\gamma\sigma^2}{(1+\gamma)^2}\left[1 - \gamma\left(\frac{\sigma}{\sigma_v}\right)^2\right] \geq 0.$$

[5] Similar expressions can be derived in the general case. In that case, m_t no longer corresponds to the common expected price level, since some public information was revealed, without the fundamental becoming common knowledge. Lagged price and output levels therefore also enter into the equation to account for the unanticipated current effects of past monetary shocks.

The variance of output is increasing in γ. The effect of γ on the variance of inflation is ambiguous: If γ is large, inflation volatility is mostly due to informational shocks, in which case, reducing γ reduces the volatility of inflation. If γ is small, inflation volatility is mostly due to noise in the monetary policy process, and reducing γ increases the short-run response of prices to monetary shocks. Finally, we observe that the model generally predicts a positive correlation between output and unexpected inflation, unless $\alpha = 0$, i.e. when there is no element of private information.

Full information revelation after one period precludes the discussion of any meaningful dynamics and persistence of shocks, since prices fully adjust after one period. To understand the effects of the informational parameters on adjustment dynamics, I therefore turn to the case, where $T = 2$.

3.3.2 $T = 2$

Let b denote the effect of a past monetary shock on the current value of θ, or

$$\Theta_t = M_t + \sigma \begin{bmatrix} 1 & b \\ 0 & 1 \end{bmatrix} E_t$$

In this case,

$$\left(BB'\right)^{-1} = \begin{bmatrix} 1 & -b \\ -b & 1+b^2 \end{bmatrix}$$

and

$$\left[I_T + \gamma \left(BB'\right)^{-1}\right]^{-1} = \frac{1}{(1+\gamma)^2 + b^2\gamma} \begin{bmatrix} 1+\gamma+\gamma b^2 & \gamma b \\ \gamma b & 1+\gamma \end{bmatrix}$$

Substituting into (3.17) and (3.18), we find

$$
\begin{aligned}
p_t - m_t &= \sigma\left\{\frac{1+\gamma+\gamma b^2}{(1+\gamma)^2+b^2\gamma}\varepsilon_t + \left[1-\frac{\gamma^2}{(1+\gamma)^2+b^2\gamma}\right]b\varepsilon_{t-1}\right\} \\
&\quad +\frac{\sigma^2}{\sigma_v}\gamma\left\{\frac{1+\gamma+\gamma b^2}{(1+\gamma)^2+b^2\gamma}v_t + \frac{\gamma}{(1+\gamma)^2+b^2\gamma}bv_{t-1}\right\} \qquad (3.22) \\
y_t &= \sigma\left\{\frac{(1+\gamma)\gamma}{(1+\gamma)^2+b^2\gamma}\varepsilon_t + \frac{\gamma^2}{(1+\gamma)^2+b^2\gamma}b\varepsilon_{t-1}\right\} \\
&\quad -\frac{\sigma^2}{\sigma_v}\gamma\left\{\frac{1+\gamma+\gamma b^2}{(1+\gamma)^2+b^2\gamma}v_t + \frac{\gamma}{(1+\gamma)^2+b^2\gamma}bv_{t-1}\right\} \qquad (3.23)
\end{aligned}
$$

It is straight-forward to show that

$$
\begin{aligned}
\frac{d}{d\gamma}\left(\frac{(1+\gamma)\gamma}{(1+\gamma)^2+b^2\gamma}\right) &= \frac{(1+\gamma)^2+\gamma^2 b^2}{\left[(1+\gamma)^2+b^2\gamma\right]^2} > 0 \\
\frac{d}{d\gamma}\left(\frac{\gamma^2}{(1+\gamma)^2+b^2\gamma}\right) &= \frac{2(1+\gamma)+\gamma b^2}{\left[(1+\gamma)^2+b^2\gamma\right]^2} > 0
\end{aligned}
$$

A decrease in γ increases the response of prices to both current and past monetary shock. On the other hand, the coefficients on v_t and v_{t-1} are increasing in γ so that reducing higher-order uncertainty also leads to a reduction of the effects of current and past informational shocks on output. However, as we have observed in the case where $T = 1$, a reduction in σ_v may actually lead to an increase in the effect of informational noise on output, if public information is very diffuse. We conclude that the results that were highlighted before apply to both current and past monetary and informational shocks.

3.3.3 Large T

In this section, I numerically solve for the impulse response function in the case of an example where $T = 30$ is set sufficiently large, so that by the time a monetary shock becomes common knowledge, it has almost entirely been factored into pricing decisions. I follow Woodford in the specification of the monetary policy process, assuming that

$$\Delta\theta_t = \rho\Delta\theta_{t-1} + \sigma\varepsilon_t \qquad (3.24)$$

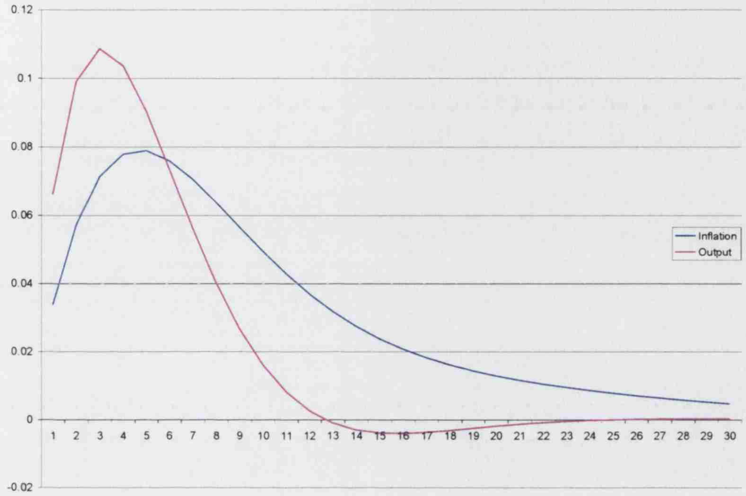

Figure 3-1: Monetary Shocks, $(5, 10)$

i.e. the first difference of nominal GDP follows an AR(1)-process. I fix $r = 0.85$ (the value used by Woodford) and $\rho = 0.9$. I then vary the informational parameters to illustrate the comparative statics effects that were identified above. The figures below plot impulse responses to monetary and informational shocks for values of (σ_u, σ_v) of $(5, 10)$, $(3.33, 20)$, $(5, 1)$ and $(5, 100)$. The first two pairs of parameters yield a value of $\gamma = 62.5$, and hence illustrate the effects of compositional changes that leave the degree of higher-order uncertainty unchanged, the last two pairs illustrate the effect of changes in σ_v, when compared with the first. We observe that the reduction in higher-order uncertainty associated with improved public information leads to faster price adjustment, we also observe that the exposure to informational noise is non-monotonic, and tends to be largest for intermediate values of σ_v.

As was already observed by Woodford, it comes as a property of any incomplete information model of price adjustment that output peaks before inflation does, since it takes time for a monetary policy shock not only to become knowledge, but common knowledge among price-setters. The above graphs illustrate that this feature, which is in line with empirical VAR estimations, for example by Christiano, Eichenbaum and Evans [6], is robust to changes in the

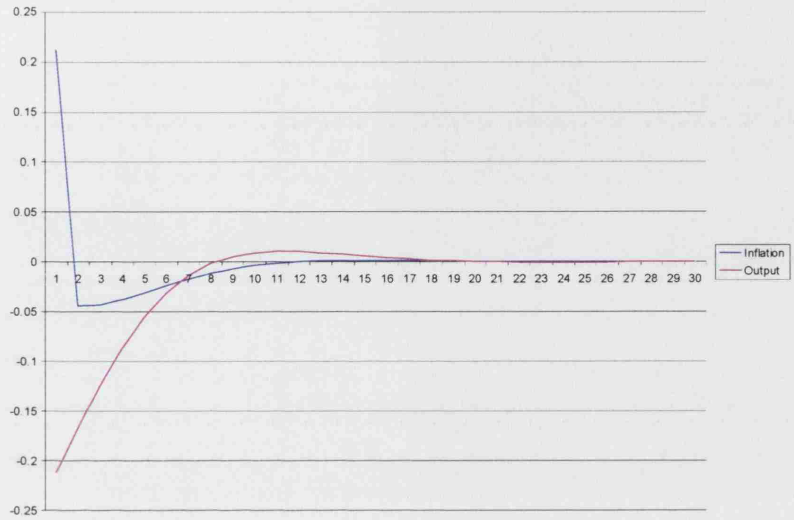

Figure 3-2: Informational Shocks, $(5, 10)$

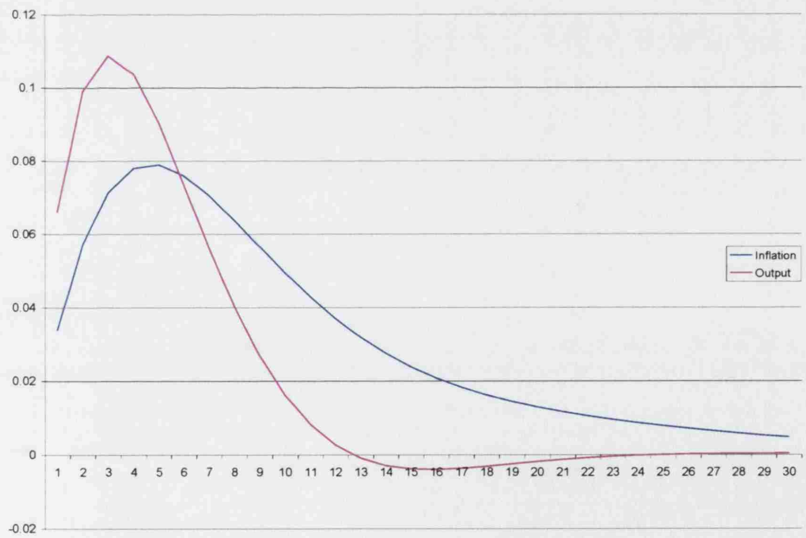

Figure 3-3: Monetary Shocks, $(3.33, 20)$

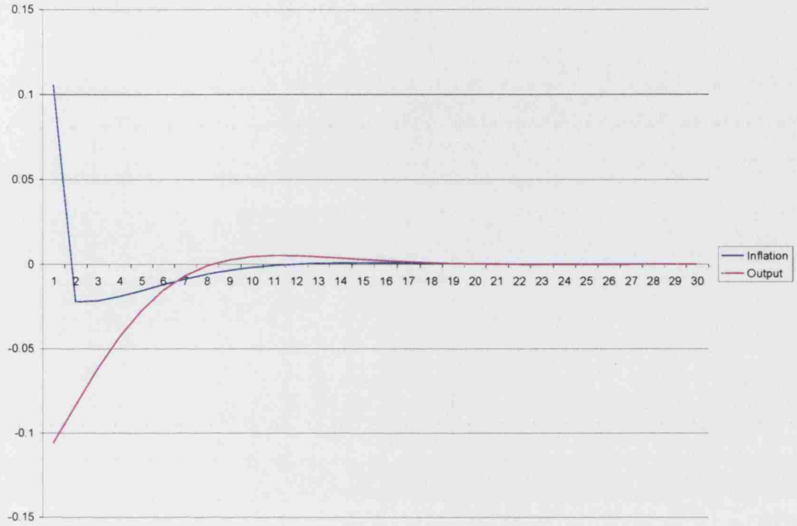

Figure 3-4: Informational Shocks, $(3.33, 20)$

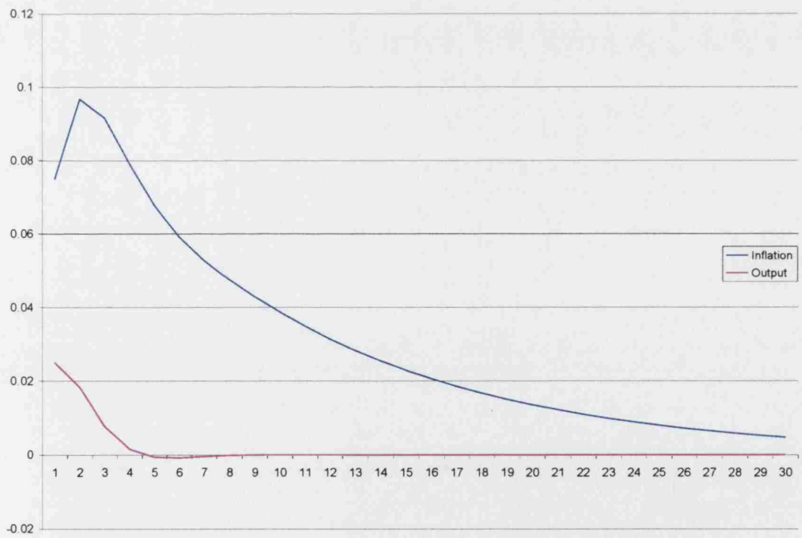

Figure 3-5: Monetary Shocks, $(5, 1)$

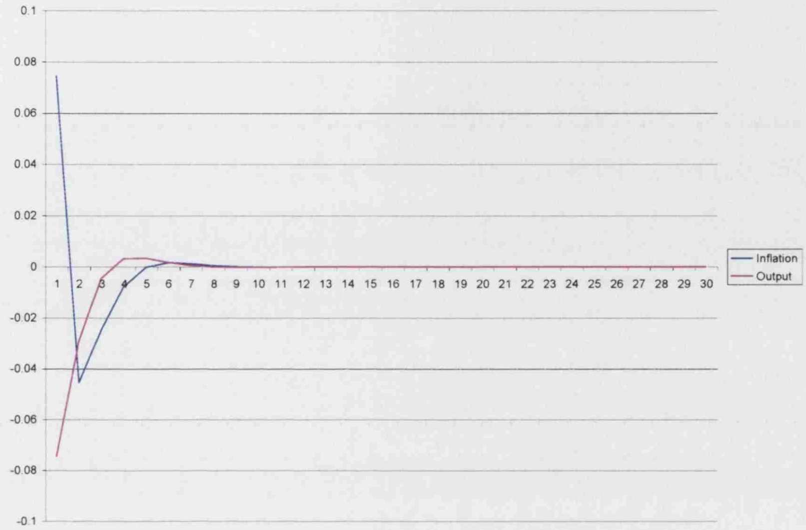

Figure 3-6: Informational Shocks, $(5, 1)$

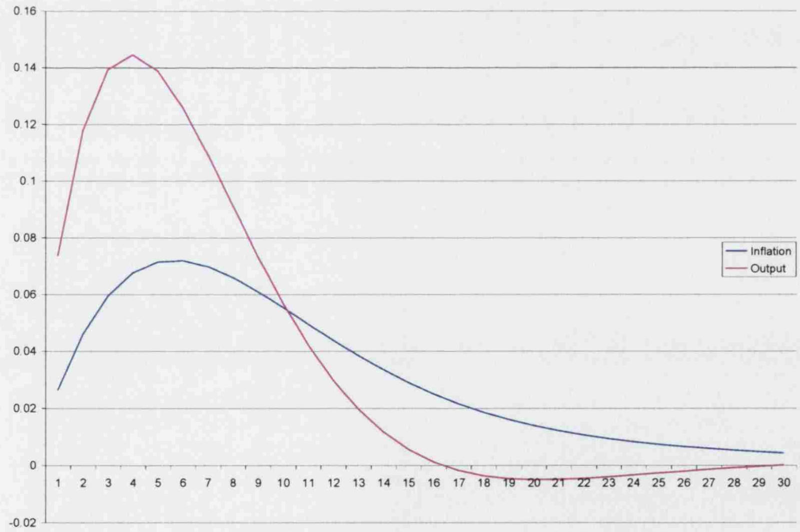

Figure 3-7: Monetary Shock, $(5, 100)$

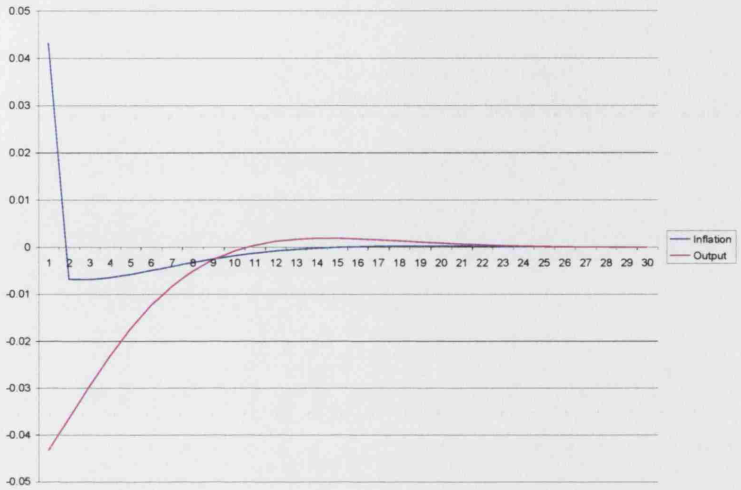

Figure 3-8: Informational Shock, $(5, 100)$

information structure. As a quantitatively testable description of nominal adjustment following a monetary shock, the present model remains incomplete, since it abstracts from monetary transmission channels other than incomplete information, as well as abstracting from other shocks. Note however, that the delays in price adjustment, and hence the lead of output effects before price effects is entirely based on learning. One should thus expect that to the extent that other shocks are also subject to higher-order uncertainty, the impulse responses they generate will exhibit similar features.

3.4 Supply Shocks

So far, the analysis focused on the effects of monetary shocks, and incomplete information about the monetary policy process. As noted by Woodford [16], restricting the analysis to a unique source of shocks highlights the fact that in contrast with Lucas' original model, the existence of an additional source of noise is not necessary to generate incomplete nominal adjustment; rather the presence of higher-order uncertainty along with strategic complementarities is sufficient to generate output effects; moreover, higher-order uncertainty leads to substantial persistence. The model, however, can easily be augmented to allow for supply shocks, as well as higher-order uncertainty about the latter, which leads to some interesting additional insights.

I therefore augment the model by assuming that the level of potential output, $\overline{y}_t$, follows a linear stochastic process, again allowing for some finite degree of integration. In this case, (3.1) continues to hold, if we adjust y_t for fluctuations in potential output, i.e. the price-setting equation becomes

$$p_t^i = E_t^i(p_t) + (1-r) E_t^i(y_t - \overline{y}_t) \tag{3.25}$$

The model can be solved using the same techniques as in section 2. Averaging (3.25), then substituting forward to express p_t and y_t as functions of the exogenous precesses θ_t and $\overline{y}_t$ yields

$$p_t = (1-r)\sum_{s=0}^{\infty} r^s \left[\overline{E}_t^{(s+1)}(\theta_t) - \overline{E}_t^{(s+1)}(\overline{y}_t)\right] \tag{3.26}$$

$$y_t = (1-r)\sum_{s=0}^{\infty} r^s \left[\theta_t - \overline{E}_t^{(s+1)}(\theta_t)\right] + (1-r)\sum_{s=0}^{\infty} r^s \overline{E}_t^{(s+1)}(\overline{y}_t) \tag{3.27}$$

To solve (3.26) and (3.27) explicitly for impulse responses to monetary, real and informational shocks, on can again proceed along the same lines as section 2.

To interpret (3.26) and (3.27), note that $\theta_t - \overline{y}_t$ is the market-clearing price level that would prevail under common knowledge of the two processes. The realized price level p_t is therefore a *weighted average of average higher-order expectations of the full-information market-clearing price.* y_t can be decomposed into a component due to monetary shocks and a weighted average of higher-order expectations concerning the potential output level. The first component reflects the incomplete nominal adjustment. The second term is new, and is due to fluctuations in potential GDP that do not become common knowledge. Because of higher-order uncertainty, real GDP responds only sluggishly to variations in potential GDP. Moreover, the stronger the coordination motive, the stronger the delays in adjustment. The analysis thus suggests that higher-order uncertainty may possibly account for the persistence of shocks other than monetary shocks; such as technology shocks. In the extreme case, where r is close to 1 and factor-specific higher-order uncertainty is important, output potentially remains far from the potential level for a long time.

Since we can separate the effects of θ_t and $\overline{y}_t$, assuming an information structure as above for each of the two components leads to identical impulse response functions as previously.[6] Real disturbances have the same, but opposite effect on prices as they have on output, however, the output gap $y_t - \overline{y}_t$ responds to real shocks as it does to nominal disturbances. *Most importantly, the factor-specific degree of higher-order uncertainty determines to what extent a disturbance has persistent effects on prices and output.* I conclude this discussion by observing that, although higher-order uncertainty may well explain the *persistence* of shocks, it is unable to account for *amplification*, at least in the short run: Since higher-order expectations are much more sluggish in responding to current information than first-order expectations are, the effects of a supply shock on output are dampened rather than amplified.

3.5 Some Preliminary Thoughts on Monetary Policy

The previous discussion has highlighted the impact of the information structure on the output-inflation trade-off. In this section, I study the welfare effects of information provision by the Central Bank. A recurring theme in the scientific debate about the optimal conduct of monetary policy is the role of Central Bank transparency, i.e. the desirability of supplying the private sector with precise information about (*i*) the objectives of monetary policy, (*ii*) the macroeconomic data on which the central bank bases its decisions, and (*iii*) the actions taken by the central bank. The main argument in favor of monetary transparency is based on the need to *monitor*: The better the information provided, the easier it is to evaluate the Central Bank's behavior ex post and analyze whether the policy objectives have been met. Transparency is thus necessary to monitor whether or not the Central Bank adheres to implicit or explicit rules that govern the Principal Agent relationship between the society and the central bank. This

[6] More generally, if the processes for θ_t and $\overline{y}_t$ are correlated (for example when the monetary authority in part responds to its own estimate of the output gap), the techniques of lemma 1 and 2 can be used to derive higher-order expectations about the full-information market-clearing price.

is emphasized in particular in the context of inflation targeting, where an independent central bank retains full authority over its policy actions, and is held accountable for them.[7]

In this section, I discuss the effects of transparent monetary policy in the present model. Some of the information that price-setting firms have access to comes from the central bank, and hence, an immediate implication of the previous results is that the information the central bank provides to the public not only has a role for monitoring purposes, but also has a direct influence on coordination among price-setters, and hence the inflation-output trade-off. This *coordination effect* of information provision is the focus of this section. Formally, I embed the present model into a simple model of monetary policy a la Barro and Gordon [2], taking the aggregate demand process as in part under the central bank's control. The analysis is far from being exhaustive, the main purpose being that of showing that by augmenting the previous model by a formalization of the central bank's objective, we have a natural framework in which to formally study the role the interplay between monetary transparency and the market's information structure.

To fully understand the role of transparency, it is important to tie this notion to parameters of the information structure. Here, I adopt the view that a higher degree of transparency means a higher degree of "common knowledge"; within the model, this is achieved by a better provision of public information. This view is also taken by Morris and Shin [11]. The motivation behind this use of the term "transparency" lies in the fact that public information coordinates the market's first- and higher-order expectations, whereas private information doesn't (see Hellwig [7] for results linking the information structure to higher-order uncertainty).

3.5.1 Transparency as an Implicit Commitment Device

In this section, I study the role of transparency in an environment, where the central bank cannot commit to a particular policy rule. The monitoring role of transparency in the commitment literature was already highlighted; here I show that in an environment characterized by discretion, transparency can act as an implicit commitment device: Even if the central bank

[7] see Bernanke and Mishkin [3] and Svensson [15] for discussions of inflation targeting that emphasize the monitoring role of transparency.

is discretionary and tries to use monetary policy to stimulate output, the provision of precise public information decreases the effectiveness of pro-active monetary policy, and hence serves as an implicit commitment device that reduces the inflationary bias. However, improved public information also comes at a cost, since a reduction of the variance of public information raises the macro-economic exposure to informational shocks, at least in environments with a high degree of informational noise; and the higher the pressure put on the central bank to stimulate output, the more it has an incentive to commit to the provision of precise public information.

To illustrate this point formally, consider the previous model, with $T = 1$ for simplicity. Suppose that a discretionary central bank controls the growth rate of nominal GDP with some noise; i.e. θ_t satisfies

$$\theta_t = \theta_{t-1} + \mu_t + \varepsilon_t \tag{3.28}$$

where μ_t is under the central bank's control, and ε_t is a monetary policy shock. In each period, the central bank has a target for the growth of θ and for real aggregate output, i.e. it minimizes with respect to $\{\mu_t\}_0^\infty$ the loss function

$$\sum_{t=0}^{\infty} \beta^t L(\mu_t)$$

where $\beta < 1$ is the discount rate, and the per period loss function is

$$L(\mu_t) = E_t (\theta_t - \theta_{t-1})^2 + bE_t (y_t - y^*)^2,$$

taking as given the private expectations about the present and future conduct of monetary policy. The output target y^* may be different from the potential output level, which is normalized to 0. The use of a money growth target instead of an inflation target in the loss function is made for convenience, it eliminates any dynamic effects resulting from the choice of μ_t. An inflation target would lead to the same results, but is technically more involved since the effects of discretionary monetary policy on inflation are spread over two periods: immediately through unexpected inflation in the current period, and once the information about the growth of nominal GDP is commonly available (i.e. the following period), through higher anticipated inflation. Here, I abstract from this additional complication.

Price-setters form expectations μ_t^e about the central bank's course of action, and set prices according to the noisy information available about the realization of θ_t. Going along the same lines as before, the output gap then satisfies

$$y_t = \frac{\gamma}{1+\gamma}\sigma\left(\mu_t - \mu_t^e + \varepsilon_t\right) - \frac{\sigma^2}{\sigma_v}\frac{\gamma}{1+\gamma}v_t \tag{3.29}$$

As a function of μ_t, the loss function can thus be rewritten as

$$\begin{aligned} L\left(\mu_t\right) &= \mu_t^2 + \sigma^2 + \frac{b\gamma^2}{(1+\gamma)^2}\left[(\mu_t - \mu_t^e)^2 + \sigma^2 + \frac{\sigma^4}{\sigma_v^2}\right] \\ &\quad - \frac{2b\gamma}{1+\gamma}\left(\mu_t - \mu_t^e\right)y^* + by^{*2} \end{aligned} \tag{3.30}$$

The first-order condition with respect to μ_t is

$$\mu_t + \frac{b\gamma^2}{(1+\gamma)^2}\left(\mu_t - \mu_t^e\right) - \frac{b\gamma}{1+\gamma}y^* = 0 \tag{3.31}$$

which, together with the rational expectations hypothesis that $\mu_t^e = \mu_t$, yields

$$\mu_t = \mu_t^e = \frac{b\gamma}{1+\gamma}y^* \tag{3.32}$$

Substituting into $L\left(\mu_t\right)$ yields an expression for the ex ante expected loss:

$$\begin{aligned} EL &= \left[1 + \frac{b\gamma^2}{(1+\gamma)^2}\right]by^{*2} \\ &\quad + \left[1 + \frac{b\gamma^2}{(1+\gamma)^2}\right]\sigma^2 \\ &\quad + \frac{b\gamma^2}{(1+\gamma)^2}\frac{\sigma^4}{\sigma_v^2} \end{aligned} \tag{3.33}$$

Hence, the inflationary bias is increasing in γ, the degree of higher-order uncertainty about θ_t. The expected loss can be decomposed into three components: The first measures the cost due to the inflationary bias, the second measures the cost due to monetary shocks, and the last component is due to informational noise. We observe that the first two components are decreasing in γ (and hence in σ_v^2), but changes in σ_v^2 have ambiguous effects on the loss due to informational noise: If σ_v^2 is large, the increase of exposure to informational noise that results from a reduction in σ_v^2 dominates the reduction of higher-order uncertainty, while for small values of σ_v^2, the opposite is true. Maximizing EL with respect to σ_v^2 then necessarily leads to a corner solution: Ideally, the central bank would want to set σ_v^2 as close to zero as possible; i.e. provide very precise public information (be very transparent): this eliminates higher-order uncertainty, and therefore the effects of monetary shocks on output as well as the inflationary bias. In practice, it may well be impossible to pursue such an information policy; rather, there is some lower bound on σ_v^2. In that case, it may be optimal to provide no public information at all, to insure the market against informational risk, while at the same time accepting more higher-order uncertainty. At what lower bound on σ_v^2 it becomes optimal to commit to transparent provision of information depends on the output target: The more biased the output target is, the more the central bank has an incentive to provide precise public information.

3.5.2 Policy Targets and the Information Structure

We can use the insights of this model to discuss the impact of the monetary regime on the information structure, and hence the output-inflation trade-off. The literature on monitoring monetary policy usually advocates that policy should be conducted within a framework that provides well-specified targets. The beneficial effects of such a framework and of policy targets with respect to the information structure are easily understood: The framework, as well as the targets, reduce higher-order uncertainty about the central bank's objectives, and hence about the course of its policy conduct. Formally, specific targets eliminate higher-order uncertainty about the values of y^* and b in the central banker's objective function, and hence reduce higher-order uncertainty about the resulting policy variable θ_t.

In addition, the targets themselves act to coordinate expectations about the targeted variables; within the model, they act as a public signal about policy. A similar role of coordinating expectations is played by published forecasts. Svensson's interpretation of inflation targets as *inflation forecast* targets (i.e. the central bank should design its policy so that its forecast of inflation is consistent with the target, cf. Svensson [15]) captures precisely this idea: to the extent that forecasts are unbiased, they act as a public signal about inflation. Monetary regimes typically differ about what variable is targeted, and hence also about the degree of higher-order uncertainty about prices. In the terminology of our model, a regime that targets money growth reduces higher-order uncertainty about θ_t, an inflation target affects higher-order uncertainty about prices. As was observed in (3.26), the inflation-output trade-off depends on higher-order uncertainty about the "full-information market-clearing price-level"; since the inflation target provides a public signal about the latter, one would conjecture that an inflation target is more beneficial in terms of its informational effect than a money growth target.

3.5.3 The Signaling Role of Monetary Policy

Finally, the model points to the role that central bank transparency may have in stabilizing output following supply shocks: The analysis in the previous section has highlighted the possibility that in an environment characterized by a high degree of strategic complementarities, adverse supply shocks can have highly persistent output effects, if there is higher-order uncertainty; in other words, even if everyone privately believes that potential output is higher, output remains depressed because of low *higher-order* expectations. In this case, transparency about supply shocks may also be beneficial, since the creation of common expectations about output reduces the persistence of the effects of real shocks.

The provision of public information about supply shocks can come through two channels: First, such information may come from the provision of public forecasts of potential output or the output gap. Second, even if such forecasts are not public, some information becomes available, if the central bank conditions monetary policy on its own estimate of the supply shock. If higher-order uncertainty is small, the "surprise" effect of monetary policy on output is small; nevertheless, the central bank may want to condition its policy on its estimates of potential output, if this increases welfare by coordinating expectations about potential output.[8]

[8] This idea mirrors results in Angeletos, Hellwig and Pavan [1], who study the informational role of policy choices in a global coordination game with multiple equilibria under common knowledge. In their environment, the information conveyed by the policy choice enables the market to coordinate on one of multiple equilibria, and this multiplicity is the root cause of the policy traps discussed in that paper. Here, the preferences of the central bank are aligned with those of the price-setters, and hence inducing better coordination will be beneficial from the central bank's point of view.

3.6 Concluding Remarks

Building on Woodford [16], this paper has developed a model of monetary business cycles, in which higher-order uncertainty about the fundamental driving processes, coupled with strategic complementarities between price-setters leads to potentially long adjustment delays for prices after monetary shocks and hence to important short-run effects on output. The main motivation of the analysis was to discuss the effect of the information structure on the inflation-output trade-off; specifically, the model showed how price stickiness becomes more important, the more relevant higher-order uncertainty is (quantifiable as a function of the information structure and of the degree of strategic complementarities). The provision of precise public information thus accelerates nominal adjustment, but potentially comes at the cost of a higher exposure to informational noise. Finally, embedding the nominal adjustment model in a simple version of the monetary policy model a la Barro and Gordon [2] shows how the provision of precise public information may serve as an implicit commitment device against inflationary biases.

The results in this paper have several positive, normative, and empirical implications. The comparative statics results rely more on the context of strategic complementarities with incomplete information than on the specific environment studied. The conclusions about the role of the information structure therefore should be expected to extend to other contexts of decision-making with strategic complementarities. Nor are the effects of higher-order uncertainty tied to the nature of the shock; the model suggests that higher-order uncertainty can lead to persistence of *any kind of shock.* Woodford's insight about monetary shocks may therefore be helpful in understanding persistence in other environments as well, for example of technology shocks in an RBC model.

Among the normative aspects of the analysis, I have highlighted the role of transparency as providing commitment against inflationary biases; I should also mention the negative impact that the degree of competition has on higher-order uncertainty: The higher the elasticity of substitution between products is, the more important the strategic complementarities between prices are, and the slower the nominal adjustment after a monetary shock. The positive effects of increasing competition are therefore in part offset by a negative welfare effect due to higher persistence of shocks.

The paper furthermore leads to potentially testable empirical implications: In particular, it points to the effects of the monetary policy regime on the information structure, and raises the question whether there is an empirical link between the way monetary policy is conducted and the inflation-output trade-off. The informational parameters can potentially be estimated from time-series data, which leads to the question whether the changes in information processing over the last 20 years have changed the inflation-output trade-off, or have otherwise altered the transmission channels of monetary policy. There seems to be at least informal evidence about how the underlying parameters have shifted: Morris and Shin [11], for example suggest that changes in the use of the media have reduced informational noise, but have also raised the public information component in financial markets (lowered α), while the conventional wisdom on product market liberalization suggests an increase in r.[9] Furthermore, evidence in Stock and Watson [14] suggests that the variance of underlying shocks has decreased; what remains unclear is how these reductions in the variance of shocks compare to the reduction in informational noise, i.e. the ratio $\frac{\sigma^2}{\Sigma}$. How these changes have altered the inflation-output trade-off and the exposure to informational noise remains a priori ambiguous.

Estimating the informational parameters (and possibly testing for their changes) is also of interest with respect to the recent debate on the decline of US output volatility (see Blanchard and Simon, [4], or Stock and Watson, [14], for an overview). These papers try to determine to what extent changes in the conduct of monetary policy have contributed to stabilize output growth, and the theoretical contribution of this paper suggests that the monetary policy changes of the 1980's, to the extent that they have altered the information structure, have influenced output volatility not only through the redefinition of policy objectives, but also through their direct effect on the structural parameters.

It should be noted that, as a descriptive model of monetary business cycles, the present model is highly simplified, and relies on information as the unique transmission channel for shocks. The analysis relies on informational assumptions, which, although more complex than

[9] In his account of Greenspan's activity at the Fed, Woodward [18] suggests based on discussions with the Fed chairman that starting in the mid-90's, increased competition prevented firms from responding to a loosening of monetary policy by raising prices. While I haven't found any similar evidence in the academic literature, the statement seems to be broadly consistent with the view that an increase in r led to an increase in γ during the 90's.

Woodford's, are very simplistic. The methodology, however, can easily be adapted to more complex information structures that involve a gradual learning of the process, or a shift from private to public information over time. The model relied on (i) the linearity of best responses, which by forward substitution led to an expression of strategic variables as weighted sums of higher-order expectations about the underlying fundamental processes, and (ii) on the derivation of average expectations out of the information structure. Average expectations can easily be computed for any kind of environment, by first deriving a filtering equation like (3.8) for the fundamental process, and then using the filtering equation to relate average expectations about the fundamentals to the fundamental process itself; iteration to higher orders then completes the procedure.

Due to its flexibility, the present model of higher-order uncertainty might therefore be useful as a vehicle for studying the role of the information structure in various other dynamic contexts, starting with a more exhaustive analysis of the role of transparency in the conduct of monetary policy. Another empirically appealing extension might be to combine the analysis of incomplete information with sticky prices a la Calvo [5]. As was discussed before, the incomplete information model can replicate the finding of VAR estimations that following a monetary shock, output peaks prior to inflation; however the model cannot account for persistent effects on output and inflation beyond the point at which the initial shock becomes common knowledge. Combining the incomplete information with some forms of price or investment rigidities might therefore lead to a further increase of the persistence of inflation. A combination of incomplete information with rigid price adjustment might also be helpful for a theoretical understanding of the insights drawn from the "new Keynesian" models, where the forward-looking nature of pricing decisions relies on a strong inter-temporal coordination of expectations. Whether such coordination of expectations remains feasible in the presence of informational differences is a yet unresolved question.

3.7 Bibliography

[1] G.-M. Angeletos, C. Hellwig, and A. Pavan, "Coordination and Policy Traps in Currency Crises", MIT and Northwestern, 2002.

[2] R. Barro and D. Gordon, "A Positive Theory of Monetary Policy in a Natural Rate Model", *J. Polit. Econ.* **91** (1983), 589-610.

[3] B. Bernanke and F. Mishkin, "Inflation Targeting: A new Framework for Monetary Policy", *J. Econ. Perspect.* **11** (1997), 97-116.

[4] O. Blanchard and J. Simon, "The Long and Large Decline in U.S. Output Volatility", MIT, 2000.

[5] G. Calvo, "Staggered Prices in a Utility Maximizing Framework", *J. Monet. Econ.* **12** (1983), 383-398.

[6] L. Christiano, M. Eichenbaum and C. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy", *NBER working paper* **8403** (2001).

[7] C. Hellwig, "Public Information, Private Information, and the Multiplicity of Equilibria in Coordination Games", *J. Econ. Theory*, forthcoming.

[8] F. Kydland and E. Prescott, "Rules rather than Discretion: the Inconsistency of Optimal Plans", *J. Polit. Econ.* **85** (1977), 473-491.

[9] R. Lucas, Jr., "Expectations and the Neutrality of Money", *J. Econ. Theory*, **4** (1972), 103-124.

[10] S. Morris and H. S. Shin, "Global Games: Theory and Applications", *in* "Advances in Economics and Econometrics, the Eighth World Congress" (M. Dewatripont, L. Hansen, and S. Turnovsky, Eds.), Cambridge University Press, forthcoming.

[11] S. Morris and H. S. Shin, "Social Value of Public Information", *Amer. Econ. Rev.*, forthcoming.

[12] E. Phelps, "Introduction: The New Microeconomics in Employment and Inflation Theory", *in* "Microeconomic Foundations of Employment and Inflation Theory", New York: Norton, 1970.

[13] C. Sims, "Rational Inattention", Princeton University, 2001.

[14] J. Stock and M. Watson, "Has the Business Cycle Changed and Why?", Princeton University, 2002.

[15] L. Svensson, "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets", *Europ. Econ. Rev.* **41** (1997), 1111-1146.

[16] M. Woodford, "Imperfect Common Knowledge and the Effects of Monetary Policy",

NBER working paper **8673** (2001).

[17] M. Woodford, "Optimizing Models with Nominal Rigidities", Princeton University, 2001.

[18] R. Woodward, "Maestro: Greenspan's Fed and the American Boom", New York: Simon and Schuster, 2000.