ESSAYS ON MONEY AND INFORMATION

Luis Araujo

A DISSERTATION

in

ECONOMICS

Presented to the Faculties of the University of Pennsylvania in Partial

Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2002

Randall Wright

Supervisor of Dissertation

Steven Matthews

Graduate Group Chairperson

To my father,
Hevelton Marcelino Araujo.

ACKNOWLEDGMENTS

I would like to thank my advisor Randall Wright for his encouragement from the very beginning to the last stages of my work in this dissertation. I am also very grateful to the members of my committee, George Mailath and Jan Eeckhout. Their help was of great importance for the development of my work.

I have a great debt with my friends and colleagues in the Department of Economics at UPenn. In particular, I would like to thank Braz Camargo, my co-author in the second chapter of this dissertation. I also would like to thank Giancarlo Gasha, Makoto Hanazono, Ping He, Tetsuya Maruyama, Brett Norwood, Sergio Parreiras, Samuel Pessoa, Silvia Pessoa, Dan Silverman, Yasuo Terajima and Clara Vega. Especially, I am very grateful with Brandon Weber. Our friendship started in the first year of my graduate studies at UPenn. I am sure it will last for many years to come.

I enjoyed very much the company of the Brazilian friends I made throughout my years in Philadelphia: Eduardo and Katia, Samuel and Silvia, Marcelo and Poliana, Gidi and Ana Claudia, Bruno and Simone, Fernando and Rucelia, Andre and Juciene, Braz, Eduardo, Francisco and Sergio.

I am very grateful to my brothers Alexandre and Junior, and to my mother, Lina Araujo.

The final years of my graduate studies I shared with the company of Christiane. Her love and patience were and are essential in my life.

To Capes and the University of Pennsylvania I am thankful for all the financial support during my years of study.

ABSTRACT

ESSAYS ON MONEY AND INFORMATION

Luis Araujo

Randall Wright

Social Norms and Money

In an economy where there is no double coincidence of wants and no record-keeping of past transactions, money is usually seen as the only mechanism that can support exchange. In this paper we show that, as long as the population is finite and agents are sufficiently patient, a social norm establishing gift-exchange can substitute for money. However, for a given discount factor, population growth leads to the breakdown of the social norm. Additionally, increases in the degree of specialization in the economy can also eventually undermine the social norm equilibrium. By contrast, a monetary equilibrium exists independent of the population size or specialization. We conclude that, while social norms can support efficient exchange in economies with small populations and limited specialization, money is essential in economies without these features.

Monetary Equilibrium with Decentralized Trade and Learning

We consider an environment where trade is decentralized and agents only obtain information about the state of the economy (the monetary regime) through private histories. We characterize the dynamics of the market under different regimes and show that changes in the degree of information transmission in the economy reduce the bank's willingness to choose a soft monetary regime. We also study the intertemporal consistency of the bank's behavior. We show that in order for patient banks to choose a tight regime they must not only be concerned with building a good reputation, but also with maintaining their good reputation by behaving in such a way as to differentiate themselves from impatient banks.

Contents

Chapter 1

Social Norms and Money

1.1 Introduction

The reasons for the existence of money and how it emerged in the economy are fundamental questions in monetary theory but only recently have they been addressed in a systematic way. In Kiyotaki and Wright (KW) (1989) and the literature that followed, money exists to play the role of a medium of exchange. However, there are other mechanisms (or institutions) in society that can play the same role. Kocherlakota (1998a ,1998b) shows that every allocation that can be attained with money can be attained with memory as long as there is perfect and public record-keeping of past transactions. We will show that a social norm establishing gift-exchange can also be a substitute for money in some situations.[1]

Memory is excluded, by construction, from the environment considered by KW. Despite departing from the Walrasian economy in the sense that they make the exchange process non-trivial, they still preserve the anonymous market as the locus of trade. Hence since agents are indistinguishable from one another at all times, there is no scope for exploiting the history of past transactions. This is why there is an interesting role for money. Our objective is not to reject the idea that money actually plays the role of a medium of exchange.

[1]The reference paper on norms, though not related to monetary theory, is Kandori (1992). See also Ellison (1994) and Okuno-Fujiwara and Postlewaite (1995). Moreover, a recent paper by Corbae et alli (2001) studies how equilibrium with money, gift-exchange and memory can emerge in an environment with endogenous matching.

This is an obvious feature of modern societies and we believe that the KW environment captures in a natural and simple way the fact that money helps to overcome frictions in the exchange process. Our concern is the study of the conditions under which alternative ways of addressing the exchange problem may or may not be viable. By considering this issue, we can have better insight into the reasons money emerged as the dominant mechanism. The main result in our paper establishes that, while exchange can be supported by a social norm in small societies, money has an essential role[2] when the population is large.

The outline of this chapter is as follows. In the next section, we will first discuss the conditions under which memory can be used in the KW environment when there is perfect and public record-keeping. However, our focus is the analysis of an equilibrium with gift-exchange without record-keeping. The main result of section 1.2 establishes that money is not essential in small populations. In section 1.3 we discuss how population growth and the degree of specialization in the economy affects this outcome. We show that when population grows the social norm equilibrium eventually breaks down. Moreover, an increment in the degree of specialization has a similar effect. In section 1.4 we show the existence of a monetary equilibrium in the KW environment with a finite number of agents (previous analysis of the model always assume a continuum of agents). We prove that money can be valued as medium of exchange independent of the size of the population. In section 1.5 we conclude with a discussion of our results.

---

[2]Following the literature, we say that money is essential if exists desirable allocations that can be achieved with money and not without money.

1.2 The Model

1.2.1 Memory

Consider an economy with a continuum of infinitely-lived agents of k different types, and with k distinct indivisible goods.[3] Following KW (1989) assume that a type j agent only consumes good j, which yields utility u, and produces good $j+1$ (modulo k) at a cost of c (with $u > c$). Types are recognized. At every date, agents enter the exchange sector and are pairwise matched under a uniform random matching technology; i.e., there is a probability $\frac{1}{k}$ that agent type j will meet, for example, an agent of type m. Finally, agents discount the future with a discount factor β. Note that, whenever $k > 2$, there is no double-coincidence of wants meeting in this economy.

Although we are mainly interested in environments without memory, we begin with the case of perfect memory as a benchmark for what follows. Suppose that this economy has a record-keeping device that keeps track of all its transactions, with this device being common knowledge across agents. In other words, every agent in the economy can observe the history of transactions of every other agent. Under these assumptions consider an exchange rule stating the following (denote it rule 1): every time an agent meets someone who likes his good, he produces the good for him, as long as everyone has done so in the past. If any agent deviates, the economy reverts to autarky.[4] Note that autarky is always an equilibrium in this economy. But we can show that, for some parameter values, the above rule is an equilibrium that makes every agent better off than in autarky. Clearly, this rule is efficient.

Lemma 1.1. *If $\beta \geq \frac{ck}{ck+u-c}$, the strategy profile associated with rule 1 constitutes a subgame*

[3] The assumption of indivisible goods is made only for simplicity, similar results hold in the divisible goods case.

[4] A more formal way of stating this rule is: If agent type j meets an agent type $j+1$(modulo k), he will produce one unit of the good and give it to $j+1$. If he meets an agent type $j-1$(modulo k), the latter will produce one unit of the good and give it to j. In any other situation, there is no exchange at all. Finally, if any agent deviates from this rule the economy reverts to autarky forever.

perfect Nash equilibrium.

Proof. Let V_c be the expected payoff of an agent on the equilibrium path and let $V_a = 0$ be the value function in autarky. We need to determine if agent j produces for agent $j+1$ (modulo k) when they meet. Since the environment is stationary, we only need to check a one-shot deviation. If he produces the good he receives $-c + \beta V_c$, and if he does not produce he receives $V_a = 0$. Thus to follow rule 1 is optimal whenever $-c + \beta V_c \geq 0$. One can calculate:

$$V_c = \frac{1}{k}(u + \beta V_c) + \frac{1}{k}(-c + \beta V_c) + \frac{k-2}{k}(0 + \beta V_c) \Longleftrightarrow V_c = \frac{1}{(1-\beta)} \frac{(u-c)}{k}$$

Hence, rule 1 is a subgame perfect Nash equilibrium if and only if:

$$\beta \geq \frac{ck}{ck + u - c}$$

For any history off the equilibrium path, since the economy reverted to autarky, each agent receives zero utility. If he deviates and produces a good he will obtain $-c$. Hence, there are no deviations off the equilibrium path. □

A natural question that arises from the previous lemma is whether it is reasonable to imagine such a stringent rule. After all, if society has full information regarding past transactions, why not retaliate against only the deviating agent, instead of the whole community? With this in mind, the following result shows that the same exchange pattern can be generated with a more reasonable rule (rule 2): every time an agent meets another who likes his good he produces the good for him, as long as everyone has done so in the past. If any agent deviates, he doesn't receive or give goods anymore.

Lemma 1.2. *If* $\beta \geq \frac{ck}{ck+u-c}$, *the strategy profile associated with rule 2 constitutes a subgame perfect Nash equilibrium.*

Proof. The proof is a straightforward modification of lemma 1.1 and hence is omitted. □

We can also modify the rule that implements this efficient outcome by allowing the possibility of forgiveness, i.e., that punishment occurs for only a finite period of time. The rule (rule 3) in this case is: every time an agent meets another who likes his good he produces the good for him, as long as everyone has done so in the past. If any agent deviates, he doesn't receive or give goods during T periods. A similar result obtains:[5]

Lemma 1.3. *For every* $T \geq 1$, *if* $\beta(1-\beta^T) \geq \frac{ck}{ck+u-c}$, *the strategy profile associated with rule 3 constitutes a subgame perfect Nash equilibrium.*

Proof. Now, if the agent deviates he receives $0+\beta V_f$, where V_f indicates the life-time utility after deviating under the punishment and forgiving norm. We have that $V_f = 0 + \beta^T V_c$. Hence, he will follow rule 3 whenever $\beta(1-\beta^T) \geq \frac{ck}{ck+u-c}$. Clearly, there will be no deviations off the equilibrium path. □

The objective of the analysis so far is to introduce the idea that equilibrium strategies can be seen as the product of rules in a society, and the three lemmas above try to capture the diversity of rules that can arise. Our results are basically the same as those obtained by Kocherlakota (1998a). The problem with this approach is that the gift-exchange equilibrium depends upon perfect and public record-keeping of past transactions, which is a very strong

[5]Note that the minimum discount factor needed to guarantee cooperation in lemma 3 is bigger than in lemmas 1 and 2, since now the cost of a deviation is smaller. Of course, when T goes to infinity, the range of β supporting rule 3 coincides with the range of β supporting rules 1 and 2.

assumption. In the next model we show that record-keeping is not needed, as long as population is finite and agents are sufficiently patient.

1.2.2 Social Norms

In what follows we are going to use the notion of a contagious equilibrium developed by Kandori (1992). Under a uniform random matching technology, Kandori shows that cooperation can be obtained in a one-shot prisoner's dilemma with a social norm establishing that: "....a single defection by a member means the end of the whole community trust, and a player who sees dishonest behavior starts cheating all of his opponents. As a result, defection spreads like an epidemic and cooperation in the whole community breaks down." (1992, page 69). The contagious defection differs from the rules of punishment previously considered. First, without public record-keeping of transactions, agents only realize that a deviation happened when they meet a deviant agent. Moreover, punishment cannot be personalized. Upon seeing a defection, an agent punishes the whole community, and this is the reason that such a pattern can be well represented by the idea of a social norm.

Consider the standard KW environment with one change: there is a finite number of agents (but as many as we wish). Moreover, from now on we are going to assume that there is no public record-keeping of past transactions in the economy, each agent can only observe his own private history. We consider first a special case where each agent in the economy is of a different type. In the next section we relax this assumption. Consider the following social norm: "every time an agent meets another who likes his good he produces the good, as long as everyone has done so in the past *for him.* If at a meeting an agent can produce a good that the other agent likes but fails to do so, neither agent will ever produce again."[6]

[6] More formally, we have: If agent j meets agent $j+1$(module N), he will produce one unit of the good and give it to $j+1$. If he meets agent $j-1$(module N), the latter will produce one unit of the good and give it to j. In any other situation, there is no exchange at all. Finally, if at a meeting an agent can produce a good that the other agent likes but fails to do so, neither agent will ever produce again.

Proposition 1.1. *For every N, there exists a discount factor β' (depending on N) such that, for every $\beta \geq \beta'$, the strategy profile corresponding to the social norm constitutes a sequential equilibrium.*

Proof. Let V_n be the expected payoff on the equilibrium path and let V_b be the expected payoff of an agent who triggers the contagious punishment. We need to determine if an agent produces a commodity when he meets another that likes his good. If he follows the social norm and produces the good he receives $-c + \beta V_n$. Otherwise he receives $0 + V_b$. He will not deviate whenever $\beta V_n - V_b \geq c$. V_b is calculated in appendix A. Since the results are similar in flavor for N odd or even, we can focus on the case where N is even. We obtain:

$$V_n = \frac{1}{N-1}(u - c + 2\beta V_n) + \frac{N-3}{N-1}(\beta V_n) \Longleftrightarrow V_n = \frac{u-c}{(1-\beta)(N-1)}$$

$$V_b = \sum_{t=1}^{\infty} \beta^t P(T > t) \frac{u}{(N-1)},$$

where $P(T > t)$ indicates the probability that the agent who deviated first is not reached by the contagious process in period t. An agent will not deviate whenever:

$$\beta V_n - V_b \geq c \Longleftrightarrow \beta \frac{u-c}{(N-1)(1-\beta)} - \sum_{t=1}^{\infty} \frac{\beta^t P(T > t) u}{(N-1)} \geq c \tag{1.1}$$

When t goes to infinity $P(T \leq t)$ goes to one (see appendix A). Hence, as β converges to 1 the expression on the left goes to infinity. So, there exists a discount factor β' (the value of β that solves the equation $\beta \frac{u-c}{(N-1)(1-\beta)} - \sum_{t=1}^{\infty} \frac{\beta^t P(T>t) u}{(N-1)} = c$) such that, for all $\beta \geq \beta'$, the above inequality is satisfied, implying that our strategy profile is a Nash equilibrium. It remains to prove that this equilibrium is sequential. Under the assumption that each agent is of a different type, this is straightforward. First, even if the agent who originally deviated wanted to slow down the contagious process, he can't because he already deviated with the only agent he can affect (which is the one that likes his good). Second, any other agent, after being contaminated by the contagious process can obtain at most 0 utility since there are no more goods available to him in this economy. Hence, he will never produce. As a result, after any possible history off the equilibrium path, the agents have incentive to follow the prescription of the social norm and transmit the punishment. □

We reiterate, the history of economy-wide transactions is not used to implement the social norm, only the personal history. However, since it takes more time to punish a deviant in this environment, as compared to one with memory, the discount factor needed to support the equilibrium is higher.

Proposition 1.1 establishes conditions under which a social norm can implement gift-exchange. However, it can also determine when this norm cannot be implemented. For a fixed β, the value of N is the key factor to be considered. Remember that, under our assumption that each agent is of a different type, N indicates at the same time population size and degree of specialization (as measured by the number of types) in the economy. In the next section we will show that population size is the key variable conducive to the breakdown of the social norm. When the population increases, defection requires a longer time to contaminate all agents. If the discount factor does not change, we reach a point where agents no longer have an incentive to follow the social norm and the equilibrium breaks down.

Nonetheless, we will show that specialization also affects the social norm equilibrium. The intuition is the following: suppose the economy is such that there are k types of agents, with $\frac{N}{k}$ agents per type (N being even or odd depending on the value of k). In this case, an increment in the value of k leads to a reduction in the future payoff, both on and off the equilibrium path. This increment has a similar effect as a reduction in the value of β, which creates more incentives for an agent to deviate in order to avoid the cost of cooperating today. An increment in the value of k also affects the speed of the contagion process. Since interactions which involve exchange become less frequent when k increases, the contagion process tends to be slower as compared to an economy where k is small.

1.3 Specialization

Up to this point we have been using as a baseline the KW (1989) model, where the description of an agent type embodies both a specification of preferences over goods and a production technology. Now (and in the next section), we consider a slightly change in the environment that allows the separation between the effect of population size and degree of specialization over the social norm equilibrium.

In particular, consider an economy with N agents (for simplicity, assume N is even) and K distinct indivisible goods. As in the previous section, agents can only observe their private histories. Each agent derives utility u from consumption of k goods, although which goods vary from agent to agent. At every meeting each agent produces one unit of a commodity, drawn randomly from the set of all commodities. An agent does not consume his own good. Moreover, we assume that there is no double-coincidence of wants meeting.[7] Agents are meet pairwise under a uniform random matching technology, and the distinct meetings that can come out of each match can be summarized in the following way:

	i likes $j's$ good	i does not like $j's$ good
j likes $i's$ good	0	σ
j does not like $i's$ good	σ	$1-2\sigma$

Single-coincidence meetings happen with probability σ, and it captures the extent to which commodities and tastes are differentiated. We will use σ as a measure of the degree

[7] This assumption restricts the production technology. It means that, in the joint distribution from which goods are drawn at every meeting, there is zero probability of an extraction corresponding to a double-coincidence of wants. We are using this assumption because it makes easier the calculation of payoffs after a deviation. Moreover, it keeps an structure similar to the model analyzed before.

of specialization in the economy. If the economy is very specialized, σ is low, i.e., there are many goods but agents only like a small fraction of them. We can also motivate the use of σ in the following way: suppose agents specialize in production, for example, agents that used to produce ten distinct goods now produce only one good. This implies that the probability of meetings where an agent likes the goods produced by the other goes down, and this is exactly what happens when σ decreases.

Consider the following social norm, similar to the norm in the previous section: "If an agent meets another who likes his good he produces the good as long as everyone has done so in the past *for him.* If at a meeting an agent can produce a good that the other agent likes but fails to do so, neither agent will ever produce again."

We can use a reasoning very similar to the one in Proposition 1.1 to prove that the strategy profile associated with this social norm constitutes a Nash Equilibrium. After a deviation from the equilibrium path, the finiteness of the population guarantees that eventually every agent will not cooperate with probability one. Hence, for a sufficiently patient agent, there is no incentive to deviate from the equilibrium path.

Formally, if an agent follows the equilibrium he obtains:

$$V_c = \frac{\sigma(u-c)}{(1-\beta)}$$

If an agent deviates from the equilibrium, he obtains (the calculation of V_d is done in appendix B):[8]

$$V_d = \sum_{t=1}^{\infty} \beta^t e_1 A^t \pi \sigma u$$

[8]To calculate V_d, we calculate first a closed form expression for the contagion process in this environment. It turns out that the process analyzed by Kandori (1992) can be seen as a particular case of the one analyzed here.

where:

$$\begin{aligned}
e_1 &= (1, 0, 0, ..., 0), \text{ N-dimensional.} \\
A &= \{(a_{ij})\}_{NxN}, \text{ where } (a_{ij}) = \Pr(n_{t+1} = j \mid n_t = i) \\
\pi &= (\pi_i), \text{ N-dimensional, where } \pi_i = \Pr(\text{defector meets a cooperator} \mid n_t = i) \\
n_t &= \text{number of defectors at time } t
\end{aligned}$$

V_d can also be rewritten as (see appendix B):

$$V_d = \sum_{t=0}^{\infty} \beta^t e_1 A^t \pi\sigma u - \sigma u = e_1(I - \beta A)^{-1}\pi\sigma u - \sigma u$$

An agent does not deviate as long as $-c + \beta V_c \geq V_d$, or:

$$-c + \frac{\beta\sigma(u - c)}{(1 - \beta)} \geq e_1(I - \beta A)^{-1}\pi\sigma u - \sigma u \tag{1.2}$$

When $\beta \to 1$, the expression on the left-hand side goes to infinity while the expression on the right-hand side is finite (see appendix B). Hence, as long as agents are sufficiently patient, there are no deviations from the social norm. This reasoning holds for any fixed population size N. However, when N goes to infinity, the following proposition can be proved:

Proposition 1.2. *For a fixed discount factor β, there exists a value of N (say N') such that, for all $N > N'$ the social norm cannot be sustained as a Nash Equilibrium.*

The proof is contained in the appendix C and it is similar to a result in Kandori (1992, page 78). Proposition 1.2 holds for any value of σ, i.e., the social norm breaks down when population grows, independent of the degree of specialization in the economy. In what follows we will change our focus and see what happens when σ changes, for a given value of N.

In order to see the effects of a change in the degree of specialization over the social norm equilibrium (if it reinforces the equilibrium or if it weakens it), we need to calculate $\text{sign}\left(\frac{\partial(\beta V_c - V_d)}{\partial\sigma}\right)$. If it is positive, increasing the degree of specialization in the economy weakens the equilibrium we have been considering so far. We do not have the analytical solution for this derivative but we analyzed the value of $\text{sign}\left(\frac{\partial(\beta V_c - V_d)}{\partial\sigma}\right)$under distinct population sizes and for various values of β and σ. We considered $N = 6$, 14; $\beta = 0.1$, 0.3, 0.5, 0.7, 0.9 and $\sigma = 0.1$, 0.3, 0.5. In all cases, we obtained $\frac{\partial(\beta V_c - V_d)}{\partial\sigma} > 0$, which supports the idea that specialization of the economy reduces the incentives to follow the equilibrium.

In the previous section we suggested that specialization can affect the social norm equilibrium in two distinct ways. First, the reduction of single-coincidence of wants meetings reduces the gains from trade over time $(u - c)$, having an effect similar to a reduction in the value of β, which creates more incentives for an agent to deviate. Another possible effect is that it can make the contagion process slower. It is important to distinguish between these components because while the former affects payoffs both on and off the equilibrium path, the latter only affects off equilibrium payoffs. Put it in another way, the first effect implies that specialization is always going to affect any equilibrium where agents are trading present cost for future benefits (be it a social norm equilibrium, a monetary equilibrium or an equilibrium using memory). Now, since specialization reduces the speed of the contagion process, we have an additional effect which is specific to the social norm equilibrium.

We calculate $\text{sign}\left(\frac{\partial V_d}{\partial\sigma}\right)$ and $\text{sign}\left(\frac{\partial V_d/\sigma}{\partial\sigma}\right)$ for the same range of parameters for N, β and σ as before. Notice that $\text{sign}\left(\frac{\partial V_d/\sigma}{\partial\sigma}\right)$ only considers the effect of σ on V_d in reducing the speed of the contagion process. We found that V_d is increasing in σ (fixing β), while $\frac{V_d}{\sigma}$ is decreasing in σ. We can interpret this result in the following way: when the economy gets more specialized the probability that a defector spreads the contagion (which happens when he meets a cooperator in a meeting where the latter likes his good) goes down, and this increases V_d. Nevertheless, the probability that a defector meets an agent that still follows the norm and has the good he likes also goes down, which reduces his expected payoff. This latter effect dominates the former leading to an overall reduction on V_d. However, V_c also goes down in such a way that $\beta V_c - V_d$ decreases. In synthesis, specialization of the economy reduces the incentives of the agents to follow the social norm.

We can conclude that, in societies where the economy is either too large or too specialized, it is unlikely to observe exchange based on a social norm.

1.4 Money

In this section we introduce money in the economy and study the conditions under which monetary equilibrium exists. We assume no record keeping (memory) and we do not take into account the possibility that agents use a social norm. This study is in the same spirit of KW (1993) but we assume a finite population. Previous analysis of this type of model always use a continuum. We consider the same environment as in the discussion of specialization. Again, agents can only observe their private histories, they differ in their preferences but have the same (random) production technology and are matched pairwise by a uniform random matching process. Moreover, σ is the probability that a meeting results in a single-coincidence of wants. We also allow a possibly positive probability of a double-coincidence of wants, δ (our environment is exactly the same as in the previous section if we set $\delta = 0$).

Following the literature, money is indivisible and agents can store at most one unit. This implies that the only potential monetary exchange involves one unit of money for one unit of a good. Money enters the economy in the following way. In period one, M units are randomly distributed among the agents. Agents are aware of the quantity of money in the economy, therefore they can calculate the expected payoff under a monetary equilibrium.

In equilibrium an agent's expected value from a match depends on whether or not he is holding money. When an agent (for example, agent i) does not have money there is a probability σ that he is matched with another with a preference for his good. In this case, with probability $\frac{M}{N-1}$ the other agent is holding money and so in a monetary equilibrium agent i produces the good (at a cost c) and obtains a continuation payoff of βV_1. With probability $\frac{N-1-M}{N-1}$ the other party has no money and there is no trade. There is a probability σ that i meets an agent which has the good he likes and probability $(1-2\sigma-\delta)$ that he faces a no-coincidence of wants meeting. In both cases, there is no trade and the continuation payoff is βV_0. Finally, there is a probability δ of a double-coincidence of wants meeting, in which case, the pair simply exchange goods.[9] This reasoning provides the expected value of not holding money. By a similar reasoning, we can obtain the expected value of holding money.

Let V_0 denote the expected payoff of an agent without money and V_1 the expected payoff of an agent with one unit of money. Then, the above reasoning yields:

$$
\begin{aligned}
V_0 \quad = \quad & \sigma[\frac{M}{N-1}(-c+\beta V_1) + \frac{(N-1-M)}{N-1}\beta V_0] + \sigma\beta V_0 + \\
& \delta[\frac{M}{N-1}(u-c+\beta V_0) + \frac{(N-1-M)}{N-1}(u-c+\beta V_0)] + (1-2\sigma-\delta)\beta V_0
\end{aligned}
$$

[9] When an agent with money meet another without money in a double-coincidence of wants, we cannot simply assume that agents exchange goods. Another possibility is that they exchange one unit of good for one unit of money. However, in an environment similar to the one considered here, Rupert, Schindler and Wright (2000) proved that the unique subgame perfect equilibrium in pure strategies involves a barter exchange.

$$
\begin{aligned}
V_1 &= \sigma[\frac{(M-1)}{N-1}\beta V_1 + \frac{(N-M)}{N-1}(u+\beta V_0)] + \sigma\beta V_1 + \\
&\quad \delta[\frac{(M-1)}{N-1}(u-c+\beta V_1) + \frac{(N-M)}{N-1}(u-c+\beta V_1)] + (1-2\sigma-\delta)\beta V_1
\end{aligned}
$$

Notice that since we assumed random production types, one does not need to know who (which types) have money, which makes it easier to calculate the expected payoff of an agent.

A monetary equilibrium exists if and only if the cost of production plus the discounted expected value of holding money exceeds the discounted expected value of not holding money; i.e., if and only if $-c + \beta V_1 \geq \beta V_0$. Letting $\gamma = \frac{M}{N}$, and substituting into the equations for V_0 and V_1, we can (after some algebraic manipulation) provide the necessary and sufficient conditions for the equilibrium in terms of the primitives of the model. Proposition 1.3 describes these conditions. Corollary 1.1 establishes that for any population size there exists a discount factor $\beta' < 1$ such that, for all $\beta \geq \beta'$ a monetary equilibrium exists.[10].

Proposition 1.3. *A monetary equilibrium exists if and only if:*

$$
\beta \geq \frac{1}{1 + (1-\gamma)\sigma\frac{(u-c)}{c}\frac{N}{N-1}} \tag{1.3}
$$

Proof. For a monetary equilibrium, we must have $-c + \beta V_1 \geq \beta V_0$. From the expressions for V_0 and V_1, this condition is equivalent to the condition in (3). □

[10] The proof of corollary 1 establishes that β' is increasing in N. The intuition for this relationship is as follows. When an agent holds money he hopes to find agents without money in order to obtain the good he likes. The probability that he will find an agent without money is $1 - \frac{M-1}{N-1} = 1 - \frac{\gamma N-1}{N-1}$. Now, $N > N'$ implies $1 - \frac{\gamma N-1}{N-1} < 1 - \frac{\gamma N'-1}{N'-1}$, i.e., this probability is decreasing with N. As a result, when the population is large agents must be more patient as compared to small populations in order to accept money.

Corollary 1.1. *For any N, a sufficient condition for a monetary equilibrium is* $\beta \geq \frac{1}{1+(1-\gamma)\sigma\frac{(u-c)}{c}}$

Proposition 1.3 shows that when agents are sufficiently patient there exists a monetary equilibrium, regardless of the size of population. This is not the case for a social norm equilibrium. If the population is large agents have incentive to deviate from the social norm because it takes a long time until they are reached by the contagion process. However, in a monetary equilibrium agents must obtain money if they want to consume in a single-coincidence of wants meeting. Hence, as long as the probability of a single-coincidence of wants is not too small, a monetary equilibrium always exist. Finally, since agents can only observe their private histories, i.e, there is no public record of past transactions, we cannot use memory in order to implement exchange. This reasoning leads to the main result of the paper, which can be stated as follows: in an economy with a large population there are allocations that only money can achieve, hence money is essential as a medium of exchange.

Finally, notice that the social norm equilibrium (when it exists) is more efficient than a monetary equilibrium. The reason is very simple: in a social norm equilibrium, there exists no coincidence of wants meeting in which exchange does not occur. In other words, the social norm can implement the best possible allocation. In a monetary equilibrium this allocation is not attainable, since a necessary condition for exchange to take place in single-coincidence meetings is that the agent demanding goods must have money.[11]

[11] The inneficiency of the monetary equilibrium reflects our assumption that agents are randomly pairwise matched. Corbae et alli (2001) show that money is efficient as a medium of exchange in an environment with endogenous random matching.

1.5 Conclusion

The objective of this chapter is to show that social norms can support exchange in the Kiyotaki-Wright environment if the population is not too large and agents are sufficiently patient. The importance of this result is that it gives a possible explanation for the dominant role of money in modern economies (as compared to social norms). As long as agents believe that money is valuable as a medium of exchange, we can support a monetary equilibrium, independent of the size of the population. Hence, even though social norms played a role in the past sustaining the exchange process, for example in village economies,[12] they cannot support exchange in modern economies.

In general, the basic result obtained here is robust to some modifications of the environment. That is, as long as the population is finite and agents are sufficiently patient, gift-exchange can be supported in the KW environment as a Nash equilibrium. Notwithstanding, the fact that we could obtain a simple expression for the contagion process (and the fact that the incentives to follow the norm off the equilibrium path are trivially satisfied such that the equilibrium is also sequential) depends on the assumption that each agent specializes in the production and consumption of only one good. This is the reason why we assumed a model with full specialization in section 1.2.

In sections 1.3 and 1.4 we considered the same model[13], an economy where heterogeneity among agents is reflected only in preferences and not in production technology, and conclude that population growth leads to the breakdown of the social norm. However, a monetary equilibrium exists independent of the population size. Moreover, since there is no public record of past transactions, memory cannot be use as an exchange device. This leads to the

[12]A good reference in this direction is Landa (1994), which gives a historical analysis of gift-exchange.

[13]Note that the introduction of a positive probability of a double-coincidence of wants (δ) in section 4 is completely irrelevant for the result in proposition 3, since this result does not depend on δ. Hence, we can just set $\delta = 0$ if we want the models in section 3 and 4 to be exactly equal.

conclusion that money is essential as a medium of exchange in a highly populated economy. Specialization, by reducing the value of future benefits and the speed of the contagion process, also increases the incentives for agents to deviate. Hence, even if population size is fixed, the social norm equilibrium can collapse if the degree of specialization in the economy increases.

The implications here are distinct from Kocherlakota (1998a,1998b). The conclusion from Kocherlakota's work is that if the cost of keeping track of the past is too high, memory cannot substitute for money, and money becomes essential. Our conclusion is that money is essential not due to the complexity of the past but to the complexity of the present, to the fact that is hard to punish defectors in a world that is either too large in terms of population or too specialized.

1.6 Appendix A

Suppose that some agent (for example, agent i) deviated from the social norm equilibrium. We want to calculate the probability distribution of the random variable T_i="number of periods until agent i is affected by the contagious process" since, after this period, he receives zero as a payoff. In the model we are analyzing here the distribution of T_i coincides with the distribution of T="number of periods until all the population is contaminated by the contagious process". This happens for the following reason: the only way an agent can be affected by the contagious process is when he meets another that was supposed to give him a good but didn't. Moreover since each agent is specialized in the production of one good, at each period there is only one agent that can transmit the defection, namely, the last agent to be contaminated. In a more general environment, agents could also transmit the defection through comunication. However, this does not happen in a complete specialized environment. Since there is only one agent that can produce the good another agent likes, if an agent sees a defection he knows that his future payoff can be at most zero, independent of the transmission of defection through comunication or not. Therefore, we cannot count on comunication as a source of transmission of the contagious punishment. Finally, the agent who gives the good to agent i is the last one to be contaminated in the population, which implies the coincidence of these two distributions. So, from now on, we restrict attention to the calculation of the distribution of T.

Consider a population of size N, where N is even. The results are similar if N is odd. After a deviation by agent i (without loss of generality, suppose that i=1) we can summarize the relevant states of the world by vectors $S_i \in \{0,1\}^N$, where the j^{th} position in the vector indicates the j^{th} agent, and the subscript i indicates that agents 1 trough i are not cooperating.

First, we calculate the probability of moving from state i to state j, where $i, j = 2, 3, ..., N-1$. Note that we cannot move backwards, that is, we cannot move from i to j, where $j < i$. Moreover, since only one agent can transmit the contamination at each period, we cannot move as well from i to j, where $j > i+1$. To obtain the probabilities P_{ij} (Probability of moving from i to j), we need to know what are all the possible partitions of the population in pairwise matches. By induction, we can find that:

$$\#(N) = (N-1)\#(N-2)$$

where $\#(N)$ indicates the number of possible partitions of the population across matches. At each state i, the probability of moving to state $i+1$ next period is equal to the number of partitions where i meets $i+1$ divided by the total of partitions. But, for a population of size N, the number of partitions where i meets $i+1$ is exactly equal to the number of partitions at $(N-2)$. Looking to the above formulas, we have (for $i \neq N$ and $i \neq 1$):

$$P_{ii} = \frac{N-2}{N-1}$$
$$P_{ii+1} = \frac{1}{N-1}$$

Finally, since the state S_N is absorbing, P_{NN}=1. We want to obtain $P(T = t)$, for every t. In order to see how this can be done, consider the following diagram:

$t = 1$	S_2										
$t = 2$	S_2	S_3									
$t = 3$	S_2	S_3	S_4								
...	...	...	...	...							
$t = N-1$	S_2	S_3	S_4	...	...	...	...	...	...	S_{N-1}	S_N
...	...	...	...	...	...	...	...	...	...	...	...
$t = t$	S_2	S_3	S_4	...	...	...	...	...	...	S_{N-1}	S_N

Notice that to find $P(T = t)$ we need to consider all the possible paths that go from S_2 at $t = 1$ to S_N at $T = t$ and then calculate the probability of each path. Let's consider a generic path P. The number of movements to the right that P needs to make in order to reach S_N at $T = t$ is equal to $N - 2$. The number of down movements that P needs to make is equal to $t - (N - 1)$. This result holds for any size of population and any number of periods (as long as $t \geq N - 1$. For $t < N - 1$, $P(T = t) = 0$). Moreover, since P_{ii} is the same for every i ($i \neq N$ and $i \neq 1$), P_{ii+1} is the same for every i ($i \neq N$), the probability of each path is equal to:

$$\left(\frac{1}{N-1}\right)^{N-2} \left(\frac{N-2}{N-1}\right)^{t-(N-1)}$$

Hence, we only need to calculate the number of paths. For every path, the last movement is always to the right. So, we can find the total number of path by calculating the number of different ways of combining $t - 2$ objects, the objects being of two different types: type 1 (movements to the right) and type 2 (down movements). This number is equal to:

$$\binom{t-2}{t-(N-1)} = \frac{(t-2)!}{(N-3)![t-(N-1)]!}$$

Thus, we have the following result:

$$P(T=t) = \binom{t-2}{t-(N-1)} \left(\frac{1}{N-1}\right)^{N-2} \left(\frac{N-2}{N-1}\right)^{t-(N-1)}$$

From the previous result, we can calculate the expected payoff of an agent after he deviates. It is equal to:

$$\sum_{t=1}^{\infty} \beta^t P(T>t) \frac{u}{(N-1)}$$

Agent i does not deviate whenever:

$$\beta V_c - V_n \geq c \Longleftrightarrow \beta \frac{u-c}{(1-\beta)(N-1)} \text{ - } \sum_{t=1}^{\infty} \beta^t P(T>t) \frac{u}{(N-1)} \geq c$$

We will show that, when t goes to infinity, $P[T \leq t]$ goes to one. First, we obtain $\lim_{t\to\infty} P[T = t]$.

$$\lim_{t\to\infty} \binom{t-2}{t-(N-1)} \left(\frac{1}{N-1}\right)^{N-2} \left(\frac{N-2}{N-1}\right)^{t-(N-1)}$$

$$\left(\frac{1}{N-1}\right)^{N-2} \left(\frac{1}{(N-3)!}\right) \lim_{t\to\infty} (t-2)(t-3)...(t-(N-2)) \left(\frac{N-2}{N-1}\right)^{t-(N-1)}$$

$$\left(\frac{1}{N-1}\right)^{N-2} \left(\frac{1}{(N-3)!}\right) \lim_{t\to\infty} (t-2) \left(\frac{N-2}{N-1}\right)^{\frac{t-(N-1)}{N-3}} . \lim_{t\to\infty} (t-(N-2)) \left(\frac{N-2}{N-1}\right)^{\frac{t-(N-1)}{N-3}}$$

$$\left(\frac{1}{N-1}\right)^{N-2} \left(\frac{1}{(N-3)!}\right) \prod_{i=4}^{N} \lim_{t\to\infty} \frac{t-(i-2)}{\left(\frac{N-1}{N-2}\right)^{\frac{t-(N-1)}{N-3}}} = 0.$$

Hence

$$\lim_{t\to\infty}\binom{t-2}{t-(N-1)}\left(\frac{1}{N-1}\right)^{N-2}\left(\frac{N-2}{N-1}\right)^{t-(N-1)}=0\Longrightarrow\lim_{t\longrightarrow\infty}P[T=t]=0.$$

Since $P[T\leq t]=1-P[T>t]$, we have

$$\lim_{t\to\infty}P[T\leq t]=1.$$

1.7 Appendix B

Let D_t be equal to the number of defectors at time t. We want to calculate the matrix $A=(a_{ij})$, where $a_{ij}=P(D_{t+1}=j\mid D_t=i)$. First, note that since the number of defectors is non-decreasing over time, $a_{ij}=0$ whenever $j<i$. Moreover, if $j\geq i$ and $(j-i)>\min(D_t,N-D_t)$, $a_{ij}=0$ since in order to increase the number of defectors by $(j-i)$ it is necessary to have at least $(j-i)$ defectors and cooperators. Another situation where it is easy to identify the value of a_{ij} is when i=1. In this case, after a first deviation of a player, in the next period there will be exactly 2 defectors in the economy. That is, $a_{1j}=0$ if $j\neq 2$ and $a_{1j}=1$ if $j=2$. The non-trivial case is when $i\neq 1$, $j\geq i$ and $(j-i)\leq\min(D_t,N-D_t)$. In this case, we have the following lemma:

Lemma 1.4. *If $i\neq 1$, $j\geq i$ and $(j-i)\leq\min(D_t,N-D_t)$, the probability of moving from i to j defectors (a_{ij}) is equal to:*

$$a_{ij}=\sum_{\tau=0}^{\max(\tau|m-2\tau\geq j-i)}A(\tau)\binom{m-2\tau}{j-i}\sigma^{(j-i)}(1-\sigma)^{(m-2\tau)-(j-i)}\frac{S(2\tau)S(M-(m-2\tau))}{S(N)}$$

where:

$A(\tau)=\binom{M}{m-2\tau}\binom{m}{m-2\tau}(m-2\tau)!$

$m=Min(D_t,N-D_t)$

$M=\max(D_t,N-D_t)$

$S(x)$=number of different ways to make $\frac{x}{2}$ pairs out of x pairs.

$S(0) = 1$

Proof. In order to generate $(j - i)$ defectors, we need to have at least $(j - i)$ meetings between defectors and cooperators. We have to consider situations where the number of meetings between cooperators and defectors is above $(j - i)$ since not always a cooperator turns into a defector upon meeting a defector. More specifically we need to consider up to m meetings between a cooperator and a defector, since this is the maximum possible number of meetings of this kind.

Fix a number of meetings $m - 2\tau$, for some $\tau \geq 0, \tau \leq \max(\tau \mid m - 2\tau \geq j - i)$. The number of possible distinct partitions in groups of two (where one agent is a cooperator and another is a defector) with size $m - 2\tau$ is equal to $\binom{M}{m-2\tau}\binom{m}{m-2\tau}(m - 2\tau)!$. Now, for each of these partitions, the chance of the number of defectors increase to $(j - i)$ follows a binomial distribution with parameter σ, and is equal to $\binom{m-2\tau}{j-i}\sigma^{(j-i)}(1 - \sigma)^{[(m-2\tau)-(j-i)]}$.Now, since we are fixing attention on only $m - 2\tau$ meetings between a defector and a cooperator, the number of possible ways of choosing pairs with the rest of the population where cooperator meets cooperator and defector meets defector is equal to $S(2\tau)S(M-(m-2\tau))$. By varying τ for all its possible values, we obtain the number of possible partitions of the population in pairs such that the number of defectors increase to $(j - i)$, and dividing this number by $S(N)$ gives the probability we are looking for. $\square$

Notice that the contagion process described by Kandori (1989, 1992) can be seen as a particular case of the one analyzed here, which happens when a cooperator always starts to defect after meeting a defector, i.e., when $\sigma = 1$. Now, we can calculate the payoff obtained by an agent who deviates from the proposed social norm equilibrium. First, an agent only has incentive to deviate when he meets another agent that has a preference for his good in a single-coincidence meeting. Hence, the overall payoff after a deviation is equal to:

$$V_d = \sum_{t=1}^{\infty} \beta^t e_1 A^t \pi \sigma u$$

V_d can also be rewritten as:

$$V_d = \sum_{t=0}^{\infty} \beta^t e_1 A^t \pi \sigma u - \sigma u$$

If we replace the last column of A by zeros (this replacement is justified because $D_t = N$ is the absorbing state and the N^{th} element of π is zero) the following equality holds (where A' is the matrix obtained from A after the replacement):

$$(I - \beta A)^{-1}\pi = \sum_{t=0}^{\infty} \beta^t A^t \pi = \sum_{t=0}^{\infty} \beta^t A'^t \pi = (I - \beta A')^{-1}\pi$$

Now, since the number of defectors never declines A' is upper-triangular and so is $(I-A')$. The determinant of an upper-triangular matrix is the product of its diagonal elements, which are all strictly positive for $(I - A')$. Hence, $(I - \beta A')^{-1}\pi < \infty \Longleftrightarrow (I - \beta A)^{-1}\pi < \infty$. We can rewrite V_d as:

$$V_d = e_1(I - \beta A)^{-1}\pi\sigma u - \sigma u$$

An agent will not deviate as long as:

$$-c + \beta V_c \geq V_d$$

or:

$$-c + \frac{\beta\sigma(u - c)}{(1 - \beta)} \geq e_1(I - \beta A)^{-1}\pi\sigma u - \sigma u$$

When $\beta \to 1$, the expression on the left-hand side goes to infinity while the expression on the right-hand side is finite. Hence, as long as agents are sufficiently patient, there will be no deviations from the social norm equilibrium.

1.8 Appendix C

Proof of Proposition 2 (based on Kandori, 1992, page 78, proposition 3):

Suppose that an agent wants to deviate from the social norm equilibrium, at a given point in time (say T). Let V_c and V_d be the continuation payoffs for this agent when he follows and when he deviates from the equilibrium. Let I_t be the set of players whose behavior at $T+t$ is affected by the defection at time T. If player i is the defector at T, $I_1 = \{i, \mu(i,T))\}$ and for $t > 1$, $I_t = \{j; j \in I_{t-1}$ or $j = \mu(k, t-1)$ for some $k \in I_{t-1}\}$, where $\mu(i,t)$ indicates the agent who matches with i at date t. The number of players in I_t is maximal when every member of I_t is matched with a member of $N - I_t$ and the latter is supposed to receive a good from the former but did not receive (so that he turns into a defector in the next period). Hence, an upper bound of $\#I_t$ is $\min\{2^t, N\}$. So, $\min\frac{\{2^t,N\}}{N-1}$ is an upper bound of the probability for the agent who deviated at T to meet a defector at $T+t$. This reasoning implies the following inequality:

$$V_c - V_d \leq \sum_{t=1}^{\infty} \beta^t \min \frac{\{2^{t-1}, N\}}{N-1} u \longrightarrow 0 \text{ when } N \longrightarrow \infty$$

Finally, we have:

$$\lim_{N \longrightarrow \infty} \beta V_c - V_d = 0 < c,$$

which implies that there exists a value of N (say N') such that, for every $N > N'$, the social norm cannot be supported as a Nash Equilibrium.

Chapter 2

Monetary Equilibrium with Decentralized Trade and Learning

(with Braz Camargo)

2.1 Introduction

In most economies the right to print fiat money is a government monopoly. Agreement about the propriety of this monopoly appears widespread. Even a fierce defender of free market systems like Milton Friedman has argued that the very nature of fiat money calls for a government role. Friedman has said, "The technical monopoly character of a pure fiduciary currency makes essential the setting of some external limit on its amount...governmental responsibility for the monetary system has of course been long and widely recognized" (1959, page 8). The idea that the stability of a monetary economy must rely on the exogenous regulation from the government is based on the belief that private agents do not have enough incentives to provide a stable value for money. According to this view, private agents do not set an external limit on the amount of money in circulation and end up printing more money than agents expect, i.e., they overissue. Ritter (1995), for example, argues that only the government is large and patient enough to resist this temptation. Though it is true that the role of the government as an external provider of monetary stability cannot be underestimated, we believe that a more comprehensive theory on the determinants of monetary stability needs to take into account the evolution of the society's ability to prevent overissue.

We construct a model that takes into account both exogenous and endogenous factors affecting the provision of monetary stability. Exogenous factors correspond to the bank's (the agent in control of money creation) behavior regarding the supply of money. This behavior is mainly determined by the bank's patience, which is an exogenous parameter. Endogenous factors correspond to the society's ability to observe the bank's behavior. We describe this ability by assuming that, even though agents do not know the bank's actual action, they can learn about it over time, from the history of the transactions in the market.

We obtain, in an environment where learning only happens slowly over time, that an impatient bank can exploit the agent's misinformation and overissue while maintaining the value of money in the short-run. Agents eventually realize the bank's actual behavior and monetary trade breaks down. However, it takes time until a complete breakdown of trade happens. In this environment, monetary stability depends upon the bank's patience, hence exogenous factors play a key role.

Now consider an environment where agents accumulate a lot of information about the bank's decision in a short period of time. In this situation, agents are able to efficiently monitor the bank's behavior and impatient banks prefer not to overissue in order to avoid the breakdown of monetary trade.

Society's ability to gather information and learn about the state of the economy changes over time. In modern economies, dissemination of information about decisions made by banks or the government is much faster than in the past. Therefore the necessity and adequacy of imposing stringent controls over the creation of fiat-money needs to be reevaluated. It may be true that government intervention played a central role in the past, regulating money creation in a poorly informed economy. However, this may not be necessary in modern economies.

We describe the market in our economy using a random search model (see Kiyotaki and Wright (1993)). These models analyze how money can be endogenously valued as a medium of exchange due to the presence of frictions in the trading process. It is shown that in a completely decentralized economy, where agents are pairwise and randomly matched, trade is constrained by the double coincidence of wants. Money emerges as a medium of exchange that helps overcome these trade constraints. Usually, it is assumed that agents know the amount of money in circulation. In our paper we relax this assumption and consider an economy where information with respect to the actual value of money is subject to the same constraint as the exchange process itself, i.e., agents only obtain information from the meetings in which they participate. In this sense, our model lies in the same tradition as Wolinsky (1990), which also analyzes information transmission in pairwise meetings. We believe that environments with decentralized trade constitute a natural framework to address incomplete information in a market economy. The same technology that restricts and enables trade in these settings can be used to describe the transmission of information throughout the economy. There is no need to adopt a particular information structure as an additional primitive of the model.[1]

In order to solve for the agent's problem in our environment we take into account what decisions an agent makes after every possible history. We solve for the agent's optimal decision rule by describing our problem in terms of a two-armed bandit (see Rothschild (1974) and Banks and Sundaram (1992)). We obtain that maximizing agents adopt a cut-off rule. There exists a threshold such that, if an agent has a belief that the bank overissues

[1]There is a growing literature of search models dealing with incomplete information regarding the amount of money in circulation. Katzman, Kennan and Wallace (2002), for example, assume that only a subset of agents are informed about the exact amount of money in the economy. See also Wallace (1997). For a model outside the search literature discussing how the presence of different information structures affects trade, see Jones and Manuelli (2001).

greater than this threshold, he does not accept money and stays in autarky. Alternatively, he accepts money and goes to the market. Another interesting feature of our solution is that an agent may stay in the market even when his flow payoff is smaller than the payoff of leaving and moving to autarky. The reason is that, by staying in the market, an agent accumulates more information and form better beliefs about the bank's behavior.

We also describe the agent's behavior when time goes to infinity. We show that, under a soft monetary regime (where the bank issues more money than agents expect) the measure of agents accepting money converges to zero. Alternatively, under a tight monetary regime a positive measure of agents always accept money from the bank. The intuition for this result runs as follows. Since an agent takes into account his whole private history when making a decision, as time goes on the effect of additional histories over his belief reduces. Hence, the length of a "bad history" (a history that leads an agent to leave the market and move to autarky) necessary to dominate long "good histories" increases over time. In particular, sufficiently long good histories cannot be entirely dominated by bad histories.

In the first part of this chapter we assume that the bank, after making a decision on how much money to issue, cannot change its choice. This assumption simplifies our problem but has the drawback of precluding an analysis of the bank's actual incentives to follow the same policy over time. In section 6 we consider the case where the bank can change its behavior at any point in time. We show that there is an equilibrium where a sufficiently patient bank always chooses a low amount of money and a sufficiently impatient bank always overissue, as long as there is a small probability that the bank suffers an exogenous shock that affects its preferences. In the absence of small preference shocks, a patient bank have incentives to not overissue only until agents become sufficiently convinced that it will never overissue in the future. At this point, the effect of additional histories over the posterior becomes very small, so that even after observing many agents with money over time, the belief that the bank is patient still remains close to one. In this situation the bank's incentive to issue a low amount of money breaks down, and there is overissue. Since agents are forward looking they anticipate this and the monetary equilibrium collapses. However, if there is an exogenous probability that the bank can become impatient, the agent's posterior is bounded away from one. In this situation, a sufficiently patient bank never overissues in an attempt to distinguish itself from an impatient bank.[2]

We can summarize our discussion in the following way. In an economy where agents learn about the environment from their private histories, an increment in the amount of information obtained through meetings in the market is beneficial in the sense that it precludes impatient banks to overissue. However, when the bank is allowed to change its decision at any point in time, we also need to introduce some degree of exogenously driven uncertainty about the bank's type if we want to support an equilibrium where patient banks never overissue.

The outline of this chapter is as follows. In the next section, we describe the physical and informational environment. In section 2.3 we discuss the agent's problem. In section 2.4 we describe the dynamics of the economy under distinct monetary regimes. Section 2.5 analyzes the bank's problem. Section 2.6 describes what happens when we allow the bank to change its behavior in every period. Section 2.7 concludes.

2.2 Environment

2.2.1 Market and Autarky

Consider an economy with a [0,1] continuum of infinitely lived agents and K indivisible goods. Each agent derives direct utility u from the consumption of only one of the K

[2]This reasoning is directly related to Mailath and Samuelson (1998). See also Mailath and Samuelson (2001).

goods, and the distribution of tastes across agents is uniform. In the first period each agent is endowed with one unit of a good. An agent's endowment does not give any direct utility but it can be used either as an input to produce consumption goods or as a medium of exchange in the market. The production of consumption goods takes time so that an agent cannot produce and at the same time go to the market. We say that an agent is in autarky if he decides to produce, and we set the overall utility obtained in autarky in every period equal to A. Agent's discount the future at a rate β.

In the economy there also exists an infinitely lived large agent, with a technology to print fiat money and a technology to store goods over time, but with no production technology. This agent is denoted bank. The bank derives utility from seigniorage, the revenue raised from money issue. We capture the idea of seigniorage by assuming that the bank's utility comes from the amount of goods it stores during a period, and goods are obtained from agents, in exchange for money. More precisely, in every period the bank proposes to a fraction m of the population entering in the market an exchange of money for goods. Agents do not know which value of m was chosen by the bank. To simplify our analysis we assume that money is indivisible and all agents (with the bank's exception) can store at most one unit of it. This implies that the only possible exchange is one unit of money for one unit of good. We control for the effects of this restriction in a similar way as in Katzman, Kennan and Wallace (2002), comparing our results to the case where agents are completely informed about the value of m. In the end of this section we analyze in detail the implications of the indivisibility of goods and money. Finally, the bank discounts the future at a rate δ.

The market is organized as follows. We have K distinct sectors, each one specialized in the exchange of one good. Agents can identify the sector but inside each sector they are pairwise matched under a uniform random matching technology. Hence, if an agent wants money he goes to the sector which trades his endowment and search for an agent with money. By the other way, if he has money he goes to the sector that trades the good he likes and search for an agent with the good. Upon consuming the good, the agent obtains utility u and receives a new endowment. In other words, an agent must consume in order to obtain a new endowment.[3]

We assume that in every period any agent that goes to the market faces n rounds of meetings ($n > 1$). Agents do not discount between meetings in a period. [4] After n rounds, if an agent ends up with money, he can come back to the bank and exchange it for one unit of the good he likes, i.e., the bank has to redeem the notes. The number of meetings n is a key parameter in our model. It is used as a measure of the degree of information transmission in the economy. Changes in the value of n affect the amount of information an agent can obtain, which in turn affects the agents' and the bank's behavior.

We can summarize the decision problem faced by an agent in this economy as follows. In the beginning of every period each agent decides between staying in autarky or going to the market. If he stays in autarky he produces consumption goods and obtains utility A. If he goes to the market, he first has the opportunity of going to the bank in an attempt to exchange his endowment for money. After that, he goes to the market and faces n rounds of meetings. After these meetings, agents with money can go back to the bank and exchange their money for the good they like. In the next period they face the same problem again.

2.2.2 Information

In our environment the frictions of the exchange process are a consequence of the fact that trade is decentralized. In what follows we assume that the transmission of information in

[3]We adopt this particular market structure just for simplicity. As we show in the next section, it delivers a simple expression for the agent's value functions as compared to the ones in a more general model of search and money. The results of the paper do not change under a more general specification.

[4]This assumption gives a simple expression for agent's flow expected payoff, but is not particularly important for any of the results in the paper.

the economy is subject to the same constraint as trade itself. In other words, agents can only gather information through their private histories. The main decision faced by each agent is whether to accept money and participate in trade or not, a decision which crucially depends upon the agent's expectation about the value of m chosen by the bank. In order to form this expectation agents will look at their personal experiences in the market.

We assume that the bank can only choose between two levels of m, m_H and m_L, with $m_H > m_L$. This assumption keeps the environment simple and delivers the main result we are pursuing, namely, the description of an economy where the bank faces a trade-off between present and future gains when deciding which value of m to choose. A high value of m gives a high flow of utility in the short-run but over time a large fraction of agents leave the market and stay in autarky after realizing (through their histories) that there is too much money in circulation.

In Section 5 we discuss in more detail the problem faced by the bank. Until there we just assume that there is a parameter $\theta_0 \in (0,1)$, which is common knowledge across agents, that indicates the probability the bank sets $m = m_H$. In this case, we say the bank is of type H (high type). Alternatively, we say the economy is under a soft monetary regime. Analogously, a bank is of type L (low type) if it sets $m = m_L$ (and the economy is said to be under a tight monetary regime). We restrict attention first to the case where a bank's type do not change over time: a bank type j will choose m_j in every period, for $j = H, L$. This assumption simplifies the problem faced by the agents. However, it precludes an analysis of the bank's intertemporal behavior. In section 6 we relax this assumption and consider the case where the bank can change its choice at any point in time.

Agents update beliefs from their private histories. A private history at t is the record of all meetings in which an agent participated up to that point. This record includes the meeting of an agent with the bank, the money or good holdings of the agent and of all his matches and what happened in each of these meetings. However, since an agent is only interested in forming a belief with respect to the bank's type, the relevant piece of information is only the record of his (the agent) money holding in the beginning of every period and the money holdings of his partners over time. Moreover, notice that any history only includes data while the agent is in the market, since an agent does not meet anyone when he is in autarky. In the next section we show that when an agent leaves to autarky his optimal decision is to stay in autarky forever. Hence, we can summarize the set of relevant histories up to date t by the set $\{0,1\}^{(t-1)(n+1)}$, where $(t-1)(n+1)$ reflects the fact that in every period when an agent is in the market he meets the bank and n other agents. For example, if $t = 2$ and $n = 3$, the vector $(0, 1, 1, 0)$ indicates a history where the agent did not receive money from the bank, met agents with money in his first two meetings, and met an agent without money in his last meeting in period 1.

Let $\theta(h^t)$ indicate the belief the bank is of type H, given a history h^t. Let $c(h^t)$ indicate the cardinality of the history h^t, i.e., $c(h^t) = \sum_{j=1}^{(t-1)(n+1)} h_j$, where $h_j \in \{0,1\}$. From Bayes rule, we have

$$\theta(h^t) = \frac{m_H^{c(h^t)}(1-m_H)^{(t-1)(n+1)-c(h^t)}\theta_0}{m_H^{c(h^t)}(1-m_H)^{(t-1)(n+1)-c(h^t)}\theta_0 + m_L^{c(h^t)}(1-m_L)^{(t-1)(n+1)-c(h^t)}(1-\theta_0)}$$

$$\theta(h^1) = \theta_0$$

An agent's belief depends only upon the cardinality of the history. Since the fraction of money in the economy does not change over time, it does not matter the order in which agents with money are met, only the total number of these agents. From now on, we use $\theta(h^t)$ and $\theta(c(h^t))$ interchangeably. Moreover, notice that $\theta(h^t)$ is an increasing function of $c(h^t)$, i.e., if an agent meets more people with money over a period of time, he increases his posterior that the bank is of a high type.

2.2.3 The Role of Indivisibility

As stated before, we assume throughout the paper that goods and money are indivisible and agents can only trade one unit of money for one unit of good. This assumption greatly simplifies the agent's Bayesian updating with respect to the bank's type since we can summarize all the relevant information in a history by a sequence of zeros and ones. In the case of divisible goods (or divisible money) we would have to specify and solve a bargaining problem for every meeting at every point in time, after every pair of private histories up to that point. Moreover, a history would have to include the decisions of the agent's match in every previous meeting. We avoid all this complexity when money and goods are indivisible.

Katzman, Kennan and Wallace (2002) deal with bargaining issues in a search environment with incomplete information and divisible goods. In their set up information is exogenously given to a fraction of agents in every period. In our model information comes endogenously, from the agent's private histories. Wallace (1997) also discusses the effects of money changes under incomplete information. He considers a model with two periods where information about the amount of money in the economy is not known in the first period but it is fully revealed to agents in the second period. However, Wallace (1997, page 1304) states that: "In contrast to the assumption that the realized change in the quantity of money is revealed to everyone with a one-period lag, the natural assumption is that it is never revealed." This is exactly what happens in our model. Agents learn about the actual amount of money in circulation only in the long-run, after sufficiently long and informative histories.

Models with divisible goods have the advantage of being able to analyze the impact of changes in the amount of money in two dimensions: an extensive-margin, which refers to the effect of increments in the amount of money over the number of trade-meetings (given by $m(1-m)$); and an intensive-margin, which refers to its effect over the quantity of goods traded in a given meeting. This analysis cannot be done with an indivisible goods-indivisible money model. In these models, there exists only an extensive margin and the agent's decision is either to trade or not to trade. Therefore, we need to make some assumptions about the possible values of m_L and m_H that are different from the ones in models with divisible goods. These assumptions are to ensure that the issues we want to address are meaningful in an environment with indivisibility.

First, we are interested in the study of the long-run costs and short-run benefits faced by a bank who issues more money than agents expect. We capture the long-run costs by assuming that a choice of m_H implies a level of liquidity in the economy inconsistent with the existence of a monetary equilibrium. In this case, if an agent becomes convinced that the bank's choice is m_H, he decides not to trade and stays in autarky. Otherwise he keeps trading until he has the chance of making a decision again. Assumption 2 in the next section reflects this idea. Second, also as a result of our focus on the banks' incentive to overissue, we restrict attention to the region where increases in the amount of money reduce the agent's expected payoff. Since we have no intensive-margin, in order to reduce the expected payoff, we assume that an increment in the amount of money affects the extensive-margin by reducing the number of trade meetings. That is the content of assumption 1.

2.3 Individual Behavior

We consider first the case of complete information, when the fraction of money in the economy in a given period is known (equal to m). Let V_j^i indicate the value function of an agent with j units of money right before his i^{th} meeting, with $i = 1, ..., n$. Note that $V_0^{n+1}(m) = 0$ and $V_1^{n+1}(m) = u$ given our assumption that the bank has to redeem the notes. We have

$$\begin{aligned} V_0^i(m) &= mV_1^{i+1}(m) + (1-m)V_0^{i+1}(m) \\ V_1^i(m) &= mV_1^{i+1}(m) + (1-m)(u + V_0^{i+1}(m)) \\ i &= 1, 2, 3, ..., n. \end{aligned}$$

For example, an agent with money (right before his i^{th} meeting) has a probability m of meeting another agent with money, in which case he does not obtain any good and moves to the next meeting with money. With probability $(1-m)$ he meets another agent without money, in which case he obtains utility u and moves to the next meeting without money. We can solve this problem recursively and obtain the value functions of an agent right before his first meeting in the market. We have

$$\begin{aligned} V_0(m) &= (n-1)m(1-m)u + mu \\ V_1(m) &= (n-1)m(1-m)u + u. \end{aligned}$$

Notice that an agent that goes to the market with money has a greater payoff than an agent with endowment. Hence, when an agent decides to go to the market he always wants to pass first in the bank in an attempt to exchange goods for money. An agent's expected payoff upon going to the market $V(m)$ is equal to $mV_1(m) + (1-m)V_0(m)$ [5], or

$$V(m) = nm(1-m)u + mu.$$

We interpret the expected payoff $V(m)$ as follows. The first component $(nm(1-m)u)$ indicates the flow payoff an agent expects from the transactions he is going to face in the market, while the second (mu) indicates the expected gain of the interaction with the bank after n periods of transactions in the market. This interpretation is important since it

[5]Notice that we are assuming that agents make their decision of going to the market before knowing if they are going to receive money from the bank or not.

separates between the gain coming from the use of money as a medium of exchange in the market and the gain coming from the redemption of the notes in the bank. Our assumption that the bank has to redeem the notes is made only for simplicity, and in what follows we want to make sure that this is not driving any of the results in the paper. We can eliminate the effect of the redemption over the agent's expected payoff by simply redefining (abusing slightly the notation) $V(m)$ as $V(m) - mu$, that is, $V(m) = nm(1 - m)u$.

We stated before that the overall gain in autarky in every period is equal to A. We now establish some microfoundations on how the value of A is obtained, so that the utility of consuming a good in autarky can be compared to the utility of consuming a good in the market (u). More specifically, we assume that the opportunity cost of each round of meetings in the market is equal to a unit of good produced in autarky, with the utility of each good produced in autarky being equal to a ($u > a$). This implies that $A = na$. Without loss of generality we can normalize $V(m)$ and A as $\frac{V(m)}{n}$ and $\frac{A}{n}$, respectively. We have $\widehat{V}(m) = m(1 - m)u$ and $\widehat{A} = a$. This normalization is useful when we make comparisons of the expected value for distinct choices of n, and when $n \to \infty$. In particular note that the difference between the flow payoffs in the market and in autarky is independent of n, the number of non-bank meetings. Therefore, when we study the effect of changes in n on the agent's optimal decision, we are looking to the informational effects only, not to the real effects that this change has over the flow gains.

In what follows we are interested in the study of a situation where the agent's expected payoff of going to the market goes down when the fraction of people with money in the economy (m) increases. Therefore, we restrict attention to the region of parameters where $\frac{\partial \widehat{V}(m)}{\partial m} \leq 0$. This leads to the following assumption (Assumption 1)

$$m_H > m_L \geq \frac{1}{2}.$$

Assumption 2 states that the values of m_H and m_L correspond, respectively, to nonexistence and existence of monetary equilibrium under full-information. This full information scenario constitutes the benchmark of our model. We have

$$m_H[(1-m_H)] < \frac{a}{u} \leq m_L[(1-m_L)].$$

We now consider the incomplete information case, where agents do not know the exact value of m. From now we are going to take n as fixed, unless otherwise stated. The decision problem of an agent is as follows: in every period t, he has to decide between going to the market or staying in autarky. [6] Since the only information he accumulates comes from his private history, each possible private history up to date t represents an information set upon which he can base his decision. An agent's strategy is then a sequence $S = \{s_t\}$ with $s_t : \widehat{H}^t \to \{M, A\}$, where

$$\widehat{H}^t = \{h^t \in \{0,1\}^{(t-1)(M+1)} \mid s^i_\tau(h^\tau) = M, \text{ for all } \tau < t\}.$$

Here M stands for going to the market and A stands for autarky. In the definition of S we are already taking into account the fact (which is proved below) that when an agent decides to stay in autarky, his optimal decision is to stay in autarky forever.

The agents' problem is an example of a two-armed bandit (see Rothschild (1974)). Our model considers a particular case where one of the arms (autarky) is deterministic and the other (market) is stochastic. In this situation, if autarky is an optimal choice at a point in time, it is also optimal in all subsequent periods, i.e., autarky is an absorbing state. The reason is the following. Consider an agent who stays in the market for $(t-1)$ periods, facing a history h^t. In the beginning of period t, he has to decide between staying in autarky or going to the market one more period. If he chooses autarky he receives a flow payoff of a, but he does not receive any additional information with respect to the bank's type. Hence, if

[6] When an agent is in the market, he also has to make trade decisions. However, since it is a dominant action to always accept money from the bank and always accept money in the market, we are not going to formally include these decisions as part of the agent's problem.

the agent's posterior at t is such that the optimal choice is autarky, this belief is unchanged in period $t+1$, and the agent chooses autarky again. We now turn to the problem of when an agent decides for the first time to go to autarky.

At any point in time, the decision of an agent to move to autarky depends only on his posterior, which indicates the probability that the bank is of type H, given the amount of information accumulated up to that point. Denote a generic posterior by θ. Given θ, agent i computes his flow payoff $\widehat{V}^E(\theta) = \theta\widehat{V}(m_H) + (1-\theta)\widehat{V}(m_L)$ and the distribution of next period's posteriors, $\theta^{(1)}(\theta)$. The elements in the support of this distribution are given by:

$$\theta^{(1)}(c \mid \theta) = \frac{m_H^c(1-m_H)^{(n+1)-c}\theta}{m_H^c(1-m_H)^{(n+1)-c}\theta + m_L^c(1-m_L)^{(n+1)-c}(1-\theta)}$$

where c is the cardinality of the history in period t, $c = 0, ..., n+1$. The probability of $\theta^{(1)}(c \mid \theta)$ is equal to the probability of a high type bank times the probability of histories with cardinality c under this regime plus the probability of a low type bank times the probability of histories with cardinality c in this other regime. We have

$$\Pr(\theta^{(1)}(c \mid \theta)) = \theta \Pr_H(N = c) + (1-\theta)\Pr_L(N = c) \qquad (2.1)$$

where

$$\Pr_H(N = c) = \binom{n+1}{c} m_H^c(1-m_H)^{(n+1)-c},$$

$$\Pr_L(N = c) = \binom{n+1}{c} m_L^c(1-m_L)^{(n+1)-c}.$$

We can write the agent's value function as

$$v(\theta) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta E v(\theta^{(1)}(\theta))\right\}.$$

If an agent moves to autarky his payoff is $\frac{a}{1-\beta}$. If he stays in the market, he receives a flow payoff of $\widehat{V}^E(\theta)$ and faces the same decision problem in the next period with the probability of posteriors given by (2.1). The following proposition can be proved:

Proposition 2.1. *$\upsilon(\theta)$ has the following properties:*

(i) $\upsilon(\theta)$ is decreasing and convex in θ;

(ii) $\upsilon(0) = \frac{\widehat{V}(m_L)}{1-\beta}$, and $\upsilon(1) = \frac{a}{1-\beta}$.

Proof. See appendix. □

From Assumption 2, $\widehat{V}^E(\theta)$ is strictly decreasing in θ. Since $\upsilon(\theta)$ is decreasing in θ, we can conclude that $\widehat{V}^E(\theta) + \beta E\upsilon(\theta^{(1)}(\theta))$ is strictly decreasing in θ. Moreover,

$$\widehat{V}^E(0) + \beta E\upsilon(\theta^{(1)}(0)) = \frac{\widehat{V}(m_L)}{1-\beta} > \frac{a}{1-\beta}$$

and

$$\widehat{V}^E(1) + \beta E\upsilon(\theta^{(1)}(1)) = \widehat{V}(m_H) + \frac{\beta a}{1-\beta} < \frac{a}{1-\beta}.$$

This allow us to conclude that there exists a unique $\theta^G \in (0,1)$ satisfying:

$$\widehat{V}^E(\theta^G) + \beta E\upsilon(\theta^{(1)}(\theta^G)) = \frac{a}{1-\beta}.$$

Hence, for every θ such that $\theta \leq \theta^G$, since $\widehat{V}^E(\theta) + \beta E\upsilon(\theta^{(1)}(\theta)) \geq \frac{a}{1-\beta}$, the agent goes to the market. Otherwise, he stays in autarky. The agent's optimal strategy is then given by $S = \{s_t\}$ with

$$\begin{aligned} s_t(h^t) &= M \text{ if } \theta(h^t) \leq \theta^G \\ s_t(h^t) &= A \text{ if } \theta(h^t) > \theta^G. \end{aligned}$$

Behaving optimally implies that after observing a history with a sufficiently high cardinality, an agent leaves the market, goes to autarky and stays in autarky forever. An implication of the agent's decision rule is that, even when there is a low type bank in the economy and the market is objectively better than autarky, an optimizing agent may stay in the market only a finite number of times and already feel confident and informed enough to conclude that the market is not good for trade.

Note that an optimal strategy involves not only the comparison between the agent's expected flow payoff from entering the market and the flow payoff from staying in autarky, but also the fact that by entering the market the agent obtains additional information about the bank's type. To make this point more clear consider the optimal strategy of a myopic agent. In this case his decision to enter or not in the economy does not take into account any gains from experimentation and depends solely on the comparison between the flow payoffs. He would enter as long as

$$\theta\frac{\widehat{V}(m_H)}{1-\beta} + (1-\theta)\frac{\widehat{V}(m_L)}{1-\beta} \geq \frac{a}{1-\beta}$$

Let $\overline{\theta}$ be the unique value of θ such that the left hand side of the above expression is equal to the right hand side. We can show that $\theta^G > \overline{\theta}$, irrespective of the value of n. Therefore, even when the flow expected payoff in the market is smaller than in autarky, the benefits of obtaining additional information through experimentation may induce the agent to stay in the market.

Proposition 2.2. *Let $\overline{\theta}$ be such that $\widehat{V}^E(\overline{\theta}) = a$. Then $\theta^G > \overline{\theta}$.*

Proof. See appendix. □

2.4 Aggregate Behavior

We now compute the measure $\mu_t(m)$ of agents staying in the market up to period t, given m. Rewriting equation (1) we can express the agent's optimal decision in terms of the cardinality of his private history as

$$\begin{aligned} s_t(h^t) &= M \text{ if } c(h^t) \leq \alpha(t-1)(n+1)+\gamma \\ s_t(h^t) &= A \text{ otherwise,} \end{aligned}$$

where

$$\alpha = \frac{\ln\left(\frac{1-m_L}{1-m_H}\right)}{\ln\left(\frac{1-m_L}{1-m_H}.\frac{m_H}{m_L}\right)} \quad \text{and} \quad \gamma = \frac{\ln\left(\frac{\theta^G(1-\theta_0)}{\theta_0(1-\theta^G)}\right)}{\ln\left(\frac{1-m_L}{1-m_H}.\frac{m_H}{m_L}\right)}.$$

An agent stays in the market until period t as long as he faces a history h^t such that $c(h^\tau) \leq \alpha(\tau-1)(n+1)+\gamma$ for all $\tau \leq t$. From now on we assume that $\theta_0 < \theta^G$, so that all agents enter in the economy in the first period.

First note, from the assumptions in our model, that meetings are independent across periods and inside each period they follow a binomial distribution with parameter $m \in \{m_L, m_H\}$. Consider an agent that enters in the market in period 1. The probability that in the beginning of period 2 he enters in the market again is given by the probability that he faces histories with cardinality less than or equal to $\alpha(n+1)+\gamma$. This is given by

$$\sum_{c_1 \leq \lfloor \alpha(n+1)+\gamma \rfloor} \binom{n+1}{c_1} m^{c_1}(1-m)^{(n+1)-c_1}$$

and is equal to $\mu_2(m)$. Now, consider an agent that enters in the market in period 2. The probability that he enters in the market in period 3 is equal to the probability that he faces a history with cardinality less than or equal to $2\alpha(n+1)+\gamma$, given that his history up to period 2 has cardinality equal to c_1 with $c_1 \leq \alpha(n+1)+\gamma$. This conditional probability is given by

$$\sum_{c_2 \leq \lfloor 2\alpha(n+1)+\gamma \rfloor - c_1} \binom{n+1}{c_2} m^{c_2}(1-m)^{(n+1)-c_2}.$$

We can obtain $\mu(3)$ by multiplying the above conditional probability by the probability that an agent faces a history of cardinality c_1 in his first period, and summing over all histories satisfying $c_1 \leq \alpha(n+1)+\gamma$. We have

$$\sum_{\substack{c_2 \leq \lfloor 2\alpha(n+1)+\gamma \rfloor - c_1 \\ c_1 \leq \lfloor \alpha(n+1)+\gamma \rfloor}} \binom{n+1}{c_1}\binom{n+1}{c_2} m^{c_1+c_2}(1-m)^{2(n+1)-(c_1+c_2)}.$$

Proceeding in a similar way, we can calculate the measure of agents entering in the economy for every value of t:

$$\mu_t(m) = \sum_{(c_1,\ldots,c_{t-1}) \in C_{t-1}} \binom{n+1}{c_1}\ldots\binom{n+1}{c_{t-1}} m^{c_1+\ldots+c_{t-1}}(1-m)^{(t-1)(n+1)-(c_1+\ldots+c_{t-1})},$$

where

$$C_{t-1} = \{(c_1,...,c_{t-1}) \mid c_\tau \leq \lfloor \alpha(\tau-1)(n+1)+\gamma \rfloor - c_1 - ... - c_{\tau-1} \text{ for } \tau = 1,...t-1\}$$

The following proposition can be proved:

Proposition 2.3. *For all t:*

(i) $\mu_{t+1}(m_i) \leq \mu_t(m_i)$, $i = L, H$.

(ii) $\mu_t(m_H) - \mu_{t+1}(m_H) \geq \mu_t(m_L) - \mu_{t+1}(m_L)$. In particular, $\mu_t(m_L) \geq \mu_t(m_H)$.

Moreover, for every γ there exists a $\bar{t}$ such that, for all $t > \bar{t}$, the above inequalities are strict.

Proof. See appendix. □

Therefore, independent of the monetary regime, the measure of agents engaged in market activities decreases over time. In an environment where agents use private histories to update their beliefs with respect to the bank's type, there is always a positive probability that some agents face histories inducing them to believe that there is too much money in the economy, even under a tight monetary regime where $m = m_L$. However, the measure of agents leaving the market under a soft monetary regime is larger than the measure under a tight regime. Agents leave the market only after collecting enough evidence that lead them to believe that the regime is actually of a high type. Since these evidences appear with higher probability in a high type regime, after a while more agents start to leave the economy in this case.

The results so far reflect the dynamics of our environment. It is also important to describe the long-run features of the economy under the two regimes. The following proposition deals with these issues.

Proposition 2.4. $\lim_{t\to\infty} \mu_t(m_H) = 0$ *and* $\lim_{t\to\infty} \mu_t(m_L) = \mu_L > 0$.

Proof. See appendix. □

Observe that we only need to prove that $\mu_t(m_H) \to 0$ as $t \to \infty$. This is so, because from (ii) of Proposition 2.3 we have that the exit rate of the market in the soft monetary regime is strictly higher than in the tight monetary regime. Therefore we cannot have $\mu_t(m_L) \to 0$.

Another way of seeing that $\mu_t(m_L) \to \mu_L > 0$ is given by Banks and Sundaram (1992). They prove (theorem 5.1) that in any denumerable-armed bandit problem with independent arms, if at any point in time an arm is selected by an optimal strategy, then there exists at least one type of this arm with the property that, conditional on the arm's type being this particular one, then this arm will remain, with non-zero probability, an optimal choice forever. However, our hypotheses imply that the 'market' arm is an optimal choice at the first period. So the above result must hold. Since we can show that $\mu_t(m_H) \to 0$, then in must be the case that $\mu_t(m_L)$ converges to a positive number.

In the long-run, the measure of agents accepting money from the high type bank goes to zero, while the measure of agents accepting money from the low type bank converges to a positive number. Therefore, even though it is true that a bank overissuing money can survive in the economy in the short-run, he gradually dies out over time. However, if a bank does not overissue, a positive measure of agents is always engaged in trade activities.

2.5 Bank's Behavior

The bank obtains utility from seigniorage, described here by the flow of goods it stores over time. We say that a bank stores a good whenever it exchanges one unit of money for one unit of good, and keeps the good while the agent stays in the market. For each unit of a good in storage during a period the bank receives a utility ν. If the bank always give money to a fraction m of agents entering in the market its utility is equal to

$$U(m,\delta) = \sum_{t=1}^{\infty} \delta^{t-1} \mu_t(m) m v$$

We will assume that the bank's discount factor δ is private information. It is determined by a draw from a random variable with cdf $F(\delta)$, for $\delta \in [0,1]$. The bank knows exactly which value of δ was chosen but agents in the economy only know the distribution $F(\delta)$. Every agent uses $F(\delta)$ to form a prior θ_0 with respect to the bank's behavior.

Propositions 2.3 and 2.4 describe how $\mu_t(m)$ evolves over time for $m = m_L$ and $m = m_H$. The next result shows how the bank's decision is affected by these dynamics. We obtain that, while patient banks choose a tight monetary regime, impatient banks choose a soft regime.

Proposition 2.5. *There exists a unique* $\overline{\delta} \in (0,1)$ *such that* $U(m_L, \overline{\delta}) = U(m_H, \overline{\delta})$. *Moreover, if* $\delta > \overline{\delta}$ *then* $U(m_L, \overline{\delta}) > U(m_H, \overline{\delta})$ *and if* $\delta < \overline{\delta}$, $U(m_L, \overline{\delta}) < U(m_H, \overline{\delta})$.

Proof. See appendix. □

In what follows we construct an equilibrium where the bank's behavior is described by proposition 2.5 and the agent's behavior corresponds to the optimal decision rule obtained in Section 2.3. First we describe our environment in terms of a Bayesian game. A Bayesian game consists of a list $\{(I, A, \Theta, u), \xi\}$ where I is the set of players, A is the set of actions, Θ is a set indicating the possible types that each player can have, u are the payoffs, and

ξ is a probability distribution over Θ. In our model, the players are the bank and a $[0,1]$ continuum of agents. The action for the bank is either m_L or m_H, while for each agent it is either to stay in autarky (A) or go to the market (M). The set of types for the bank is the interval $[0,1]$, the set of possible discount factors. Agents are ex-ante homogeneous, hence they have only one type, described by their initial prior θ_0. The probability distribution over the types is given by the cdf $F(\delta)$ and by the prior θ_0. The payoff of the bank is given by $\sum_{t=1}^{\infty} \delta^{t-1}\mu_t(m)mv$ and the payoff of the agents is given by $\upsilon(\theta)$, for each agent

Proposition 2.6. *Let $\theta_0 = F(\overline{\delta}) < \theta^G$, and consider the strategy profile $\Delta = (S^*, m^*)$, where $S^* = \{s_t^*\}_{t=1}^{\infty}$, with $s_t^* : \widehat{H}^t \to \{M, A\}$ given by*

$$\begin{aligned} s_t(h^t) &= M \;\; \text{if} \;\; \theta(h^t) \leq \theta^G \\ s_t(h^t) &= A \;\; \text{if} \;\; \theta(h^t) > \theta^G. \end{aligned}$$

and $m^*(\delta) : [0,1] \to \{m_L, m_H\}$ is such that

$$\begin{aligned} m(\delta) &= m_L \text{ if } \delta \geq \overline{\delta} \\ m(\delta) &= m_H \text{ otherwise.} \end{aligned}$$

Then Δ is a Bayesian Nash Equilibrium.

Proof. From Proposition 2.5 we know that each type of bank is maximizing utility given the agents' behavior. Moreover, agents update their beliefs with respect to the bank's type from Bayes rule and, given these beliefs, they maximize utility at any point in time. □

This proposition describes one possible equilibrium in our model. If $\theta_0 = 1$, there is another equilibrium where agents never accept money (the bank's choice of m is irrelevant in this case). However, our assumption that $\theta_0 < \theta^G$ rules it out. Nevertheless, we can have many degenerate equilibria where money does not circulate for some or all periods of time. This can happen if an agent believes that nobody is going to the market in a given period or if he believes that everybody who is going to the market is not going to accept money. For example, all agents can decide that they are not going to enter in the market at t=1 (for any of the above reasons), or they can decide to wait for a number of periods before entering. Another type of coordination equilibria is one where the agents agree to stay in the market up to a certain period T (they may decide to leave earlier, if they get pessimistic enough). This is possible since the bank must make an once and for all decision about the level of m at $t = 1$. If T is big, then a sufficiently patient bank will have the incentive to choose $m = m_L$, even knowing that after a finite number of periods all agents leave the economy no matter their beliefs.

However, we believe that the equilibrium described in proposition 2.6 is more appealing than other equilibria. The reason for this lies in our assumption that the agent's expected payoff does not depend on the size of the market, only on the ratio of buyers and sellers. Hence, whenever a group of agents (no matter how small, but of positive measure) decide to go to the market and accept money—their beliefs must, of course, be smaller than an appropriate threshold belief θ^G—all other agents with a posterior smaller than or equal to θ^G will also find profitable to enter the market. So, whenever the agents and the bank believe that a positive measure of agents is going to the market in every period, proposition 2.6 describes their behavior.

Proposition 2.6 describes how a utility maximizing bank acts in an environment where agents can only form beliefs from their private histories. A crucial aspect is that the bank has a degree of freedom to choose m which does not exist when there is full information. If agents know with certainty which bank is acting in the economy, the only equilibrium involving monetary trade has $m = m_L$. However, with partial information, the bank can overissue money and still operate in the economy over time.

Economic History provides various examples on how monetary regimes are subjected to instability associated with the bank's temptation to overissue. One of them is Tullock's (1957) description of the development of paper money in China:

"All of the governments in China between 1100 and 1500 succumbed to this temptation (overissue), and their monetary histories have a strong family resemblance. In each there was a period of inflation, usually quite a long one. Except in the case of the Southern Sung dynasty, which was conquered by the Mongols before the evolution was completed, the use of paper money was, in each case, eventually abandoned. This abandonment of the use of paper money in China is the most interesting feature of the history of paper money in China..." (pages 395 and 396)

However, the fact that incomplete information can bring instability does not necessarily implies that every economy where agents face uncertainty is necessarily unstable. As long as the agent in charge of printing money is patient trade will be supported over time. Ritter (1995), for example, suggests that the government is the only agent sufficiently patient[7] to make the transition from a barter economy to a fiat money economy possible. Private agents are either too small or too impatient to resist the temptation to overissue, which leads to the breakdown of a monetary system based on private money.

The idea that governments are more able than private agents to resist the temptation of overissue can also be seen in the American Legislation regarding money issue. According to Timberlake (1987), "...since the abolition of the operational Gold Standard in the early 1930s, the federal government through its agency, the Federal Reserve System, has been almost the sole creator of the monetary base....No money of any significant amount can be created today without some sanction or act of the Federal Reserve System." Timberlake

[7]Ritter also assumes that the government must be sufficiently large in order for the transition from barter to fiat money to be possible.

continues: "This condition has encouraged the notion that government is a necessary, or at least desirable, regulator of any monetary system - that without government involvement any monetary system quickly degenerates into chaos" (pages 437 and 438).

We believe that considerations about who (if any) is the most suited agent to print fiat money cannot be addressed without a further reference to the environment where such decisions are made. It may be true that governments care more about the welfare of future generations and hence are more patient. However, governments have been part of our society since ancient times and, as Ritter (1995) points out, the widespread use of purely fiat money is a 20th century development. What our model suggests is that there is another element which is crucial for a better understanding of the development and consolidation of monetary regimes, namely, the society's ability to monitor the bank's behavior.

In what follows we show that if the number of non-bank meetings (n) within periods is large, the incentives to overissue decrease. When an agent has more meetings, he obtains more information with respect to the state of the economy. Hence, in general, agents form more precise beliefs about the bank's behavior. In other words, under a soft monetary regime, the average value of $\theta(h^t)$ increases and, under a tight monetary regime, it decreases. We are going to show that, for a given discount factor, the difference $U(m_L, \delta) - U(m_H, \delta)$ increases as n increases, which gives more incentives for a bank to choose m_L. In particular, a bank who was indifferent between m_L and m_H, for a given value of n, prefers to choose m_L if n becomes sufficiently large.

Up to this point we have taken n to be fixed, but from now on we need to consider the case where n is allowed to assume any value in $\mathbf{N}$, the set of natural numbers. This brings in additional complications to our environment since now the value function should depend in both θ and n. Being n a discrete variable, it is better if we instead view the agent's value function as an infinite sequence $\{v_n\}$ of value functions depending on θ alone.

From the results in proposition 2.1 we know that, for all $n \in N$, υ_n is a decreasing function of θ with $\upsilon_n(0) = \frac{\widehat{V}(m_L)}{1-\beta}$ and $\upsilon_n(1) = \frac{a}{1-\beta}$. Another consequence of letting n vary is that the threshold belief θ^G also depends on n. This leads to the issue of how θ^G behaves as n increases, in particular, what happens when n goes to infinity. One feature we would expect to happen is that θ^G is bounded away from one when n goes to infinity. In what follows we show that this is indeed the case. Before giving a formal statement of this fact, let us argue informally why it must be true.

Consider the hypothetical case where $n = \infty$. If an agent has prior θ and decides to enter the economy, his flow payoff does not depend on n and is equal to

$$\widehat{V}^E(\theta) = [\theta m_H(1-m_H) + (1-\theta)m_L(1-m_L)]u.$$

Moreover, he is going to learn the type of the bank he is dealing with. Given that his prior is θ, with probability θ he learns that the bank is of the high type and with probability $1-\theta$ he learns that the bank is of the low type. Therefore, his overall payoff from entering the economy is

$$\widehat{V}^E(\theta) + \beta\left[\theta\frac{a}{1-\beta} + (1-\theta)\frac{\widehat{V}(m_L)}{1-\beta}\right],$$

from which we can conclude that the agent's value function in this case is

$$W(\theta) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta\left[\theta\frac{a}{1-\beta} + (1-\theta)\frac{\widehat{V}(m_L)}{1-\beta}\right]\right\}.$$

Now let $\widetilde{\theta}$ be such that $\upsilon_n(\widetilde{\theta}) = \frac{a}{1-\beta}$. Since $\widehat{V}^E(1) < a$ it is easy to see that $\widetilde{\theta}$ must be smaller than 1. But $\widetilde{\theta}$ is just the equivalent of θ^G when $n = \infty$. So it is natural to expect that $\theta^G(n)$ will not converge to 1 when $n \to \infty$, since we expect the agent's problem to become increasingly similar to the $n = \infty$ problem. The following results justify this intuition.

Proposition 2.7. $\{\upsilon_n\}$ *is a non-decreasing sequence that converges uniformly to* $W(\theta)$ *where*

$$W(\theta) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta\left[\theta\frac{a}{1-\beta} + (1-\theta)\frac{\widehat{V}(m_L)}{1-\beta}\right]\right\}$$

Proof. See appendix. □

Corollary 2.1. *$\exists \rho < 1$ such that $\theta^G(n) \leq \rho$ for all n.*

Proof. See appendix. □

We are now in conditions to describe how changes in the value of n affects the bank's incentives to choose a monetary regime. We obtain that, unless for very impatient banks, if the number of meetings in every period is sufficiently large the bank is always going to choose a tight regime and is not going to overissue.

Proposition 2.8. *For every $\delta > \frac{m_H - m_L}{m_H}$, there exists $n(\delta)$ such that, for all $n \geq n(\delta)$, $U(m_L, \delta, n) > U(m_H, \delta, n)$.*

Proof. See appendix. □

The lower bound in this proposition is due to our assumption that all agents enter in the economy in the first period. Therefore, a very impatient bank prefers to choose m_H if $m_H > \frac{m_L}{1-\delta}$, i.e., $\delta > \frac{m_H - m_L}{m_H}$.

2.6 Building and Maintaining a Reputation

We now relax the assumption that the bank can only make an once and for all decision about the value of m and ask whether there is an equilibrium where a sufficiently patient bank chooses m_L in every period and a sufficiently impatient bank chooses m_H in every period. It turns out that the existence of such an equilibrium depends not only on the patient bank's concern in building a good reputation, but it also depends on its concern in maintaining this good reputation over time. The discussion in this section is similar to the analysis in Mailath and Samuelson (1998), starting with the interpretation of reputation as the agent's posterior about the bank behavior.

First, it is important to make clear the difference between building a good reputation and maintaining it over time. In the previous section, we only captured the incentives to build a good reputation. Since patient banks care about the long-run acceptability of its money, the best choice is to set $m = m_L$, so that the bank has an improving reputation in the society. The bank can only make an once and for all decision, hence there is no sense in which we can talk about incentives to maintain a good reputation over time. This does not happen when we allow the bank to change its decision in every period. In this new set up, as times goes on, since agents are continually revising their beliefs about the bank's type, the bank is facing a different distribution of its reputation in every period, and this affects its future behavior. It is reasonable to expect that a very patient bank chooses m_L for a large number of periods not only to improve its reputation, but also to maintain a good reputation once it is acquired. However, when the number of periods becomes large, if a bank always chooses $m = m_L$, the distribution of the agents' posteriors about its type converges in probability to 0. This implies that, after a sufficiently long period of time, agents that stay in the market become virtually convinced that the bank never overissues. Moreover, further experience in the market practically does not affect their beliefs. Even after facing further histories with high cardinality, an agent remains basically convinced that the bank is choosing m_L. At this point the bank did such a good job in the past in building its reputation that now there is no more need to care about it. As a consequence, the equilibrium where a patient bank always chooses m_L collapses.

The above reasoning leads to the conclusion that in order to have an equilibrium with no overissue we cannot reach a point where the bank's decision exerts no significant effect over the distribution of its reputation. In what follows we are going to slightly modify the environment in order to address this issue.

Up to this point we have treated the bank as a black box. The only relevant variable was its discount factor δ. We now add structure to the bank by assuming that its ownership can change and the bank's discount rate is given by that of its current owner. More specifically, there is a probability λ that the owner is replaced in every period. The new owner discount factor is determined by a draw from the discount factor's pool and the agents' common prior about the new owner is $\phi \in [0,1]$, with ϕ not necessarily equal to θ_0. Moreover, this change of ownership is not observed by the agents in the economy. In this new setup agents' beliefs never reach a point where the bank's actions have a negligible influence on its future reputation. After observing a very negative history an agent attributes some probability to the event that the bank has changed its preferences. This gives a patient bank the right incentives to not overissue.

A strategy for the bank is now a sequence $\Delta = \{m_t\}$ of contingent plans where $m_t \in \{f : [0,1] \times \{m_L, m_H\}^{t-1} \longrightarrow \{m_L, m_H\}\}$. At period t, the bank's action depends on its type and on all of its previous actions. Implicit in this description is the assumption that changes of ownership do not erase the bank's memory. Since we are mainly interested in establishing that the bank's behavior described in section 2.5 is dynamically consistent, we restrict attention only to strategies where patient banks always choose m_L and impatient banks always choose m_H, irrespective of its past history. Formally, the class of strategies that we are going to consider is $M = \{m_t\}$ with $m_t : [0,1] \times \{m_L, m_H\}^{t-1} \longrightarrow \{m_L, m_H\}$ given by (for some $\widetilde{\delta}$)

$$m_t(\delta, .) = \begin{cases} m_L \text{ if } \delta \geq \widetilde{\delta} \\ m_H \text{ if } \delta < \widetilde{\delta} \end{cases} \tag{2.2}$$

As before, a bank is of the low type if its discount factor is in $[\widetilde{\delta}, 1]$, otherwise it is of the high type.

Our first objective is to determine the agents' optimal response to this particular strategy. For this we need to describe how agents update beliefs when the bank's strategy is given by (2.2). Let h_{t-1} indicate the history during period $t-1$ and let $\theta(h^{t-1})$ be the prior probability that the bank is impatient (in which case it chooses m_H), given the history h^{t-1}. We have:

$$\theta(h^t) = (1-\lambda)\frac{m_H^{c(h_t)}(1-m_H)^{(n+1)-c(h_t)}\theta(h^{t-1})}{m_H^{c(h_t)}(1-m_H)^{(n+1)-c(h_t)}\theta(h^{t-1})+m_L^{c(h_t)}(1-m_L)^{(n+1)-c(h_t)}(1-\theta(h^{t-1}))} + \lambda\phi \quad (2.3)$$

Observe that at all points in time posteriors are restricted to the interval $[\lambda\phi, 1-\lambda+\lambda\phi]$.

The agent's decision problem now involves 3 variables, the belief θ, the probability of replacement λ and the common prior about a new owner ϕ. His value function, taking the number of per-period meetings fixed at some n, is given by

$$\upsilon_n(\theta,\lambda,\phi) = \max\left\{a+\beta\upsilon_n((1-\lambda)\theta+\lambda\phi,\lambda,\phi), \widehat{V}_n^E(\theta)+\beta E\upsilon_n(\theta^{(1)}(\theta,\lambda,\phi),\lambda,\phi)\right\},$$

where $\theta^{(1)}(\theta,\lambda,\phi)$ is, as before, the distribution of the next period's beliefs (given θ,λ,ϕ and n). Unlike the previous environment, autarky need not be absorbing. At any period in time, an agent with belief θ has to decide between going to the market or staying in autarky. If he stays in autarky, his next period's belief will be $(1-\lambda)\theta+\lambda\phi$, which can be either greater or smaller than θ. Hence we cannot assume, from the start, that autarky is an absorbing state. We obtain, however, that for ϕ sufficiently close to 1, this will be the case. First, we prove the following result

Proposition 2.9. *For all $n \in N$, $\upsilon_n(\theta,\lambda,\phi)$ is a decreasing function of θ.*

Proof. See appendix. □

We now show that, for $\phi = 1$, there exists a unique θ which makes agents indifferent between autarky and market. First, note that if $\theta^{'} \in [0, 1]$ has this property, we have

$$\upsilon(\theta^{'}, \lambda, 1) = a + \beta\upsilon((1-\lambda)\theta^{'} + \lambda, \lambda, 1).$$

Since $(1-\lambda)\theta^{'} + \lambda \geq \theta^{'}$, and because $\upsilon_n(\theta, \lambda, \phi)$ is decreasing in θ

$$\upsilon(\theta^{'}, \lambda, 1) \leq a + \beta\upsilon(\theta^{'}, \lambda, 1) \Rightarrow \upsilon(\theta^{'}, \lambda, 1) \leq \frac{a}{1-\beta}.$$

However, we must have $\upsilon(\theta, \lambda, \phi) \geq \frac{a}{1-\beta}$ for all choices of θ, λ and ϕ. Hence $\upsilon(\theta^{'}, \lambda, 1) = \frac{a}{1-\beta}$. Now, take any $\theta^{''} > \theta^{'}$ and suppose that the agent is indifferent between the market and autarky if his belief is $\theta = \theta^{''}$. We must have

$$\widehat{V}^E(\theta^{''}) + \beta E\upsilon(\theta^{(1)}(\theta^{''}, \lambda, 1), \lambda, 1) = a + \beta\upsilon((1-\lambda)\theta^{''} + \lambda, \lambda, 1) = \upsilon(\theta^{''}, \lambda, 1) = \frac{a}{1-\beta}$$

Therefore, since $\widehat{V}^E$ is strictly decreasing and (see the proof of proposition (2.9))

$E\upsilon(\theta^{(1)}(\theta, \lambda, \phi), \lambda, \phi)$ is decreasing in θ,

$$\begin{aligned} \frac{a}{1-\beta} &= \widehat{V}^E(\theta^{''}) + \beta E\upsilon(\theta^{(1)}(\theta^{''}, \lambda, 1), \lambda, 1) < \widehat{V}^E(\theta^{'}) \\ +\beta E\upsilon(\theta^{(1)}(\theta^{'}, \lambda, 1), \lambda, 1) &= \frac{a}{1-\beta}, \end{aligned}$$

a contradiction. Similarly, we can show that for $\theta^{''} < \theta^{'}$ we cannot have the agent indifferent.

So, when $\phi = 1$ the equation

$$a + \beta\upsilon((1-\lambda)\theta + \lambda, \lambda, \phi) = \widehat{V}^E(\theta) + \beta E\upsilon(\theta^{(1)}(\theta, \lambda, \phi), \lambda, \phi) \tag{2.4}$$

has at most one solution. However, when $\theta \in \{0,1\}$,

$$\upsilon((1-\lambda)\theta+\lambda,\lambda,\phi) = E\upsilon(\theta^{(1)}(\theta,\lambda,\phi),\lambda,\phi),$$

irrespective of the values of λ and ϕ. So, since $a > \widehat{V}^E(1)$ and $a < \widehat{V}^E(0)$, we can conclude that (2.4) has a unique solution in $(0,1)$, which we denote by $\theta^G(\lambda,1)$. Reducing λ if necessary, this unique solution can be made to lie in the interval $[\lambda\theta_0, 1)$ (just remember, from proposition (2.9), that $\upsilon(\theta,\lambda,\phi)$ is continuous), which is the relevant interval when $\phi = 1$.

Therefore, at least when $\phi = 1$, the agent's optimal decision is the same as in the previous sections: go the market if and only if his belief θ about the bank is smaller than or equal to a certain threshold belief $\theta^G(\lambda,1)$. Moreover, since once in autarky the agent's next period belief about the bank only gets more pessimistic, we have that autarky is an absorbing state.

Now suppose that $\phi < 1$. Denote by $\theta^G(\lambda,\phi)$ a posterior that makes the agent indifferent between market and autarky. By continuity we know that $\theta^G(\lambda,\phi)$ converges to $\theta^G(\lambda,1)$ as ϕ goes to 1.[8] Since $\theta^G(\theta,1) < 1$, it must be the case that there exists a $\overline{\phi} < 1$ with the property that if $\phi \geq \overline{\phi}$, then $\theta^G(\lambda,\phi) \leq \phi$. Therefore, if we restrict ϕ to the interval $[\overline{\phi}, 1)$, we have that $(1-\lambda)\theta^G(\lambda,\phi) + \lambda\phi \geq \theta^G(\lambda,\phi)$. We can then replicate the argument used to prove that (2.4) has a unique solution when $\phi = 1$, to conclude that indeed it has

[8]Let $\{\phi_n\}$, with $\phi_n \in [0,1]$, be such that $\phi_n \to 1$ and suppose, by contradiction, that $\{\theta^G(\lambda,\phi_n)\}$ does not converge to $\theta^G(\lambda,1)$. This means that there is an $\epsilon > 0$ such that for all $n \in N$ there exists an $n' \geq n$ with the property that $|\theta^G(\lambda,\phi_{n'}) - \theta^G(\lambda,1)| > \epsilon$. So we can construct a subsequence $\{\theta^G(\lambda,\phi_{n_k})\}$ such that for all $k \in N$, $|\theta^G(\lambda,\phi_{n_k}) - \theta^G(\lambda,1)| > \epsilon$. Since $\{\theta^G(\lambda,\phi_{n_k})\}$ is bounded, it has a converging subsequence, and this subsequence cannot converge to $\theta^G(\lambda,1)$. Without loss of generality we can take this subsequence as the original one. Denote by α its limit. However, we know that

$$a + \upsilon((1-\lambda)\theta^G(\lambda,\phi_n)+\lambda,\lambda,\phi_n) = \widehat{V}^E(\theta^G(\lambda,\phi_n)) + \beta E\upsilon(\theta^{(1)}(\theta^G(\lambda,\phi_n),\lambda,\phi_n),\lambda,\phi_n),$$

and so, by the continuity of υ, we must have

$$a + \upsilon((1-\lambda)\alpha+\lambda,\lambda,1) = \widehat{V}^E(\alpha) + \beta E\upsilon(\theta^{(1)}(\alpha,\lambda,1),\lambda,1),$$

a contradiction.

a unique solution as long as we consider $\phi \geq \overline{\phi}$. Moreover, autarky is absorbing, and the reason is exactly the same as when $\phi = 1$. One final remark is that, once more because υ is continuous, if we take λ to be sufficiently small, $\theta^G(\lambda, \phi) > \theta_0$, the initial common prior in this economy. So, we are sure that this perturbed environment is meaningful for small values of the perturbation parameter. From now on, we are going to assume that $\phi \geq \overline{\phi}$.

Therefore, we can write down the agent's optimal strategy as $S^* = \{s_t^*\}$ with

$$\begin{aligned} s_t^*(h^t) &= M \ if \ \theta(h^t) \leq \theta^G(\lambda, \phi), \text{and} \\ s_t^*(h^t) &= A \ \text{ if } \theta(h^t) > \theta^G(\lambda, \phi) \end{aligned} \tag{2.5}$$

for all t and $h^t \in \widehat{H}^t$.

Let θ_{t-1} denote the agent's prior at the beginning of period $t-1$. Simple algebra shows that if an agent had a prior $\theta_{t-1} \leq \theta^G(\lambda, \phi)$ and faced a history h_{t-1} during period $t-1$, he enters in the economy in period t if and only if $c(h_{t-1}) \leq \alpha(n+1) + \gamma(\theta_{t-1})$, where α is the same as before and[9]

$$\gamma = \frac{\ln\left(\frac{\widetilde{\theta}^G(\lambda,\phi)(1-\theta_{t-1})}{\theta_{t-1}(1-\widetilde{\theta}^G(\lambda,\phi))}\right)}{\ln\left(\frac{1-m_L}{1-m_H}\cdot\frac{m_H}{m_L}\right)} \quad \text{with } \widetilde{\theta}^G(\lambda, \phi) = \frac{\theta^G(\lambda, \phi) - \lambda\theta_0}{1-\lambda}.$$

The discussion above shows that, if agents believe that patient banks always choose m_L and impatient banks always choose m_H, their best behavior is to follow a cut-off rule with the cut-off value given by $\theta^G(\lambda, \phi)$. We now check if the bank's optimal choice is indeed consistent with the agent's beliefs. First we consider the patient bank's case. For this, let $\nu_t(\theta, m)$ be the probability that if an agent enters the market with a belief θ, and the bank chooses $m \in \{m_L, m_H\}$ in the present period and m_L thereafter, this agent will enter in the economy t periods in the future. The following is true

[9]Note that $\theta_0 < \theta^G(\lambda, \phi)$ hence $\widetilde{\theta}^G(\lambda, \phi)$ is well-defined.

Lemma 2.1.

(i) $\nu_t(\theta, m_L) \geq \nu_t(\theta, m_H)$, *for all* $\theta \in (0,1)$ *and for all* $t \geq 1$.

(ii) $\nu_1(\theta, m_L) - \nu_1(\theta, m_H)$ *increases with* θ *in* $[0, \theta^G(\lambda, \phi)]$.

Proof. See appendix. □

If the low type bank faces an agent with prior θ and does a one-shot deviation, the probability that this agent will enter in the market in the future decreases. The net gain from a deviation depends upon θ. In particular, when θ diminishes, the bank's incentive to deviate becomes stronger. If an agent is really confident that the bank is of a low type, the probability that he enters the market in the future, even after experiencing bad histories, is high. In the extreme case where $\theta = 0$ and either λ or ϕ is equal to 0, the agent always enter the market, that is, $\nu_t(\theta, m_L) = \nu_t(\theta, m_H) = 1$ for all t. In this situation, the bank's choice has no effect over the agent's posteriors, and the bank's dominant action is m_H in every period. This case does not happen in our environment since λ and ϕ are positive and θ is bounded below by $\lambda\phi$, so that the bank always has a trade-off between present and future gains when faced with the choice of m. Therefore, we expect that a sufficiently patient bank prefers to choose m_L in order to increase its future gains. This result is proved in the proposition below. Proposition 2.10 also considers the incentives of the impatient bank to choose m_H. This bank has to compare the present loss incurred by a deviation with the future gains coming from an increase in reputation. It is straightforward to show that an impatient bank never have incentives to choose m_L if we take δ to be sufficiently small.

Proposition 2.10. *Let* $\Delta' = (M^*, S^*)$ *be the strategy profile where* S^* *is the agent's strategy* *defined in (2.5) and* M^* *is the bank's strategy defined as (for* $\delta_1 \leq \delta_2$*)*

$$\begin{aligned} m_t &: [0,\delta_1] \cup [\delta_2, 1] \times \{m_L, m_H\}^{t-1} \longrightarrow \{m_L, m_H\} \textit{ such that} \\ m_t(\delta, .) &= \begin{cases} m_L \textit{ if } \delta \geq \delta_2 \\ m_H \textit{ if } \delta \leq \delta_1 \end{cases} \end{aligned}$$

There exists $n_1 \in N$ *such that this strategy profile, along with the Bayesian updating rule given by (2.3), constitutes a sequential equilibrium if the number* n *of non-bank meetings per period is greater or equal than* n_1*.*

Proof. The only thing we have to prove is that given an appropriate choice of δ_1 and δ_2, the bank will have no incentives to deviate after any history of its play. Let us first deal with the patient bank. We are going to do so by looking at the bank's incentives to do one-shot deviations after any possible histories of its play. Given any such history, we can group the agents in the economy according to their present beliefs. Then, within each of these groups we can determine whether a one-shot deviation pays for the bank. Take the case where $\theta = \lambda\theta_0$. Agents with this belief are the most optimistic about the bank. According to (i) of lemma (2.1), a one-shot deviation involves a trade-off for the bank: A present gain of $m_H - m_L$ (normalized by the size of the group of agents with this belief θ), and a loss of payoff in all future periods (due to a decrease in its reputation). This loss is bounded below by $\delta[\nu_1(\theta, m_L) - \nu_1(\theta, m_H)]$. We want to show that with a convenient choice of n and δ we can make this lower bound bigger than $m_H - m_L$. From the proof of proposition (2.8) we have that $\lim_n \nu_1(\theta, m_L) = 1$ and $\lim_n \nu_1(\theta, m_H) = 0$ for all $\theta \in (0, \theta^G(\lambda)]$. Let then ϵ be any (small) positive number such that

$$\delta_2 = \frac{m_H - m_L}{m_L - \epsilon(m_L + m_H)} < 1.$$

We know that such an ϵ exists since $m_H - m_L < 1/2$ and $m_L \geq 1/2$. We also know that $\exists n_1 \in \mathbf{N}$ such that if $n \geq n_1$, then $\nu_1(\theta, m_L) > 1 - \epsilon$ and $\nu_1(\theta, m_H) < \epsilon$. So, if $n \geq n_1$ we have that

$$\delta[\nu_1(\theta, m_L) - \nu_1(\theta, m_H)] \geq \delta(m_L - \epsilon(m_L + m_H)) \geq m_H - m_L$$

whenever $\delta \geq \delta_2$. Hence, even when facing the most optimistic agents that are possible, a one-shot deviation doesn't pay for the bank if $n \geq n_1$ and $\delta \geq \delta_2$. Now consider the case where $\theta > \lambda\theta_0$. According to (*ii*) of lemma (2.1), we know that the bank's loss in the period following its one-shot deviation will be at least as high as in the case where $\theta = \lambda\theta_0$. So, a one-shot deviation will never pay for the bank as long as $n \geq n_1$ and $\delta \geq \delta_2$. The case of the impatient bank is simpler. Let $\eta_t(\theta, m)$ be the probability that if an agent enters the market with a belief θ, and the bank chooses $m \in \{m_L, m_H\}$ in the present period and m_H thereafter, he enters in the economy t periods in the future. From the proof of (*i*) of lemma (2.1) we can see that $\eta_t(\theta, m_L) \geq \eta_t(\theta, m_H)$ for all t and all relevant θ. We can interpret this in a way similar to the way we did in the previous paragraph. At any period, the impatient bank has to consider the tradeoffs involved in a one-shot deviation. But unlike the previous case, the tradeoff is between a loss in the present period and a gain in all subsequent periods due to an increase in the bank's reputation. However, since the bank is impatient, it's discount factor is $\delta \leq \delta_1$, we can always make the future gains as small as we want by simply making δ_1 sufficiently close to zero. Therefore, $\exists\delta_1$ such that a one-shot deviation will never pay for an impatient bank. This proves our proposition. □

□

2.7 Conclusion

This chapter constitutes an effort to address in a formal way the determinants of monetary stability in an economy where information about the value of money is decentralized. In particular, we considered the bank's incentive to follow a tight monetary regime in an environment where agents only learn about the amount of money in circulation from their personal experiences in the market. We obtained that the bank's temptation to overissue is limited in two different ways. First, it depends on the bank's commitment in maintaining the long-run value of money, which in turn depends on the bank's patience. Second, it depends on the society's ability to monitor the bank's behavior, as measured by the number of transactions an agent faces in a given period.

In an economy where agents face few transactions, the presence of external controls over the amount of money is crucial. However, in economies where there is a rapid transmission of information about the state of the economy, the government intervention is less necessary. In this case, the society itself is able to monitor the behavior of the bank.

Society's ability to obtain information changes over time. In modern economies, information about decisions made by the bank is much more accessible than in the past. Therefore we believe that it is necessary to make a reevaluation of the adequacy of imposing stringent controls over the creation of fiat-money. It may be true that government intervention was necessary in the past, in order to regulate money creation in a poorly informed economy. However, such intervention may not be necessary in modern economies.

Finally, we discussed under what conditions a patient bank is willing to follow the same policy and not overissue even if it is free to change m over time. Not surprisingly, we obtained that in the initial periods the bank's concern in building a good reputation guarantees that it will not overissue. As time goes on and a good reputation is achieved, the incentives for a patient bank to choose m_L are driven by the necessity of maintaining this good reputation. The interesting result is that in the case where the bank's type is known with certainty, the bank does such a good job in building and maintaining a reputation that after some point in time agents basically ignore their private histories and take for granted that the bank will never overissue. This eventually causes the bank to start overissuing.

However, as long as agents face a small probability that the bank's type may change and this change may not be observed, a sufficiently patient bank always chooses m_L in order to maintain a good reputation and separate itself from the impatient banks.

2.8 Appendix

Proposition 2.1 *$\upsilon(\theta)$ has the following properties:*

(i) $\upsilon(\theta)$ is decreasing and convex in θ;

(ii) $\upsilon(0) = \frac{\widehat{V}(m_L)}{1-\beta}$, and $\upsilon(1) = \frac{a}{1-\beta}$.

Proof. We first prove (ii). If $\theta = 0$ then $\theta^{(1)}(\theta)$ is the degenerate distribution putting weight one in $\theta = 0$. Hence

$$\upsilon(0) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(0) + \beta\upsilon(0)\right\}.$$

Moreover, since $\widehat{V}^E(0) = \widehat{V}(m_L) > a$, we can conclude that $\upsilon(0) = \frac{\widehat{V}(m_L)}{1-\beta}$. If $\theta = 1$, then $\theta^{(1)}(\theta)$ is again degenerate, but now it puts all the weight in $\theta = 1$. However, $\widehat{V}^E(1) = \widehat{V}(m_H) < a$, and so we have that $\upsilon(1) = \frac{a}{1-\beta}$. Here we are implicitly assuming that $\upsilon(\theta)$ is well-defined, a fact that is established in the proof of (i).

(i) Let $T : \mathcal{C}[0,1] \to \mathcal{C}[0,1]$ be the mapping given by:

$$Tf(\theta) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta E f(\theta^{(1)}(\theta))\right\}$$

where $\mathcal{C}[0,1] = \{f : [0,1] \to \mathbf{R} \mid f \in C^0\}$. It is easy to see that Tf is indeed continuous because $\theta^{(1)}(\theta)$ is a discrete probability distribution and the Bayesian updating rule is continuous on the prior. Observe now that if $f, g \in \mathcal{C}[0,1]$, with $f \leq g$, then $Ef(\theta^{(1)}(\theta)) \leq Eg(\theta^{(1)}(\theta))$, and so $Tf \leq Tg$. Therefore, T is monotonic. Also, since

$$\max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta E(f+c)(\theta^{(1)}(\theta))\right\} \leq \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta E f(\theta^{(1)}(\theta))\right\} + \beta c,$$

we have that $T(f+c)(\theta) \leq Tf(\theta) + \beta c$ for any $\theta \in [0,1]$ and any $f \in \mathcal{C}[0,1]$. Therefore, the Blackwell conditions are satisfied and T is a contraction. This means that it has a unique fixed point in $\mathcal{C}[0,1]$, which is exactly $\upsilon(\theta)$.

If we show that $\theta < \theta'$ implies that $\theta^{(1)}(\theta')$ first order stochastically dominates $\theta^{(1)}(\theta)$, we then have that if f is decreasing in θ, $Ef(\theta^{(1)}(\theta))$ is also decreasing in θ. To see this, remember that

$$\theta^{(1)}(c \mid \theta) = \frac{m_H^c(1-m_H)^{(n+1)-c}\theta}{m_H^c(1-m_H)^{(n+1)-c}\theta + m_L^c(1-m_L)^{(n+1)-c}(1-\theta)}.$$

Then, for all $x \in [0,1]$,

$$\Pr(\theta^{(1)}(\theta) \leq x) = \Pr(S_{n+1}(\theta) \leq \phi(x,\theta)),$$

where $S_{n+1}(\theta)$ is a random variable indicating the number of successes in $n+1$ trials and

$$\phi(x,\theta) = \max\{c \in \{0, ...n+1\} \mid \theta^{(1)}(c \mid \theta) \leq x\}.$$

Note that $\Pr(S_{n+1}(\theta) = c) = \Pr(\theta^{(1)}(c \mid \theta))$ as given by (2). Since $\theta^{(1)}(c \mid \theta)$ is strictly increasing both in θ and c, $\phi(x,\theta)$ is decreasing in θ. Therefore, for $\theta < \theta'$,

$$\Pr(S_{n+1}(\theta') \leq \phi(x,\theta')) \leq \Pr(S_{n+1}(\theta) \leq \phi(x,\theta')) \leq \Pr(S_{n+1}(\theta) \leq \phi(x,\theta))$$

where the first inequality comes from $\theta < \theta'$ and $\Pr_H(N \leq c) < \Pr_L(N \leq c)$. Hence, for all $x \in [0,1]$:

$$\Pr(\theta^{(1)}(\theta') \leq x) \leq \Pr(\theta^{(1)}(\theta) \leq x),$$

which is the desired stochastic dominance result. Since $\widehat{V}^E(\theta)$ is also decreasing in θ , the map T takes decreasing functions into decreasing functions. Because $\{f \in \mathcal{C}[0,1] \mid f \text{ is decreasing}\}$ is a closed subset of $\mathcal{C}[0,1]$, we can conclude that the unique fixed point of T, $\upsilon(\theta)$, must be decreasing in θ.

We still have to establish the convexity of v. However, $\{f \in \mathcal{C}[0,1] \mid f \text{ is convex}\}$ is also a closed subset of $\mathcal{C}[0,1]$. Since the intersection of 2 closed sets is also closed, the only thing left to do is to show that the map T defined above takes convex functions into convex functions. The proof of this fact is included in the proof of (i) of Lemma 3.1 in Banks and Sundaran (1992), and so we are not including it here. See also footnote 2 in their paper for this respect. $\square$

Proposition 2.2 *Let $\overline{\theta}$ be such that $\widehat{V}^E(\overline{\theta}) = a$. Then $\theta^G > \overline{\theta}$.*

Proof. Let $f(\theta)$ be given by

$$f(\theta) = \frac{\widehat{V}^E(\theta)}{1-\beta}.$$

Then

$$Ef(\theta^{(1)}(\theta)) = \frac{\widehat{V}^E(\theta)}{1-\beta},$$

since $E(\theta^{(1)}(\theta)) = \theta$ (because the agents are doing Bayesian updating) and $\frac{\widehat{V}^E(\theta)}{1-\beta}$ is a linear function of θ. Therefore we must have:

$$Tf(\theta) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta Ef(\theta^{(1)}(\theta))\right\} = \max\left\{\frac{a}{1-\beta}, \frac{\widehat{V}^E(\theta)}{1-\beta}\right\} \geq f(\theta),$$

where T is the map considered in Proposition 2.1. Since T is monotone, $T^{n-1}f(\theta) \leq T^n f(\theta)$, for all $n \in N$, and for all $\theta \in [0,1]$. Consequently, $f(\theta) \leq \lim_{n\to\infty} T^n f(\theta)$. However, T is a contraction and so $T^n f$ converges uniformly (in the sup norm) to v, it's unique fixed point. But uniform convergence implies pointwise convergence, and so $f(\theta) \leq v(\theta)$ for all $\theta \in [0,1]$.

From the definition of θ^G, and from the fact that $f \leq v$, we know that:

$$\frac{a}{1-\beta} = \widehat{V}^E(\theta^G) + \beta Ev(\theta^{(1)}(\theta^G)) \geq \widehat{V}^E(\theta^G) + \beta Ef(\theta^{(1)}(\theta^G)) = \frac{\widehat{V}^E(\theta^G)}{1-\beta}$$

Since $\widehat{V}^E(\overline{\theta}) = a$ and $\widehat{V}^E(\theta)$ is decreasing in θ, we must then have that $\overline{\theta} \leq \theta^G$. Observe now that

$$Tf(\theta) = \max\left\{\frac{a}{1-\beta}, f(\theta)\right\} > f(\theta)$$

for all $\theta \in (\overline{\theta}, 1]$. Since $\exists \theta \in$supp$(\theta^{(1)}(\overline{\theta}))$ such that $\theta > \overline{\theta}$, we then have that

$$ETf(\theta^{(1)}(\overline{\theta})) > Ef(\theta^{(1)}(\overline{\theta})) = f(\overline{\theta}).$$

Therefore

$$\widehat{V}^E(\overline{\theta}) + \beta ETf(\theta^{(1)}(\overline{\theta})) > \widehat{V}^E(\overline{\theta}) + \beta f(\overline{\theta}) = \frac{\widehat{V}^E(\overline{\theta})}{1-\beta} = \frac{a}{1-\beta},$$

from which we can conclude that:

$$T^2 f(\overline{\theta}) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\overline{\theta}) + \beta Ef(\theta^{(1)}(\overline{\theta}))\right\} > \frac{\widehat{V}^E(\overline{\theta})}{1-\beta} = f(\overline{\theta}).$$

Consequently $f(\overline{\theta}) < \lim_{n\to\infty} T^n f(\overline{\theta}) = \upsilon(\overline{\theta})$. Hence:

$$\upsilon(\overline{\theta}) > f(\overline{\theta}) = \frac{a}{1-\beta} = \upsilon(\theta^G) \Longrightarrow \theta^G > \overline{\theta},$$

which ends our proof. □

Proposition 2.3 *For all t, we have:*

(i) $\mu_{t+1}(m_i) \leq \mu_t(m_i)$, $i = L, H$;

(ii) $\mu_t(m_H) - \mu_{t+1}(m_H) \geq \mu_t(m_L) - \mu_{t+1}(m_L)$. In particular $\mu_t(m_L) \geq \mu_t(m_H)$.

(iii) Moreover, for every γ there exists a $\overline{t}$ such that, for all $t > \overline{t}$, the above inequalities are strict.

Proof. Note that (i) is immediate, since autarky is an absorbing state. So we just have to prove (ii) and (iii).

(ii) We know that for $m \in \{m_L, m_H\}$, and for all $t \geq 2$,

$$\mu_t(m) = \sum_{(c_1,\ldots,c_{t-1})\in C_{t-1}} B_{t-1} m^{c_1+\cdots+c_{t-1}} (1-m)^{(t-1)(n+1)-c_1-\cdots-c_{t-1}},$$

where

$$B_{t-1} = \binom{n+1}{c_1} \cdots \binom{n+1}{c_{t-1}},$$

and

$$C_{t-1} = \{(c_1, \ldots, c_{t-1}) \mid c_\tau \leq \lfloor \tau\alpha(n+1) + \gamma \rfloor : \text{for} : \tau = 1, \ldots, t-1\}$$

is the set of feasible cardinalities. The fraction of individuals leaving the economy between periods t and $t+1$ is given by $\mu_t(m) - \mu_{t+1}(m) = L_t(m)$. Remembering that $\mu_1(m) = 1$, we can show that for all $t \geq 1$,

$$L_t(m) = \sum_{(c_1,\ldots,c_t)\in\overline{C}_t} B_t m^{c_1+\cdots+c_t}(1-m)^{t(n+1)-c_1-\cdots-c_t},$$

where $\overline{C}_1 = \{c_1 : c_1 \geq \lfloor \alpha(n+1) + \gamma \rfloor + 1\}$ and

$$\overline{C}_t = \{(c_1, \ldots, c_t) \mid (c_1, \ldots, c_{t-1}) \in C_{t-1},\ c_t \geq \lfloor t\alpha(n+1) + \gamma \rfloor - c_1 - \cdots - c_{t-1} + 1\},$$

for $t \geq 2$. Therefore,

$$L_t(m_L) - L_t(m_H) =$$
$$\sum_{(c_1,\ldots,c_t)\in\overline{C}_t} B_t m_L^{c_1+\cdots+c_t}(1-m_H)^{t(n+1)-c_1-\cdots-c_t}$$
$$\left[\left(\frac{1-m_L}{1-m_H}\right)^{t(n+1)-c_1-\cdots-c_t} - \left(\frac{m_H}{m_L}\right)^{c_1+\cdots+c_t}\right].$$

Observe now that

$$\left(\frac{1-m_L}{1-m_H}\right)^{t(n+1)-c_1-\cdots-c_t} < \left(\frac{m_H}{m_L}\right)^{c_1+\cdots+c_t} \qquad (2.6)$$

if, and only if, $c_1 + \cdots + c_t > \alpha t(n+1)$. But $(c_1, \ldots, c_t) \in \overline{C}_t$ implies that

$$c_1 + \cdots + c_t \geq \lfloor t\alpha(n+1) + \gamma \rfloor + 1 \geq t\alpha(n+1),$$

and so (2.6) is true for all $(c_1, \ldots, c_t) \in \overline{C}_t$. Consequently $L_t(m_L) - L_t(m_H) \leq 0$, since $\overline{C}_t$ might be empty (if γ is high, for example). But

$$L_t(m_L) - L_t(m_H) = \mu_t(m_L) - \mu_{t+1}(m_L) - (\mu_t(m_H) - \mu_{t+1}(m_H)),$$

and so we can conclude that for all$t \geq 1$

$$\mu_t(m_L) - \mu_{t+1}(m_L) \leq \mu_t(m_H) - \mu_{t+1}(m_H).$$

(*iii*) According to the proof of (*ii*), we just need to show that $\overline{C}_t$ will be non-empty if t is sufficiently large. Let $\bar{t}$ be the smallest positive integer such that $\bar{t}(n+1) > \lfloor \bar{t}\alpha(n+1)+\gamma \rfloor$. Since γ is finite for every choice of the initial belief θ_0 (as long as $\theta_0 > 0$), such a $\bar{t}$ will always exist. By definition, if $t < \bar{t}$, then $t(n+1) \leq \lfloor \bar{t}\alpha(n+1)+\gamma \rfloor$, and so $(c_1, \dots, c_{\bar{t}-1}) \in C_{\bar{t}-1}$ if we let $c_\tau = (n+1)$, $\tau = 1, \dots, \bar{t}-1$. But if we also let $c_{\bar{t}} = (n+1)$, then $c_1 + \cdots + c_{\bar{t}} = \bar{t}(n+1)$, and so $(c_1, \dots, c_{\bar{t}}) \in \overline{C}_{\bar{t}}$. Therefore $\overline{C}_t$ is not empty if $t = \bar{t}$. It is now easy to see that if $\overline{C}_t$ is not empty, then $\overline{C}_{t'}$ will be non-empty for all $t' > t$. □

Proposition 2.4 $\lim_{t\to\infty} \mu_t(m_H) = 0$ *and* $\lim_{t\to\infty} \mu_t(m_L) = \mu_L > 0$.

Proof. First note that

$$\binom{n+1}{c_1} \cdots \binom{n+1}{c_{t-1}} \leq \binom{(t-1)(n+1)}{c_1 + \cdots + c_{t-1}},$$

and that

$$C_{t-1} \subset \{(c_1, \dots, c_{t-1}) \mid c_1 + \cdots + c_{t-1} \leq \lfloor (t-1)\alpha(n+1) + \gamma \rfloor\}.$$

Hence

$$\begin{aligned} \mu_t(m_H) &\leq \sum_{k=0}^{\lfloor (t-1)\alpha(n+1)+\gamma \rfloor} \binom{(t-1)(n+1)}{k} m_H^k (1-m_H)^{(t-1)(n+1)-k} \\ &= \Pr\{S_{(t-1)(n+1)}(m_H) \leq \lfloor (t-1)\alpha(n+1) + \gamma \rfloor\}, \end{aligned}$$

where $S_{(t-1)(n+1)}(m_H)$ denotes the number o successes of $(t-1)(n+1)$ Bernoulli trials when the probability of success is m_H. By the law of large numbers for the Binomial distribution (see Feller, chapter 6), we know that for all $\epsilon > 0$,

$$\lim_{t\to\infty} \Pr\{S_{(t-1)(n+1)}(m_H) < (t-1)(n+1)[m_H - \epsilon]\} = 0. \qquad (2.7)$$

But, since $\alpha < m_H$, we know that $\exists \bar{t}$ such that if $t \geq \bar{t}$, then $\lfloor (t-1)\alpha(n+1) + \gamma \rfloor < (t-1)(n+1)m_H$. Let $\epsilon = m_H - (\lfloor (\bar{t}-1)\alpha(n+1) + \gamma \rfloor)/(\bar{t}-1)(n+1) > 0$. Then, for any $t \geq \bar{t}$,

$$\Pr\{S_{(t-1)(n+1)}(m_H) \leq \lfloor (t-1)\alpha(n+1) + \gamma \rfloor\} \leq$$

$$\Pr\{S_{(t-1)(n+1)}(m_H) \leq (t-1)(n+1)[m_H - \epsilon]\},$$

and we can apply (2.7) to conclude that $\mu_t(m_H) \to 0$ as t goes to infinity.

For the asymptotic behavior of $\mu_t(m_L)$, observe that $\{\mu_t(m_L)\}$ is a bounded and decreasing sequence. Therefore it must have a limit. The above result, together with (*ii*) of proposition **2.3**, implies that this limit must be strictly positive. $\square$

Proposition 2.5 *There exists a unique* $\overline{\delta} \in [0,1)$ *such that* $U(m_L, \overline{\delta}) = U(m_H, \overline{\delta})$. *Moreover, if* $\delta > \overline{\delta}$ *then* $U(m_L, \overline{\delta}) > U(m_H, \overline{\delta})$ *and if* $\delta < \overline{\delta}$, $U(m_L, \overline{\delta}) < U(m_H, \overline{\delta})$

Proof. Let $F(\delta)$ be given by

$$F(\delta) = \frac{1}{v}\left[U(m_L, \delta) - U(m_H, \delta)\right] = \sum_{t=1}^{\infty} \delta^{t-1} d_t,$$

where $d_t = \mu_t(m_L)m_L - \mu_t(m_H)m_H$. We want to show that $\exists \underline{\delta} \in [0,1)$ such that $F(\delta) < 0$ if $\delta < \underline{\delta}$, and $F(\delta) > 0$ if $\delta > \underline{\delta}$. First note that

$$F^{(k)}(\delta) = \sum_{t=k+1}^{\infty} (t-1)(t-2)\dots(t-k)\delta^{t-k-1} d_t.$$

Since $\mu_t(m_L)$ decreases (monotonically) to some $\mu_L > 0$ and $\mu_t(m_H)$ decreases (monotonically) to zero, we have that $\exists t' \geq k+1$ such that if $t \geq t'$, then $d_t \geq \frac{1}{4}\mu_L m_L$. Therefore,

$$F^{(k)}(\delta) \geq \sum_{t=k+1}^{t'-1} (t-1)\dots(t-k)\delta^{t-k-1}d_t + \frac{\mu_L}{4}\sum_{t=t'}^{\infty}(t-1)\dots(t-k)\delta^{t-k-1},$$

and so, since the first term after the above inequality is finite for all δ, we can conclude that $\forall k \geq 0$,

$$\lim_{\delta \to 1^-} F^{(k)}(\delta) = +\infty.$$

Let now $\bar{t}$ be the smallest integer such that $d_t > 0$ for all $t \geq \bar{t}$. Such a t exists because of the asymptotic behavior of the population measures in the 2 regimes. Since

$$F^{(\bar{t}-1)}(\delta) = \sum_{t=\bar{t}}^{\infty}(t-1)\dots(t-\bar{t}+1)\delta^{t-\bar{t}}d_t,$$

we have that $F^{(\bar{t}-1)}(\delta) > 0$ for all $\delta > 0$. If $\bar{t} = 1$, we are done, just set $\underline{\delta} = 0$.

So, suppose that $\bar{t} > 1$. Now observe that

$$\begin{aligned} F^{(\bar{t}-2)}(\delta) &= \sum_{t=\bar{t}-1}^{\infty} (t-1)\dots(t-\bar{t}+2)\delta^{t-\bar{t}+1}d_t \\ &= (\bar{t}-2)\dots 2d_{\bar{t}-1} + \sum_{t=\bar{t}}^{\infty}(t-1)\dots(t-\bar{t}+2)d_t, \end{aligned}$$

and so $F^{(\bar{t}-2)}(0) \leq 0$. Since $F^{(\bar{t}-2)}(.)$ is strictly increasing and $F^{(\bar{t}-2)}(1^-) = +\infty$, we have that there exists a unique δ_1 such that $F^{(\bar{t}-2)}(\delta) < 0$ if, and only if, $\delta < \delta_1$. If $\bar{t} = 2$, we are done, just set $\underline{\delta} = \delta_1$.

So, suppose now that $\bar{t} > 2$. The same reasoning as above shows that $F^{(\bar{t}-3)}(0) < 0$. Since $F^{(\bar{t}-3)}(\delta)$ decreases until δ_1, and after this point it increases strictly to $+\infty$, we have that there is a unique $\delta_2 > \delta_1$ such that $F^{(\bar{t}-3)}(\delta) < 0$ if, and only if $\delta < \delta_2$. If $\bar{t} = 3$, we are again done. If $\bar{t} > 3$, we can continue with this process. Since $\bar{t}$ is finite, we eventually reach an end to it. □

Proposition 2.7 *The sequence $\{\upsilon_n\}$ of value functions converges uniformly to*

$$\mathsf{V}(\theta) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta\left[\theta\frac{a}{1-\beta} + (1-\theta)\frac{\widehat{V}^E(m_L)}{1-\beta}\right]\right\}$$

Proof. Let $\mathcal{S}$ be the set given by

$$\mathcal{S} = \prod_{i=1}^{\infty} \mathcal{C}[0,1],$$

the infinite cartesian product of the set $\mathcal{C}[0,1]$. Then $\{\upsilon_n\}$ is a fixed point of the map $\Gamma : \mathcal{S} \to \mathcal{S}$ defined by $(\Gamma f)_n(\theta) = Tf_n(\theta)$; that is,

$$(\Gamma f)_n(\theta) = \max\left\{\frac{a}{1-\beta}, \widehat{V}_n^E(\theta) + \beta E f_n(\theta^1(\theta, n))\right\}.$$

The proof is quite long, so we give an outline of it before getting into the details. The first thing we do is define a topology on $\mathcal{S}$ such that it becomes a complete topological (vector) space with a countable basis. Since $\mathcal{S}$ has a countable basis, it is metrizable. We exhibit a metric on $\mathcal{S}$ that is compatible with the topology we defined on it and show that Γ is a contraction with respect to this metric. Therefore Γ has a unique fixed point, $\{\upsilon_n\}$. Then we show that the set $\mathcal{X} = \{f \in \mathcal{S} : \{f_n\} \text{ is uniformly convergent}\}$ is a closed subset of $\mathcal{S}$, and that Γ maps $\mathcal{X}$ into itself. This completes the proof.

(1) Let $p_n(f) = \| f_n \|_{\sup}$, where $\| \,.\, \|_{\sup}$ is the sup norm in $\mathcal{C}[0,1]$. Then p_n is a semi-norm on $\mathcal{S}$, and the collection $\mathcal{P} = \{p_n, n \in \mathbf{N}\}$ is a separating family of semi-norms on $\mathcal{S}$. Indeed, if $p_n(f) = 0$ for all n, then $f_n = 0$ for all n, which implies that $f = 0$. For each $p_n \in \mathcal{P}$ and each $k \in \mathbf{N}$, set

$$B(p_n, k) = \left\{f \in \mathcal{S} : p_n(f) < \frac{1}{k}\right\},$$

and let $\mathcal{B}$ be the collection of all finite intersections of sets of the above form. Then (see Rudin, Functional Analysis, Theorem 1.37), $\mathcal{B}$ is a local base for a vector topology τ on $\mathcal{S}$

($\mathcal{S}$ is a vector space with the vector operations defined in the usual way). From now on, $\mathcal{S}$ will denote the topological vector space $(\mathcal{S}, \tau)$ just defined. Because $\mathcal{P}$ is countable, $\mathcal{B}$ is countable. Hence $\mathcal{S}$ is a topological vector space with a countable local base, and so it is metrizable (see Rudin, Functional Analysis, Theorem 1.24); i.e., there exists a metric d on $\mathcal{S}$ such that the topology induced by d on $\mathcal{S}$ coincides with τ (we say that d is compatible with τ).

To show that $\mathcal{S}$ is complete, suppose that $\{f^n\}$ is a Cauchy sequence in $\mathcal{S}$. This means that $\forall B \in \mathcal{B}$, $\exists N \in \mathbf{N}$ such that if $m, n \geq N$, then $f^m - f^n \in B$. In particular, given $n \in \mathbf{N}$, we have that $\forall k \in \mathbf{N}$, $\exists N$ such that if $m, m^{'} \geq N$, then

$$f^m - f^{m^{'}} \in B(p_n, k) \Rightarrow \| f_n^m - f_n^{m^{'}} \|_{\text{sup}} < \frac{1}{k}.$$

This implies that $\{f_n^m\}$ is a Cauchy sequence in $\mathcal{C}[0,1]$ for all $n \in \mathbf{N}$. Since this space is complete, $\{f_n^m\}$ is convergent in $\mathcal{C}[0,1]$ for all $n \in \mathbf{N}$. Let f_n be its limit. We want to show that $\{f^m\}$ converges to $f = (f_1, \ldots, f_n, \ldots)$ in $\mathcal{S}$. For this, let $V \in \mathcal{B}$. We know that $\exists p_{n_1}, \ldots, p_{n_j} \in \mathcal{P}$ and $\exists k_1, \ldots, k_j \in \mathbf{N}$ such that $V = \cap_{i=1}^{j} B(p_{n_i}, k_i)$. Since $f_n^m \to f_n$ for all $n \in \mathbf{N}$, $\exists N_i \in \mathbf{N}$ such that if $m \geq N_i$, then $\| f_{n_i}^m - f_{n_i} \|_{\text{sup}} < \frac{1}{k_i}$ for $i = 1, \ldots, j$. If we let $N = \max\{N_1, \ldots, N_j\}$, $m \geq N$ implies that $f^m - f \in V$. Consequently $f^m \to f$ in $\mathcal{S}$, and so $\mathcal{S}$ is complete.

(2) Let $d : \mathcal{S} \times \mathcal{S} \to \mathbf{R}_+$ be such that

$$d(f, g) = \max_n \frac{c_n p_n(f - g)}{1 + p_n(f - g)},$$

where $\{c_n\}$ is some strictly positive sequence of real numbers such that c_n converges to zero. Then one can show (see Rudin, Functional Analysis, Remark 1.38) that d is a metric on S that is compatible with the vector topology τ we defined in **(1)**. We will now establish that Γ, the map we defined at the very beginning, is a contraction. For this, suppose, by contradiction, that there is no $\delta < 1$ such that $d(\Gamma f, \Gamma g) \leq \delta d(f,g)$. So, for all $n \in \mathbf{N}$, $\exists j(n) \in \mathbf{N}$ such that

$$\begin{aligned} \frac{c_{j(n)} \parallel Tf_{j(n)} - Tg_{j(n)} \parallel_{:\sup}}{1+ \parallel Tf_{j(n)} - Tg_{j(n)} \parallel_{:\sup}} &> \left(1 - \frac{1}{n}\right) \max_k \frac{c_k \parallel f_k - g_k \parallel_{:\sup}}{1+ \parallel f_k - g_k \parallel_{:\sup}} \\ &\geq \left(1 - \frac{1}{n}\right) \frac{c_{j(n)} \parallel f_{j(n)} - g_{j(n)} \parallel_{:\sup}}{1+ \parallel f_{j(n)} - g_{j(n)} \parallel_{:\sup}}. \end{aligned}$$

This implies that

$$\begin{aligned} &\frac{\parallel Tf_{j(n)} - Tg_{j(n)} \parallel_{\sup} - \beta \parallel f_{j(n)} - g_{j(n)} \parallel_{\sup}}{d_{j(n)}} :> \\ &\left(1 - \frac{1}{n} - \beta\right) \frac{\parallel f_{j(n)} - g_{j(n)} \parallel_{:\sup}}{d_{j(n)}} \\ &- \frac{1}{n} \underbrace{\frac{\parallel Tf_{j(n)} - Tg_{j(n)} \parallel_{:\sup} \parallel f_{j(n)} - g_{j(n)} \parallel_{:\sup}}{d_{j(n)}}}_{e_{j(n)}}, \end{aligned}$$

where $d_j = (1+ \parallel Tf_j - Tg_j \parallel_{:\sup})(1+ \parallel f_j - g_j \parallel_{:\sup})$. But $\{e_{j(n)}\}$ is a bounded sequence, and so $\frac{1}{n} e_{j(n)}$ must converge to zero. Since $1 - \frac{1}{n} - \beta$ converges to $1 - \beta > 0$, we then have that if n is sufficiently large, the right-hand side of the above inequality will be positive. Hence, if we take n large enough,

$$\parallel Tf_{j(n)} - Tg_{j(n)} \parallel_{:\sup} > \beta \parallel f_{j(n)} - g_{j(n)} \parallel_{:\sup},$$

which contradicts the fact that $\forall j \in \mathbf{N}$ the map $f_j \mapsto (\Gamma f)_j = Tf_j$ is a contraction of modulus β. Consequently Γ is a contraction in $\mathcal{S}$, which allows us to conclude that $\{v_n\}$ is its unique fixed point.

(3) Let now $\mathcal{X} = \{f \in \mathcal{S} : \{f_n\} \text{ is uniformly convergent}\}$ and suppose that $\{f^m\}$ is a sequence in $\mathcal{X}$ such that $f^m \to f$ in $\mathcal{S}$. We want to show that $f \in \mathcal{X}$; that is, that $\{f_n\}$ is uniformly convergent. We have seen in (**1**) that if $\{f^m\}$ is convergent in $\mathcal{S}$, then $\{f_n^m\}$ is uniformly convergent for all n, and it must converge to f_n. Now observe that $\forall \theta \in [0,1]$,

$$f_n(\theta) - f_{n'}(\theta) = f_n(\theta) - f_n^m(\theta) + f_n^m(\theta) - f_{n'}^m(\theta) + f_{n'}^m(\theta) - f_{n'}(\theta),$$

where the choice of m is arbitrary, so that

$$|f_n(\theta) - f_{n'}(\theta)| \leq \| f_n - f_n^m \|_{\sup} + \| f_n^m - f_{n'}^m \|_{\sup} + \| f_{n'}^m - f_{n'} \|_{\sup} .$$

Take $\epsilon > 0$. Since $\{f_q^m\}$ is uniformly convergent for all m by hypothesis, we know that $\exists N$ such that if $n, n^{'} \geq N$, then $\| f_n^m - f_{n'}^m \|_{\sup} < \frac{\epsilon}{3}$. Take then $n, n^{'}$ greater than N. Because $\{f_n^m\}$ and $\{f_{n'}^m\}$ converge uniformly to f_n and $f_{n'}$, respectively, there is $m_0(n, n^{'}) \in \mathbf{N}$ such that if $m \geq m_0$, $\| f_n - f_n^m \|_{\sup} < \frac{\epsilon}{3}$ and $\| f_{n'}^m - f_{n'} \|_{\sup} < \frac{\epsilon}{3}$. If we now take $m \geq m_0$, we can conclude that $|f_n(\theta) - f_{n'}(\theta)| < \epsilon$ for all $\theta \in [0, 1]$; that is, $\| f_n - f_{n'} \|_{\sup} < \epsilon$. Consequently $\{f_n\}$ is Cauchy, and so uniformly convergent. This proves that $\mathcal{X}$ is indeed a closed subset of $\mathcal{S}$.

(4) To finish, we want to show that if $\{f_n\}$ is uniformly convergent, then $\{Tf_n\}$ will also be. For this, let f be the uniform limit of $\{f_n\}$, and let

$$Tf(\theta) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta[\theta f(1) + (1-\theta)f(0)]\right\},$$

where $\widehat{V}^E$ is the uniform limit of $\widehat{V}_n^E$. Remember that

$$\widehat{V}_n^E(\theta) = \frac{1}{n}[\theta V^E(m_H) + (1-\theta)V^E(m_L)],$$

with $V^E(m) = nm(1-m)u + mu$, and so

$$\widehat{V}^E(\theta) = \theta m_H(1-m_H)u + (1-\theta)m_L(1-m_L)u.$$

Since $\max\{g, h\} - \max\{m, n\} \leq \max\{g - m, h - n\}$, we have that

$$\begin{aligned} Tf_n(\theta) - Tf(\theta) \quad &\leq \quad \max\Big\{0, \widehat{V}_n^E(\theta) - \widehat{V}^E(\theta) : + \\ &\qquad \beta[Ef_n(\theta^1(\theta, n)) - \theta f(1) - (1-\theta)f(0)]\Big\}. \end{aligned}$$

First observe that $|\widehat{V}_n^E(\theta) - \widehat{V}^E(\theta)| \leq \| \widehat{V}_n^E - \widehat{V}^E \|_{:\sup}$ for all $\theta \in [0, 1]$, and that $\| \widehat{V}_n^E -$

$\widehat{V}^E \|_{:\sup} \to 0$ by construction. Now note that

$$Ef_n(\theta^1(\theta,n)) - [\theta f(1) + (1-\theta)f(0)] :=$$

$$Ef_n(\theta^1(\theta,n)) - Ef(\theta^1(\theta,n)) :+$$

$$Ef(\theta^1(\theta,n)) - [\theta f(1) + (1-\theta)f(0)].$$

However,

$$|Ef_n(\theta^1(\theta,n)) - Ef(\theta^1(\theta,n))| = |E(f_n - f)(\theta^1(\theta,n))| \leq \| f_n - f \|_{:\sup},$$

and $\| f_n - f \|_{:\sup} \to 0$ by hypothesis. So we're left with the term

$$Ef(\theta^1(\theta,n)) - [\theta f(1) + (1-\theta)f(0)].$$

If we can show that $Ef(\theta^1(\theta,n))$ converges uniformly to $\theta f(1) + (1-\theta)f(0)$ we are done.

For this, note first that we can assume that the sequence $\{f_n\}$ is such that

$$\theta f_n(1) + (1-\theta)f_n(0) \geq Ef_n(\theta^1(\theta,n))$$

for all $n \in \mathbf{N}$ and all $\theta \in [0,1]$ (see **(5)** below). But we have seen above that the uniform convergence of f_n to f implies that $Ef_n(\theta^1(\theta,n)) - Ef(\theta^1(\theta,n))$ converges uniformly to zero. Moreover, it is easy to see that $\theta f_n(1) + (1-\theta)f_n(0)$ converges uniformly to $\theta f(1) + (1-\theta)f(0)$. Therefore, $\exists \overline{n}$ such that if $n \geq \overline{n}$, then

$$\theta f(1) + (1-\theta)f(0) \geq Ef(\theta^1(\theta,n))$$

for all $\theta \in [0,1]$. If we let $h_n(\theta) = Ef(\theta^1(\theta,n))$ and $h(\theta) = \theta f(1) + (1-\theta)f(0)$, we have thus established that if $n \geq \overline{n}$, $h_n(\theta) \leq h(\theta)$ for all $\theta \in [0,1]$. Since $\theta^1(\theta,n)$ converges in distribution to $\theta\delta(1) + (1-\theta)\delta(0)$, where $\delta(x)$ is the probability distribution putting all mass on $x \in \mathbf{R}$, we know that if g is any continuous function, then $Eg(\theta^1(\theta,n)) \to$

$\theta g(1) + (1-\theta)g(0)$ pointwisely. Hence $h_n(\theta)$ also converges pointwisely to $h(\theta)$ in $[0,1]$. So, for all $\theta \in [0,1]$ there is a subsequence of $\{h_n(\theta)\}$ that increases monotonically to $h(\theta)$. This allows us to construct o subsequence $\{h_{n_k}\}$ of $\{h_n\}$ with the property that $h_{n_k}(\theta)$ increases monotonically to $h(\theta)$ for $\forall\theta \in D$, where D is a dense subset of the $[0,1]$ interval—just apply a standard diagonal method to obtain a subsequence $\{h_{n_k}\}$ of $\{h_n\}$ with the desired property on the (countable) set of rational numbers in $[0,1]$. But all h_n are continuous functions of θ, and so $\{h_{n_k}\}$ must be monotonically increasing in all of $[0,1]$. We can then conclude that $\{h_{n_k}\}$ converges uniformly to h, since every sequence of functions that converges pointwisely and monotonically in a compact set must be uniformly convergent (see Rudin, Principles of Mathematical Analysis, Theorem 7.13).

In fact, this reasoning implies that any subsequence of $\{h_n\}$ has a uniformly convergent subsequence. Let $E = \{h_n\}$. Then E is a set of continuous functions defined on a compact set with the property that every sequence in E has a uniformly convergent subsequence. This implies that E is equicontinuous. Therefore $\{h_n\}$ is an equicontinuous sequence of functions defined on a compact set that converges pointwisely. This allows us to conclude (see Rudin, Principles of Mathematical Analysis, pg. 168) that $\{h_n\}$ converges uniformly.

(5) The only thing left to do is to prove the claim made in **(4)** that we can assume f_n to satisfy

$$\theta f_n(1) + (1-\theta) f_n(0) \geq E f_n(\theta^1(\theta, n)) \quad (1)$$

for every $n \in \mathbf{N}$ and for all $\theta \in [0,1]$. Notice first of all that the set $\mathcal{Y} = \{f \in \mathcal{S} \mid f$ satisfies (1) for all $n \in \mathbf{N}\}$ is a closed subset of $\mathcal{S}$: remember that if $\{f^m\}$ is a convergent sequence in $\mathcal{S}$, the topology we defined on $\mathcal{S}$ implies that $\{f_n^m\}$ is uniformly convergent for all n; a standard $\frac{\epsilon}{3}$ argument then allows us to prove this fact. So we just have to prove that T maps

$\mathcal{Y}$ into itself. Since

$$Tf_n(x) = \max\left\{\frac{a}{1-\beta}, \widehat{V}_n^E(x) + \beta E f_n(\theta^{(1)}(x,n))\right\}$$

for $x \in \{0,1\}$, we have that

$$\begin{aligned} \theta Tf_n(1) + (1-\theta)Tf_n(0) &\geq \max\left\{\frac{a}{1-\beta}, \widehat{V}_n^E(\theta) + \beta[\theta f_n(1) + (1-\theta)f_n(0)]\right\} \\ &\geq \max\left\{\frac{a}{1-\beta}, \widehat{V}_n^E(\theta) + \beta E f_n(\theta^1(\theta,n))\right\} \\ &= Tf_n(\theta), \end{aligned}$$

by hypothesis. So, if we let $H(\theta) = \theta Tf_n(1) + (1-\theta)Tf_n(0)$,

$$EH(\theta^1(\theta,n)) = \theta Tf_n(1) + (1-\theta)Tf_n(0) \geq ETf_n(\theta^1(\theta,n)),$$

which shows that $\mathcal{Y}$ is indeed mapped into itself by T. One last remark is that $\mathcal{Y}$ is not empty, even when we restrict ourselves to the subset of $f \in \mathcal{S}$ where f_n is non-increasing in θ. Just take $f_n =$ constant, for example.

The argument in **(1)**–**(5)** shows that the sequence $\{v_n\}$ of value functions is uniformly convergent. From **(4)** we can conclude that its uniform limit must be

$$\mathsf{V}(\theta) = \max\left\{\frac{a}{1-\beta}, \widehat{V}^E(\theta) + \beta\left[\theta\frac{a}{1-\beta} + (1-\theta)\frac{\widehat{V}^E(m_L)}{1-\beta}\right]\right\}.$$

□

Corollary 2.1 $\exists\rho' < 1$ *such that* $\theta^G(n) \leq \rho'$ *for all* n.

Proof. Suppose, by contradiction, that $\theta^G(n) > \rho$, with $\rho > \widetilde{\theta}$, for infinitely many $n's$. Since $\theta^G(n) < 1$ for all n, there exists a converging subsequence of $\{\theta^G(n)\}$ with the property that its limit is $\theta \geq \rho$. Denote it by $\theta^G(n_k)$. Now observe that:

$$v_{n_k}(\theta^G(n_k)) = \widehat{V}^E(\theta^G(n_k)) + \beta E v_{n_k}(\theta^{(1)}(\theta^G(n_k), n_k)) = \frac{a}{1-\beta}.$$

Since $\{\upsilon_n\}$ converges uniformly to $\mathsf{V}(\theta)$, we have that

$$\lim_{k\to\infty} \upsilon_{n_k}(\theta^G(n_k)) = \mathsf{V}(\theta).$$

Here we are using the fact that if $\{f_n\}$ is uniformly convergent to f and $\{x_n\}$ is a sequence of points that converges to x, then $\lim_{n\to\infty} f_n(x_n) = f(x)$. Because $\{h_n\}$, with $h_n(\theta) = E\upsilon_n(\theta^{(1)}(\theta, n))$, converges uniformly to $h(\theta) = \theta\mathsf{V}(1) + (1-\theta)\mathsf{V}(0)$, we also have that

$$\lim_{n\to\infty} \upsilon_{n_k}(\theta^{(1)}(\theta^G(n_k), n_k)) = \theta\mathsf{V}(1) + (1-\theta)\mathsf{V}(0)$$

Hence, since $\theta \geq \rho > \widetilde{\theta}$, we have

$$\frac{a}{1-\beta} = \upsilon_{n_k}(\theta^G(n_k)) \to \widehat{V}^E(\theta) + \beta\left[\theta\frac{a}{1-\beta} + (1-\theta)\frac{\widehat{V}(m_L)}{1-\beta}\right] < \frac{a}{1-\beta},$$

a contradiction. Therefore, $\exists\rho < 1$ and $\exists\overline{n}$ such that if $n > \overline{n}$, then $\theta^G(n) \leq \rho$. Now let $\rho' = \max\{\theta^G(1), \ldots, \theta^G(\overline{n}-1), \rho\}$. Since we know that $\theta^G(n) < 1$ for all $n \in \mathbf{N}$, we have that $\rho' < 1$ and $\theta^G(n) \leq \rho'$ for all n. □

Proposition 2.8 *For every $\delta > (m_H - m_L)/m_H$, there exists $n(\delta)$ such that, for all $n \geq n(\delta)$, $U(m_L, \delta, n) > U(m_H, \delta, n)$.*

Proof. We first need to prove that $\forall t \geq 2$, $\mu_t(m_H, n) \to 0$ and $\mu_t(m_L, n) \to 1$ as $n \to \infty$.

(1) We know, from the proof of proposition 2.4, that

$$\mu_t(m_H) \leq \Pr\{S_{(t-1)(n+1)}(m_H) \leq \lfloor (t-1)\alpha(n+1) + \gamma(\theta^G(n))\rfloor\},$$

where, as we have seen before,

$$\gamma(\theta^G(n)) = \frac{\ln\left(\frac{\theta^G(n)(1-\theta_0)}{\theta_0(1-\theta^G(n))}\right)}{\ln\left(\frac{1-m_L}{1-m_H}\cdot\frac{m_H}{m_L}\right)},$$

and $S_{(t-1)(n+1)}(m_H)$ denotes the number of successes in $(t-1)(n+1)$ Binomial trials with probability of success being given by m_H. Since, by proposition 2.2 and corollary 2.1,

$\overline{\theta} < \theta^G(n) < \rho$ for all $n \in \mathbf{N}$, where $\rho < 1$, we have that $\{\gamma(\theta^G(n))\}$ is bounded. Therefore, by applying the law of large numbers for the Binomial distribution in exactly the same way as we did in the proof of proposition 2.4, we can conclude that, indeed, $\mu_t(m_H, n) \to 0$ as $n \to \infty$ for all $t \geq 2$.

(2) To prove that $\mu_t(m_L, n) \to 1$ as $n \to \infty$ for all $t \geq 2$, we need a different argument. From the proof of proposition 2.3 we know that for all $t \geq 1$,

$$\mu_t(m_L, n) - \mu_{t+1}(m_L, n) = \sum_{(c_1,\ldots,c_t)\in\overline{C}_t} \binom{n+1}{c_1} \cdots \binom{n+1}{c_t} m_L^{c_1+..+c_t}(1-m_L)^{t(n+1)-\ldots c_t},$$

where $\overline{C}_t = \{(c_1, \ldots, c_t) \mid (c_1, \ldots, c_{t-1}) \in C_{t-1},\ c_t \geq \lfloor t\alpha(n+1) + \gamma(\theta^G(n)) - c_1 - \cdots - c_{t-1}\rfloor + 1\}$. Therefore,

$$\begin{aligned}
\mu_t(m_L, n) - \mu_{t+1}(m_L, n) &\leq \sum_{\lfloor t\alpha(n+1)+\gamma(\theta^G(n))\rfloor+1}^{t(n+1)} \binom{t(n+1)}{c} m_L^c(1-m_L)^{t(n+1)-c} \\
&= \Pr\{S_{t(n+1)}(m_L) \geq \lfloor t\alpha(n+1) + \gamma(\theta^G(n))\rfloor + 1\} \\
&\leq \Pr\{S_{t(n+1)}(m_L) \geq t\alpha(n+1) + \gamma(\theta^G(n))\}.
\end{aligned}$$

Since $m_L < \alpha$, we can once more apply the law of large numbers for the Binomial distribution, and obtain that for all $t \geq 1$, $\mu_t(m_L, n) - \mu_{t+1}(m_L, n) \to 0$ as $n \to \infty$. Since $\mu_1(m_L, n) = 1$ for all n, we can then conclude that $\mu_t(m_L, n) \to 1$ for all $t \geq 2$.

Now we're ready to prove our claim. For this, suppose $\delta > (m_H - m_L)/m_H$ and let

$$\epsilon = \frac{m_L}{1-\delta} - m_H > 0.$$

Since $\mu_2(m_H, n) \to 0$, $\exists n_1(\delta)$ such that if $n \geq n_1(\delta)$, then $\mu_2(m_H, n) \leq \frac{\epsilon(1-\delta)}{4m_H\delta}$. Moreover, since $\mu_{t+1}(m_H, n) \leq \mu_t(m_H, n)$ for all t and n, we have, in fact, that if $n \geq n_1(\delta)$, then

$\mu_t(m_H, n) \le \frac{\epsilon(1-\delta)}{4m_H\delta}$ for all $t \ge 2$. Therefore,

$$U(m_H, \delta, n) = \sum_{t=1}^{\infty} \delta^{t-1} m_H \mu_t(m_H, n) v < \left(m_H + \frac{\epsilon}{4}\right) v$$

whenever $n \ge n_1(\delta)$. We now need to find an appropriate lower bound for $U(m_L, \delta, n)$. For this, let $n_2(\delta, t)$ be such that if $n \ge n_2(\delta, t)$, then $\mu_t(m_L, n) > 1 - \frac{\epsilon}{4Nm_L\delta^{t-1}}$, where N is such that

$$\frac{1-\delta^{N+1}}{1-\delta} m_L - m_H - \frac{\epsilon}{2} > 0.$$

We know that such an N exists by hypothesis. The existence of an $n_2(\delta, t)$ with the desired property is guaranteed by (2). Now let $n_2(\delta) = \max\{n_2(\delta, t) \mid t = 2, \ldots, N\}$. If $n \ge n_2(\delta)$,

$$U(m_L, \delta, n) = \sum_{t=1}^{\infty} \delta^{t-1} m_L \mu_t(m_L, n) v > \left(\frac{m_L}{1-\delta}(1-\delta^{N+1}) - \frac{\epsilon}{4}\right) v.$$

Hence, if $n \ge n(\delta)$, with $n(\delta) = \max\{n_1(\delta), n_2(\delta)\}$,

$$U(m_L, \delta, n) - U(m_H, \delta, n) = \left(\frac{1-\delta^{N+1}}{1-\delta} m_L - m_H - \frac{\epsilon}{2}\right) v > 0.$$

□

Proposition 2.9 *For all $n \in N$, $\upsilon_n(\theta, \lambda, \phi)$ is a decreasing function of θ.*

Proof. We have to show that the mapping $T : \mathcal{C}[0,1]^3 \to \mathcal{C}[0,1]^3$ given by

$$Tf(\theta, \lambda, \phi) = \max\left\{a + f((1-\lambda)\theta + \lambda\phi, \lambda, \phi), \widehat{V}_n^E(\theta) + \beta Ef(\theta^{(1)}(\theta, \lambda, \phi), \lambda, \phi)\right\}$$

is a contraction, and sends the set $\mathcal{D} = \{f \in \mathcal{C}[0,1]^3 \mid f : \text{is:non-increasing:in:}\theta\}$ into itself. For the sake of brevity, we are omitting n from now on. That T is a contraction follows immediately from the fact that

$$\begin{aligned} Tf(\theta, \lambda, \phi) - Tg(\theta, \lambda, \phi) : \quad \le \quad & \beta \max\{f((1-\lambda)\theta + \lambda\phi, \lambda, \phi) - g((1-\lambda)\theta + \lambda\phi, \lambda, \phi), \\ & Ef(\theta^{(1)}(\theta, \lambda, \phi), \lambda, \phi) - Eg(\theta^{(1)}(\theta, \lambda, \phi), \lambda, \phi)\} \end{aligned}$$

and

$$|Ef(\theta^{(1)}(\theta,\lambda),\lambda) - Eg(\theta^{(1)}(\theta,\lambda),\lambda)| \leq \| f-g \|_{:\sup},$$

where $\| \, . \, \|_{:\sup}$ is the sup-norm on $\mathcal{C}[0,1]^3$. To show that T sends $\mathcal{D}$ into itself we have to check that for all $\lambda \in [0,1]$, $\theta^{(1)}(\theta',\lambda,\phi)$ first order stochastically dominates $\theta^{(1)}(\theta,\lambda,\phi)$ when $\theta' > \theta$. If we look at the proof of proposition 1, the facts needed to show that $\theta^{(1)}(\theta',0,\phi)$ first order stochastically dominates $\theta^{(1)}(\theta,0,\phi)$ when $\theta' > \theta$, were that the next period's beliefs are an increasing function of θ, the present belief, and of the number of meetings in the present period. But these 2 facts hold for any value of λ and ϕ, and so the desired result holds. Consequently, $v(\theta,\lambda,\phi)$ is decreasing in θ, since the maximum between two functions that are decreasing in θ is also decreasing in θ. □

Lemma 2.1 *(i)* $\nu_t(\theta,m_L) \geq \nu_t(\theta,m_H)$, for all $\theta \in (0,1)$ and for all $t \geq 1$

(ii) $\nu_1(\theta,m_L) - \nu_1(\theta,m_H)$ increases with θ in $[0,\theta^G(\lambda)]$.

Proof. The first thing that we need to do is obtain an expression for the probabilities $\nu_t(\theta,m)$, $m \in \{m_L,m_H\}$. Following the reasoning presented in the text before the statement of this lemma, it is easy to see that

$$\nu_1(\theta,m) = \sum_{\{c \leq \lfloor \alpha(n+1)+\gamma(\theta) \rfloor\}} \binom{n+1}{c} m^c(1-m)^{n+1-c} = \Pr\{\theta^1(\theta,\lambda) \leq \theta^G(\lambda) \mid m\}.$$

Now let $\Omega = \{S : |S| < \infty : \text{and} : S \subset [0,1]\}$, the set of all finite subsets of $[0,1]$, and consider the correspondence $\chi : \Omega \rightarrow \Omega$ such that

$$\chi(\Omega) = \bigcup_{\theta \in S} \text{supp}\{\theta^{(1)}(\theta,\lambda)\}.$$

The correspondence χ maps any finite set of beliefs S into the set of next period's beliefs that can arise from the elements of S. Then we have that

$$\nu_2(\theta,m) = \sum_{\tilde{\theta} \in \chi(\{\theta\}) \cap [0,\theta^G]} \sum_{\{c \leq \lfloor \alpha(n+1)+\gamma(\tilde{\theta}) \rfloor\}} \binom{n+1}{c} m_L^c(1-m_L)^{n+1-c} \Pr\{\theta^{(1)}(\theta,\lambda) = \tilde{\theta} | m\},$$

where $\Pr\{\theta^{(1)}(\theta,\lambda)=\tilde{\theta}|m\}$ is the probability that the next period's belief is $\tilde{\theta}\in\chi(\{\theta\})$, given that the bank chooses $m\in\{m_L,m_H\}$ in the present period. The rationale for this expression is almost identical to the one given in section 3 to justify the expressions for $\mu_t(m)$. Anyway, we are going to present it here: If in the next period ($t=1$) the agent has a prior $\tilde{\theta}\in\chi(\{\theta\})\cap[0,\theta^G]$, so that he decides to enter the economy, his probability of staying in the economy for one more period ($t=2$) is

$$\sum_{\{c\le\lfloor\alpha(n+1)+\gamma(\tilde{\theta})\rfloor\}}\binom{n+1}{c}m_L^c(1-m_L)^{n+1-c}=p(\tilde{\theta}),$$

since the bank is assumed to be choosing m_L from $t=1$ on. To obtain $\nu_2(\theta,m)$ we must multiply $p(\tilde{\theta})$ by the probability that $\tilde{\theta}$ happens, and then sum over all possible beliefs in $t=1$ for which the agent decides to stay in the economy.

If we let

$$\mathcal{R}^{(n)}(\{\theta\})=\chi(\mathcal{R}^{(n-1)}(\{\theta\}))\cap[0,\theta^G],$$

with $\mathcal{R}^{(0)}(\{\theta\})=\{\theta\}$, the above reasoning in fact shows that for all $t\ge 2$,

$$\nu_t(\theta,m)=\sum_{\tilde{\theta}\in\mathcal{R}^{(t-1)}(\{\theta\})}\sum_{\{c\le\lfloor\alpha(n+1)+\gamma(\tilde{\theta})\rfloor\}}\binom{n+1}{c}m_L^c(1-m_L)^{n+1-c}\Pr\{\theta^{(t-1)}(\theta,\lambda)=\tilde{\theta}|m\}$$

where $\Pr\{\theta^{(t-1)}(\theta,\lambda)=\tilde{\theta}|m\}$ is the probability that $t-1$ periods from now on the agent's belief is $\tilde{\theta}$, given that the bank chooses $m\in\{m_L,m_H\}$ in the present period and m_L thereafter.

Now we are ready to prove (i) and (ii).

(i) We will first show that if $\theta_1>\theta_2$, then $\theta^{(t)}(\theta_1,\lambda)|m$ first order stochastically dominates $\theta^{(t)}(\theta_2,\lambda)|m$ for all t and $m\in\{m_L,m_H\}$. For this, note that the statement is true when $t=1$ (see proof of proposition 2.9). Suppose then, by induction, that it is true for some

$k \in \mathbf{N}$. Since

$$\Pr\{\theta^{(k+1)}(\theta,\lambda) = x|m\} = \sum_{\theta' \in \mathcal{R}^{(1)}(\{\theta\})} \Pr\{\theta^{(k)}(\theta',\lambda) = x|m_L\} \Pr\{\theta^{(1)}(\theta,\lambda) = \theta'|m\}$$

we have that

$$\begin{aligned} \Pr\{\theta^{(k+1)}(\theta,\lambda) \le x|m\} &= \sum_{\tilde{\theta}\in[0,x]} \sum_{\theta' \in \mathcal{R}^{(1)}(\{\theta\})} \Pr\{\theta^{(k)}(\theta',\lambda) = x|m_L\} \Pr\{\theta^{(1)}(\theta,\lambda) = \theta'|m\} \\ &= \sum_{\theta' \in \mathcal{R}^{(1)}(\{\theta\})} \Pr\{\theta^{(k)}(\theta',\lambda) \le x|m_L\} \Pr\{\theta^{(1)}(\theta,\lambda) = \theta'|m\}. \end{aligned}$$

But, by our induction hypothesis, we know that $\Pr\{\theta^{(k)}(\theta',\lambda) \le x|m_L\}$ is a decreasing function of θ'. Therefore, since the result is true for $t = 1$, we have that $\Pr\{\theta^{(k+1)}(\theta',\lambda) \le x|m\}$ is also decreasing in θ. So, by induction, $\Pr\{\theta^{(k)}(\theta',\lambda) \le x|m\}$ is decreasing in θ for all t (and all x).

Now observe that $\gamma(\theta)$ decreases when θ increases, and so

$$\Psi(\tilde{\theta}) = \sum_{\{c \le \lfloor \alpha(n+1)+\gamma(\tilde{\theta})\rfloor\}} \binom{n+1}{c} m_L^c (1-m_L)^{n+1-c}$$

is decreasing in $\tilde{\theta}$ (as the set of feasible cardinalities gets smaller). Moreover, we have that

$$\nu_t(\theta,m) = \sum_{\tilde{\theta}\in\mathcal{R}^{(t-1)}(\{\theta\})} \Psi(\tilde{\theta}) \Pr\{\theta^{(t-1)}(\theta,\lambda) = \tilde{\theta}|m\} = E[\Psi(\theta^{(t-1)}(\theta,\lambda)|m)].$$

This, together with what we established in the above paragraph, allows us to conclude that $\nu_t(\theta,m)$ must be decreasing in θ. Also observe that $\theta^{(1)}(\theta,\lambda)|m_H$ first order stochastically dominates $\theta^{(1)}(\theta,\lambda)|m_L$. In fact,

$$\Pr\{\theta^{(1)}(\theta,\lambda) \le x|m\} = \sum_{\{c \le \lfloor \alpha(n+1)+\gamma(x,\theta)\rfloor\}} \binom{n+1}{c} m^c (1-m)^{n+1-c} = f(m,x),$$

for $m \in \{m_L, m_H\}$, and from the proof of proposition (2.3), we have that $\partial f/\partial m < 0$. Note that this also proves that $\nu_1(\theta,m_L) \ge \nu_1(\theta,m_H)$. To finish, note that for all $t \ge 2$,

$$\nu_t(\theta,m) = \sum_{\theta' \in \mathcal{R}^{(1)}(\{\theta\})} \nu_{t-1}(\theta',m_L) \Pr\{\theta^{(1)}(\theta,\lambda) = \theta'|m\}.$$

Therefore, from the fact that $\nu_{t-1}(\theta', m_L)$ is decreasing in θ' and $\theta^{(1)}(\theta, \lambda)|m_L$ is first order stochastically dominated by $\theta^{(1)}(\theta, \lambda)|m_H$, we can conclude that $\nu_t(\theta, m_L) \geq \nu_t(\theta, m_H)$ for all $t \geq 2$. This proves (i).

(ii) Note that

$$\begin{aligned}\nu_1(\theta, m_L) - \nu_1(\theta, m_H) &= \sum_{\{c \leq \lfloor \alpha(n+1)+\gamma(\theta) \rfloor\}} \binom{n+1}{c} m(c) \\ \text{where } m(c) &= \left[m_L^c(1-m_L)^{n+1-c} - m_H^c(1-m_H)^{n+1-c}\right]\end{aligned}$$

Note also that

$$\tilde{\theta}^G(\lambda) = \frac{\theta^G(\lambda) - \lambda\theta_0}{1-\lambda} > \frac{\theta^G(\lambda) - \lambda\theta^G}{1-\lambda} = \theta^G(\lambda),$$

because $\theta_0 < \theta^G(\lambda)$. Hence, for any $\theta \in [0, \theta^G(\lambda)]$ we have that $\gamma(\theta)$ is positive. Moreover, we know that $\gamma(\theta)$ is decreasing in θ. Since $m_L^c(1-m_L)^{n+1-c} < m_H^c(1-m_H)^{n+1-c}$ if, and only if $c > \alpha(n+1)$, we can then conclude that if θ increases, $\nu_1(\theta, m_L) - \nu_1(\theta, m_H)$ must decrease. $\square$

Bibliography

[1] Banks, J. and Sundaran, R. (1992) Denumerable-Armed Bandits, Econometrica, 60, 1071-1096

[2] Corbae, D.; Temzelides, T. and Wright, R. (2001) Matching and Money, Working Paper.

[3] Ellison, G. (1994) Cooperation in the Prisoner's Dilemma with Anonymous Random Matching, Review of Economic Studies, 61, 567-588.

[4] Feller, W. (1971) An Introduction to Probability Theory and its Applications, John Willey and Sons.

[5] Friedman, M. (1959) A Program for Monetary Stability, Fordham University Press.

[6] Jones, L. and Manuelli, R. (2001) Volatile Policy and Private Information: The Case of Monetary Shocks, Journal of Economic Theory, 99 265-*296.*

[7] Kandori, M. (1989) Social Norms and Community Enforcement, CARESS Working Paper, (89-14), University of Pennsylvania.

[8] Kandori, M. (1992) Social Norms and Community Enforcement, Review of Economic Studies, 59, 63-80.

[9] Katzman, B.; Kennan, J. and Wallace, N. (2002) Output and Price Level Effects of Monetary Uncertainty in a Matching Model, forthcoming Journal of Economic Theory.

[10] Kiyotaki, N., Wright, R. (1989) On Money as a Medium of Exchange, Journal of Political Economy, 97, 927-954.

[11] Kiyotaki, N., Wright, R. (1993) A Search-Theoretic Approach to Monetary Economics, American Economic Review, 83(1), 63-77.

[12] Kocherlakota, N. (1998a) The Technological Role of Fiat Money, Federal Reserve Bank of Minneapolis Quarterly Review, Vol. 22(3), 2-10.

[13] Kocherlakota, N. (1998b) Money is Memory, Journal of Economic Theory, Vol.81(2), 232-51.

[14] Landa, J. T. (1994) The Enigma of the Kula-Ring: Gift-Exchange and Primitive Law and Order, in: Trust, Ethnicity and Identity: Beyond the New Institutional Economics of Ethnic Trading Networks, Contract Law and Gift-Exchange, Landa, J.T., Economics, Cognition and Society Series, Ann Arbor: University of Michigan Press, 141-70.

[15] Mailath, G. and Samuelson, L. (1998) Your Reputation Is Who You're Not, Not Who You'd Like To Be, CARESS Working Paper, University of Pennsylvania.

[16] Mailath, G. and Samuelson, L. (2001) Who wants a Good Reputation?, Review of Economic Studies, 68, 415-441.

[17] Okuno-Fujiwara, M. and Postlewaite, A. (1995) Social Norms and Random Matching Games, Games and Economic Behavior, Vol.09(1), 79-109.

[18] Ritter, J. (1995) The Transition from Barter to Fiat-Money, American Economic Review 85(1), 134-149.

[19] Rothschild, M. (1974) A two-armed Bandit Theory of Market-Pricing, Journal of Economic Theory 9, 185-202

[20] Rudin, W. (1991) Functional Analysis, Second Edition, McGraw-Hill.

[21] Rudin, W. (1976) Principles of Mathematical Analysis, Third Edition, McGraw-Hill.

[22] Rupert, P., Schindler, M., Wright, R. (2000) Generalized Search-Theoretic Models of Monetary Exchange, forthcoming, Journal of Monetary Economics.

[23] Timberlake, R. (1987) Private Production of Scrip-Money in the Isolated Community, Journal of Money, Credit and Banking, 19(4), 437-447.

[24] Tullock, G. (1957) Paper Money-A Cycle in Cathay, Economic History Review IX, 393-407.

[25] Wallace, N. (1997) Short-run and Long-run Effects of Changes in Money in a Random Matching Model, Journal of Political Economy, 105, 1293-1307.

[26] Wolinsky, A. (1990) Information Revelation in a Market with Pairwise Meetings, Econometrica, 58, 1-23.