Monetary Equilibrium with Decentralized Trade and Learning*

Luis Araujo and Braz Camargo
University of Pennsylvania

November 16, 2001

Abstract

We consider an environment where trade is decentralized and agents only obtain information about the state of the economy (the amount of money in circulation) through their personal experiences in the market. We describe the individual agent decision as a two-armed bandit problem and characterize the dynamics of the market under different regimes. We study how the bank's (the agent in charge of money creation) incentives to choose a regime depend on the degree of information transmission in the economy. The main result is that, under rapid dissemination of information, society's ability to monitor the level of money in circulation puts a strong constraint on the bank's willingness to issue too much money. We also study the intertemporal consistency of the bank's behavior and find that a patient bank must not only be concerned with building a good reputation, but also with continuously behaving so as to separate itself from impatient banks in order to maintain its good reputation.

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.

---

[*] We are especially grateful to Randall Wright, George Mailath and Jan Eeckhout. We also benefited from discussions with Andrew Postlewaite, Brett Norwood, Dan Silverman and Brandon Weber.

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. In particular, if we believe that the government is less tempted to overissue than private agents, it is desirable to have government monopoly or regulation over the amount of money in circulation.

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 will prefer not to overissue in order to avoid the breakdown of monetary trade. It still may be the case that very impatient banks prefer to overissue. However, money from such banks will last only few periods in circulation.

Society's ability to gather information and learn about the state of the economy changes over time. In modern economies, the 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 the government is more patient than private agents and that it played a central role in the past, regulating money creation in a poorly informed economy. However, such intervention may not be necessary in modern economies.

We describe the market in our economy using a random search model (see Kiyotaki and Wright (1989,1993). 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 a double coincidence of wants problem. 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. Therefore, there are no informational frictions with respect to the actual value of money. In our paper we will relax this last 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. In other words, agents only obtain information from the meetings in which they participate, that is, from their private histories. In this sense, our model lies in the same tradition as Wolinsky (1990), which also analyzes information transmission in pairwise meetings. We believe that search models of money constitute a natural framework to address incomplete information in a market economy. The same technology that restricts and enables trade in these settings, i.e., random and pairwise meetings, 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 have to take into account what decisions an agent makes after every possible history. Though potentially this is a complex task, it turns out that we can solve for the agent's optimal decision rule by reinterpreting our problem in terms of a two-armed bandit problem, which is standard in the experimentation and learning literature (see Rothschild (1974) and Banks and Sundaram (1992)). We obtain that maximizing agents will adopt a cut-off rule. There exists a threshold with the property that, if an agent has a belief that the bank overissues greater than this threshold, he does not accept money and moves to autarky. Alternatively, if his belief is less than or equal to the threshold, he accepts money and stays in 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 described the agent's behavior in the limit, 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

---

[1] There is a growing literature of search models dealing with incomplete information regarding the amount of money in circulation. For example, Cavalcanti, Erosa and Temzelides (1999) assume that the history of transactions of a subset of agents in the economy (banks) is public. Katzman, Kennan and Wallace (2001) 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 (2000).

measure of agents will 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 the beliefs 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.

Throughout the paper 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 will have incentives to not overissue only until agents become sufficiently convinced that he will never overissue in the future, i.e., until the agents posterior probability that the bank is patient is very close to one. At this point, the effect of 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. However, if there is a continual exogenous probability that the bank can become impatient, agent's posterior will be bounded away from one. In this situation, a sufficiently patient bank will never overissue in an attempt to distinguish itself from the impatient bank.[2]

We can summarize this 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 his 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 the paper is as follows. In the next section, we describe the physical and informational environment. In section 3 we discuss the agent's problem. In section 4 we describe the dynamics of the economy under distinct monetary regimes. Section 5 analyzes the bank's problem and gives an example. Section 6 describes what happens when we allow the bank to change his behavior in every period. Section 7 concludes the paper.

---

[2] This reasoning follows Mailath and Samuelson (1998) and (2001).

2 Environment

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 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 he 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 (2001), comparing our results to the case where agents are completely informed about the value of m. In the end of this section we will analyze in greater detail the implications of the indivisibility of goods and money. Finally, the bank discounts the future at a rate δ.

Our objective in this paper is to study the interaction of agents and the bank under incomplete information, with a focus on the bank's decision of how much money to issue. We do not attempt to analyze here how the bank uses the resources he obtains from agents. For this reason we believe that the idea of a bank as a goods-holder suffices for our purposes.

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. [3] 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,

---

[3] The random matching assumption inside a sector can be interpreted as the cost associated with waiting in line to buy a good, for example.

if he has money he goes to the sector that trades the good he likes and search for an agent that has 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.

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. Notice that, since there is no discounting between meetings, even if an agent has the option of leaving the market earlier in order to redeem his note in the bank, he prefers not to do so.

The number of meetings n is a key parameter in our model. It will be 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. However, the number of meetings also increases the agent's flow payoff (more meetings imply more opportunities of consumption). Therefore, whenever we study changes in n, in order to focus on the informational effect, we assume that any increment in the flow payoff is associated with an increment in the gain from autarky A.

We can summarize the decision problem faced by an agent in this economy in the following way: in the beginning of every period each agent decides between going to autarky or moving to the market. If he moves to autarky he can produce consumption goods and obtain utility A. If he goes to the market, he first has the possibility of passing in 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 have the opportunity to 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 Information

In a search environment, the frictions of the trading process are captured by the assumption that trade is decentralized, with agents being pairwise and randomly matched. We believe that a natural way of addressing informational frictions with respect to the value of money is to assume that the information transmission is subject to the same constraints as trade itself. In other words, agents can only gather information from the meetings in which they participate, that is, from their private histories.

In our economy the main decision faced by each agent is whether to accept money and participate

---

[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.

in trade or not. This decision crucially depends upon the agent's expectation about the value of m chosen by the bank. They will use their private histories in order to form a belief with respect to the bank's choice of m.

In our setup 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, to describe 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 will leave the market and move to autarky after realizing (through their histories) that there is too much money in circulation. We believe that the more general assumption that the bank can choose any value of m between zero and one will not bring any new results or insights. At the same time it will require the use of a more cumbersome notation.

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 will set $m = m_H$. In this case, we say the bank is of type H. Alternatively, we say the economy is under a high monetary regime. Analogously, a bank is of type L if he sets $m = m_L$ (and the economy is said to be under a low 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 his 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:

$$\begin{aligned}
\theta(h^t) &= \frac{m_H^{c(h^t)}(1-m_H)^{(t-1)(n+1)-c(h^t)}\theta_0}{m_H^{c(h^t)}(1-m_H)^{(t-1)(n+1)-c(h^t)}\theta_0 + m_L^{c(h^t)}(1-m_L)^{(t-1)(n+1)-c(h^t)}(1-\theta_0)}, \text{ for all } h^t \\
\theta(h^1) &= \theta_0 \qquad (1)
\end{aligned}$$

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.

2.3 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, a potentially complex object. We avoid all this complexity when money and goods are indivisible.

Katzman, Kennan and Wallace (2001) deal with bargaining issues in a search environment with incomplete information and divisible goods. However, in order to simplify the problem, they assume that 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 the intensive-margin, which refers to its effect over the quantity of goods traded in a given meeting. This 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 the 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 move to 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 our 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 will affect the extensive-margin by reducing the number of trade meetings. That is the content of assumption 1.

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 $i = 1, ..., n$. Let V_j^{n+1} indicate the value function of a meeting between an agent with j units of money and the bank after n rounds in the market. We have:

$$\begin{aligned} V_0^{n+1}(m) &= 0 \\ V_1^{n+1}(m) &= u \end{aligned}$$

and

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

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 following value functions:

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

Notice that an agent with money has a greater payoff than an agent with endowment. Hence, an agent going to the market will always accept an exchange of goods for money. An agent's expected payoff upon going to the market $V(m)$ is equal to $mV_1(m)$ $+(1-m)V_0(m)$. That is,

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

We stated before that the overall gain in autarky in every period is equal to A. We will establish some microfoundations on how this 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 will be 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 $\frac{A}{n}$. We have $\widehat{V}(m) = m(1-m)u + \frac{mu}{n}$ and $\widehat{A} = a$. This normalization is useful when we make comparisons of the expected value for distinct choices of n, and when $n \to \infty$.

As we discussed in the previous section, in what follows we are mainly 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 will 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 $m_H > m_L \geq \dfrac{n+1}{2n}$

Assumption 2 states that the values of m_H and m_L correspond, respectively, to non-existence and existence of monetary equilibrium under full-information. This full information scenario constitutes the benchmark of our model. We have:[5]

$$\textbf{Assumption 2} \quad m_H[(1-m_H)+\frac{1}{n}] < \frac{a}{u} \leq m_L[(1-m_L)+\frac{1}{n}]$$

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 moving to 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 moving to autarky. In the definition of S we are already taking into account the fact (which is proved below) that when an agent decides to move to autarky, his optimal decision is to stay in autarky forever. Hence, an agent is called upon to make a decision in a given period only in case he decided to stay in the market in all the previous ones.

The problem at hand is to determine the agent's optimal behavior. It turns out that the agents' problem belongs to a class of problems studied by probabilists under the general name of two-armed bandit problems: "Consider a gambler condemned to, or bent on, putting a quarter into one of several (by convention the number is usually two) slot-machines (one-armed bandits) from now on until the end of time. The *ith* machine gives a payoff of q_i with probability Π_i, and nothing with probability $1 - \Pi_i$. He wants to choose a strategy which will maximize his expected discounted earnings...However, it is assumed that, although he may have prior beliefs, the man does not know which machine offers the most favorable odds" (Rothschild (1974), pages 189 and 190). Our model considers a particular case where one of the arms (autarky) is deterministic and the other (the market) is stochastic. In this situation, if autarky is an optimal choice at a point in time, it will also be 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 going to autarky or staying in the market one more period. If he chooses autarky he receives a flow payoff of a, but he does not receive any additional

---

[5] Note that when $n \to \infty$, assumption 1 implies $m_H > m_L \geq \frac{1}{2}$ and assumption 2 implies $m_H(1-m_H) < \frac{a}{u} \leq m_L(1-m_L)$.

[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.

information with respect to the bank's type. Hence, if 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 will choose 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:

$$\Pr(\theta^{(1)}(c \mid \theta)) = \theta \Pr_H(N = c) + (1-\theta) \Pr_L(N = c) , \tag{2}$$

where

$$\begin{aligned} \Pr_H(N &= c) = \binom{n+1}{c} m_H^c(1-m_H)^{(n+1)-c} , \\ \Pr_L(N &= c) = \binom{n+1}{c} m_L^c(1-m_L)^{(n+1)-c} . \end{aligned}$$

We can write agent's value function in the following way:

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

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). The following proposition can be proved:

Proposition 1 *$\mathcal{V}(\theta)$ has the following properties:*
(i) $\mathcal{V}(\theta)$ is decreasing in θ;
(ii) $\mathcal{V}(0) = \frac{\widehat{V}(m_L)}{1-\beta}$, and $\mathcal{V}(1) = \frac{a}{1-\beta}$.

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

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

Moreover, since $\widehat{V}^E(0) = \widehat{V}(m_L) > a$, we can conclude that $\mathcal{V}(0) = \frac{\widehat{V}(m_L)}{1-\beta}$. 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 $\mathcal{V}(1) = \frac{a}{1-\beta}$. Here we were implicitly assuming that $\mathcal{V}(\theta)$ 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:

(1) If $f, g \in \mathcal{C}[0,1]$, with $f \leq g$, then $Ef(\theta^{(1)}(\theta)) \leq Eg(\theta^{(1)}(\theta))$, and so $Tf \leq Tg$.

Moreover, 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:

(2) $T(f+c)(\theta) \leq Tf(\theta) + \beta c$ for any $f \in \mathcal{C}[0,1]$.

Therefore, 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 $\mathcal{V}(\theta)$.

Moreover, if we show that $\theta < \theta'$ implies that $\theta^{(1)}(\theta')$ first order stochastically dominates $\theta^{(1)}(\theta)$, we can conclude 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, $\mathcal{V}(\theta)$, must be decreasing in θ.

□

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

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

and

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

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

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

Hence, for every θ such that $\theta \leq \theta^G$, since $\widehat{V}^E(\theta) + \beta E\mathcal{V}(\theta^{(1)}(\theta)) \geq \frac{a}{1-\beta}$, the agent stays in the market. Otherwise, he moves to 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, moves to autarky and stay in autarky forever. An interesting 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.

It is important to 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 moving to autarky, but also the fact that by entering the market the agent will obtain 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. We have:

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

From the definition of θ^G, and from the fact that $f\leq\mathcal{V}$, we know that:

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

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

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

which ends our proof.

$\square$

4 Aggregate Behavior

From the analysis of the agent's problem we can describe how the measure of people in the market evolves over time, under the low and the high monetary regimes. This is our objective in this section. More specifically, we are going to compute the probability that an agent will stay in the market up to period t, for every t. Since we are dealing with a continuum of agents we can reinterpret this probability as the measure of agents staying in the market up to that period.

Simple algebra with equation (1) allows us to write 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 will stay 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 $\tau \leq t$. Since we are interested in obtaining a non-trivial dynamic in our environment, from now on we assume that $\theta_0 < \theta^G$, so that all agents enter in the economy in the first period. The measure of agents staying in the economy up to period t (with $t > 1$) is then given by

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

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 intuition for this result is as follows. 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 will be willing to enter in the market again is given by the probability that he will face 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 will continue in the market in period 3 is equal to the probability that he will have 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

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

From these expressions we can obtain what is the probability that an agent entering in the market in period 1 will enter again in period 3. We obtain it 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 $c_1 \leq \alpha(n+1)+\gamma$. After some algebraic manipulations we have:

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

This probability describes the measure of agents that enter in the market in period 3, $\mu_3(m)$. Proceeding in a similar way, we can calculate the measure of agents entering in the economy for every value of t. The following proposition can be proved:

Proposition 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* $\mu_t(m_L) \geq \mu_t(m_H)$*.*
Moreover, for every γ *there exists a* $\bar{t}$ *such that, for all* $t > \bar{t}$*, the above inequalities are strict.*

Proof: See appendix.

Proposition 3 tells us that, independent of the monetary regime, the measure of agents engaged in market activities decreases over time. This result is a direct consequence of agents having incomplete information. In particular, 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 an agent will face a history inducing him to believe that there is too much money in the economy, even if the bank chose m_L. If we compare this situation to our benchmark of complete information, we see that the presence of uncertainty makes agents worse off. Some agents abandon trading under a low type regime, others stay in the market for more than one period under a high type regime. However, as we will see in the next section, the presence of incomplete information can make the bank better off.

Proposition 3 also shows that trade is more unstable in a high monetary regime. For a given value of γ, after a sufficiently large extent of time, the measure of agents leaving the market under a high monetary regime is strictly larger than the measure under a low monetary regime. For example, when agents have a strong prior that the bank is of a low type (which implies a large γ), an agent will leave the market only after collecting enough evidence through experimentation 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 will 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 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 3 we have that the exit rate of the market in the high monetary regime is strictly higher than in the low 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 will always be engaged in trade activities.

In the next section, we use these results to analyze the bank's behavior when it can make a one-time decision between m_L and m_H.

5 Bank's Behavior

Up to this point we assumed that agents have a common prior regarding the bank's type. In this section, we analyze the bank's behavior and give a more clear characterization of how the prior is formed. More specifically, we 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 $F(\delta)$.

As stated before, 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. We assume that for each unit of good in storage during a period the bank receives a utility v. If the bank decides to always give money to a fraction m of agents entering in the market, its utility can be written in the following way:

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

Propositions 3 and 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.

Proposition 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 are going to construct an equilibrium where the bank's behavior is described by the above proposition and the agent's behavior corresponds to the optimal decision rule obtained in Section 3. But first we characterize our environment in terms of a Bayesian game. A Bayesian game consists of a list $\{(\mathcal{I}, \mathcal{A}, \Theta, u), \xi\}$ where $\mathcal{I}$ is the set of players, $\mathcal{A}$ 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 move to autarky (A) or enter in 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 $\sum\limits_{t=1}^{\infty} \delta^{t-1} \mu_t(m) m v$ and the payoffs of the agents are given by $\mathcal{V}(\theta)$. We can then prove the following result:

Proposition 6 *Let* $0 < \theta_0 = F(\overline{\delta}) < \theta^G$, *and consider the strategy profile* $\Delta = (S^*, m^*)$, *where* $S^* = \{s_t^*\}$, *with* $s_t^* : \widehat{H}^t \to \{M, A\}$ *given by*

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

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

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

Then Δ *is a Bayesian Nash Equilibrium.*

Proof: From Proposition 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.

$\square$

This proposition describes one possible equilibrium in our model. If $\theta_0 = 1$, there is another equilibrium where the 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 the 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 will leave the economy no matter their beliefs. A third example is when the agents, after staying in the market for a number of periods, decide not to go to the market for a finite number of periods, returning to it after that.

However, we believe that the equilibrium described in proposition 6 is much more appealing than the other equilibria mentioned in the previous paragraph. 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 the 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 6 describes their behavior.

Proposition 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 will have $m = m_L$. However, with partial information, the bank can overissue money and still operate in the economy over time.

A consequence of this result is that incomplete information can bring a lot of instability to trade involving money. For the monetary system to be stable the bank must have interest in keeping money valuable as a medium of exchange over a long period of time. However, this interest depends on how the bank discounts the future, which is exogenously given. Since the economy has no mechanism that fully reveals the bank's policy, the society has to rely on the possibility that the bank is patient and care about its future earnings.

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 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

---

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

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 high monetary regime, the average value of $\theta(h^t)$ increases and, under a low 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 larger.

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 $\{\mathcal{V}_n\}$ of value functions depending on θ alone. From the results in proposition 1 we know that, for all $n \in N$, $\mathcal{V}_n$ is a decreasing function of θ with $\mathcal{V}_n(0) = \frac{\widehat{V}(m_L)}{1-\beta}$ and $\mathcal{V}_n(1) = \frac{a}{1-\beta}$. Another consequence of letting n vary is that the threshold belief θ^G will also depend 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 thing we would expect not to happen is that θ^G goes to one as 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 is $\widehat{V}^E(\theta)$ where:

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

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 will learn that the bank is of the high type and with probability $1-\theta$ he will learn 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

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

Now let $\widetilde{\theta}$ be such that $\mathcal{V}_n(\widetilde{\theta}) = \frac{a}{1-\beta}$. Since $\widehat{V}^E(1) < a$ it is easy to see that $\widetilde{\theta}$ must be less 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 7 *$\{\mathcal{V}_n\}$ converges uniformly to $\mathsf{V}(\theta)$, where*

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

Proof: See appendix.

Corollary 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:

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

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

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

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 $h_n(\theta) = E\mathcal{V}_n(\theta^{(1)}(\theta, n))$, converges uniformly to $h(\theta) = \theta\mathsf{V}(1) + (1-\theta)\mathsf{V}(0)$, we also have that

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

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

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

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.

□

We are now in conditions to prove the following proposition formalizing our very initial comparative statics argument in n.

Proposition 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., $\delta > \frac{m_H - m_L}{m_H}$.

In modern economies, the dissemination of information about decisions made by the bank or government is much faster than in the past. We can interpret this change as in increment in the number of opportunities an agent has to gather information about the state of the economy. In our model, this increment can be interpreted as an increase in the value of n. Therefore, as a result of the above proposition, we expect that in a modern society, the incentives of a bank to overissue are less pronounced than in past societies.

In what follows we consider a particular case of an economy where $m_H = 1$. This restriction allows us to obtain a closed form expression for the agent's and the bank's strategy under the Bayesian Nash Equilibrium described in proposition 6.

5.1 Example

First, we determine what are the agent's beliefs with respect to the bank's type for each possible history that he can face. It turns out, since $m_H = 1$, that for every history where at some point in time an agent meets another one without money (or the agent does not receive money from the bank), he will know with certainty that the bank is of the low type. We have:

$$\theta(h^t) = 0, \text{ for all } h^t \text{ such that } c(h^t) \neq (t-1)(n+1).$$

Hence, the only history where an agent is not sure about the bank's type is the one where in all his previous meetings he always met agents with money. In this case, his posterior will be

$$\theta(h^t) = \frac{\theta_0}{\theta_0 + m_L^{(t-1)(n+1)}(1-\theta_0)}, \text{ for } h^t \mid c(h^t) = (t-1)(n+1).$$

We can then find the value $\tau^* = (t^*-1)(n+1)$ of τ such that for all $t \leq t^*$, the agents will stay in the market. It is given by:

$$\tau^* = \max_{\tau \in \mathbf{N}} \left\{ \tau \leq \frac{\ln[\frac{\theta_0(1-\theta^G)}{\theta^G(1-\theta_0)}]}{\ln m_L} \right\}.$$

Given the beliefs, we know how agents behave after every possible history. Their optimal strategy $S^* = \{s_t^*\}$ will be such that

$$s_t(h^t) = M \text{ for } h^t \text{ such that } c(h^t) \neq (t-1)(n+1) ,$$

$$s_t(h^t) = \begin{cases} M \text{ for } h^t \text{ such that } c(h^t) = (t-1)(n+1) \text{ and } t \leq t^* \\ A \text{ for } h^t \text{ such that } c(h^t) = (t-1)(n+1) \text{ and } t > t^* \end{cases} ,$$

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

From the above strategy we can obtain how the measure of agents evolve over time for each bank's type. If the bank is of type H, the only history an agent faces is the one where he only meets agents with money. However, he decides to move to autarky only after $t > t^*$, since it takes some time until the beliefs change so as to make $\theta(h^t) > \theta^G$. This implies that

$$\mu_t(m_H) = 1, \text{ for all } t \leq t^*,$$

and

$$\mu_t(m_H) = 0, \text{ for all } t > t^*.$$

If the bank is of type L, the measure of agents that will face a history inducing them to move to autarky is equal to $m_L^{(t^*-1)(n+1)}$. All other histories in the set $\widehat{H}^t$ are such that an agent meets at least another agent without money (or he does not receive money from the bank), in which case he knows with certainty the type of the bank. Therefore,

$$\mu_t(m_L) = 1, \text{ for all } t \leq t^*,$$

and

$$\mu_t(m_L) = 1 - m_L^{(t^*-1)(n+1)}, \text{ for all } t > t^*.$$

We can see, as shown in proposition 4, that under a high type bank the measure of agents in the market goes to zero, while with a low type bank this measure remains positive.

We are now in condition to calculate the expected payoffs of a bank in each possible monetary regime. They are:

$$U(m_H, \delta) = (1 - \delta^{t^*})\frac{v}{1-\delta}$$

and

$$U(m_L, \delta) = \left[(1 - \delta^{t^*}) + \delta^{t^*}\left(1 - m_L^{(t^*-1)(n+1)}\right)\right]\frac{m_L v}{1-\delta}.$$

A bank will choose m_L if, and only if, $U(m_L, \delta) \geq U(m_H, \delta)$. This condition can be rewritten in terms of a lower bound for δ. Just observe that

$$U(m_L) \geq U(m_H) \Longleftrightarrow \delta \geq \left[\frac{1 - m_L}{1 - m_L^{(t^*-1)(n+1)+1}} \right]^{\frac{1}{t^*}}$$

We can then conclude that as long as the bank is sufficiently patient ($\delta \geq \delta^*$), it will prefer not to issue too much money, so that over time it can gain profits from a constant stream of agents entering in the market at every period. However, if the bank is impatient, even if the measure of agents in the market goes to zero when it issues too much money, it still prefers to do so, and obtain a large short-run profit.

6 Intertemporal Consistency of the Bank's Behavior

In the previous section we considered a simple model of the bank's behavior and discussed how the incentives to overissue change as a function of the degree of information in the economy. In what follows we are going to consider a distinct but related issue. We 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 follows Mailath and Samuelson (1998) and (2001), starting with the interpretation of reputation as the agent's posterior that the bank chooses m_L.

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 his future behavior. It is reasonable to expect that a very patient bank will choose 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 experiencing further histories with high cardinality, an agent will still be 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. The incentives to choose m_L in order to maintain it disappear. The bank is free to choose his one-period dominant action m_H. 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. The patient bank must have an incentive to continuously separate itself from the impatient bank and maintain a good 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 sort of a black box. It's internal structure was in no way relevant for our purposes. The only relevant variable was its discount factor δ, which was the bank's private information. We want to add some structure to it, and we're going to do so by assuming that the bank's discount rate is given by the discount factor of its current owner. Previously we were assuming that the owner was the same throughout the bank's infinite lifetime. Now we're going to assume that there is always a small positive probability λ that the owner is replaced in every period. The new owner corresponds to a new random draw of the discount factor's pool. For reasons that will become clear later on, we will assume that the agents' common prior about this new owner is $\phi \in [0, 1]$, with ϕ not necessarily equal to θ_0. Moreover, we are going to assume that this change of ownership is not observed by the agents in the economy. We will see that in this new setup the agents' beliefs will never reach a point where the bank's action will have a negligible influence on its future reputation For example, if an agent faces a very negative history he will attribute some probability to the event that the bank has changed its preferences. This will give a patient bank the right incentives to not overissue. Thus it is adequate to talk about creation and maintenance of reputation as an equilibrium phenomena. At the same time, we want the main features of the environment that we obtained so far to be maintained, and we can do so if we take λ sufficiently small.

In this new environment, since the bank is allowed to choose the amount of money it will issue in every period, the description of its strategy changes. 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 5 is dynamically consistent, we will restrict attention only to the strategy where a patient bank always chooses m_L and an impatient one always chooses m_H, irrespective of its past history. We will further restrict the set of possible discount factors to $[0, \delta_1] \cup [\delta_2, 1]$, with $0 < \delta_1 \leq \delta_2 < 1$. The reason for this is that it is possible, for intermediate discount factors, that any of the bank's optimal strategies will involve changing the amount of money issued over time. We will determine δ_1 and δ_2 later on. Formally, the bank's strategy that we are going to consider is $\mathcal{M}^* = \{m_t^*\}$ with $m_t^* : [0, \delta_1] \cup [\delta_2, 1] \times \{m_L, m_H\}^{t-1} \longrightarrow \{m_L, m_H\}$ given by:

$$m_t^*(\delta, .) = \begin{cases} m_L \text{ if } \delta \geq \delta_2 \\ m_H \text{ if } \delta \leq \delta_1 \end{cases} . \tag{3}$$

As before, a bank is of the high type if its discount factor is in $[0, \delta_1]$, otherwise it is of low type.

Our first objective is to determine the agents' optimal response to this particular strategy. For this we need to describe how the agents update their beliefs in this case; that is, when they believe that the bank's strategy is given by (3). 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:

$$\theta(h^t) = (1-\lambda) \frac{m_H^{c(h_t)}(1-m_H)^{(n+1)-c(h_t)}\theta(h^{t-1})}{m_H^{c(h_t)}(1-m_H)^{(n+1)-c(h_t)}\theta(h^{t-1}) + m_L^{c(h_t)}(1-m_L)^{(n+1)-c(h_t)}(1-\theta(h^{t-1}))} + \lambda\phi. \tag{4}$$

Observe that, since the agents' initial belief is θ_0, at all points in time all posteriors will be restricted to the interval $[\lambda\theta_0, 1-\lambda+\lambda\phi]$.

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

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

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 will see, however, that for ϕ sufficiently close to 1, this will be the case.

Proposition 9 *For all $n \in N$, $\mathcal{V}_n(\theta, \lambda, \phi)$ is a decreasing function of θ.*

Proof: See appendix.

Now take $\phi = 1$, and suppose that there exists $\theta' \in [0,1]$ such that the agent is indifferent between the market and autarky if his belief is $\theta = \theta'$. Then we have (omitting the subscript n) that

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

Since $(1-\lambda)\theta' + \lambda \geq \theta'$, and because of the above proposition,

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

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

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

since we have seen above that, if $\phi = 1$, whenever the agent is indifferent between going to the market and staying in the autarky, his value function must be equal to $\frac{a}{1-\beta}$. Therefore, since $\widehat{V}^E$ is strictly decreasing and (see the proof of proposition (9)) $E\mathcal{V}(\theta^{(1)}(\theta, \lambda, \phi), \lambda, \phi)$ is decreasing in θ,

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

a contradiction. Similarly, we can show that for no $\theta'' < \theta'$ we can have the agent indifferent between his 2 alternatives.

So, when $\phi = 1$ the equation

$$a + \beta\mathcal{V}((1-\lambda)\theta + \lambda, \lambda, \phi) = \widehat{V}^E(\theta) + \beta E\mathcal{V}(\theta^{(1)}(\theta, \lambda, \phi), \lambda, \phi) \tag{5}$$

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

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

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 (5) 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 (9), that $\mathcal{V}(\theta, \lambda, \phi)$ 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 $\phi < 1$. We have seen above that (5) has at least one solution. Denote one such solution by $\theta^G(\lambda, \phi)$. By continuity we know that $\theta^G(\lambda, \phi)$ converges to $\theta^G(\lambda, 1)$ as ϕ goes to 1 [8]. Since $\theta^G(\theta, \lambda) < 1$, it must be the case that there exists a $\overline{\phi} < 1$ with the property if $\phi \geq \overline{\phi}$, then $\theta^G(\lambda, \phi) \leq \phi$. Therefore, if we restrict ϕ to the interval $[\overline{\phi}, 1)$, we will 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 (5) had a unique solution when $\phi = 1$, to conclude that indeed it has a unique solution as long as we consider $\phi \geq \overline{\phi}$. Moreover, autarky will be absorbing, and the reason is exactly the same as when $\phi = 1$. One final remark is that, once more because $\mathcal{V}$ 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 $\phi \geq \overline{\phi}$. With this we have, as before, that the agent's optimal decision is to enter the market if his prior is less than or equal to a certain threshold belief $\theta^G(\lambda, \phi)$, and to move to autarky otherwise. Moreover, as argued above, autarky will be absorbing. 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 will enter 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]

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

---

[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 $n \in \mathbf{N}$ 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 $k \in \mathbf{N}$, $|\theta^G(\lambda, \phi_{n_k}) - \theta^G(\lambda, 1)| > \epsilon$. 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

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

and so, by the continuity of $\mathcal{V}$, we must have that

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

an absurd.

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

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

$$\begin{aligned} s_t^*(h^t) &= M \text{ if } \theta(h^t) \leq \theta^G(\lambda, \phi) \\ s_t^*(h^t) &= A \text{ if } \theta(h^t) > \theta^G(\lambda, \phi) \end{aligned} \tag{6}$$

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

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 will 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, he enters in the economy t periods in the future. The following is true:

Lemma 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: 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. Moreover, the bank's incentive to deviate gets stronger as θ diminishes. If an agent is really confident that the bank will not deviate, the probability that he enters the market in the future, even after experiencing bad histories, is high. In the extreme case where $\theta = 0$, the agent will 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 θ is bounded below by $\lambda\theta_0$, so that the bank always has a tradeoff between present and future gains when faced with the choice of m. Therefore, we expect that a sufficiently patient bank will prefer to choose m_L in order to increase its future gains. This result is proved in the proposition below. This proposition 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 will never have incentives to choose m_L if we take δ to be sufficiently small.

Proposition 10 *Let $\Delta' = (\mathcal{M}^*, S^*)$ be the strategy profile where $\mathcal{M}^*$ is the bank's strategy defined in* (3)*, and S^* is the agent's strategy defined in* (6)*. There exist $n_1 \in \mathbf{N}$, δ_1 and δ_2, $\delta_1 \leq \delta_2$, such that this strategy profile, along with the Bayesian updating rule given by* (4)*, constitutes a sequential equilibrium if the number n of non-bank meetings per period is greater or equal than n_1.*

Proof: See appendix.

7 Conclusion

This paper 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 overissue in an environment where agents only learn about the amount of money in circulation from their personal experience 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 will basically ignore their private histories and take for granted that the bank will never overissue. This is will eventually tempt 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, the bank will always choose m_L in order to maintain a good reputation and separate itself from the impatient banks.

References

[1] **J. BANKS and R. SUNDARAM**, Denumerable-Armed Bandits, *Econometrica*, **60** (1992) 1071-1096

[2] **R. CAVALCANTI, A. EROSA and T. TEMZELIDES**, Private Money and Reserve Management in a Random-Matching Model, *Journal of Political Economy*, **107** (1999), 929-945.

[3] **W. FELLER**, An Introduction to Probability Theory and its Applications, *John Willey and Sons* (1971).

[4] **L. JONES and R. MANUELLI**, Volatile Policy and Private Information: The Case of Monetary Shocks, *Working Paper* (2001).

[5] **B. KATZMAN, J. KENNAN and N. WALLACE**, Output and Price Level Effects of Monetary Uncertainty in a Matching Model, *Working Paper* (2001).

[6] **M. FRIEDMAN**, A Program for Monetary Stability, *Fordham University Press* (1959).

[7] **N. KIYOTAKI and R. WRIGHT,** On Money as a Medium of Exchange, *Journal of Political Economy*, **97**(1989), 927-954.

[8] **N. KIYOTAKI and R. WRIGHT,** A Search-Theoretic Approach to Monetary Economics, *American Economic Review* **83(1)** (1993), 63-77.

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

[10] **G. MAILATH and L. SAMUELSON**, Who wants a Good Reputation?, *Review of Economic Studies*, **68** (April 2001), 415-441.

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

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

[13] **W. RUDIN**, Functional Analysis, Second Edition, *McGraw-Hill* (1991).

[14] **W. RUDIN**, Principles of Mathematical Analysis, Third Edition, *McGraw-Hill* (1976).

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

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

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

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

8 Appendix

Proposition 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 $\bar{t}$ such that, for all $t > \bar{t}$, 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

$$B_{t-1} = \left(\begin{array}{c} n+1 \\ c_1 \end{array} \right) \cdots \left(\begin{array}{c} n+1 \\ c_{t-1} \end{array} \right),$$

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 then given by

$$\begin{aligned} \mu_t(m) - \mu_{t+1}(m) \;=\; L_t(m) \;=\; & \sum_{(c_1,\ldots,c_{t-1}) \in C_{t-1}} B_{t-1} m^{c_1+\cdots+c_{t-1}} (1-m)^{(t-1)(n+1)-c_1-\cdots-c_{t-1}} \times \\ & \left[1 - \sum_{c_t=0}^{\lfloor t\alpha(n+1)+\gamma \rfloor - c_1 - \cdots - c_{t-1}} \left(\begin{array}{c} n+1 \\ c_t \end{array} \right) m^{c_t} (1-m)^{n+1-c_t} \right]. \end{aligned}$$

We can rewrite the exit rate $L_t(m)$ as

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

where

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

Therefore,

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

Observe now that

$$\left(\frac{1-m_L}{1-m_H}\right)^{t(n+1)-c_1-\cdots-c_t} < \left(\frac{m_H}{m_L}\right)^{c_1+\cdots+c_t} \tag{7}$$

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 (7) 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 have that for all $t\geq 2$

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

If we can show that $\mu_2(m_H)\leq\mu_2(m_L)$, we are done. To see this, note that

$$\frac{d\mu_2}{dm} = \sum_{c_1=0}^{\lfloor\alpha(n+1)+\gamma\rfloor}\left(\begin{array}{c} n+1 \\ c \end{array}\right) m^{c_1-1}(1-m)^{n-c_1}[c_1-m(n+1)].$$

If for all $c_1\leq\lfloor\alpha(n+1)+\gamma\rfloor$, $c_1\leq m(n+1)$, then $d\mu_2/dm<0$, as we want. Suppose then that this is not the case; that is, $\exists\bar{c}_1\leq\lfloor\alpha(n+1)+\gamma\rfloor$ such that $\bar{c}_1>m(n+1)$. Then, because $c_1-m(n+1)$ is increasing in c_1,

$$\frac{d\mu_2}{dm}\leq\sum_{c_1=0}^{n+1}\left(\begin{array}{c} n+1 \\ c \end{array}\right) m^{c_1-1}(1-m)^{n-c_1}[c_1-m(n+1)] = \frac{1}{m(1-m)}\{E[c_1\mid m]-m(n+1)\}=0,$$

since we are adding only positive terms to $d\mu_2/dm$. So, indeed, $\mu_2(m_H)\leq\mu_2(m_L)$.

(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 $(c_1,\ldots,c_{\bar{t}-1})\in C_{\bar{t}-1}$ if we let $c_\tau=(n+1)$, $\tau=1,\ldots,\bar{t}-1$. But if we also let $c_{\bar{t}}=(n+1)$, then $c_1+\cdots+c_{\bar{t}}=\bar{t}(n+1)$, and so $(c_1,\ldots,c_{\bar{t}})\in\overline{C}_{\bar{t}}$. 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 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

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

and that

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

Hence

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

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 $\epsilon > 0$,

$$\lim_{t\to\infty} \Pr\{S_{(t-1)(n+1)}(m_H) < (t-1)(n+1)[m_H - \epsilon]\} = 0. \tag{8}$$

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 (8) to conclude that $\mu_t(m_H) \to 0$ as t goes to infinity.

As 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 **3**, implies that this limit must be strictly positive.

$\square$

Proposition 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 $d_t \geq 0$ 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 δ. If $\bar{t} = 2$, we are done, just set $\underline{\delta} = 0$.

So, suppose that $\bar{t} > 2$. 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 $F^{(\bar{t}-2)}(0) < 0$. 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 $\bar{t} = 3$, we are done, just set $\underline{\delta} = \delta_1$.

So, suppose now that $\bar{t} > 3$. 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 $\bar{t} = 4$, we are again done. If $\bar{t} > 4$, we can continue with this process. Since $\bar{t}$ is finite, we eventually reach and end to it.

Now we can prove the proposition. For this note that

$$F(0) = \frac{1}{v}[m_L - m_H] < 0,$$

and that $\lim_{\delta \to 1^-} F(\delta) = +\infty$. We proved above that $\exists \underline{\delta}$ such that if $\delta < \underline{\delta}$, $F^{'}(\delta) < 0$, and if $\delta > \underline{\delta}$, $F^{'}(\delta) > 0$. This means that F is always negative between 0 and $\underline{\delta}$, and after this point it increases strictly to $+\infty$. So there is a unique $\overline{\delta}$ with the desired properties.

$\square$

Proposition 7 *The sequence $\{\mathcal{V}_n\}$ 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 $\{\mathcal{V}_n\}$ is a fixed point of the map $\Gamma : \mathcal{S} \to \mathcal{S}$ defined by $(\Gamma f)_n(\theta) = T f_n(\theta)$; 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, $\{\mathcal{V}_n\}$. 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 $p_n(f) = \| f_n \|_{\text{sup}}$, where $\| \, . \, \|_{\text{sup}}$ 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

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

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 $\| f_{n_i}^m - f_{n_i} \|_{\sup} < \frac{1}{k_i}$ 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

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

This implies that

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

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 $\{\mathcal{V}_n\}$ 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 $\{f^m_n\}$ is uniformly convergent for all n, and it must converge to f_n. Now observe that $\forall \theta \in [0, 1]$,

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

where the choice of m is arbitrary, so that

$$|f_n(\theta) - f_{n'}(\theta)| \leq \parallel f_n - f^m_n \parallel_{\sup} + \parallel f^m_n - f^m_{n'} \parallel_{\sup} + \parallel f^m_{n'} - f_{n'} \parallel_{\sup} .$$

Take $\epsilon > 0$. Since $\{f^m_q\}$ is uniformly convergent for all m by hypothesis, we know that $\exists N$ such that if $n, n' \geq N$, then $\parallel f^m_n - f^m_{n'} \parallel_{\sup} < \frac{\epsilon}{3}$. Take then n, n' greater than N. Because $\{f^m_n\}$ and $\{f^m_{n'}\}$ converge uniformly to f_n and $f_{n'}$, respectively, there is $m_0(n, n') \in \mathbf{N}$ such that if $m \geq m_0$, $\parallel f_n - f^m_n \parallel_{\sup} < \frac{\epsilon}{3}$ and $\parallel f^m_{n'} - f_{n'} \parallel_{\sup} < \frac{\epsilon}{3}$. 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, $\parallel f_n - f_{n'} \parallel_{\sup} < \epsilon$. 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 $\widehat{V}^E_n$. Remember that

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

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

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

First observe that $|\widehat{V}_n^E(\theta) - \widehat{V}^E(\theta)| \leq \| \widehat{V}_n^E - \widehat{V}^E \|_{\text{sup}}$ for all $\theta \in [0,1]$, and that $\| \widehat{V}_n^E - \widehat{V}^E \|_{\text{sup}} \to 0$ by construction. Now note that

$$\begin{aligned} Ef_n(\theta^1(\theta,n)) - [\theta f(1) + (1-\theta)f(0)] \;=& \\ &Ef_n(\theta^1(\theta,n)) - Ef(\theta^1(\theta,n)) \;+ \\ &Ef(\theta^1(\theta,n)) - [\theta f(1) + (1-\theta)f(0)]. \end{aligned}$$

However,

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

and $\| f_n - f \|_{\text{sup}} \to 0$ 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

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

for every $n \in \mathbf{N}$ and for all $\theta \in [0,1]$. Notice first of all that the set $\mathcal{Y} = \{f \in \mathcal{S} \mid f \text{satisfies (1) for all } n \in \mathbf{N}\}$ 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

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

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

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

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 $\{\mathcal{V}_n\}$ of value functions is uniformly convergent.

From **(4)** we can conclude that its uniform limit must be

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

□

Proposition 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 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 the probability of success being given by m_H. Since, by proposition 2 and corollary 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 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. First note that

$$\mu_2(m_L, n) = \Pr\{S_{n+1}(m_L) \leq \lfloor \alpha(n+1) + \gamma(\theta^G(n))\rfloor\} \geq \Pr\{S_{n+1}(m_L) \leq \alpha(n+1) - \gamma\},$$

where $\gamma \geq 0$ is fixed. Since $m_L < \alpha$, we can again apply the law of large numbers, but now to conclude that $\mu_2(m_L, n) \to 1$ as $n \to \infty$.

From the proof of proposition 3 we know that

$$\begin{aligned}\mu_t(m_L, n) - \mu_{t+1}(m_L, n) &= L_t(m_L) \\ &= \sum_{(c_1,\ldots,c_t)\in\overline{C}_t} \begin{pmatrix} n+1 \\ c_1 \end{pmatrix} \cdots \begin{pmatrix} n+1 \\ c_t \end{pmatrix} m_L^{c_1+\cdots+c_t}(1-m_L)^{t(n+1)-c_1-\cdots-c_t},\end{aligned}$$

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,

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

Hence, by once more applying the law of large numbers for the Binomial distribution, we have that $\mu_t(m_L, n) - \mu_{t+1}(m_L, n) \to 0$ as $n \to \infty$ (since $\alpha > m_L$). Consequently, $\mu_t(m_L, n) \to 1$ 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 $\mu_t(m_H, n) \leq \frac{\epsilon(1-\delta)}{4m_H\delta}$ for all $t \geq 2$. Therefore,

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

whenever $n \geq n_1(\delta)$. 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 $n \geq n_2(\delta, t)$, 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 $n \geq n_2(\delta)$,

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

Hence, if $n \geq n(\delta)$, 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 9 *For all $n \in N$, $\mathcal{V}_n(\theta, \lambda, \phi)$ 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

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

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

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

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, $\mathcal{V}(\theta,\lambda,\phi)$ is decreasing in θ, since the maximum between two functions that are decreasing in θ is also decreasing in θ.

□

Lemma 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 $\chi : \Omega \to \Omega$ such that

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

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

$$\sum_{\{c\leq\lfloor\alpha(n+1)+\gamma(\tilde{\theta})\rfloor\}}\left(\begin{array}{c}n+1\\c\end{array}\right)m_L^c(1-m_L)^{n+1-c}=p(\tilde{\theta}),$$

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 $t\geq2$,

$$\nu_t(\theta,m)=\sum_{\tilde{\theta}\in\mathcal{R}^{(t-1)}(\{\theta\})}\sum_{\{c\leq\lfloor\alpha(n+1)+\gamma(\tilde{\theta})\rfloor\}}\left(\begin{array}{c}n+1\\c\end{array}\right)m_L^c(1-m_L)^{n+1-c}\Pr\{\theta^{(t-1)}(\theta,\lambda)=\tilde{\theta}|m\},\quad(9)$$

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 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

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

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

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

$$\Psi(\tilde{\theta}) = \sum_{\{c \leq \lfloor \alpha(n+1)+\gamma(\tilde{\theta}) \rfloor\}} \begin{pmatrix} n+1 \\ c \end{pmatrix} m_L^c (1-m_L)^{n+1-c}$$

is decreasing in $\tilde{\theta}$ (as the set of feasible cardinalities gets smaller). Moreover, from (9) 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,

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

for $m \in \{m_L, m_H\}$, and from the proof of proposition (3), we have that $\partial f / \partial m < 0$. Note that this also proves that $\nu_1(\theta, m_L) \geq \nu_1(\theta, m_H)$. To finish, note that for all $t \geq 2$,

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

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

$$\nu_1(\theta, m_L) - \nu_1(\theta, m_H) = \sum_{\{c \leq \lfloor \alpha(n+1)+\gamma(\theta) \rfloor\}} \begin{pmatrix} n+1 \\ c \end{pmatrix} \left[m_L^c (1-m_L)^{n+1-c} - m_H^c (1-m_H)^{n+1-c} \right].$$

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.

□

Proposition 10 *Let $\Delta' = (\mathcal{M}^*, S^*)$ be the strategy profile where $\mathcal{M}^*$ is the bank's strategy defined in* (3)*, and S^* is the agent's strategy defined in* (6)*. There exist $n_1 \in \mathbf{N}$, δ_1 and δ_2, $\delta_1 \leq \delta_2$, such that this strategy profile, along with the Bayesian updating rule given by* (4)*, 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 (one in which the bank played m_L up to a certain period t), 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 (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 (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 (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 (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.

□