

Endogenous Supply of Fiat Money

Luis Araujo^{*,†}

Braz Camargo[‡]

January 25, 2006

Abstract

We address whether reputation concerns can discipline the behavior of a self-interested agent who has a monopoly over the provision of fiat money. We obtain that when this agent can commit to a plan of action, there is a monetary equilibrium where it never overissues. We show, however, that such equilibrium is no longer possible when there is no commitment. This happens because the incentives this agent has to maintain a reputation for providing valuable currency disappear once its reputation is high enough. More generally, we prove that there is no monetary equilibrium where overissue happens only infrequently. We conclude by showing that imperfect memory can restore the positive result obtained in the presence of commitment.

JEL Classification: C73, D82, D83, E00

Key Words: Endogenous Money, Reputation

*Michigan State University, Department of Economics. Address: 101 Marshall Adams Hall, East Lansing, MI, 48824-1038, USA. E-mail:araujolu@msu.edu

†Fucape. Address: Avenida Fernando Ferrari, 1358, Goiabeiras, Vitória, Espírito Santo, 29075-010, Brazil.

‡The University of Western Ontario, Department of Economics. Address: Social Science Centre, London, Ontario, N6A 5C2, Canada. E-mail:bcamargo@uwo.ca

1 Introduction

Frictions in trade and lack of enforceability are necessary if money is to be valued as a medium of exchange. The standard approach to model this is to assume that trade is anonymous and decentralized and there is no record keeping. Under these assumptions, a large body of work has established the essentiality of money when its supply is exogenous. However, if the amount of money in circulation in an economy is determined by self-interested agents, a complete lack of monitoring may lead to the so-called “dynamic inconsistency” problem. If money has value, any agent with the ability to print money faces a temptation to overissue, as any deviation from a pre-specified plan of action is likely to go unnoticed. In other words, money may not be feasible in such economies.

The way the literature has dealt with this problem is to assume a form of partial record keeping: the behavior of note issuers can be publicly monitored. Using this approach, Berentsen [5] and Martin and Schreft [16] have established that fiat money issued by self-interested agents can be feasible, while Cavalcanti et. al. [6] and Cavalcanti and Wallace [7] have done the same for inside money, i.e., private circulating liabilities.¹

This solution has two shortcomings. First, it leaves open the question of how this form of partial record keeping arises in the first place. Second, it is not robust to the introduction of monitoring costs, no matter how small, because of a free rider problem. The success any technology that allows note issuers to be monitored has in disciplining their behavior depends on how many agents use it. If only a small number of agents do so, the punishment note issuers face if they overissue is insufficient to induce good behavior. However, if a large number of agents uses this technology, its effectiveness is not affected when a single agent stops using it. Therefore, if all agents have to pay a cost, no matter how small, to use it, there can be no situation where a large group of agents uses this technology in the first place. Otherwise, any such agent would have an incentive to stop using it and free ride on the social benefit it brings.

In this article we address the feasibility of fiat money when it is issued by a single self-interested agent, the bank, and its choice of money supply cannot be monitored by the other agents in the economy.² We do so in an environment where trade is decentralized and agents are anonymous and have heterogeneous preferences, so that there is a natural role for money. The absence of

¹An exception is Monnet [18], who establishes the feasibility of private fiat money without monitoring. His result, however, depends critically on the assumption that the matching process is deterministic and money is costly to produce. Moreover, as this cost decreases to zero, so does the value of money. This last result is consistent with what we obtain in our analysis.

²It is usually argued that the technology to print money has increasing returns to scale, see [8], and so is a natural monopoly.

a monitoring technology means that information can only be acquired and transmitted in trade meetings; i.e., information is decentralized. This assumption about information is a natural one in an economy with decentralized trade.

Our starting point is a simple version of the model introduced in Kiyotaki and Wright [11]. We modify it in a number of ways. First, as indicated above, the money supply is now privately determined by the bank in each period. Moreover, the bank can either be patient or impatient, and this is also its private information. Second, the other agents in the economy can now decide between staying in autarky or entering the market and transacting with the help of money. The bank's revenue from money issue in a given period is proportional to how much new currency it prints and to the number of agents who choose the market at that point in time. In particular, holding everything else constant, this revenue is higher if the bank overissues. Finally, autarky is always better than the market if money is always overissued, which happens when the bank is impatient, but the opposite is true when overissue never takes place.

In principle, the bank's choice of money supply would affect both the frequency of trade meetings (extensive margin) and the terms of trade in such meetings (intensive margin). However, since in our framework money and goods are indivisible and there is a unit bound on money holdings, the only margin that is affected by the bank's decision is the extensive one. Nevertheless, agents are still able to learn about the bank's decision from their private experience in the market. Because information is decentralized, this is the only way in which learning can take place. Hence, the environment considered is one where monitoring is private and imperfect.

We first consider the case where the bank can commit to a choice of money supply, the so-called full-commitment case. This is done by assuming that its choice of money supply in the first period is binding. We show that despite the bank's choice being its private information, there exists an equilibrium where the patient bank never overissues as long as it is sufficiently patient. The intuition for this result is simple. Suppose the patient bank indeed never overissues. Since the impatient bank always overissues, the agents in the market learn from their private experiences that they are facing the patient bank. In other words, the patient bank's reputation that it is indeed patient increases over time. This guarantees to it a steady stream of revenue from money issue, as the market is preferred when there is no overissue. Suppose, instead, that the patient bank deviates and always overissues. Its revenue from money issue increases in the short-run as a result of this. However, as time progresses, its reputation for being patient decreases, as agents in the market attribute the type of experience they face to the impatient bank. In the long-run, this leads to a decrease in the patient bank's revenue from money issue. If the patient bank cares enough about the future, this is sufficient to discourage it from deviating.

The conclusion that reputation concerns can solve the “dynamic inconsistency” problem is not new. Klein [12] considers an environment where a tradeoff between short-run gains from overissue and long-run losses due to a decrease in the reputation for providing valuable currency disciplines the behavior of note issuers. In his model, however, this tradeoff is assumed rather than derived.

We then consider the no-commitment case, where the bank may change its behavior at any point in time. In this case, a policy where the patient bank never overissues is not time-consistent. Indeed, if the patient bank were never to overissue, its reputation for being patient, and thus providing valuable currency, would increase over time. Eventually a point would be reached where all agents in the market are so convinced that the bank they face is patient that any negative experience in a given period is attributed to bad luck. At this stage, the patient bank would rather overissue. The cost of doing so, a reduction in future revenue from money issue due to a decrease in its reputation, is almost zero, while the immediate benefit is substantial.

In light of this negative result, a natural question to ask is what type of equilibria are possible in the no-commitment case. For instance, is it possible to have a monetary equilibrium where the patient bank overissues only infrequently, so that when the bank is of this type, the gain of choosing the market is always positive and does not disappear over time? We show that when the bank follows a pure strategy, the same logic that rules out the no-overissue equilibrium also rules out these other equilibria.

The above discussion suggests that private fiat money can be made feasible if a way is found to keep the “reputational” cost of overissue bounded away from zero. Motivated by this, we modify the no-commitment case to include a continuous threat to the patient bank’s reputation. More precisely, we assume that in every period a fraction of the population becomes uninformed. We show that with this form of imperfect memory, an equilibrium where the patient bank never overissues is possible; i.e., a policy of no overissue for the patient bank is time-consistent. The reason is that now a patient bank has always an incentive to look after its reputation: any time it overissues, the negative impact on its reputation will be non-negligible.

Besides the literature on private money, this work is also related to the literature that looks at reputation as a separation device.³ The paper in this second literature that is closest to ours is Mailath and Samuelson [14], who consider an environment where monitoring, like in this article, is imperfect and private. See also Moav and Neeman [17], who study the interplay between memory and reputation.

The basic setup is developed in the next section. Equilibrium is defined in Section 3. The

³Mailath and Samuelson [15] discuss how this approach differs from the more standard approach to reputation pioneered by Kreps, Milgrom, Roberts, and Wilson [13].

full-commitment case is considered in Section 4, while the no-commitment case is considered in Section 5. The modification of the no-commitment case to include imperfect memory is analyzed in Section 6. Section 7 concludes and the Appendix collects proofs and technical details that are omitted from the main text. Throughout the text, if S is metric space, $\mathcal{P}(S)$ denotes the set of Borel probability measures over S .

2 Basic Setup

Time is discrete and indexed by t . The economy has one large infinitely lived agent that we call the bank. Its discount factor δ is either zero or $\delta_p > 0$. In the first case we say the bank is *impatient*, while in the second we say the bank is *patient*. The value of δ , however, is known to the bank only. The economy is also populated by a large number of small infinitely lived agents that we describe in the paragraphs that follow. For simplicity, we refer to these small agents as *agents* only.

The economy starts in $t = 1$ with a mass one of agents, each with a prior belief θ_0 that the bank is patient. For each one of these agents the value of θ_0 is determined by an independent random draw from a p.d.f. λ such that $\lambda(\theta_0) > 0$ for all $\theta_0 \in (0, 1)$.⁴ Moreover, in every $t > 1$ each agent born in the previous period gives birth to another agent, who inherits his parent's private history. As we are going to see later on, an agent's private history determines his belief about the bank's type. Hence, this last assumption implies that any agent born after $t = 1$ starts with the same belief about δ as his parent. An agent in his first period of life is said to be *newly born*.

All agents have the same discount factor $\beta \in (0, 1)$. They also have a type that is determined when they first enter the economy. There are $K > 2$ of these types, one for each of the K types of goods that can be produced in the economy. The probability that a newly born agent is of the type $k \in \{1, \dots, K\}$ is the same in every period, $1/K$. Agents of type k can only consume a type k good, their so-called preferred good.

Production works as follows. All newly born agents receive a non-perishable endowment and make a once and for all decision between moving to autarky or entering the market. In autarky, an agent uses his endowment as an input to a production technology. In each period there are $n \geq 2$ such production possibilities, and each good produced yields utility a . In the market, an agent uses his endowment in the production of indivisible and perishable goods. An agent of type k can only produce, at a cost c per unit, a good of type $k + 1 \pmod{K}$, his so-called endowment good. Any agent in the market can hold at most one unit of either goods or money at any point in time.

⁴ All the results we obtain hold if the agents born in 1 have the same prior belief that the bank is patient and this prior is sufficiently high. The assumption that the period 1 distribution of priors has full support simplifies the proof of the main result of Section 5 (Theorem 2) though.

The bank derives utility from the consumption of all K goods, but cannot produce any of them. It has, however, the technology to print indivisible units of fiat money. These units provide no direct benefit, but can be offered in exchange for goods. More precisely, each newly born who enters the market is approached by the bank with a certain probability m , in which case he receives one unit of fiat money in exchange for one unit of his endowment good. The value of m is restricted to $\{m_L, m_H\}$, with $\frac{1}{2} \leq m_L < m_H < 1$, and is determined by the bank in each period. No agent observes this choice. If μ is the measure of newly born agents who enter the market in a given period, the bank's flow payoff from choosing m is $(\mu + \kappa)m$, where κ is infinitesimal. Hence, it is myopically optimal for the bank to choose m_H even when the measure of agents who enter the market is zero. From now on we say that the bank *overissues* when it chooses m_H .

The market is organized as follows. There are K sectors, each one specialized in the exchange of one of the K available goods. Agents can identify sectors, but inside each sector they are randomly and anonymously matched in pairs. Since $K > 2$, there are no double coincidence of wants meetings. An agent, however, can trade his endowment good for money and use money to buy his preferred good. If an agent wants money, he goes to the sector that trades his endowment good and searches for an agent with money. If he has money, he goes to the sector that trades his preferred good and searches for an agent who can produce it. When a single coincidence of wants meeting takes place, the buyer transfers his unit of money to the seller and the latter produces one unit of his endowment good for the buyer, who consumes it to obtain utility $u > c$.⁵ Any agent in the market faces n rounds of meetings per period, where n is the same number as above. Agents don't discount within a period.

Notice that we take the behavior of the agents and the bank in the market as given. It is possible, in a natural way, to model the market environment itself as a game involving the agents and the bank. This game has an equilibrium where the agents always exchange their endowment for one unit of money if approached by the bank and, as long as their discount factor is close enough to one, their behavior in the market is as described.

An implicit assumption in the above description of the market is that there is a positive measure of agents in it at any point in time. Since once in the market an agent cannot leave it, a necessary and sufficient condition for this is that a positive measure of agents enters the market in period 1. When the measure of agents in the market is zero, i.e., the market is "empty", money does not circulate and the market flow payoff is zero.

Suppose that in period t the choice of m by the bank is m_t and the fraction of agents in the

⁵This market structure is adopted for simplicity. It implies that any single coincidence of wants meeting involves a buyer with money and a seller without money, so that trade always occurs at such meetings.

market that have (one unit of) money is η_t . If the market is empty in t , $\eta_t = 0$. If the market is not empty in t , η_t depends on m_t and on the previous choices of m by the bank. In this section, however, we treat the sequences $\{m_t\}$ and $\{\eta_t\}$ as being independent of each other. Notice, though, that if $\eta_t > 0$, then it must be that $\eta_k > 0$ for all $k > t$. Moreover, if $\eta_t > 0$, then $\eta_t \in [m_L, m_H] \subseteq [\frac{1}{2}, 1)$, since the bank's choice of m is constrained to $\{m_L, m_H\}$.

Let $w_{j,t}^i$ denote the lifetime expected payoff for an agent in the market in period t with j units of money right before his i th meeting. Also, let $w_{j,t}^{n+1} = \beta w_{j,t+1}^1$ be the lifetime expected payoff for an agent that has j units of money after the last round of meetings in t . Then, if $\eta_t > 0$,

$$\begin{aligned} w_{1,t}^i &= \eta_t w_{1,t}^{i+1} + (1 - \eta_t)[u + w_{0,t}^{i+1}] \\ w_{0,t}^i &= \eta_t[w_{1,t}^{i+1} - c] + (1 - \eta_t)w_{0,t}^{i+1} \end{aligned}, \quad i \in \{0, \dots, n+1\}. \quad (1)$$

For the first equation, notice that an agent with money right before his i th meeting in t has probability η_t of meeting another agent with money, in which case no trade occurs. With probability $1 - \eta_t$ he meets an agent without money, in which case trade occurs and he obtains utility u . A similar interpretation holds for the second equation.

Solving (1) recursively, we obtain that if $\eta_t > 0$, then

$$\begin{aligned} w_{1,t}^1 &= \beta \eta_t w_{1,t+1}^1 + \beta(1 - \eta_t)w_{0,t+1}^1 + (n - 1)\eta_t(1 - \eta_t)(u - c) + (1 - \eta_t)u \\ w_{0,t}^1 &= \beta \eta_t w_{1,t+1}^1 + \beta(1 - \eta_t)w_{0,t+1}^1 + (n - 1)\eta_t(1 - \eta_t)(u - c) - \eta_t c \end{aligned}. \quad (2)$$

Observe that $w_{1,t}^1 - c - w_{0,t}^1 = (1 - \eta_t)(u - c) > 0$ in this case. Hence, a newly born agent who enters the market is always willing to accept one unit of money from the bank in exchange for his endowment good if he knows the market is not empty.

Rewrite (2) as $w_t = \beta B(\eta_t)w_{t+1} + b(\eta_t)$ for all $t \in \mathbb{N}$ such that $\eta_t > 0$, where

$$B(\eta) = \begin{pmatrix} \eta & 1 - \eta \\ \eta & 1 - \eta \end{pmatrix}, \quad w_t = \begin{pmatrix} w_{1,t}^1 \\ w_{0,t}^1 \end{pmatrix}, \quad \text{and} \quad b(\eta) = \begin{pmatrix} (n - 1)\eta(1 - \eta)(u - c) + (1 - \eta)u \\ (n - 1)\eta(1 - \eta)(u - c) - \eta c \end{pmatrix}. \quad (3)$$

Solving this last system of equations recursively, we obtain that if the market is not empty in t , then

$$w_t = b(\eta_t) + \sum_{\tau=1}^{\infty} \beta^\tau B(\eta_t) \cdots B(\eta_{t+\tau-1}) b(\eta_{t+\tau}). \quad (4)$$

If, on the other hand, the market is empty in t , then $w_t = \beta^{\underline{t}} w_{\underline{t}}$, where \underline{t} is the first period when the market is not empty and $\underline{t} = +\infty$ if the market is always empty.

Let v_t be the lifetime expected reward from entering the market in period t and $a(m)$ be the 2×1 row vector $(m \ 1 - m)$. Then,

$$v_t = m_t(w_{1,t}^1 - c) + (1 - m_t)w_{0,t}^1 = a(m_t)w_t - m_t c. \quad (5)$$

We denote the dependence of $w_{i,t}^1$, w_t , and v_t on the sequence $\{\eta_{t+k-1}\}_{k=1}^\infty$ by writing $w_{i,t}^1 = w_{i,t}^1(\{\eta_{t+k-1}\})$, $w_t = w_t(\{\eta_{t+k-1}\})$, and $v_t = v_t(\{\eta_{t+k-1}\})$. In the particular case where the market is not empty in some period t and $\eta_k = m_k = \eta \in \{m_L, m_H\}$ for all $k \geq t$, so that $v_k = v_t$ for $k \geq t$, we write $v_t = v(\eta)$.

Notice that if $\eta_t > 0$, then $w_{1,t+1}^1 - w_{0,t+1}^1 = (1 - \eta_{t+1})(u - c) + c$, since $\eta_{t+1} > 0$ as well. Therefore,

$$\begin{aligned} dv_t/d\eta_t &= \beta(w_{1,t+1}^1 - w_{0,t+1}^1) + (n-1)(1-2\eta_t)(u-c) - c - m_t(u-c) \\ &< [(1-m_t-\eta_{t+1}) + (n-1)(1-2\eta_t)](u-c) \leq 0 \end{aligned}$$

when the market is not empty in period t . Consequently, when $\eta_t > 0$, v_t is a strictly decreasing function of η_k for $k \geq t$.

Consider now the particular case where $m_t = \eta_t = m$ for all $t \in \mathbb{N}$ and denote by $w_j(m)$ the lifetime expected utility from entering the market with j units of money. Since $B(\eta)^2 = B(\eta)$, it is straightforward to see from (4) that

$$\begin{aligned} w_1(m) &= (1-\beta)^{-1}nm(1-m)(u-c) + (1-m)u - m(1-m)(u-c) \\ w_0(m) &= (1-\beta)^{-1}nm(1-m)(u-c) - mc - m(1-m)(u-c) \end{aligned}$$

Hence, in this case, $v(m) = (1-\beta)^{-1}nm(1-m)(u-c) - mc$ is the lifetime expected payoff from choosing the market in any given period. Let $v_A = (1-\beta)^{-1}na$ denote the lifetime expected payoff from choosing autarky.

Assumption 1. $v(m_H) < v_A < v(m_L)$.

Observe that the impatient bank, being myopic, will always choose m_H , whether it can commit to its period one choice of m or not. Also observe that if a positive measure of agents enters the market in period 1 and the bank's choice of money supply is the same in every period, then η_t is equal to this choice of m for all t . Consequently, Assumption 1 implies that if a positive measure of agents enters the market in the first period, then: (i) the market is always worse than autarky when the bank is impatient; (ii) the market is always better than autarky when the bank is patient and chooses m_L in every period.

To finish, notice that Assumption 1 also implies that there exists $\theta^* \in (0, 1)$ depending on β and n such that $\theta^*v(m_L) + (1-\theta^*)v(m_H) = v_A$. Since the p.d.f. λ has full support, a positive measure of agents born in period 1 have beliefs in $[\theta^*, 1]$. If this were not the case, no agents would ever want to enter the market. Indeed, we know from above that v_t is a strictly increasing function of η_k for $k \geq t$ when $\eta_t > 0$. Hence, since the market flow payoff is zero when the market is empty, the highest payoff a newly born agent can obtain if he chooses the market is $v(m_L)$.

3 Equilibrium

Let \mathcal{H}_t denote the set of possible period t histories for the bank. By definition, $\mathcal{H}_1 = \{\emptyset\}$. Implicit in the description of the bank's flow payoff given in the previous section is the assumption that in each period it can observe the measure of newly born agents who enter the market. Besides its previous choices of m , this is the only information available to the bank at any given point in time. Hence, $\mathcal{H}_t = ([0, 1] \times \{m_L, m_H\})^{t-1}$ for $t > 1$. A strategy for the bank is then a sequence $M = \{M_t\}$ of contingent plans, where M_t maps $\{0, \delta_p\} \times \mathcal{H}_t$ into $\mathcal{P}(\{m_L, m_H\})$.

When making his market/autarky decision, the only piece of information a newly born agent has is the private history he inherits from his parent. Let H_t , with arbitrary element h^t , denote the set of all possible histories for an agent born in t . By assumption, $H_1 = [0, 1]$, the set of possible prior beliefs that the bank is patient. We describe H_t for $t > 1$ in the next two paragraphs. A strategy (decision rule) for an agent born in t is then a (Borel measurable) function $d_t : H_t \rightarrow [0, 1]$, where $d_t(h^t)$ is the probability that he chooses the market given a private history h^t .

Notice that an agent's history in his first period of life is his decision together with his subsequent experience in this period. If he chooses autarky, he observes nothing (\emptyset). If, instead, he goes to the market, his experience consists of: (i) how many units of money, 0 or 1, he receives from the bank; (ii) the money holdings of his partners and the terms of trade in all his market meetings, if the market is not empty. Define a family to be the collection of all agents whose genealogy can be traced back to a particular agent born in $t = 1$. Then, the history an agent born in $t > 1$ inherits is the collection of first period histories of the generation k members of his family, where $k \leq t - 1$.

Now observe that since money and goods are indivisible and there is a unit bound on money holdings, trade is always one-to-one when it occurs. Hence, if a newly born agent chooses the market, and the market is not empty, the number $j \in \{0, \dots, n\}$ of market meetings where his partner carries one unit of money summarizes the information about the bank's type that this agent gathers in his first period of life. Let $\Pi = \{A, \emptyset\} \cup \{M, \{0, 1\} \times \{e, 0, \dots, n\}\}$, where A denotes the event that autarky is chosen, M denotes the event that the market is chosen, and e denotes the event that the market is empty. We can then conclude that $H_t = [0, 1] \times \Pi^{t-1}$ for $t > 1$. In what follows, we denote an arbitrary element of Π by $\pi = (d, \omega)$.

Identify the set of families with the unit interval and let Δ_t denote the set of (Borel measurable) functions from H_t into $[0, 1]$. A strategy profile for the agents is an equivalence class of sequences $\tau = \{\tau_t\}$, where τ_t maps $[0, 1]$ into Δ_t and two sequences $\tau_1 = \{\tau_{1,t}\}$ and $\tau_2 = \{\tau_{2,t}\}$ are considered to be the same if for all $t \in \mathbb{N}$ the functions $\tau_{1,t}$ and $\tau_{2,t}$ differ on a set of measure zero in $[0, 1]$. We interpret $\tau_t(i) \in \Delta_t$ as the strategy of the generation t member of the family labeled by $i \in [0, 1]$.

Suppose the bank uses a strategy $M = \{M_t\}$ such that $M(\delta_p) = \{M_t(\delta_p, \cdot)\}$ is pure and let $\tau = \{\tau_t\}$ be a strategy profile for the agents. Since there is no aggregate uncertainty, the measure γ_1 of agents who enter the market in period 1 is deterministic. It is also independent of the bank's type. Because the impatient bank always chooses m_H , the bank's choice of m in period 1, m_1 , is also deterministic. This choice of m together with τ_1 induces a probability measure λ_2 over H_2 such that $\lambda_2(D)$ is the fraction of agents born in $t = 2$ with private histories in (the Borel set) $D \subseteq H_2$.

Now observe that the pair (γ_1, m_1) completely determines the bank's choice of m in period 2, m_2 . If the bank is impatient, $m_2 = m_H$. If, on the other hand, the bank is patient, $m_2 = M(\gamma_1, m_1)$. Also notice, once more because there is no aggregate uncertainty, that the pair (λ_2, τ_2) completely determines the measure γ_2 of agents who enter the market in $t = 2$. Unlike γ_1 , the measure γ_2 potentially depends on the bank's type. To finish, observe that m_2 together with λ_2 determine a probability measure λ_3 over H_3 such that $\lambda_3(D)$ is the fraction of agents born in $t = 3$ that have private histories in $D \subseteq H_3$.

Continuing with this process, we obtain sequences $\{m_t(M, \tau, \delta)\}$ and $\{\gamma_t(M, \tau, \delta)\}$ such that if δ is the bank's discount factor, then $m_t(M, \tau, \delta)$ is the bank's choice of m in t and $\gamma_t(M, \tau, \delta)$ is the fraction of agents born in t who enter the market. Notice that $m_t(M, \tau, 0) \equiv m_H$. We also obtain a sequence $\{\lambda_t(M, \tau, \delta)\}$, with $\lambda_t(M, \tau, \delta) \in \mathcal{P}(H_t)$, such that: (i) $\lambda_1(M, \tau, \delta) \equiv \hat{\lambda}_1$, the probability measure over H_1 induced by the p.d.f. λ ; (ii) $\lambda_t(M, \tau, \delta)(D) = \lambda_t(D|M, \tau, \delta)$ is the fraction of agents born in t with private histories in $D \subseteq H_t$ as a function of the bank's type. The important point is that if the patient bank uses a pure strategy, then both its behavior over time and the aggregate behavior of the agents over time are deterministic. This plays a fundamental role in Section 5.

Let m_t be the choice of m in t and μ_t be the measure of agents who enter the market in t when the bank is patient. Notice the change of notation. Let ν_t be the measure of agents who enter the market in the same period when the bank is impatient. Notice that if the bank is impatient, the fraction of agents in the market that have money is either zero, when the market is empty, or m_H . When the bank is patient, however, this fraction, that we denote by α_t , can change over time even when the market is not empty. Moreover, if $t > 1$, α_t depends on μ_1 to μ_t . Indeed, $\alpha_t = 0$ if $\sum_{\tau=1}^t \mu_\tau = 0$ and

$$\alpha_t = \frac{\sum_{\tau=1}^t \mu_\tau m_\tau}{\sum_{\tau=1}^t \mu_\tau} \quad (6)$$

if $\sum_{\tau=1}^t \mu_\tau > 0$. Also notice that it is possible that there are periods when the market is empty if the bank is of one of type, but not of the other. As a consequence, the belief an agent born in t has that the bank is patient depends not only on his private history $h^t \in H_t$, but also on the

sequences $\{\mu_t\}$, $\{\nu_t\}$, and $\{m_t\}$. Denote this belief by $\theta(h^t; \{\mu_t\}, \{\nu_t\}, \{m_t\})$. When there is no risk of confusion, we omit its dependence on the sequences $\{\mu_t\}$, $\{\nu_t\}$, and $\{m_t\}$.

Let $\Omega = \{0, 1\} \times \{e, 0, \dots, n\}$ and define $X_t(\delta; \{\mu_t\}, \{\nu_t\}, \{m_t\})$ to be the random variable on Ω such that if $i \in \{0, 1\}$ and $j \in \{0, \dots, n\}$, then:

$$\begin{aligned} \Pr\{X_t(0; \{\mu_t\}, \{\nu_t\}, \{m_t\}) = (i, e)\} &= \begin{cases} m_H^i (1 - m_H)^{1-i} & \text{if } \sum_{\tau=1}^t \nu_\tau = 0 \\ 0 & \text{otherwise} \end{cases}; \\ \Pr\{X_t(0; \{\mu_t\}, \{\nu_t\}, \{m_t\}) = (i, j)\} &= \begin{cases} m_H^{i+j} (1 - m_H)^{m+1-(i+j)} & \text{if } \sum_{\tau=1}^t \nu_\tau > 0 \\ 0 & \text{otherwise} \end{cases}; \\ \Pr\{X_t(\delta_p; \{\mu_t\}, \{\nu_t\}, \{m_t\}) = (i, e)\} &= \begin{cases} m_t^i (1 - m_t)^{1-i} & \text{if } \sum_{\tau=1}^t \mu_\tau = 0 \\ 0 & \text{otherwise} \end{cases}; \\ \Pr\{X_t(\delta_p; \{\mu_t\}, \{\nu_t\}, \{m_t\}) = (i, j)\} &= \begin{cases} m_t^i (1 - m_t)^{1-i} \alpha_t^j (1 - \alpha_t)^{n-j} & \text{if } \sum_{\tau=1}^t \mu_\tau > 0 \\ 0 & \text{otherwise} \end{cases}. \end{aligned}$$

Observe that $m_t^i (1 - m_t)^{1-i}$ is the probability an agent has of receiving i units of money from the patient bank if he enters the market in t , and that if $\alpha_t > 0$, $\alpha_t^j (1 - \alpha_t)^{n-j}$ is the probability he has of meeting j agents with one unit of money in the market in the same period. When the bank is impatient, these probabilities are $m_H^i (1 - m_H)^{1-i}$ and $m_H^j (1 - m_H)^{n-j}$, respectively. As with $\theta(h^t)$, we omit the dependence of X_t on $\{\mu_t\}$, $\{\nu_t\}$, and $\{m_t\}$ when there is no chance of confusion.

The belief $\theta(h^t)$ can then be computed in the following inductive way. In the first period, $\theta(h^1) = h^1$, the prior belief that the bank is patient. Now fix $t \geq 1$ and assume that $\theta(h^t)$ is defined for all $h^t \in H_t$. Moreover, let $h^{t+1} = (h^t, \pi)$, with $\pi = (d, \omega) \in \Pi$, be an element of H_{t+1} . If $d = A$, then $\theta(h^t, \pi) = \theta(h^t)$. If $d = M$, so that $\omega \in \Omega$, then

$$\theta(h^t, \pi) = \frac{\theta(h^t) \Pr\{X_t(\delta_p) = \omega\}}{\theta(h^t) \Pr\{X_t(\delta_p) = \omega\} + (1 - \theta(h^t)) \Pr\{X_t(0) = \omega\}} \quad (7)$$

when the denominator is positive. When the denominator is zero, set $\theta(h^t, \pi) = \theta(h^t)$.

It is clear that in order to define an equilibrium where the patient bank uses a pure strategy, we have to take into account that: (i) the agents need the sequences $\{\mu_t\}$, $\{\nu_t\}$, and $\{m_t\}$ to compute the expected payoff from choosing the market, since this payoff depends on the belief that the bank is patient; (ii) the sequences $\{\mu_t\}$, $\{\nu_t\}$, and $\{m_t\}$ depend on the aggregate behavior of the agents

(together with the bank's behavior). Hence, the requirement that correct expectations about $\{\mu_t\}$, $\{\nu_t\}$, and $\{m_t\}$ are held needs to be included in such a definition. We do not require the patient bank's behavior to be sequentially rational, though. This is not an issue in the full commitment case and it emphasizes that the negative results we obtain in Section 5 are not a consequence of requiring subgame perfection. In Section 6, where we show that positive results are possible if memory is imperfect, we do require sequential rationality, though.

Definition 1. Let M be a strategy for the bank such that $M(\delta_p)$ is pure and τ be a strategy profile for the agents. Let $\Theta : \bigcup_{t=1}^{\infty} H_t \rightarrow [0, 1]$ be a belief updating rule for the agents, i.e., $\Theta(h^t)$ is the belief an agent born in t with history h^t has that the bank is patient. The list $\sigma = (M, \tau, \Theta, \{\mu_t\}, \{\nu_t\}, \{m_t\})$ is a (deterministic) equilibrium if :

- (a) The impatient bank always chooses m_H , i.e., $M_t(0, \cdot) \equiv m_H$ for all $t \in \mathbb{N}$;
- (b) $\Theta(h) = \theta(h; \{\mu_t\}, \{\nu_t\}, \{m_t\})$ for all $h \in \bigcup_{t=1}^{\infty} H_t$;
- (c) Agents hold correct expectations about $\{m_t\}$, $\{\mu_t\}$, and $\{\nu_t\}$; that is, $m_t = m_t(M, \tau, \delta_p)$, $\mu_t = \gamma_t(M, \tau, \delta_p)$, and $\nu_t = \gamma_t(M, \tau, 0)$ for all $t \in \mathbb{N}$;
- (d) The patient bank behaves optimally, i.e.,

$$\{m_t(\delta_p)\} \in \operatorname{argmax} \left\{ (1 - \delta_p) \sum_{t=1}^{\infty} \delta_p^{t-1} \mu_t(\delta_p) m_t : \{m_t\} \in \{m_L, m_H\}^{\infty} \right\};$$

- (e) The decision rules of almost all agents are optimal given the belief updating rule Θ .

Suppose σ is an equilibrium and let $\mathbb{N}_1(\sigma) = \{t : \mu_t > 0, \nu_t > 0\}$ be the set of periods where, regardless of the bank's type, a positive measure of agents enters the market in this equilibrium. Lemma 1 implies that $(\mathbb{N}_1(\sigma))^c = \mathbb{N}_0(\sigma) = \{t : \mu_t = \nu_t = 0\}$. One consequence of this result is that the market is never completely informative about the bank's type in any equilibrium.

Lemma 1. *Suppose $(M, \tau, \Theta, \{\mu_t\}, \{\nu_t\}, \{m_t\})$ is an equilibrium. Then, for all $t \in \mathbb{N}$, $\mu_t = 0$ if, and only if, $\nu_t = 0$.*

Proof: Let $\underline{t} = \inf\{t \in \mathbb{N} : \mu_t > 0 \text{ or } \nu_t > 0\}$, where the infimum of an empty set is taken to be plus infinity. If $\underline{t} = +\infty$, we are done. Suppose then that $\underline{t} < +\infty$. First observe that both $\mu_{\underline{t}}$ and $\nu_{\underline{t}}$ must be positive. Indeed, since $\mu_t = \nu_t = 0$ for $t < \underline{t}$, $\lambda_t(D|M, \tau, \delta) > 0$ if, and only if, $D \subseteq H'_t = [0, 1] \times \{A, \emptyset, A, \emptyset, \dots, A, \emptyset\}$. Moreover, if $D \subseteq H'_{\underline{t}}$, then

$$\lambda_t(D|M, \tau, 0) = \int_{\theta_0 \in D} \lambda(\theta_0) d\theta_0 = \lambda_t(D|M, \tau, \delta_p).$$

Hence, $\mu_{\underline{t}} = \nu_{\underline{t}}$, and so both must be positive by the definition of \underline{t} . Consequently, for all $t > \underline{t}$, the measure of agents in the market is greater than zero regardless of the bank's type. Because $m_L, m_H \in (0, 1)$, we then have that when $t > \underline{t}$, any event in H_t happens with positive probability when the bank is patient if, and only if, it happens with positive probability when the bank is impatient. In particular, the set of all private histories that lead an agent born in $t > \underline{t}$ to enter the market has positive probability when the bank is patient if, and only if, it has positive probability when the bank is impatient. \square

If σ is an equilibrium such that $\mathbb{N}_1(\sigma)$ is empty, we say σ is non-monetary. Such an equilibrium exists. In fact, let $\tau^n = \{\tau_t^n\}$, $M^n = \{M_t^n\}$, and $\Theta^n : \bigcup_{t=1}^{\infty} H_t \rightarrow [0, 1]$ be such that: (i) $\tau_t^n(i) \equiv A$ for all $i \in [0, 1]$ and $t \in \mathbb{N}$; (ii) $M_t^n \equiv m_H$ for all $t \in \mathbb{N}$; and (iii) $\Theta^n(h) = h^1$ for all private histories h . Then, $(M^n, \tau^n, \Theta^n, \{\mu_t\}, \{\nu_t\}, \{m_t\})$, with $\mu_t = \nu_t = 0$ and $m_t = m_H$ for all $t \in \mathbb{N}$, is an equilibrium.

When the patient bank uses a mixed strategy, the evolutions of both m_t and μ_t are no longer necessarily deterministic. Instead, they can follow a stochastic process. If this is the case, the belief a newly born agent has about the bank's type is not enough to determine his lifetime expected payoff from choosing the market. He also needs a (history dependent) conjecture about how m_t and μ_t are going to behave over time starting with his period of birth. This means that the equilibrium concept introduced above needs to be modified if one wants to consider this more general case. One exception is when the bank's behavior fails to be deterministic only off the equilibrium path, in which case the definition given above is appropriate.

4 The Full Commitment Case

In this section we assume that the bank can commit to its period 1 choice of m , that is, once it chooses the value of m in the first period, it cannot change it afterwards. This is equivalent to using the equilibrium notion introduced in the previous section, but reducing the set of strategies of the bank to $\{M^{LL}, M^{LH}, M^{HL}, M^{HH}\}$, where $M^{kl} = \{M_t^{kl}\}$, with $k, l \in \{L, H\}$, is such that

$$M_t^{kl}(\delta, \cdot) \equiv \begin{cases} m_k & \text{if } \delta = 0 \\ m_l & \text{if } \delta = \delta_p \end{cases}.$$

We show that when δ_p is close enough to 1, an equilibrium where the bank chooses M^{HL} and money circulates exists. Moreover, the sequence $\{\mu_t\}$ is bounded away from zero in this equilibrium, which implies that the subset of the population that transacts with money when the bank is patient does not die out over time.

For this, let Θ^* be such that: (i) $\Theta^*(h^1) = h^1$ for all $h^1 \in H_1$; (ii) If $h^{t+1} = (h^t, d, \omega)$, then $\Theta^*(h^{t+1}) = \Theta^*(h^t)$ when $d = A$ or $d = M$ and $\omega \in \{0, 1\} \cup \{e\}$, and

$$\Theta^*(h^{t+1}) = \frac{\Theta^*(h^t)m_L^{i+j}(1-m_L)^{n+1-i-j}}{\Theta^*(h^t)m_L^{i+j}(1-m_L)^{n+1-i-j} + (1-\Theta^*(h^t))m_H^{i+j}(1-m_H)^{n+1-i-j}} \quad (8)$$

when $d = M$ and $\omega = (i, j) = \{0, 1\} \times \{0, \dots, n\}$. Recall that e denotes the event that the market is empty. Moreover, let $\theta^*(\beta, n)$ be the value of θ such that $\theta v(m_L) + (1-\theta)v(m_H) = v_A$ and define $\tau^* = \{\tau_t^*\}$ to be such that $\tau_t^*(\cdot) \equiv d_t^*$, where

$$d_t^*(h^t) = \begin{cases} M & \text{if } \Theta^*(h^t) \geq \theta^*(\beta, n) \\ A & \text{if } \Theta^*(h^t) < \theta^*(\beta, n) \end{cases}. \quad (9)$$

To finish, let $\mu_t^* = \gamma_t(M^{HL}, \tau^*, \delta_p)$, $\nu_t^* = \gamma_t(M^{HL}, \tau^*, 0)$, and $m_t^* \equiv m_L$. Notice that $\mu_1^* = \nu_1^* > 0$.

Theorem 1. *There exists $\underline{\delta} \in (0, 1)$ such that if $\delta_p \geq \underline{\delta}$, then $\sigma^* = (M^{HL}, \tau^*, \Theta^*, \{\mu_t^*\}, \{\nu_t^*\}, \{m_t^*\})$ is an equilibrium. Moreover, the sequence $\{\mu_t^*\}$ is bounded away from zero.*

The following fact is useful in the proof of Theorem 1. If $\{x_t\}$ is a convergent sequence in the real line with limit x_∞ , then

$$\lim_{\delta \rightarrow 1} (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} x_t = x_\infty. \quad (10)$$

Proof of Theorem 1: Suppose the bank follows M^{HL} . Since under τ^* a positive measure of agents enters the market in period 1, the fraction of agents in the market with money is constant over time: equal to m_L when the bank is patient and equal to m_H when the bank is impatient. Hence, $\Theta^*(h^t) = \theta(h^t; \{\mu_t^*\}, \{\nu_t^*\}, \{m_t^*\})$. Moreover, expectations are satisfied by construction. Because $v(m_L)$ is the lifetime expected payoff from entering the market when the bank is patient and $v(m_H)$ is the same payoff when the bank is impatient, d_t^* is, for all $t \in \mathbb{N}$, an optimal decision rule for an agent born in this period. Therefore, we only need to show that m_L is indeed the optimal choice of m by the patient bank in order to establish that σ^* is an equilibrium. For this, notice that $\nu_t^* \downarrow 0$ and there exists $\underline{\mu} > 0$ such that $\mu_t^* \downarrow \underline{\mu} > 0$ by [2, Proposition 2].⁶ Since $(1-\delta_p) \sum_{t=1}^{\infty} \delta_p^{t-1} (\mu_t^* + \kappa) m_L$ is the patient bank's payoff when it chooses m_L and $(1-\delta_p) \sum_{t=1}^{\infty} \delta_p^{t-1} (\nu_t^* + \kappa) m_H$ is its payoff when it chooses m_H , the desired result is a consequence of (10). Recall our assumption that κ is infinitesimal. \square

⁶The setting in [2] is slightly different from the setting considered in this article. It is straightforward to adapt the proof of their Proposition 2 to our environment.

5 The No Commitment Case

By restricting the bank to make a once and for all decision on the value of m in period 1, we rule out any considerations about the time-consistency of its behavior. In this section we investigate what happens when the bank can change its decision of m at the beginning of every period, but this decision is still its private information. The main result we obtain is the following.

Theorem 2. *Suppose $\sigma = (M, \tau, \Theta, \{\mu_t\}, \{\nu_t\}, \{m_t\})$ is an equilibrium. Then either $\mathbb{N}_1(\sigma) = \emptyset$, i.e., σ is non-monetary, or*

$$\limsup_{t \rightarrow \infty} v_t(\{\alpha_{t+k-1}\}) = v_A.$$

Two results are needed before we can prove Theorem 2.

Lemma 2. *Let $\sigma = (M, \tau, \Theta, \{\mu_t\}, \{\nu_t\}, \{m_t\})$ be an equilibrium. Then $m_t = m_H$ for all $t \in \mathbb{N}_0(\sigma)$.*

Proof: Suppose $t \in \mathbb{N}_0(\sigma)$. Notice that the patient bank's flow payoff in this period is κm_t . Moreover, its choice of m in t does not affect future payoffs. Indeed, since the measure of agents who enter the market in t is zero, the behavior of the patient bank in this period does not affect the distribution of private histories across the agents in future periods. Hence, a profitable deviation for the patient bank is possible if $m_t = m_L$. \square

Lemma 3. *Suppose $\sigma = (M, \tau, \Theta, \{\mu_t\}, \{\nu_t\}, \{m_t\})$ is an equilibrium. Then, either $\mathbb{N}_1(\sigma) = \emptyset$ or $\mathbb{N}_1(\sigma)$ is infinite.*

Proof: Suppose $\mathbb{N}_1(\sigma)$ is non-empty, but finite; i.e., suppose there is $\bar{t} \in \mathbb{N}$ such that both $\mu_{\bar{t}}$ and $\nu_{\bar{t}}$ are positive, but $\mu_t = \nu_t = 0$ for $t > \bar{t}$. Recall that $w_{i,t}^1 = w_{i,t}^1(\{\alpha_{t+k-1}\})$ is the lifetime expected reward from entering the market in period t with i units of money when the bank is patient. Hence, if an agent born in $t \geq \bar{t}$ with belief θ enters the market, his lifetime expect payoff is

$$V_t(\theta) = \theta[m_t(w_{1,t}^1 - c) + (1 - m_t)w_{0,t}^1] + (1 - \theta)v(m_H). \quad (11)$$

Now observe that: (a) $\alpha_t = \alpha_{\bar{t}}$ for all $t > \bar{t}$; (b) $m_{\bar{t}} = m_H$, otherwise the patient bank would have a profitable deviation; (c) the measure of agents born in \bar{t} with belief in $W_{\bar{t}} = \{\theta : V_{\bar{t}}(\theta) \geq v_A\}$ is positive regardless of the bank's type. These facts imply that: (d) $V_{t+1}(\theta) = V_{\bar{t}}(\theta)$ for all $\theta \in [0, 1]$; and (e) $\alpha_{\bar{t}} > m_H$. Indeed, $w_{i,\bar{t}+1}^1 = w_{i,\bar{t}}^1$ by (a) and $m_{\bar{t}+1} = m_H$ by Lemma 2, so that (d) holds by (11) and (b). To see why (e) is also true, notice that if $\alpha_{\bar{t}} = m_H$, then (a) and (b) imply that $V_{\bar{t}}(\theta) \equiv v(m_H) < v_A$, in which case $W_{\bar{t}}$ is empty. A consequence of (e) is that the market is informative about the bank's type in period \bar{t} . Therefore, there is a positive probability that any agent who enters the market in \bar{t} ends this period with a strictly higher belief about the bank.

Since (e) also implies that $V_{\bar{t}}$ is strictly increasing in θ , the measure of agents who are born in $\bar{t} + 1$ with belief in $\{\theta : V_{\bar{t}+1}(\theta) > v_A\}$ is positive when the bank is patient by (c) and (d). This, however, implies that $\mu_{\bar{t}+1} > 0$, a contradiction. \square

For each equilibrium σ define $B_t(\sigma, \delta)$ to be the probability measure over $[0, 1]$ such that

$$B_t(\sigma, \delta)(D) = B_t(D|\sigma, \delta) = \lambda_t(\{h^t \in H_t : \theta(h^t) \in D\}|\delta). \quad (12)$$

By definition, if σ is the equilibrium under play, then $B_t(D|\sigma, \delta)$ is the measure of agents born in t with belief in D when the bank's discount factor is δ . Notice that $B_1(D|\sigma, \cdot) = \hat{\lambda}_1(D)$, the measure of agents born in 1 with a belief in D . Since $\hat{\lambda}_1$ has full support, the same is true of $B_t(\sigma, \delta)$ for $t > 1$.

Proof of Theorem 2: Suppose $\mathbb{N}_1(\sigma)$ is infinite and denote $v_t(\{\alpha_{t+k-1}\})$ by v_t . First notice that $\limsup v_t \geq v_A$, otherwise there exists $\underline{t} \in \mathbb{N}$ such that $v_t < v_A$ for $t \geq \underline{t}$, in which case $\mathbb{N}_1(\sigma)$ would be finite. Suppose then that $\limsup v_t > v_A$. We want to show that a profitable deviation is possible for the patient bank.

Step 1: Since $B_t([\theta, 1]|\sigma, \delta) > 0$ for all $t \in \mathbb{N}$ and $\theta \in [0, 1]$, it must be that $v_t \leq v_A$ for all $t \in \mathbb{N}_0(\sigma)$. Therefore, there exists $v > v_A$ such that $\bar{\mathbb{N}}_1 = \{t \in \mathbb{N}_1(\sigma) : v_t \geq v\}$ is infinite. Assume, without loss, that $\bar{\mathbb{N}}_1 = \mathbb{N}_1(\sigma)$ and let $\bar{\theta} \in (0, 1)$ be such that $\bar{\theta}v/2 + (1 - \bar{\theta})v(m_H) = 0$. Then, any agent born in $t \in \mathbb{N}_1(\sigma)$ with belief in $[\bar{\theta}, 1]$ enters the market. As a consequence, at least one of the two subsequences $\{\mu_t\}_{t \in \mathbb{N}_1(\sigma)}$ and $\{\nu_t\}_{t \in \mathbb{N}_1(\sigma)}$ is bounded away from zero. This follows from a straightforward modification of the proof of Theorem 5.1 in [4].

Step 2: We now show that $\{\alpha_t\}$ cannot converge to m_H . Suppose not. Since $a(m)B(\eta) = a(\eta)$ for all $m, \eta \in (0, 1)$, $v_t = a(m_t)b(m_t) - m_t c + \sum_{k=1}^{\infty} \beta^k a(\alpha_{t+k})b(\alpha_{t+k})$. Therefore,

$$\begin{aligned} |(v_t - a(m_t)b(m_H) + m_t c) - (v(m_H) - a(m_H)b(m_H) + m_H c)| \leq \\ \sum_{k=0}^{\infty} \beta^k |a(\alpha_{t+k})b(\alpha_{t+k}) - a(m_H)b(m_H)| \leq (1 - \beta)^{-1} \sup_{k \geq 0} |a(\alpha_{t+k})b(\alpha_{t+k}) - a(m_H)b(m_H)|, \end{aligned}$$

and so $v_t - a(m_t)b(m_H) + m_t c$ converges to $v(m_H) - a(m_H)b(m_H) + m_H c$ as $t \rightarrow \infty$. Since

$$a(m_H)b(m_H) - m_H c - [a(m_t)b(m_H) - m_t c] = (m_H - m_t)(1 - m_H)(u - c) \geq 0,$$

we can then conclude that there exists $\underline{t} \in \mathbb{N}$ such that $v_t \leq v(m_H) < v_A$ for $t \geq \underline{t}$, a contradiction. Consequently, there exist $m \in [m_L, m_H)$ and a subsequence $\{\alpha_{t_j}\}$ of $\{\alpha_t\}$ such that $\alpha_{t_j} \rightarrow m$. Let $k(t) = \max\{t' \leq t : t' \in \mathbb{N}_1(\sigma)\}$. Since, by construction, $\alpha_t = \alpha_{k(t)}$ for all $t \in \mathbb{N}$, $\alpha_{k(t_j)} \rightarrow m$ as well. Assume, without loss, that

$$\{\alpha_t\}_{t \in \mathbb{N}_1(\sigma)} \rightarrow m. \quad (13)$$

Step 3: We now show that (13) implies that $\{\mu_t\}_{t \in \mathbb{N}_1(\sigma)}$ is bounded away from zero. Suppose not. Then, $\{\nu_t\}_{t \in \mathbb{N}_1(\sigma)}$ must be bounded away from zero by Step 1. For each family $i \in [0, 1]$ let $n_t(i)$ be the number of members of this family that enter the market up to period $t-1$. By assumption, there is a positive mass of families for which $n_t(i) \uparrow \infty$ when the bank is impatient, otherwise $\{\nu_t\}_{t \in \mathbb{N}_1(\sigma)}$ converges to zero. Denote this set of families by \mathfrak{F} . Now let $\theta_t(i)$ be the belief of the member of i that is born in t and suppose that all members of i who enter the market *do not* use their initial money holdings – whether they receive money from the bank upon entering the market – to update beliefs. If $i \in \mathfrak{F}$, then (13) together with the Kolmogorov's Strong Law of Large Numbers implies that $\{\theta_t(i)\}$ converges to zero almost surely. Therefore, by [1, Lemma 1], $\{\theta_t(i)\}$ converges to zero almost surely even when the members of i who enter the market use their initial money holdings to update beliefs. However, we know that it is optimal to choose autarky if $\theta_t(i) < \theta^*(\beta, n)$. Hence, \mathfrak{F} cannot have a positive mass, a contradiction. Let $\mu > 0$ be a lower bound for $\{\mu_t\}_{t \in \mathbb{N}_1(\sigma)}$.

Step 4: Let $\mathbb{N}_1(\sigma, m_k) = \{t \in \mathbb{N}_1(\sigma) : m_t = m_k\}$, with $k \in \{L, H\}$, and suppose that $\mathbb{N}_1(\sigma, m_L)$ is finite. Because $\{\mu_t\}_{t \in \mathbb{N}_1(\sigma)}$ is bounded away from zero, $\lim_{t \rightarrow \infty} \sum_{k \in \mathbb{N}_1(\sigma) \cap \{1, \dots, t\}} \mu_k = +\infty$. Hence,

$$\alpha_t = \underbrace{\frac{\sum_{k \in \mathbb{N}_1(\sigma, m_L) \cap \{1, \dots, t\}} \mu_k}{\sum_{k \in \mathbb{N}_1(\sigma) \cap \{1, \dots, t\}} \mu_k}}_{A_t} m_L + \underbrace{\frac{\sum_{k \in \mathbb{N}_1(\sigma, m_H) \cap \{1, \dots, t\}} \mu_k}{\sum_{k \in \mathbb{N}_1(\sigma) \cap \{1, \dots, t\}} \mu_k}}_{1-A_t} m_H \rightarrow m_H,$$

since A_t converges to zero. This, however, contradicts (13). Consequently, $\mathbb{N}_1(\sigma, m_L)$ is infinite.

Step 5: Let $\overline{\mathbb{N}}$ be the one point compactification of $\mathbb{N} \cup \{\infty\}$ and denote the set of all signed measures over $\overline{\mathbb{N}}$ that have bounded variation by $ba(\overline{\mathbb{N}})$. We take $ba(\overline{\mathbb{N}})$ to be endowed with the bounded variation norm. Moreover, let $\chi_t(i)$ be the (Borel) probability measure on $\overline{\mathbb{N}}$ such that $\chi_t(i)(\{n\})$ is the probability that $n_t(i) = n$. It is easy to see that for each family i , $\chi_t(i)$ converges in norm to $\chi_\infty(i)$, where $\chi_\infty(i)(\{n\})$ is the probability that n members of i enter the market. Therefore, by Egoroff's theorem, for all $\epsilon > 0$ there exists a set \mathfrak{F}_ϵ of families such that $\chi_t(i)$ converges uniformly to $\chi_\infty(i)$ for all $i \in \mathfrak{F}_\epsilon$ and the measure of \mathfrak{F}_ϵ is at least $1 - \epsilon$. We can then assume, without loss, that $\chi_t(i)$ converges uniformly to $\chi_\infty(i)$ on a set of measure 1. Denote this set of families by \mathfrak{F} . From now on, we take the bank to be patient.

Observe that there exist $K \in \mathbb{N}$ such that $\sum_{t=K+1}^\infty \delta_p^t < \mu(m_H - m_L)/2$ and $0 < \epsilon \leq 1 - \bar{\theta}$ such that if an agent is born with belief in $[1 - \epsilon, 1]$, then the next K generations of his family always enter the market as long as their date of birth lies in $\mathbb{N}_1(\sigma)$. Just choose ϵ close enough to zero so that if $\Theta(h) \geq 1 - \epsilon$ and (h, h') denotes the history with initial segment h and final segment h' , then $\Theta((h, h')) \geq 1 - \epsilon$ for all $h' \in \Pi^K$. Also observe that if $n_t(i) \uparrow \infty$, then $\{\theta_t(i)\}_t$ converges to

one almost surely. This follows from [1, Lemma 1] and (13). Consequently, there exists \underline{N} such that if $n_t(i) > \underline{N}$, then the probability that $\theta_t(i) \geq 1 - \epsilon$ is greater than $1 - \mu(1 - \delta_p)(m_H - m_L)/2$, as almost sure convergence implies convergence in measure.

Now let t_1 be such that $\#(\mathbb{N}_1(\sigma) \cap \{1, \dots, t_1\}) = \underline{N} + 1$, where $\#(A)$ is the number of elements in a set A . We know that for all $\epsilon > 0$ there exists t_2 such that if $t \geq t_2$, then $|\chi_t(i)(A) - \chi_\infty(i)(A)| \leq \epsilon$ for all $A \in \mathcal{A} = \{\{1\}, \{2\}, \dots, \{\underline{N}\}, \{\underline{N} + 1, \dots, \infty\}\} = \mathcal{A}$ and $i \in \mathfrak{F}$. We can then assume, without loss, that there exists t_2 such that if $t > t_2$, then $\chi_t(i)(A) = \chi_\infty(i)(A)$ for all $A \in \mathcal{A}$ and $i \in \mathfrak{F}$. This implies that if $t \geq \bar{t} = \max\{t_1, t_2\}$, then for all $i \in \mathfrak{F}$ either: (i) $n_t(i)$ is constant, so that no members of i enter the market after t ; or (ii) $n_t(i) \geq \bar{N} + 1$. In other words, if $t \in \mathbb{N}_1(\sigma) \cap \{\bar{t}, \dots\}$, then almost all agents who enter the market in t have at least $\bar{N} + 1$ previous members of his family who have entered the market.

To finish, consider a period $t \in \mathbb{N}_1(\sigma, m_L)$ such that $t > \bar{t}$ and ignore κ , since it is infinitesimal. The (patient) bank's lifetime payoff from sticking to its prescribed strategy is

$$\sum_{k=1}^{\infty} \delta_p^k \mu_{t+k} m_{t+k} + \mu_t m_L.$$

If, instead, it does a one-shot deviation, its lifetime payoff is at least

$$\mu_t m_H + \left(1 - \frac{1}{2} \mu(1 - \delta_p)(m_H - m_L)\right) \sum_{k=1}^K \delta_p^k \mu_{t+k-1} m_{t+k-1}.$$

Hence, the patient bank's payoff gain from a one-shot deviation at t is no less than

$$\begin{aligned} & \mu_t(m_H - m_L) - \frac{1}{2} \mu(1 - \delta_p)(m_H - m_L) \sum_{k=1}^K \delta_p^k \mu_{t+k} m_{t+k} - \sum_{k=K+1}^{\infty} \delta_p^k \mu_{t+k} m_{t+k} \\ & \geq \mu(m_H - m_L) - \frac{1}{2} \mu(1 - \delta_p)(m_H - m_L) \sum_{k=1}^K \delta_p^k - \sum_{k=K+1}^{\infty} \delta_p^k \\ & > \mu(m_H - m_L) - \mu(m_H - m_L) > 0, \end{aligned}$$

since $\mu_t m_t < 1$ for all $t \in \mathbb{N}$. We have the desired result. \square

A straightforward consequence of Theorem 2 is that the patient bank overissues infinitely many times in any monetary equilibrium.⁷

Corollary 1. *Suppose σ is a monetary equilibrium. Then $\mathbb{N}_1(\sigma, m_H)$ is infinite.*

⁷Theorem 2 is a characterization result. We have not been able to establish the existence of a monetary equilibrium so far, but we believe that one where $\{t \in \mathbb{N} : m_t = m_L\} = \{1, \dots, T\}$, with $T \geq 2$, exists. The main difficulty is with the condition that expectations must be correct.

Proof: Suppose not. Then, there exists $t_0 \in \mathbb{N}$ and $\alpha \in [m_L, m_H)$ such that $\{\alpha_{t+t_0-1}\}$ converges from below to α . If $\mathbb{N}_1(\sigma, m_H)$ is empty, so that $t_0 = 1$, then $\alpha_t \equiv m_L$, and so $v_t(\{\alpha_{t+k-1}\}) \equiv v(m_L) > v_A$, a contradiction. Hence, it must be that $\{\alpha_{t+t_0-1}\}$ is non-constant. Let

$$v(\alpha, m_L) = a(m_L)b(m_L) - m_L c + \sum_{k=1}^{\infty} \beta^k a(\alpha)b(m_L).$$

By assumption, if $t \geq t_0$, then $v_t(\{\alpha_{t+k-1}\})$ is non-constant and converges to $v(\alpha, m_L)$ from below. Since it must be that $v_t \geq v_A$ for all $t \in \mathbb{N}_1(\sigma)$, we can then conclude that $v(\alpha, m_L) > v_A$, a contradiction as well. \square

6 Imperfect Memory

The analysis so far suggests that in order to have an equilibrium where the patient bank never overissues there must be something that prevents its reputation from increasing too much when it always chooses m_L . Putting it differently, we need a mechanism that provides the patient bank with the incentive to always invest in its reputation by never choosing m_H . With this in mind we modify our environment by assuming that for any $t \geq 2$ there is a probability $\lambda > 0$ that a newly born agent does not inherit his parent's private history. Instead, this agent's market/autarky decision is based on his prior belief that the bank is patient. For simplicity, we assume that all agents who fail to inherit their parent's histories have the same prior θ_0 . Nothing would change if this prior were to be determined by an independent random draw from the p.d.f. λ . In this way, regardless of what the patient bank does, there is always a positive measure of newly born agents for whom its reputation is not high.

The equilibrium notion we use is the one introduced in Section 3 with the added requirement that the bank has now to behave in a sequentially rational way; i.e., it must behave optimally after every history. We also have to change the definition of H_t for $t > 1$ to incorporate the fact that memory is imperfect. Now, when $t > 1$, $H_t = ([0, 1] \times \Pi^{t-1}) \cup \{N\}$, where N denotes the event that no private history is inherited.

Let $\Theta^{**} : \bigcup_{t=1}^{\infty} H_t \rightarrow [0, 1]$ be such that

$$\Theta^{**}(h) = \begin{cases} \theta_0 & \text{if } h = N \\ \Theta^*(h) & \text{otherwise} \end{cases}, \quad (14)$$

where Θ^* is the belief updating rule of Section 4, and define $\tau^{**} = \{\tau_t^{**}\}$ to be such that $\tau^{**}(\cdot) \equiv d_t^*$, where d_t^* is given by (9) with Θ^{**} in the place of Θ^* . Now let $\{\mu_t^{**}(n)\}$, $\{\nu_t^{**}(n)\}$, and $\{m_t^{**}\}$ be

such that $\mu_t^{**}(n) = \gamma_t(M^{HL}, \tau^{**}, \delta_p)$, $\nu_t^{**}(n) = \gamma_t(M^{HL}, \tau^{**}, 0)$, and $m_t^{**} \equiv m_L$, where M^{HL} is the strategy profile for the bank defined in Section 4 and we make the dependence of μ_t^{**} and ν_t^{**} on n explicit. To finish, let α and $\underline{\delta}$ be such that

$$\alpha = \frac{m_H - m_L}{m_L} \quad \text{and} \quad \underline{\delta} = \frac{\alpha}{(1 - \lambda)(\lambda + \alpha)},$$

and assume that $\theta_0 \geq \theta^*(\beta) = \max_{n \geq 2} \theta^*(\beta, n)$, where we recall that $\theta^*(\beta, n)$ is the value of θ for which $\theta v(m_L) + (1 - \theta)v(m_H) = v_A$. It is easy to see that $\theta^*(\beta) = \theta^*(\beta, 2)$.

Theorem 3. *Let $\sigma^{**} = (M^{HL}, \tau^{**}, \Theta^{**}, \{\mu_t^{**}\}, \{\nu_t^{**}\}, \{m_t^{**}\})$. Suppose $\alpha < 1/2$, $\lambda \in (\alpha, 1/2)$, and $\delta_p > \underline{\delta}$. Then, there exists $n(\alpha) \in \mathbb{N}$ such that σ^{**} is an equilibrium if $n \geq n(\alpha)$.*

Observe that $\underline{\delta} < 1/2(1 - \lambda)$ since $\alpha < \lambda$, and so $\underline{\delta} < 1$ when $\lambda < 1/2$. That α cannot be too big is expected, as it is the relative one-period gain when a patient bank deviates and chooses $m = m_H$. The bounds on λ are also intuitive. The probability that private histories are transmitted cannot be too large, otherwise the incentives for the patient bank to maintain a good reputation are not enough to prevent it from choosing m_H . This probability cannot be too small as well, for in this case the patient bank cannot benefit from a good reputation, as it disappears too quickly. Finally, notice that the assumption on θ_0 implies that under τ^{**} any newly born agent who fails to inherit a private history chooses the market.

Proof of Theorem 3: Notice, from the proof of Theorem 1 in Section 4, that for all $n \geq 2$, (i) $\Theta^{**}(h^t) = \theta(h^t; \{\mu_t^{**}\}, \{\nu_t^{**}\}, \{m_t^{**}\})$, (ii) expectations are (trivially) satisfied, and (iii) agents are behaving optimally given the bank's behavior and the behavior of the other agents. Hence, we just need to check that no one-shot deviation pays for the patient bank if n is sufficiently high. As in the proof of Theorem 2, we ignore κ since it is infinitesimal.

Fix a period k and suppose the bank is patient. Given any history $h \in \mathcal{H}_k$ for the bank, we can compute the distribution of beliefs across the agents born in k induced by h . Denote this distribution by $\lambda_k(h)$ and let $\Delta_k(\theta)$ be the bank's gain from a one-shot deviation in k if all agents born in this period have the same belief θ . The bank's gain from a one-shot deviation in k , as a function of h , is then given by

$$\Delta_k(h) = \int \Delta_k(\theta) d\lambda_k(\theta|h).$$

Given our imperfect memory assumption, we know that for each $h \in \mathcal{H}_k$ there exist a distribution $\hat{\lambda}_k(h)$ such that

$$\Delta_k(h) = (1 - \lambda) \int \Delta_k(\theta) d\hat{\lambda}_k(\theta|h) + \lambda \Delta_k(\theta_0).$$

Consider an agent born in k with belief θ . For each $t \geq 1$, denote by $\xi_{t,k}(m, \theta)$ the probability that the generation $k+t$ member of this agent's family enters the market in $k+t$ if the patient bank

chooses $m \in \{m_L, m_H\}$ in k and m_L from $k + 1$ on, and private histories are always transmitted from one generation to the next. Now let $\phi_{t,k}(m, \theta)$ be the same probability when private histories can fail to be transmitted. Then,

$$\phi_{t,k}(m, \theta) = (1 - \lambda)^t \xi_{t,k}(m, \theta) + \sum_{\tau=1}^{t-1} \lambda(1 - \lambda)^\tau \xi_{\tau, k+(t-1)-\tau}(m_L, \theta_0) + \lambda. \quad (15)$$

Indeed, with probability $(1 - \lambda)^t$ private histories are always passed from one generation to the next, while with probability $\lambda(1 - \lambda)^\tau$ the last period before $k + t$ where private histories are not passed from one generation to the following is $k + (t - 1) - \tau$, with $\tau \in \{0, \dots, t - 1\}$.

Observe $\xi_{t,k}$ is independent of k , since all newly born agents follow the same cutoff belief strategy and the market environment is stationary. Therefore, the same is true for the probabilities $\phi_{t,k}$. Because of this, we omit k from both $\xi_{t,k}$ and $\phi_{t,k}$ in what follows. Let ζ be the indicator function of $[\theta^*, 1]$. Then, for all $k \in \mathbb{N}$,

$$\begin{aligned} \Delta_k(\theta) &= (m_H - m_L)\zeta(\theta) + m_L \sum_{t=1}^{\infty} \delta_p^t [\phi_t(m_H, \theta) - \phi_t(m_L, \theta)] \\ &= (m_H - m_L)\zeta(\theta) + m_L \sum_{t=1}^{\infty} \delta_p^t (1 - \lambda)^t [\xi_t(m_H, \theta) - \xi_t(m_L, \theta)], \end{aligned}$$

where the second equality follows from (15). Now observe that for each $t \in \mathbb{N}$, $\xi_t(m_L, \theta) \geq \xi_t(m_H, \theta)$ for all $\theta \in [0, 1)$ and $\xi_t(m_L, 1) = \xi_t(m_H, 1)$. Hence, $\Delta_k(\theta) \leq \Delta_k(1)$ for all $k \in \mathbb{N}$ and $\theta \in [0, 1)$, and so, for each $k \in \mathbb{N}$,

$$\Delta_k(h) \leq \underbrace{(1 - \lambda)\Delta_k(1) + \lambda\Delta_k(\theta_0)}_{\Delta_k} \quad (16)$$

for all $h \in \mathcal{H}_k$. It is then enough to show that the right-hand side of (16) is negative for all $k \in \mathbb{N}$ when n is sufficiently large.

In the Appendix we prove that $\xi_t(m_L, \theta_0) - \xi_t(m_H, \theta_0) \leq \xi_{t+1}(m_L, \theta_0) - \xi_{t+1}(m_H, \theta_0)$ for all $t \in \mathbb{N}$; that is, the loss to the patient bank from a worse reputation increases over time. Hence,

$$\begin{aligned} \Delta_k &= \lambda \left[(m_H - m_L) + m_L \sum_{t=1}^{\infty} \delta_p^t (1 - \lambda)^t [\xi_t(m_H, \theta_0) - \xi_t(m_L, \theta_0)] \right] + (1 - \lambda)(m_H - m_L) \\ &\leq m_L \left[\alpha - \frac{\delta_p \lambda (1 - \lambda)}{1 - \delta_p (1 - \lambda)} \Delta \xi_1 \right] = m_L \Psi(\delta_p, \lambda, \alpha). \end{aligned}$$

where $\Delta \xi_1 = \xi_1(m_L, \theta_0) - \xi_1(m_H, \theta_0)$. To finish, notice that Ψ is strictly increasing in δ_p and that $\Psi(\underline{\delta}, \lambda, \alpha) = \alpha[1 - \Delta \xi_1]$. Because $\lim_{n \rightarrow \infty} \Delta \xi_1 = 1$ by the Strong Law of Large numbers, we can then conclude that there exists $n(\alpha)$ such that $\Delta_k < 0$ for all $k \in \mathbb{N}$ when $n \geq n(\alpha)$ and $\delta_p > \underline{\delta}$. \square

7 Conclusion

This work contributes to the literature on endogenous money. It addresses the feasibility of fiat money when its supply is determined by a single self-interested agent (the bank). This is done in an environment where trade is decentralized and agents are anonymous and have heterogeneous preferences, so that money is essential. We depart from previous work by assuming that: (i) there is uncertainty about the bank's preferences, so that there is a role for reputation; (ii) there is no technology that allows the bank to be publicly monitored, so that information is decentralized and its flow is constrained by the same technology that hinders trade. The main feature of our model is that the bank faces a trade-off between short-run gains from overissue and long-run losses due to a decrease in its reputation for providing valuable currency.

We show that if the patient bank can commit to a choice of money supply, then a monetary equilibrium where it does not overissue exists as long as it is sufficiently patient. This equilibrium, however, is not time-consistent when the patient bank cannot commit to a plan of action. In fact, we show that in the absence of such commitment, the only monetary equilibria possible have the patient bank destroying the value of money over time through repeated overissue. The reason for this negative result is that if the patient bank behaves in a way to distinguish itself from the impatient bank long enough, its reputation eventually becomes insensitive to its choice of money supply. At this point, the patient bank has an incentive to overissue, since this has a marginal impact on its future revenue from money issue. Agents anticipate this, and as a consequence any monetary equilibrium where the patient bank does not drive the value of money to zero breaks down. To summarize, the failure of reputation to discipline the behavior of the patient bank is because it works too well.

Following this insight, we show that if memory is imperfect, the patient bank's incentive to maintain a good reputation never disappears, and so a monetary equilibrium with no overissue is possible. The shortcoming of this approach is that even though reasonable, it is somewhat ad-hoc. An interesting question is whether there are other mechanisms that can discipline the bank's behavior and preclude overissue. In this regard, we believe that there are two possible directions for study. First, it may be that inconvertibility is at the root of the overissue problem. As Friedman and Schwartz point out, "historically, producers of money have established confidence by promising convertibility into some dominant money, generally, specie. Many examples can be cited of fairly long-continued and successful producers of private moneys convertible into specie" [9, p. 45].⁸ The second alternative, which we are currently investigating, is to introduce competition among

⁸The reason why convertibility may work together with reputation is that its failure can act as a signal that the bank is impatient, thus disciplining the behavior of the patient bank.

money issuers. According to von Hayek, one of its main advocates, “convertibility is a safeguard necessary to impose upon a monopolist, but unnecessary with competing suppliers who cannot maintain themselves in the business unless they provide money at least as advantageous to the user as anybody else” [10, p. 111].

References

- [1] M. Aoyagi, Mutual Observability and the Convergence of Actions in a Multi-Armed Bandit Model, *J. Econ. Theory* 82 (1998) 405-424.
- [2] L. Araujo, B. Camargo, Information, Learning, and the Stability of Fiat Money, *J. Monet. Econ.* forthcoming, 2005
- [3] L. Araujo, B. Camargo, Monitoring versus Reputation in Monetary Economies, unpublished manuscript, 2005.
- [4] J. Banks, R. Sundaram, Denumerable-Armed Bandits, *Econometrica* 60 (1992) 1071-1096.
- [5] A. Berentsen, Time-Consistent Private Supply of Outside Paper Money, *Europ. Econ. Rev.* forthcoming, 2005
- [6] R. Cavalcanti, A. Erosa, T. Temzelides, Private Money and Reserve Management in a Random-Matching Model, *J. Polit. Economy* 107 (1999) 929-945.
- [7] R. Cavalcanti, N. Wallace, A Model of Private Banknote Issue, *Rev. Econ. Dynam.* 2 (1999) 104-136.
- [8] S. Fischer, Friedman Versus Hayek on Private Money, *J. Monet. Econ.* 17 (1986) 433-439.
- [9] M. Friedman, A. Schwartz, Has Government Any Role in Money?, *J. Monet. Econ.* 17 (1986) 37-62.
- [10] F. Hayek, Denationalisation of Money - The Argument Refined, The Institute of Economic Affairs, London, 1990.
- [11] N. Kiyotaki, R. Wright, A Search-Theoretic Approach to Monetary Economics, *Amer. Econ. Rev.* 83 (1993) 63-77.
- [12] B. Klein, The Competitive Supply of Money, *J. Money, Credit, Banking* 6 (1974) 423-453.

- [13] D. Kreps, P. Milgrom, D. Roberts, R. Wilson, Rational Cooperation in the Finitely Repeated Prisoner's Dilemma, *J. Econ. Theory* 27 (1982) 245-252.
- [14] G. Mailath, L. Samuelson, Your Reputation Is Who You're Not, Not Who You'd Like To Be, CARESS Working Paper #98-11, University of Pennsylvania, 1998.
- [15] G. Mailath, L. Samuelson, Who Wants a Good Reputation, *Rev. Econ. Stud.* 68 (2001) 415-441.
- [16] A. Martin, S. Schreft Currency Competition: A Partial Vindication of Hayek, *J. Monet. Econ.* forthcoming, 2005.
- [17] O. Moav, Z. Neeman, The Quality of Information and Incentives for Effort, unpublished manuscript, 2005.
- [18] C. Monnet, Private vs. Public Money, *Int. Econ. Rev.* forthcoming, 2005

Appendix

Lemma 4. $\xi_t(m_L, \theta_0) - \xi_t(m_H, \theta_0) \leq \xi_{t+1}(m_L, \theta_0) - \xi_{t+1}(m_H, \theta_0)$ for all $t \in \mathbb{N}$

Proof: Let $c(h^t)$ denote the number of meetings with money in the private history h^t . It is straightforward to show that an agent born in t with private history h^t chooses the market if, and only if, $c(h^t) \leq \alpha(t - k)(n + 1) + \gamma$, where α is a positive constant depending on m_L and m_H , and γ is a non-negative constant depending on θ_0 . Therefore, $\xi_t(m, \theta_0)$ is given by

$$\sum_{(c_1, \dots, c_t) \in C_t} \underbrace{\binom{n+1}{c_1} \dots \binom{n+1}{c_t}}_{B_t(c_1, \dots, c_t)} m^{c_1} (1-m)^{n+1-c_1} m_L^{c_2+\dots+c_t} (1-m_L)^{(t-1)(n+1)-c_2-\dots-c_t},$$

where $C_t = \{(c_1, \dots, c_t) : c_\tau \leq \lfloor \alpha\tau(n+1) + \gamma \rfloor - c_1 - \dots - c_{\tau-1}, \tau = 1, \dots, t-1\}$. Here $\lfloor x \rfloor$ is, by definition, the greatest integer smaller than x .

Now let $L_t(m) = \xi_t(m, \theta_0) - \xi_{t+1}(m, \theta_0)$. Then

$$L_t(m) = \sum_{(c_1, \dots, c_{t+1}) \in D_{t+1}} B_t(c_1, \dots, c_t) m^{c_1} (1-m)^{n+1-c_1} \underbrace{m_L^{c_2+\dots+c_{t+1}} (1-m_L)^{t(n+1)-c_2-\dots-c_{t+1}}}_{m_L(c_2, \dots, c_{t+1})},$$

where $D_{t+1} = \{(c_1, \dots, c_{t+1}) : (c_1, \dots, c_t) \in C_t, c_{t+1} \geq \lfloor \alpha(t+1)(n+1) + \gamma \rfloor - c_1 - \dots - c_t + 1\}$. Hence $L_t(m_L) - L_t(m_H) = \xi_t - \xi_{t+1} - (\hat{\xi}_t - \hat{\xi}_{t+1})$ is equal to

$$\sum_{(c_1, \dots, c_{t+1}) \in D_{t+1}} B_{t+1}(c_1, \dots, c_{t+1}) m_L(c_2, \dots, c_{t+1}) [m_L^{c_1} (1-m_L)^{n+1-c_1} - m_H^{c_1} (1-m_H)^{n+1-c_1}].$$

To finish observe that the term in brackets in the above expression is non-positive if, and only if, $c_1 \leq \alpha(n+1)$. Therefore $L_t(m_L) \leq L_t(m_H)$, the desired result. \square