Private Supply of Fiat Money

Luis Araujo*,†
Braz Camargo‡

June 2, 2005

Abstract

The question of whether private money is viable in the absence of any external control is an old one. We address it in an economy with decentralized trade and information where a self-interested agent, the bank, has the monopoly over the provision of fiat money and is not limited on how much it can issue over time. The assumption of decentralized information is a crucial one and sets this work apart from all previous investigations in the literature. We argue below that the latter is a necessary assumption if one wants to capture the absence of any external limits on the bank's behavior. We first show that if the bank can commit to a choice of money supply, a monetary equilibrium with no overissue exists. This equilibrium, however, is not time-consistent, and so does not survive when no commitment is possible for the bank. In fact, we show that with no commitment, the only monetary equilibria possible have overissue happening infinitely many times. We finish by showing how this time-consistency problem can be solved.

JEL Classification: D82, D83, E00

Key Words: Private Money Supply, Time-Consistency, Reputation

1 Introduction

In most economies the right to print fiat money is a government monopoly. Agreement about the propriety of this monopoly is widespread. Even a fierce defender of free market systems like Milton Friedman has argued that the very nature of fiduciary money calls for a governmental role.

*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 Saturnino Rangel Mauro 245, Vitória, Espírito Santo, 29062-030, Brazil.

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

Friedman has said: "The technical monopoly character of a pure fiduciary currency makes essential the setting of some external limit on its amount" (1959, p. 4). Underlying Friedman's view is the belief that private agents do not have enough incentives to provide a stable value for money: "Something like a moderately stable monetary framework seems an essential prerequisite for the efective operation of a private market economy. It is dubious that the market can by itself provide such a framework. Hence, the function of providing one is an essential governmental function" (1959, p. 8).

In this paper we evaluate Friedman's point. We consider whether a stable currency can be provided in an economy with a private monopoly in the supply of fiat money when there are no external limits on its supply. We do so in an environment with decentralized trade and asymmetric preferences, elements that lead to a natural role for money as a medium of exchange. We also assume that information about the behavior of the note issuer is decentralized and its flow is restricted by the same technology that hinders trade. This is a key assumption in our model and sets it apart from previous work, where the existence of some exogenously given public information about the behavior of the note issuer(s) is assumed.[1] In our opinion, the availability of any such information can only be justified by the existence of some institution put in place to monitor, and as a consequence limit, the note issuer's behavior.

We take a simple version of the model introduced in Kiyotaki and Wright (1993) as a starting point for our analysis.[2] We modify it by allowing the money supply to be determined in every period by a self-interested individual, the bank, and by making this choice its private information. The bank derives utility from money issue and is either patient or not. The other agents in the economy have two options. They either stay in autarky or go to the market, where they trade with the help of money. The bank's revenue from money issue in any given period is proportional to the number of agents that enter the market at that point in time. The only way agents learn about the bank's choice is through their private experiences in the market. We assume that autarky is better than the market if money is always overissued, which the impatient banks does, while the opposite is true when overissue never takes place.

[1]Berentsen (2005) shows that in the presence of a monopoly in the supply of fiat money a monetary equilibrium is possible when the monopolist's history is public. Cavalcanti and Wallace (1999a,b) and Cavalcanti, Erosa, and Temzelides (1999) consider an economy where a subset of agents is allowed to print fiduciary money. They show that this arrangement is an equilibrium as long as the history of these agents is commonly known. See also Martin and Schreft (2005).

[2]One may object that the indivisibilities in money and goods together with the restriction on money holdings that are present in the Kiyotaki-Wright environment are themselves an external limit on the behavior of the note issuer. These, however, are technological restrictions, not institutional ones. Besides, as we are going to see, they do not prevent overissue, unlike public monitoring.

We first consider the case where the patient bank can commit to a choice of money supply. This is done by assuming that its choice of money supply in the first period is binding. We refer to this situation as the full commitment case. We show that despite the bank's choice being its private information, there exists an equilibrium where the patient bank never overissues. 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 agents prefer the market 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 meetings in the market are informative enough, this drop in reputation is sufficiently fast to discourage the patient bank from deviating.

This result is not new. Klein (1974) also considers an environment where the trade-off between a short-run gain from overissue and a long term loss due to the negative impact on the bank's reputation is present. In his model, however, the relation between reputation and the amount of money in circulation is assumed rather then derived.

We then look at the no-commitment scenario, where the patient bank may change its behavior at any point in time. In this case a policy where it never overissues is not time-consistent. Indeed, if the patient bank does not overissue, its reputation increases over time. Eventually a point is reached where all agents in the market are so convinced that the bank is patient that any negative experience they face in a given period is attributed to bad luck. At this point 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 fact, we establish, through a similar argument, the stronger result that with no commitment the only monetary equilibria possible have the patient bank overissuing infinitely many times.

It is interesting to compare the above argument with the argument put forth by Friedman in defense of his point. He states that "So long as the fiduciary currency has a value greater than its cost of production - which under conditions can be compressed close to the cost of the paper in which it is printed - any individual issuer has an incentive to issue additional amounts" (1959, p. 7). As pointed out by Klein (1974), Friedman's argument fails to take into consideration the adverse effect in the revenue from money issue due to a decrease in the reputation of the note issuer. This is acknowledged by Friedman and Schwartz (1986). According to them, "Klein's theoretical case, resting on the necessity for a producer of money to establish confidence in his money is impeccable, and has received wide acceptance" (1986, p. 45). However, they also object that "It is not clear that his argument can be carried over to a 'pure fiduciary' currency" (1986, p. 45).

We formalize the above objection in the context of a monopolistic supply of money by showing that the reputational effect is not enough to prevent overissue when commitment is not possible. Since a patient bank can do such a good job in convincing the agents in the market that it is indeed patient, eventually its cost of overissuing, including the reputational cost, becomes smaller than the benefit. Putting it in a different way, overissue takes place because even though the patient bank has an incentive to build a good reputation, it has no incentives to maintain it once it is good enough.

The above reasoning makes it clear that an arrangement with private money and no external regulations is not viable unless the reputational cost of issuing additional units of money remains positive over time. Motivated by this, we extend our model in order to build a continuous threat for the patient bank's reputation. More precisely, we assume that in every period a fraction of the population is replaced by new agents that are uninformed about the bank's current behavior.[3] In this scenario, we show that there is an equilibrium where the policy of never overissuing is time consistent for the patient bank. The reason is that a patient bank always has an incentive to gain a good reputation with the new agents that enter the economy.

Our work is also related to the literature on the time-inconsistency of monetary policy.[4] There are two key distinctions, however. First, since this literature is focused on the optimality of monetary policy, and not on the feasibility of a monetary regime, money is introduced through a cash-in-advance type of constraint. In our case, the role of money is motivated by the decentralized nature of trade. Second, these models assume some form of public information on the past behavior of the note issuer while, for the reasons given above, we don't do so in our model.

The paper is structured as follows. In the next section we introduce the basic setup. The full-commitment case is considered in Section 3, while the no-commitment case and its extension are considered in Section 4. Section 5 concludes and the Appendix collects the proofs omitted from the main text.

2 Basic Setup

The model we consider is a modification of the one from Araujo and Camargo (AC hereafter) (2005). Time is discrete and indexed by $t \in \mathbb{N} \cup \{0\}$. The economy starts in $t = 0$ with one large infinitely lived agent that we call the bank. There are two possible discount factors δ for the bank,

[3] A similar setting has been considered by Mailath and Samuelson (1998), but in a completely different context.

[4] See, for example, Calvo (1978), Chari and Kehoe (1990), and Chang (1998).

either zero or $\delta_p > 0$. The value of δ is determined in $t = 0$ and is known to the bank only. We say the bank is *patient* when $\delta = \delta_p$ and *impatient* when $\delta = 0$. The probability of a patient bank is $\theta_0 \in (0, 1)$. The economy is also populated by a large number of infinitely lived small agents that we describe next. To make the exposition simpler, we refer to the small agents as *agents* only.

In $t = 1$ a continuum of mass one of agents enters the economy. Moreover, for any $t \geq 2$, each agent born in $t-1$ gives birth to another agent, so that the mass of agents increases by one in every period. For any $t \geq 1$ an agent born in that period is referred to as *newly born*. Beginning in $t = 2$, a newly born agent inherits his parent's private history. We show below that an agent's private history is what determines his belief about the type of bank he faces. Hence this last assumption implies that any agent born after $t = 1$ starts with the same belief about δ as his parent. All agents have the same discount factor $\beta \in (0, 1)$.

Besides being newly born or mature, an agent has a type that is determined in his first period of life. There are $K > 2$ of them, each one corresponding to one of the K possible types of goods that can be produced in the economy. At every date t, the distribution of types across the newly born population is uniform. In other words, for any type $k \in \{1, \ldots, K\}$, the fraction of newly born agents that is of this type is $\frac{1}{K}$. Agents of type k can only consume a type k good, their so-called preferred good.

Production works as follows. Each newly born agent receives a non-perishable endowment and makes a once and for all decision. He either goes to autarky or enters the market. When in autarky, an agent uses his endowment as an input to a production technology. There are n production possibilities in each period, and each good produced yields utility $a > 0$. In the market, an agent uses his endowment in the production of indivisible goods. An agent of type k can only produce, at zero cost, a good of type $k+1 (\bmod K)$, his so-called endowment good. Each agent can hold at most one unit of either goods or money at any point in time. In what follows we assume that when indifferent between the market and autarky, a newly born agent chooses the market.

The bank derives utility from the consumption of all K goods in the economy, 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, we assume that in every period after $t = 0$ the bank approaches a fraction m of the newly born agents who decided to enter the market in that period and exchanges one unit of fiat money for one unit of their corresponding endowment good. If μ is the measure of agents entering the market in a given period, the bank's flow payoff from choosing m is μm. The value of m is restricted to the set $\{m_L, m_H\}$, with $m_L < m_H$, but no agent in the economy observes the bank's choice. When the bank chooses m_H we say it *overissues*. We assume that when indifferent between m_L and m_H a bank always chooses m_H. Notice that this rules out equilibria where the bank uses a mixed strategy.

The market is organized as follows. There are K distinct sectors, each one specialized in the exchange of one of the K possible goods. Agents can identify sectors, but inside each one of them they are pairwise matched under an uniform random matching technology. 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. For example, 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 with it. As soon as an agent obtains one unit of his preferred good, he consumes it and obtains utility $u > a$. After that, he produces one more unit of his endowment good, that can be used for further trading. Any agent going to the market faces $n \geq 2$ rounds of meetings, where n is the same as above. We assume that agents don't discount within a trading period. [5]

Notice that despite the cost of production being zero for the agents, gift-giving is not an equilibrium in this environment. The reason for this is the market structure adopted. In fact, if gift-giving were to be an equilibrium, all agents would like to stay in the market where their preferred good is traded. This, however, rules out single-coincide of wants meetings, making gift-giving impossible.

An implicit assumption in the above description of the market is that there is always a positive measure of agents in it at any point in time. Since once in the market an agent cannot leave it, a sufficient condition for having a positive measure of agents in the market at all points in time is that a positive measure of agents enters it in period 1. In the next section we compute the market flow payoffs when it has a positive measure of agents in it. When the measure of agents in the market is zero (the market is "empty"), the market flow payoff is zero.

Throughout this paper we make the following assumption about m_L and m_H.

Assumption 1: $m_H(1 - m_H) < a/u < m_L(1 - m_L)$.

In the next section we show that this implies that if the bank always makes the same choice of m, the market is always worse than autarky if $m = m_H$ and always better when $m = m_L$.

[5] Notice that we took 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. As the next subsection makes clear, this game has an equilibrium where the agents always exchange their endowment for one unit of money if approached by the bank and their behavior in the market (sectors to visit and trading decisions) is as above. We omit the above details for the sake of brevity since they are of little relevance for the issues here considered.

3 The Full-Commitment Case

In this section we assume that the patient bank can commit to its choice of m in period 1; that is, once it chooses the value of m in $t = 1$, it cannot change it afterwards. Notice that an impatient bank always overissues, whether it has the ability to commit to a choice of m or not. Therefore, in the full commitment case, the amount of money introduced in the market by the bank is the same in all periods. In particular, the fraction of agents with (one unit of) money in the market is the same at all points in time.

We first describe the bank's payoffs and strategies. Since an impatient bank always overissues we consider the patient bank only. A strategy for the latter is just a choice of m in $t = 1$. The measure of agents entering the market on a given period depends on the aggregate behavior of the agents born in this period. Since the bank's decision is its private information, this measure also depends on the patient bank's presumed choice of m in $t = 1$. Indeed, its expected choice of m affects how agents use their market experience to make inferences about the type of bank they face. For example, if the agents believe that the patient bank behaves in the same way as the impatient one, the market is completely uninformative. Since the number n of market meetings per unit of time also affects this inference process, the above measure is a function of n as well. Let m_p denote the presumed choice of m by the patient bank. Suppose the bank is patient and denote by $\mu_t(m, n|m_p)$ the measure of newly born agents that enter the market in t as a function of n and its actual choice of m. The patient bank's lifetime utility as a function of m and n is then

$$U(m, m_p, n) = \sum_{t=1}^{\infty} \delta_p^{t-1} \mu_t(m, n|m_p) m.$$

We now consider the agents' payoffs and strategies. Suppose that the fraction of agents with money in the market is m and let $w_i(m)$ denote the expected lifetime utility to a newly born agent if he enters the market with $i \in \{0, 1\}$ units of money. This payoff is independent of t since the market environment is stationary in the full commitment case. It is possible to show, see AC (2005) for the details, that

$$\begin{aligned} w_0(m) &= (1-\beta)^{-1}[(n-1)m(1-m)u + \beta m(1-m)u], \\ w_1(m) &= (1-\beta)^{-1}[(n-1)m(1-m)u + (1-m)u - \beta(1-m)^2 u] \end{aligned}$$

Notice that $w_1(m) = w_0(m) + (1-m)u > w_0(m)$, and so an agent is always willing to exchange one unit of his endowment good for one unit of money. A newly born agent's lifetime expected payoff $w(m)$ from going to the market as a function of m is then equal to

$$w(m) = (1-m)w_0(m) + m w_1(m) = (1-\beta)^{-1} n m (1-m) u.$$

Since the flow payoff from staying in autarky is na, the flow payoff gain or loss from going to the market is equal to

$$(1-\beta)w(m) - na = nm(1-m)u - na.$$

From now on we look at the flow payoff gain or loss per unit of market meetings,

$$\left[(1-\beta)w(m) - na\right]/n = m(1-m)u - a.$$

We denote the above quantity by $v(m)$. This normalization is useful because it emphasizes that a newly born agent's decision between the market and autarky is independent of n, the number of per period market meetings. Assumption 1 at the end of the previous section then implies that: (i) If agents knew that money is always overissued, they would never enter the market no matter their date of birth; (ii) The opposite would take place if agents knew that money is never overissued.

We need to assume that if the patient bank chooses m_L the ex-ante expected payoff from entering the market is not smaller than the payoff from choosing autarky. Otherwise the only equilibrium possible with full commitment is the uninteresting non-monetary equilibrium where no agent ever enters the market and the bank always chooses m_H. Notice that this non-monetary equilibrium always exists, whether there is full commitment or not.

Assumption 2: $\theta_0 v(m_H) + (1-\theta_0)v(m_L) \geq 0$.

Now that we know the full information flow payoffs, we can describe the agents' strategies and payoffs in the incomplete information case, the case of interest. All agents are assumed to use Bayes rule to update their beliefs about the type of bank they face. The only piece of information available to newly born agents when they make their market/autarky decision is the private history they inherit from their parents. Moreover, the only thing that matters for this decision is the expected lifetime value of entering the market compared with the lifetime value of staying in autarky, zero. Consequently, given the presumed behavior of the patient bank, these private histories can be summarized by $\theta(h)$, the posterior belief that the bank faced is patient given a private history h.[6] Notice that if the patient bank chooses m_L, $\theta(h)$ is also the belief that overissue never occurs. In period 1, when private histories are empty, $\theta(h) = \theta_0$.

Let A stand for autarky and M for market. Given our tie-breaking assumption, a behavioral strategy for an agent born in $k \in \mathbb{N}$ is a function $\tau : [0,1] \to \{A, M\}$ where $\tau(\theta)$ is the agent's

[6] At this point it may seem more natural to summarize private histories by the posterior belief that overissue does not take place. This, however, is no longer possible when the patient bank is allowed to change its behavior at every point in time. On the other hand, as we show in the next section, it is always possible to summarize private histories by the posterior belief that the bank faced is patient, even when the choice of m can be changed in every period. Therefore, to keep the notation uniform, we adopt the latter identification.

choice when θ is his belief that the bank is patient. From now on we always use θ to denote such beliefs. The payoff to such a strategy, as a function of m_p, is

$$V(\tau,\theta,m_p)=\begin{cases}(1-\beta)^{-1}\{\theta I_{\{m_L\}}(m_p)v(m_L)+[1-\theta I_{\{m_L\}}(m_p)]v(m_H)\} & \text{if} \quad \tau(\theta)=M\\ 0 & \text{if} \quad \tau(\theta)=A\end{cases},$$

where $I_{\{m_L\}}$ is the indicator function of the set $\{m_L\}$. An equilibrium in this environment is then a strategy profile where:

(a) The patient bank is best responding to all agents;

(b) Each agent is best responding to the patient bank and the market/autarky decisions of all other agents.

As mentioned above, a non-monetary equilibrium always exists. In the remaining part of this section we show that under certain conditions an equilibrium where the patient bank never overissues also exists. As will be seen, these conditions have a very natural interpretation in terms of the incentives the patient bank has to build a reputation that it is indeed patient. By reputation we mean the distribution of beliefs θ across the agents that are in the market. For this let $\theta^*(a)$ be given by

$$\theta^*(a)=\frac{a-m_H(1-m_H)u}{m_L(1-m_L)u-m_H(1-m_H)u},$$

and define $\tau^*:[0,1]\to\{A,M\}$ to be such that

$$\tau^*(\theta)=\begin{cases}M & \text{if } \theta\geq\theta^*(a)\\ A & \text{otherwise}\end{cases}. \tag{1}$$

Consider now the strategy profile where $m=m_L$ and all newly born agents behave according to τ^*. Denote it by σ^*.

Theorem 1. *Suppose $\delta_p>(m_H-m_L)/m_H$. There exists $n_0\in\mathbb{N}$ such that σ^* is an equilibrium when $n\geq n_0$. Moreover, σ^* is the unique strategy profile that can be an equilibrium where a positive measure of agents enters the market in all periods.*

Notice that if there is a positive measure of agents in the market at all points in time, the tie breaking assumption about the agents' behavior implies that τ^* is the unique optimal decision rule for any newly born agent. What happens then if n is large? Loosely speaking, the market is, on average, very informative about the bank's behavior. Hence, after $t=1$, the newly born agents are, again on average, better informed when they make their market/autarky decisions. Consequently, if the patient bank deviates and chooses m_H, agents soon become very convinced that the bank they face is impatient. In other words, its reputation for being patient (and thus not overissuing) decreases quickly. Therefore, even if the patient bank's revenues from money issue are initially bigger than when it chooses m_L, these revenues decrease fast. On the other hand, if the patient bank chooses m_L, agents quickly learn that the bank they face is indeed the patient one. This guarantees the patient bank a steady stream of revenue from money issue.[7] Therefore, for high n, the patient bank prefers m_L despite the initially lower revenues; that is, if the market transmits information about m fast enough, the incentives for the patient bank to develop a reputation for being patient are sufficient to induce it to choose m_L.

Let us be precise about how agents learn the value of m from their market experience. In what follows, a family is the collection of all agents whose genealogy can be traced to a particular agent born in $t = 1$. Consider one such family and refer to its member born in $t = k$ as the generation k member. Since money and goods are indivisible, they can only be exchanged on a one-to-one basis. Moreover, all agents hold either one unit of money or one unit of their endowment good at any point in time. Therefore, for any $t \geq 1$, the only relevant piece of information for the generation t member is the record of money holdings of the partners of the generation k members in their respective first periods of life, where $k \in \{1, \ldots, t-1\}$. This includes their meetings with the government, where exchanging their endowment good for one unit of money is interpreted as the government having one unit of money.

Suppose now that the patient bank is expected to choose m_L. If a newly born agent goes to autarky, his belief θ that the bank is patient does not change, as he receives no new information about the value of m. If, instead, he enters the market, his updated belief is

$$B^n(c,\theta) = \frac{\theta m_L^c (1-m_L)^{(n+1)-c}}{\theta m_L^c (1-m_L)^{(n+1)-c} + (1-\theta) m_H^c (1-m_H)^{(n+1)-c}}, \tag{2}$$

where $c \in \{0, \ldots, n+1\}$ is the number of meetings with money he faces. This updated belief is transferred to his son. By the consistency of Bayes estimates for the binomial distribution, see De Groot (1970), we know that: (i) If the true choice of m is m_H, then with probability one $B^n(c,\theta) \to 0$ as $n \to \infty$ for all $\theta \in (0,1)$; (ii) If the true choice of m is m_L, then with probability one $B^n(c,\theta) \to 1$ as $n \to \infty$ for all $\theta \in (0,1)$. In particular, an agent's market experience in his first period of life is very informative about the true value of m when n is large.

Let $\Omega = \{A, M\} \times \{0, \ldots, n+1\}$. An element of Ω contains the relevant part of a newly born agent's private history in any given period. If he goes to autarky (A), he observes nothing. For practical purposes this is equivalent to always observing zero meetings with money regardless of the bank's type. If, instead, he enters the market, he observes $c \in \{0, \ldots, n+1\}$ meetings with

[7] It is possible to show, see AC (2005), that for fixed n the measure of agents entering the market when $m = m_H$ converges to zero, while the same measure is bounded away from zero when $m = m_L$.

(one unit) of money. The set of period t histories for a newly born agent is then $\Omega^t = \times_{\tau=1}^{t-1}\Omega$ if $t > 1$ and $\Omega^t = \emptyset$ if $t = 1$. Suppose then that the bank is patient (and is still expected to choose m_L). For each $h^t \in \Omega^t$ it is possible to construct the belief $\theta(h^t)$ an agent born in t with h^t has that the bank he faces is patient. Given this, we can construct, for each $t \in \mathbb{N}$, the random variable $\theta_t(m, m_L) : \Omega^\infty \to [0, 1]$, where $\Omega^\infty = \times_{t=1}^{\infty}\Omega$, describing the distribution of period t beliefs for an agent born in t as a function of the patient bank's actual choice of m. The details of how this is done can be found in Easley and Kiefer (1988). Since there is no aggregate uncertainty, the probability any agent born in t has of inheriting a private history $h^t \in \Omega^t$ is equal to the fraction of newly born agents with this particular history. Therefore $\mu_t(m, n|m_L) = \Pr\{\theta_t(m, m_L) \geq \theta^*(a)\}$ is the measure of agents that enter the market in t as a function of m and n.

Proof of Theorem 1: Suppose the patient bank chooses m_L, so that the ex-ante probability of $m = m_L$ is θ_0. We know that if a positive measure of agents enters the market in period 1, then τ^* is the unique best response for any agent born in this and in all subsequent periods. Since Assumption 2 implies that $\theta_0 > \theta^*(a)$, all agents born in $t = 1$ enter the market if they follow τ^*, no matter the bank's type. Hence, under σ^* all agents are best responding to the other agents and the patient bank.

Suppose now that all agents follow τ^*. We know from the last paragraph that $\mu_1(m_L, n|m_L) = \mu_1(m_H, n|m_L) = 1$ for all $n \in \mathbb{N}$. We also know, from the reasoning preceding this proof and the fact that $\theta^*(a) \in (0, 1)$, that $\mu_t(m_L, n|m_L) \to 1$ and $\mu_t(m_H, n|m_L) \to 0$ as $n \to \infty$ for all $t > 1$. Let $V_m : [0, 1]^\infty \to \mathbb{R}$ be such that

$$V_m(x_1, x_2, \ldots) = V_m(\{x_t\}) = \sum_{t=1}^{\infty} \delta_p^{t-1} x_t m,$$

where $m \in (0, 1)$ is a fixed constant. It is easy to see that V_m is, for each $m \in (0, 1)$, a continuous function when $[0, 1]^\infty$ is endowed with the product topology. Consequently, $U(m_H, m_L, n) = V_{m_H}(\{\mu_t(m_H, n|m_L)\}) \to m_H$ and $U(m_L, m_L, n) = V_{m_L}(\{\mu_t(m_L, n|m_L)\}) \to (1 - \delta_p)^{-1} m_L$ as $n \to \infty$. Since $\delta_p > (m_H - m_L)/m_H$ by assumption, there exists $n_0 \in \mathbb{N}$ such that $U(m_H, m_L, n) < U(m_L, m_L, n)$ if $n \geq n_0$. In other words, the best response for the patient bank is to indeed choose m_L.

The second statement follows from the first paragraph together with the fact that the only equilibrium possible where the patient bank chooses m_H is the non-monetary one. $\square$

4 The No-Commitment Case

By restricting the patient 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 doing so, we are able to focus on the trade-off between the bank's patience and the agents ability to monitor its choice. In this section we investigate what happens when the bank is allowed to change its decision of m at the beginning of every period, but this decision is still its private information. Our interest is in whether there is an equilibrium in this new environment where the patient bank always chooses m_L. In what follows we show that no such equilibrium exists. In fact, we establish the much stronger result that in any pure strategy equilibrium where money circulates, in the long-run all agents become indifferent between the market and autarky when the bank is patient. In particular, the patient bank must overissue infinitely many times in any such equilibrium. We also show that if we modify our environment in a certain natural way, a positive answer is possible.

4.1 Bad News

We begin with the negative result. When the patient bank is allowed to choose m in every period, its strategy set changes. A strategy for it is now a sequence $\{M_t\}$ of contingent plans, where M_t is a function from $\{m_L, m_H\}^{t-1}$ into $\mathcal{P}\{m_L, m_H\}$, the set of all probability measures on $\{m_L, m_H\}$. From now until the end of this subsection we restrict attention to pure strategies. An important feature of any such strategy is that it uniquely determines the sequence $\{m_t\}$ of choices of m by the patient bank.[8]

As in the full commitment case, given any pure strategy for the patient bank agents use their market experience to learn about the type of bank they face. The set of period t histories for a newly born agent is the same as in the previous section. However, for reasons that become clear below, we represent it in a different way. Let $\Delta = \{A, M\} \times \{0, 1\} \times \{0, \ldots, n\}$. We interpret an element of Δ in the following way: the first component is the market/autarky decision, the second is the number of meetings with the bank, and the third is the number of meetings with money in the first period of life. If an agent chooses autarky, we assume that the number of meetings with the bank and the number of meetings with money is zero regardless of the bank's type. The set of period t histories for an agent born in t can then be represented by Δ^t.

The updating of beliefs is now more complex. The reasons are two. First, the fraction of agents with money in the market when the bank is patient is no longer necessarily equal to m_L at all points in time. Second, for any $k > 1$, unless $m_t = m_L$ or $m_t = m_H$ for all $t \geq 1$, this fraction

[8] Allowing behavioral strategies for the patient bank introduces substantial complications, since in this case the sequence of choices of m is no longer deterministic.

depends on the measures of agents that entered the market from $t = 1$ to $t = k - 1$.

Let μ_t denote the measure of agents entering the market in t when the bank is patient, ν_t denote the same measure when the bank is impatient, and suppose $\{m_t\}$ is the sequence of choices of m by the patient bank. We know that the impatient bank always chooses m_H. Hence, in this case the fraction of agents in the market with money is always m_H as long as there is a positive measure of them in there. On the other hand, if the bank is patient, the period t fraction of agents in the market with money is

$$\eta_t = \frac{\sum_{\tau=1}^{t} \mu_\tau m_\tau}{\sum_{t=1}^{\tau} \mu_\tau}$$

as long as the denominator is positive. The case where the denominator is zero corresponds to the situation where at least until t the measure of agents in the market is zero. For convenience we set $\eta_t = 0$ in the latter case.

Let $\theta(h^t; \{\mu_t\}, \{\nu_t\}, \{m_t\})$ denote the posterior belief that the bank is patient given a private history $h^t \in \Omega^t$ and the sequences $\{\mu_t\}$, $\{\nu_t\}$, and $\{m_t\}$. When there is no risk of confusion, the dependence of this belief on these sequences is omitted. It can be computed in the following inductive way:

1. $\theta(h^1) = \theta_0$, where h^1 is the empty history;

2. Suppose $\theta(h^t)$ is defined for all $h^t \in \Omega^t$, where $t \geq 1$ is fixed. Let $h^{t+1} = (h^t, \omega)$, with $\omega = (\omega_1, \omega_2, \omega_3) \in \{A, M\} \times \{0, 1\} \times \{0, \ldots, n\}$, be an element of Δ^{t+1}. If $\omega_1 = A$ or $\sum_{\tau=1}^{t+1} \mu_\tau = \sum_{\tau=1}^{t+1} \nu_\tau = 0$, then $\theta(h^{t+1}) = \theta(h^t, \omega) = \theta(h^t)$. The second contingency corresponds to the case where the measure of agents in the market in $t + 1$ is zero no matter the bank's type. If $\omega_1 = M$, then: (i) $\theta(h^t, \omega) = 1$ if $\sum_{\tau=1}^{t+1} \mu_\tau > 0$ and $\sum_{\tau=1}^{t+1} \nu_\tau = 0$; (ii) $\theta(h^t, \omega) = 0$ if $\sum_{\tau=1}^{t+1} \mu_\tau = 0$ and $\sum_{\tau=1}^{t+1} \nu_\tau > 0$; (iii)

$$\theta(h^t, \omega) = \frac{\theta(h^t) m_{t+1}^{\omega_2} (1 - m_{t+1})^{1-\omega_2} \eta_t^{\omega_3} (1 - \eta_t)^{n-\omega_3}}{\theta(h^t) m_{t+1}^{\omega_2} (1 - m_{t+1})^{1-\omega_2} \eta_t^{\omega_3} (1 - \eta_t)^{n-\omega_3} + (1 - \theta(h^t)) m_H^{\omega_2+\omega_3} (1 - m_H)^{n+1-\omega_2-\omega_3}}$$

if $\sum_{\tau=1}^{t+1} \mu_\tau > 0$ and $\sum_{\tau=1}^{t+1} \nu_\tau > 0$. The first contingency corresponds to the case where the measure of agents in the market in t is positive if the bank is patient and zero when it is impatient. The second contingency corresponds to the exact opposite. The third contingency corresponds to the case where the measure of agents in the market in t is greater than zero regardless of the bank's type.

A strategy for a newly born agent in t is still a map $\tau : [0, 1] \to \{A, M\}$ such that $\tau(\theta)$ is this agent's market/autarky decision as a function of his belief θ that the bank is patient.[9] In order to define an equilibrium in the no-commitment case we need to take into account that: (1) The agents

[9] Recall that the tie-breaking assumption about the agents' behavior rules out mixed strategies.

need the sequences $\{\mu_t\}$ and $\{\nu_t\}$ to update their beliefs; (2) The sequences $\{\mu_t\}$ and $\{\nu_t\}$ depend on the agents' behavior together with their belief updating rule. Hence, we need to enlarge the equilibrium notion used in Section 3 to include the requirement that correct expectations about $\{\mu_t\}$ and $\{\nu_t\}$ are held. Accordingly, an equilibrium is a tripe $(\sigma, \{\mu_t\}, \{\nu_t\})$, where σ is a pure strategy profile for the patient bank and the agents, that satisfies the following properties:

(a) Agents update beliefs according to $\theta(h^t; \{\mu_t\}, \{\nu_t\}, \{m_t\})$;

(b) The patient bank optimally chooses the value of m in each period given $\{\mu_t\}$;

(c) Given the sequence $\{m_t\}$ of choices of m by the patient bank and the sequences $\{\mu_t\}$ and $\{\nu_t\}$, each agent's market/autarky decision in his first period of life is optimal;

(d) The agents and the patient bank hold correct expectations. In other words, let $\widehat{\mu}_t$ be the measure of agents that actually enter the market in t when σ is followed, the agents update beliefs according to (a), and the bank is patient. Define $\widehat{\nu}_t$ to be the corresponding fraction when the bank is impatient. Then $\widehat{\mu}_t = \mu_t$ and $\widehat{\nu}_t = \nu_t$ for all $t \in \mathbb{N}$.

Lemma 1. *Suppose $(\sigma, \{\mu_t\}, \{\nu_t\})$ is an equilibrium. Then $\mu_t = 0$ if, and only if, $\nu_t = 0$.*

Proof: Let $\underline{t} = \inf\{t \in \mathbb{N} : \mu_t > 0 \text{ or } \nu_t > 0\}$. 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. In fact, suppose, by contradiction, that $\mu_{\underline{t}} > 0$ but $\nu_{\underline{t}} = 0$. The other possibility is dealt with in exactly the same way. Since $\mu_t = \nu_t = 0$ for $t < \underline{t}$, all agents born in $\underline{t}$ have the same belief about the bank's type, θ_0, whether the bank is patient or not. Hence they should all make the same decision in $\underline{t}$. This, however, implies that $\mu_{\underline{t}} = \nu_{\underline{t}}$, a contradiction. Now observe that from $\underline{t}$ on the measure of agents in the market is always positive. Therefore, any event in Ω^t, with $t > \underline{t}$, happens with positive probability when the bank is patient if, and only if, it happens with positive probability when the bank is impatient.[10] In particular, the set of all private histories in $t > \underline{t}$ that lead an agent born in 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. □

Let $\mathbb{N}_0 = \{t \in \mathbb{N} : \mu_t > 0 \text{ and } \nu_t > 0\}$ be the set of points in time where a positive measure of agents enter the market regardless of the bank's type. From the above lemma we know that $(\mathbb{N}_0)^c = \{t \in \mathbb{N} : \mu_t = \nu_t = 0\}$. From Section 3 we know that if in a given period the fraction of agents in the market with money is m, then the market flow payoff is $v(m) = m(1-m)u - a$. Hence, if the bank is patient, the lifetime expected payoff to an agent born in t from entering the market is $\sum_{\tau=1}^{\infty} \beta^{\tau-1} v(\eta_{t+\tau-1})$. Recall that we set $\eta_t = 0$ when the measure of agents in the market

[10] This relies on the fact that both m_L and m_H belong to $(0, 1)$.

is zero. In this case $v(\eta_t) = -a$, which indeed is the flow payoff loss from entering the market when it is empty.

Theorem 2. *Suppose* $(\sigma, \{\mu_t\}, \{\nu_t\})$ *is an equilibrium. Then either* $\mathbb{N}_0 = \emptyset$ *or* $\mathbb{N}_0$ *is infinite, but*

$$\lim_{t\to\infty} \sum_{\tau=1}^{\infty} \beta^{\tau-1} v(\eta_{t+\tau-1}) = 0.$$

The first alternative corresponds to the non-monetary equilibrium where the patient bank always overissues and no agents ever enter the market. Therefore, in any equilibrium where money circulates, the lifetime value of entering the market in t converges to zero as t increases when the bank is patient. Below we show that this implies, in particular, that the patient bank must overissue infinitely many times in any such equilibrium.

The intuition for the result that there is no equilibrium where the patient bank never overissues is simple. From the previous section we know that if the patient bank never overissues, then the measure of agents entering the market is positive for all $t \geq 1$. Moreover, the posterior belief of any agent in the market eventually becomes so close to one that even if in one period he faces the most negative experience possible, he remains very convinced that the bank is patient. When this happens, the patient bank looses its incentives to choose m_L. Indeed, if it deviates and chooses m_H, the impact on the distribution of beliefs across the agents in the market is negligible. Hence, its loss of future revenue from money issue, due to a reduction in the measure of newly born agents entering the market in subsequent periods, is negligible as well because of discounting. The present gain from overissue is, however, bounded away from zero. In other words, once its reputation for being patient is high enough, the patient bank's cost of overissue, reflected in lower future revenues from money issue due to a worse reputation, gets smaller then the immediate benefit. A similar, if more involved argument, works in the more general case that we now establish.

Proof of Theorem 2: We first show that $\mathbb{N}_0$ is either empty of infinite. For this suppose, by contradiction, that $\mathbb{N}_0$ is non-empty but finite and let $\bar{t}$ be the last period where μ_t is positive. Consider now an agent born in $\bar{t}$ that enters the market and denote his belief by θ^*. By Lemma 1 there is a positive measure of agents in the market in $\bar{t}$ whether the bank is patient or not. Moreover, $\eta_{\bar{t}} \geq m_H$. If $\eta_{\bar{t}} = m_H$, the market is uninformative in $\bar{t}$, and so the above agent's belief at $\bar{t}+1$ is θ^* independently of his market experience. On the other hand, the market is informative in $\bar{t}$ when $\eta_{\bar{t}} > m_H$. In particular, there is a positive probability that the above agent's belief in $\bar{t}+1$ is greater than θ^*. Consequently the probability that his son is born in $\bar{t}+1$ with a belief not smaller than θ^* is positive in both cases. Now observe that since $\mu_t = 0$ for all $t > \bar{t}$, $\eta_t = \eta_{\bar{t}}$ for all $t > \bar{t}$. Therefore, the expected payoff to any agent born in $\bar{t}+1$ with a belief at least equal to θ^* is no less than

$$\theta^* \sum_{t=1}^{\infty} \beta^{t-1} v(\eta_{t+\bar{t}}) + (1-\theta^*)(1-\beta)^{-1} v(m_H) = (1-\beta)^{-1}[\theta^* v(\eta_{\bar{t}}) + (1-\theta^*) v(m_H)].$$

The last quantity, however, is the lifetime expected payoff from entering the market to an agent born in $\bar{t}$ with belief θ^*, and so is non-negative by assumption. Consequently a positive measure of agents enters the market in $\bar{t}+1$, a contradiction.

From now on we assume that $\mathbb{N}_0$ is infinite. Notice first that if $\alpha_t = \sum_{\tau=1}^{\infty} \beta^{\tau-1} v(\eta_{t+\tau-1}) \leq 0$, then

$$\theta \sum_{\tau=1}^{\infty} \beta^{\tau-1} v(\eta_{t+\tau-1}) + (1-\theta)(1-\beta)^{-1} v(m_H) < 0$$

for all $\theta \in (0,1)$, and so no agent born in t wants to enter the market.[11] Therefore $\alpha_t > 0$ for all $t \in \mathbb{N}_0$. Let $\hat{k}(t) = \max\{t' \in \mathbb{N} : t' \leq t \text{ and } \mu_{t'} > 0\}$ be the last period before t where a positive measure of agents entered the market. We know that this definition is correct by Lemma 1. Now observe that $\eta_t = \eta_{\hat{k}(t)}$ for all $t \in \mathbb{N}$, and so the same holds for α_t. We can then conclude that there exists $\underline{t} \in \mathbb{N}$ such that $\alpha_t > 0$ for all $t \geq \underline{t}$.

We need to show that $\{\alpha_t\}$ converges to zero. Suppose not. Since $\{\alpha_t\}$ is bounded, the previous paragraph shows that it has a subsequence converging to some $\alpha > 0$. As will become clear in what follows we can assume, without loss, that $\{\alpha_t\}$ itself converges to α. Now let $\overline{\theta}$ be such that $\overline{\theta}\alpha/2 + (1-\overline{\theta})v(m_H) = 0$. Notice that $\overline{\theta} \in (0,1)$. By assumption there exists $t_0 \in \mathbb{N}$ such that

$$\overline{\theta} \sum_{\tau=t}^{\infty} \beta^{\tau-1} v(\eta_\tau) + (1-\overline{\theta})(1-\beta)^{-1} v(m_H) > (1-\beta)^{-1}[\overline{\theta}\alpha/2 + (1-\overline{\theta}) v(m_H)] = 0$$

for $t \geq t_0$. Therefore, any agent born in $t \geq t_0$ with $\theta \geq \overline{\theta}$ always enters the market.

Let t_i denote the i^{th} element of $\mathbb{N}_0$ in increasing order and consider the sequence $\{\lambda_k\}$ such that $\lambda_k = \eta_{t_k}$. We claim that $\{\lambda_k\}$ cannot converge to m_H. To see why observe that $\eta_t = \lambda_{\hat{k}(t)}$ for all $t \in \mathbb{N}$, and so $\{\eta_t\}$ converges to m_H if $\{\lambda_k\}$ does. Moreover,

$$|\alpha_t - (1-\beta)^{-1} v(m_H)| = \left| \sum_{\tau=1}^{\infty} \beta^{\tau-1} v(\eta_{t+\tau-1}) - (1-\beta)^{-1} v(m_H) \right| \leq \sum_{\tau=1}^{\infty} \beta^{\tau-1} |v(\eta_{t+\tau-1}) - v(m_H)|,$$

and so $\{\alpha_t\}$ converges to $(1-\beta)^{-1} v(m_H) < 0$ if $\{\lambda_k\}$ converges to m_H, a contradiction. Consequently $\{\lambda_k\}$ has a subsequence $\{\lambda_{k_l}\}$ that converges to some $m \in (m_H, 1]$.

Given $\{\mu_t\}$ and $\{\nu_t\}$ we can construct, for each $t \in \mathbb{N}$, a random variable $\theta_t(\delta) : \Omega^\infty \to [0,1]$ describing the distribution of period t beliefs across the population as a function of the bank's type.

[11] Lemma 1 implies that in all periods agents are born with a belief in $(0,1)$.

Once again see Easley and Kiefer (1988) for the details. By the definition of equilibrium we have that $\mu_t \geq \underline{\mu}_t = \Pr\{\theta_t(\delta_p) \geq \overline{\theta}\}$ and $\nu_t \geq \underline{\nu}_t = \Pr\{\theta_t(0) \geq \overline{\theta}\}$ for $t \geq t_0$. An argument similar to the one employed in Banks and Sundaram (1992) in the proof of their Theorem 5.1 allows us to show that at least one of the two sequences $\{\underline{\mu}_t\}$ and $\{\underline{\nu}_t\}$ is bounded away from zero.

Let $E \subset \Omega^\infty$ be the event where the market is chosen for all $t \in \mathbb{N}_0$ such that $t = t_{k_l}$ for some k_l in the index set of the subsequence $\{\lambda_{k_l}\}$. This event has positive probability since a positive measure of agents enters the market in all $t \in \mathbb{N}_0$. Conditional on E, $\theta_t(0)$ converges to 0 with probability one by Kolmogorov's Strong Law of Large Numbers. See Camargo (2005) for a proof of this. Now let $F \subset \Omega^\infty$ be the event where the market is chosen for all $t \in \mathbb{N}_0$. Then, conditional on F, $\theta_t(0)$ also converges to 0 almost surely. See Lemma 1 in Aoyagi (1998). Hence it must be that $\underline{\nu}_t \to 0$, and so we can conclude that there exists $\underline{\mu}$ such that $\underline{\mu}_t \geq \underline{\mu}$ for t sufficiently large.

A consequence of the last paragraph is that μ_t is positive and bounded away from zero when t is sufficiently large. We show that this implies that $m_t = m_L$ for infinitely many $t \in \mathbb{N}_0$. Suppose not and let $\mathbb{N}_1 = \{t \in \mathbb{N}_0 : m_t = m_H\}$, so that $(\mathbb{N}_1)^c$ is finite. Since $\sum_{\tau=1}^{t_k} \mu_\tau \to \infty$ as $k \to \infty$, the term (A) below converges to zero (and so (B) converges to one). Therefore

$$\lambda_k = \underbrace{\frac{\sum_{\tau \in \mathbb{N}_0 \cap (\mathbb{N}_1)^c \cap \{1,\ldots,t_k\}} \mu_\tau}{\sum_{\tau \in \mathbb{N}_0 \cap \{1,\ldots,t_k\}} \mu_\tau}}_{(A)} m_L + \underbrace{\frac{\sum_{\tau \in \mathbb{N}_0 \cap \mathbb{N}_1 \cap \{1,\ldots,t_k\}} \mu_\tau}{\sum_{\tau \in \mathbb{N}_0 \cap \{1,\ldots,t_k\}} \mu_\tau}}_{(B)} m_H \to m_H,$$

a contradiction.

We can now finish the proof. For this let $\epsilon < \underline{\mu}(m_H - m_L)$. We know that there exists $t_1 \in \mathbb{N}$ such that $\sum_{t=t_1}^{\infty} \delta_p^{t-1} < \epsilon/2$. Given $\{\mu_t\}$ and $\{\nu_t\}$ we can compute, for each $t \in \mathbb{N}$, the cutoff belief θ_t that makes an agent born in t indifferent between the market and autarky.[12] We know that $\theta_t \leq \overline{\theta} < 1$ for $t \geq t_0$. By the same argument as above, it is possible to show that conditional on the event F, $\theta_t(\delta_p)$ converges to one almost surely, and so in measure. Therefore there exists $t_2 \geq t_0$ with the following property: Regardless of how the patient bank behaves in t, when $t \geq t_2$ the probability that $\theta_{t'}(\delta_p) \geq \theta_{t'}$ is greater than $1 - \epsilon(1-\delta_p)/2$ for all $t' \in \{t, t+1, \ldots, t+t_1-1\}$. Since $m_t = m_L$ for infinitely many $t \in \mathbb{N}_0$, there exists $t' \geq t_2$ such that $m_{t'} = m_L$. In t' the lifetime payoff to the patient bank from sticking to its prescribed strategy is

$$\sum_{t=2}^{\infty} \delta_p^{t-1} \mu_{t'+t-1} m_{t'+t-1} + \mu_{t'} m_L.$$

[12] It may be that for some values of t an agent born in this period never wants to enter the market regardless of his belief. In this case we set $\theta_t = 1$.

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

$$\mu_{t'} m_H + \Big(1 - \frac{\epsilon}{2}(1-\delta_p)\Big) \sum_{t=2}^{t_1 - 1} \delta_p^{t-1} \mu_{t'+t-1} m_{t'+t-1}.$$

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

$$\begin{aligned} \mu_t(m_H - m_L) - \frac{\epsilon}{2}(1-\delta_p) \sum_{t=2}^{t_1-1} \delta_p^{t-1} \mu_{t'+t-1} m_{t'+t-1} - \sum_{t=t_1}^{\infty} \delta_p^{t-1} \mu_{t'+t-1} m_{t'+t-1} & \\ \geq \ \underline{\mu}(m_H - m_L) - \frac{\epsilon}{2}(1-\delta_p) \sum_{t=2}^{t_1-1} \delta_p^{t-1} - \sum_{t=t_1}^{\infty} \delta_p^{t-1} & \\ > \ \underline{\mu}(m_H - m_L) - \frac{\epsilon}{2} - \frac{\epsilon}{2}, & \end{aligned}$$

since $\mu_t m_t < 1$ for all $t \in \mathbb{N}$. Therefore the above one-shot deviation is profitable, and so we cannot have $\alpha > 0$. We thus have the desired result. □

The above theorem is a characterization result. The question of whether there is an equilibrium with $\mathbb{N}_0$ infinite, although interesting, is beyond the scope of this paper.

Corollary 1. *Suppose* $\mathbb{N}_0$ *is infinite. Then* $\mathbb{N}_1 = \{t \in \mathbb{N}_0 : m_t = m_H\}$ *is infinite. Therefore in any equilibrium where money circulates the patient bank overissues infinitely many times.*

Proof: Suppose not. Following the convention introduced in the proof of the previous theorem, there exists $\overline{k} \in \mathbb{N}$ such that if $k \geq \overline{k}$, then $m_{t_k} = m_L$; that is, after the first $\overline{k}$ periods in $\mathbb{N}_0$, the patient bank always chooses m_L. Hence $\{\lambda_k\}$ is strictly increasing after $\overline{k}$, and so must converge to some $m > m_H$. Consequently $\{\eta_t\}$ also converges to m. This, however, implies that α_t is positive and bounded away from zero for t sufficiently large, a contradiction. □

4.2 Good News

The previous subsection suggests that 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. Put in another way, 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. In this way, no matter 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 is still the one introduced in the previous subsection.

In what follows we show that there are conditions under which an equilibrium where the patient bank never overissues exists in this modified environment. For this let $\mathrm{M}^* = \{M_t^*\}$ be the strategy for the patient bank where it chooses always m_L, no matter how it behaved previously. Moreover, let α and $\underline{\delta}$ be such that

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

Theorem 3. *Let σ^{**} be the strategy profile where the patient bank follows M^* and all agents follow τ^* given by (1). Suppose that $\alpha < 1/2$, $\lambda \in (\alpha, 1/2)$, and $\delta_p > \underline{\delta}$. There exists $n(\alpha) \in \mathbb{N}$ and two sequences $\{\mu_t^{**}(n)\}$ and $\{\nu_t^{**}(n)\}$ depending on n such that $(\sigma^{**}, \{\mu_t^{**}(n)\}, \{\nu_t^{**}(n)\})$ 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 if the patient bank follows M^*, agents update their beliefs in the same way as they do in the full commitment case when the patient bank chooses m_L. In particular, the belief updating is independent of the sequences $\{\mu_t\}$ and $\{\nu_t\}$.

Proof of Theorem 3: For each $n \in \mathbb{N}$, let $\mu_k^{**}(n)$ ($\nu_k^{**}(n)$) be the measure of agents entering the market in k when $m_L(m_H)$ is always chosen and agents follow τ^* and update beliefs according to (2). Since $\mu_1^{**}(n) = \nu_1^{**}(n)$ for all $n \in \mathbb{N}$, there is a positive measure of agents in the market at all points in time. Hence $\theta(h^t; \{\mu_t^{**}(n)\}, \{\nu_t^{**}(n)\})$ coincides with the belief updating given by (2). Moreover, expectations are trivially satisfied. That each agent is best responding when the patient bank follows M^* and all other agents follow τ^* has been established in the proof of Theorem 1. Therefore we just need to show that no one-shot deviation ever pays for the patient bank.

First observe that given any on- or off-the-equilibrium-path history of length $k-1$ ($k \geq 1$) for the patient bank, we can compute the distribution of beliefs across the agents when they all play according to τ^* and update their beliefs as described above. Group the newly born agents by the beliefs they have. For all $k \in \mathbb{N}$ there are finitely many of these groups and all have a positive mass. We can then determine within each one of these groups whether a one-shot deviation pays for the patient bank or not. A sufficient condition is that it does not pay in all of them.

Consider an agent born in k that enters the market with belief θ. Denote by $\xi_{t,k}(m,\theta)$ the probability that the generation $k+t$ member of this agent's family, with $t \geq 1$, enters the market in period $k+t$ in case the patient bank chooses $m \in \{m_L, m_H\}$ in k and m_L from $k+1$ on. From the definition of τ^* and the fact that $B^n(c,\theta)$ given by (2) is increasing in θ, we have that $\xi_{t,k}$ is weakly increasing in θ. Hence, a one-shot deviation for the patient bank is most profitable in the most optimistic group of agents, as it entails the smallest loss of future revenue from a decreased reputation. In particular, if we show that a one-shot deviation does not pay in an arbitrary k even in the case where a fraction $1-\lambda$ of the newly born agents have $\theta = 1$ and the remaining ones have $\theta = \theta_0$, we are done. In what follows we consider this case.

First notice that $\xi_{t,k}$ is independent of k, since all newly born agents follow the same strategy and update beliefs in the same way. Because of this we omit k from $\xi_{t,k}$ in the rest of this proof. Notice also that in any $k \geq 1$ a one-shot deviation by the patient bank has no impact on the behavior of future generations in the family of an agent born in k with $\theta = 1$. Hence, the patient bank's net gain with the $\theta = 1$ group from a one-shot deviation is $(1-\lambda)(m_H - m_L)$. Now let $H_t(m)$, with $m \in \{m_L, m_H\}$ and $t \geq 1$, be such that

$$H_t(m) = (1-\lambda)^t \xi_t(m, \theta_0) + \sum_{\tau=1}^{t-1} \lambda(1-\lambda)^\tau \xi_\tau(m_L, \theta_0) + \lambda.$$

If an agent is born in k with a belief θ_0 that the bank is patient, then $H_t(m)$ is the probability that the generation $k+t$ member of his family enters the market when the patient bank chooses m in k and m_L from $k+1$ on. 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 $k+t$ that private histories are not passed from one generation to the following is $k+(t-1)-\tau$, where $\tau \in \{0, \ldots, t-1\}$. The net loss from a one-shot deviation by the patient bank is then equal to

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

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

$$\Delta \geq (1-\lambda)(m_L - m_H) + \lambda\left[(m_L - m_H) + \frac{\delta_p(1-\lambda)}{1-\delta_p(1-\lambda)} m_L \Delta\xi_1\right] = m_L \underbrace{\left[\frac{\delta_p \lambda(1-\lambda)}{1-\delta_p(1-\lambda)} \Delta\xi_1 - \alpha\right]}_{\Psi(\delta_p, \lambda, \alpha)},$$

where $\Delta\xi_1 = \xi_1(m_L, \theta_0) - \xi_1(m_H, \theta_0)$. To finish notice that Ψ is increasing in δ_p and that $\Psi(\underline{\delta}, \lambda, \alpha) = \alpha[\Delta\xi_1 - 1]$. Because $\lim_{n\to\infty} \Delta\xi_1 = 1$, we can then conclude there exists $n(\alpha)$ such that $\Delta > 0$ when $n \geq n(\alpha)$ and $\delta_p > \underline{\delta}$. □

5 Conclusion

This paper considers whether a stable monetary regime is possible in an economy with a monopoly in the supply of fiat money when there are no external limits on its supply. The environment we consider has two key characteristics. First, money is essential as a medium of exchange. Second, information is decentralized and its flow is constrained by the same technology that hinders trade. In our view, any public information about the bank's choice of money supply can only be justified by the existence of some institution put in place to monitor it and thus limit its behavior. Therefore, the second characteristic captures the absence of any form of control, direct or indirect, on the bank's behavior. An interesting question, not pursued in this article, is whether the existence of a monitoring institution can be made endogenous in this type of framework.

The main feature of our model is the existence of a tradeoff between the immediate gain from overissue and the long-run loss in revenue from money issue due to a decrease in the bank' reputation. We show that if commitment to a choice of money supply is possible, then a monetary equilibrium where the patient bank does not overissue exists under certain conditions. The market must transmit information about the bank's behavior fast enough. This equilibrium, however, is not time-consistent when no such commitment is possible for the bank. In fact, the only equilibria where money circulates have the patient bank overissuing infinitely often. The reason for this negative result is that once a bank's reputation is very high, it becomes insensitive to the choice of money supply. Hence, even though a patient bank has an incentive to build a good reputation by not overissuing, the incentive to maintain it disappears once its reputation is high enough.

Following the above insight, we show that if information transmission is not perfect, the patient bank's incentive to maintain a good reputation never disappears, and so a stable monetary equilibrium is possible. The shortcoming of this approach to solve the time-consistency problem is that even though reasonable, it is somewhat ad-hoc. An interesting alternative, and the object of current research, is to introduce oligopolistic competition among note issuers.

References

[1] M. Aoyagi, "Mutual Observability and the Convergence of Actions in a Multi-Armed Bandit Model", *Journal of Economic Theory*, **82** (1998), 405-424.

[2] L. Araujo and B. Camargo, "Information, Learning, and the Stability of Fiat Money", Working Paper, 2005.

[3] J. Banks and R. Sundarm, “Denumerable-Armed Bandits”, *Econometrica*, **60** (1992), 1071-1096.

[4] A. Berentsen, “Time-Consistent Private Supply of Outside Paper Money”, Forthcoming *European Economic Review*, 2005.

[5] G. Calvo, “On the Time Consistency of Optimal Policy in a Monetary Economy”, *Econometrica*, **46** (1978), 1411-1428.

[6] B. Camargo, “Learning in Society - Characterization”, Working Paper, 2005.

[7] 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.

[8] R. Cavalcanti and N. Wallace, “A Model of Private Banknote Issue”, *Review of Economic Dynamics* , **2** (1999a), 104-136.

[9] R. Cavalcanti and N. Wallace, “Inside and Outside Money as Alternative Media of Exchange”, *Journal of Money, Credit and Banking*, **31** (1999b), 443-457.

[10] R. Chang, “Credible Monetary Policy in an Infinite Horizon Model: Recursive Approaches”, *Journal of Economic Theory*, **81** (1998), 431-461.

[11] V.V. Chari and P. Kehoe, “Sustainable Plans”, *Journal of Political Economy*, **98** (1990), 783-202.

[12] D. Easley and N. Kiefer, “Controlling a Stochastic Process with Unkown Parameters”, *Econometrica*, **56** (1988), 1045-1064.

[13] M. Friedman, *A Program for Monetary Stability*, Fordham University Press, 1959.

[14] M. Friedman and A. Schwartz, “Has Government Any Role in Money?”, *Journal of Monetary Economics*, **17** (1986), 37-62.

[15] M. DeGroot, *Optimal Statistical Decisions*, McGraw-Hill, New York, 1970.

[16] B. Katzman, J. Kennan, and N. Wallace, “Output and Price Level Effects of Monetary Uncertainty in a Matching Model”, *Journal of Economic Theory* **108** (2003), 217-255.

[17] N. Kiyotaki and R. Wright, “A Search-Theoretic Approach to Monetary Economics”, *American Economic Review* **83** (1993), 63-77.

[18] B. Klein, "The Competitive Supply of Money", *Journal of Money, Credit and Banking*, **6** (1974), 423-453.

[19] 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.

[20] A. Martin and S. Schreft, "Currency Competition: A Partial Vindication of Hayek", Forthcoming *Journal of Monetary Economics*, 2005.

Appendix

Lemma 2. $\xi_t(m_L, \theta_0) - \xi_t(m_H, \theta_0) \le \xi_{t+1}(m_L, \theta_0) - \xi_{t+1}(m_H, \theta_0)$ *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, $c(h^t) \le \alpha(t-k)(n+1) + \gamma$, 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

$$\sum_{(c_1,\ldots,c_t)\in C_t} \underbrace{\begin{pmatrix} n+1 \\ c_1 \end{pmatrix} \cdots \begin{pmatrix} n+1 \\ c_n \end{pmatrix}}_{B_t(c_1,\ldots,c_t)} m^{c_1}(1-m)^{n+1-c} m_L^{c_2+\cdots+c_t}(1-m_L)^{(t-1)(n+1)-c_2-\cdots-c_t},$$

where $C_t = \{(c_1, \ldots, c_t) : c_\tau \le \lfloor \alpha\tau(n+1) + \gamma \rfloor - c_1 - \cdots - c_{\tau-1},\ \tau = 1, \ldots, t-1\}$. 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

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

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

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

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