

Escola de
ECONOMIA
de São Paulo

**Textos para
Discussão**

222

Outubro
de 2009



**MONETARY EQUILIBRIUM WITH DECENTRALIZED
TRADE AND LEARNING**

**LUIS ARAUJO
BRAZ CAMARGO**



Os artigos dos *Textos para Discussão da Escola de Economia de São Paulo da Fundação Getúlio Vargas* são de inteira responsabilidade dos autores e não refletem necessariamente a opinião da FGV-EESP. É permitida a reprodução total ou parcial dos artigos, desde que creditada a fonte.

Escola de Economia de São Paulo da Fundação Getúlio Vargas FGV-EESP
www.fgvsp.br/economia

Monetary Equilibrium with Decentralized Trade and Learning

Luis Araujo - Michigan State University*
Braz Camargo - University of Western Ontario

March 22, 2005

Abstract

This paper analyzes the stability of monetary regimes in an economy where fiat money is endogenously created by the government, information about its value is imperfect, and learning is decentralized. We show that monetary stability depends crucially on the speed of information transmission in the economy. Our model generates a dynamic on the acceptability of fiat money that resembles historical accounts of the rise and eventual collapse of overissued paper money. It also provides an explanation of the fact that, despite its obvious advantages, the widespread use of fiat money is only a recent development.

JEL Classification: D82, D83, E00

Key Words: Monetary Stability, Endogenous Money, Decentralized Trade, Learning.

“All of the governments in China between 1100 and 1500 succumbed to this temptation, and their monetary histories have a strong family resemblance. In each there was a period of inflation, usually quite a long one. Except in the case of the Southern Sung dynasty, which was conquered by the Mongols before the evolution was completed, the use of paper money was, in each case, eventually abandoned. This abandonment of the use of paper money in China is the most interesting feature of the history of paper money in China.” (Gordon Tullock, 1957)

*We are specially grateful to Randall Wright and George Mailath. We also benefited from conversations with Jan Eeckhout, Brett Norwood, Dan Silverman and Brandon Weber. Seminar participants at Cornell, Georgetown, Michigan State, Penn, Purdue, Texas-Austin, the 2002 NSF Decentralization Conference, the 2002 SED Meeting and the Federal Reserve Bank of Cleveland provided many comments and suggestions. Corresponding author: Luis Araujo, 101 Marshall-Adams Hall, Department of Economics, Michigan State University, East Lansing, MI 48824-1038. E-mail: araujolu@msu.edu.

1 Introduction

In recent years, a large body of work has been devoted to the study of economies where the value of money arises endogenously. Most of the literature, following Kyotaki and Wright (KW hereafter) (1989, 1993), takes as a starting point the assumption that the amount of money in circulation, be it known or not, fixed or changing over time, is exogenously given. Less attention has been given to the questions that arise when we allow the quantity of money to be endogenous; i.e., when we make some of the agents inside the economy responsible for money creation. Of particular importance is the emergence and stability of fiat money regimes. One of the few contributions in this front is Ritter (1995), where the transition from a barter to a fiat money economy is analyzed.¹ He considers an economy where a coalition of agents, that he identifies as the government, is allowed to issue money, and shows that in order for this transition to take place, the size and the patience of this coalition must be large. Patience is important because the government must care about the future if money is to have value. Size plays a role as it allows the government to internalize the costs of overissue.

While providing a framework where the emergence of fiat money occurs endogenously, Ritter does not address the concomitant issue of its stability, i.e., whether money remains in circulation in the long-run or not. In his framework, once fiat money is introduced, it always stays in the economy. There is, however, varied evidence, as we can see in the above excerpt from Tullock (1957), that paper money issued by governments was, in the past, subject to much instability. Yang (1952), in his study of money in ancient China also describes a succession of failures in the transition to paper money due to instabilities related to overissue. A similar pattern of overissue and consequent abandonment of paper money appears throughout the U.S. history, as can be seen in Galbraith's (1975) account of the monetary experience in the Massachusetts Bay Colony in 1690 and in various American colonies during the mid-18th century. The rise and the collapse of paper money seems to be such a common phenomenon throughout history that, according to Friedman and Schwartz (1986), a continuous and widespread use of fiat money is only a twentieth century development.

In what follows, we build a model where monetary stability depends on both exogenous and endogenous factors. We take a simplified version of KW (1989) as our starting point. The main difference from KW is that the amount m of money in circulation is determined by a self-interested agent, the government, and is not known in advance by the other agents in the economy. We refer to the value of m as the monetary regime. These agents can react against the government by not accepting money if they think its value is low. Agents can obtain information about m from the trade meetings in which they participate; that is, from their private histories. Hence, the

¹See also Sik Kim (2001) and Norwood (2003).

same technology that governs trade and makes money essential in our environment, random and anonymous pairwise meetings, is also used to describe the transmission of information. Related papers in the literature always assume that there is a time when the quantity of money in circulation is revealed to all agents.²

The exogenous factors affecting the stability of money are the government's patience and the speed of information transmission. The endogenous factor comes from the ability the agents have to over time learn the nature of the monetary regime. In an environment where learning happens slowly, an impatient government can exploit the agents' misinformation and overissue, while maintaining the value of money in the short-run. Agents eventually realize the government's actual behavior and monetary trade breaks down. However, it takes time until a complete breakdown of trade happens, a result that matches with Tullock's observation on the history of paper money in China. In this case, monetary stability is only feasible with a patient government. On the contrary, when agents accumulate a lot of information in a short period of time, even impatient governments prefer not to overissue in order to avoid the breakdown of monetary trade.

Our model then provides an informational rationale for the late emergence of fiat money. Societies' ability to gather information and learn about the state of the economy increased over time. In modern economies, the dissemination of information is much faster than in the past, and so a government's temptation to overissue should be less pronounced. Therefore, the late widespread implementation of fiat money is not necessarily a result of exogenous factors like an increase in the government size and patience. It can, instead, be the result of an increase in the society's ability to monitor the behavior of the money issuer.

The paper is structured as follows. In the next section the environment is described. In Section 3 the agents' problem is solved. Section 4 describes the dynamics of the economy under distinct monetary regimes and solves the government's problem. Section 5 discusses how changes in the degree of information transmission affect the behavior of the government. Section 6 deals with existence and characterization of equilibria of the game between the agents and the government. In Section 7 some remarks about our modeling assumptions are made. Section 8 concludes. An appendix collects all proofs.

²That is the case, for example, in Wallace (1997) and Katzman, Kennan, and Wallace (2003). Outside the money literature, Wolinsky (1990) is the closest to our paper in how the transmission of information is modeled. A key difference, and something that plays a central role in our analysis, is that we parameterize the speed of information transmission in our environment by allowing the number of meetings per unit of time to change.

2 The Model

We have a discrete time economy populated with one large infinitely lived agent, that we call the government, and a $[0, 1]$ continuum of small infinitely lived agents. The discount factor of the government, denoted by δ , is its private information and is determined by a draw from a distribution function F with p.d.f. f and support $[0, 1]$. The discount factor of each small agent is an independent draw from a distribution function G with a p.d.f. g . We assume the support of G is $[\underline{\beta}, \bar{\beta}]$, with $1 < \underline{\beta} < \bar{\beta} < 1$.

Throughout the paper we refer to the small agents simply as *agents*. They are of $K > 2$ different types, each one corresponding to one of the K possible types of goods that can be produced in this economy. Goods are indivisible and perishable, i.e., they only last for one period. An agent of type $k \in \{1, \dots, K\}$ only derives utility from consuming a good of type k , with each unit consumed yielding utility $u > 0$. Moreover, a type k agent can only be endowed with a type $k + 1 \pmod{K}$ good, that we call his endowment good. The distribution of agent types across the population is uniform. In other words, for any type k , the fraction of the population that is of this type is $\frac{1}{K}$. At the beginning of every period, all agents receive one unit of their endowment good. They then have two alternatives, either go to the market, where trade is possible, or stay in autarky. If an agent stays in autarky, he receives some flow utility \bar{a} . We discuss \bar{a} below.

The government can neither produce nor consume any of the K goods that are available in the economy. It has, however, the technology to print indivisible fiat money and store goods over time, with the latter being the source of its utility. With this, we aim at capturing the idea that the government derives utility from seigniorage, i.e., the revenue from money issue. We are going to be precise about how the government derives utility at the end of this section. The way the government obtains goods for storage is the following: In any given period, after the agents have made their market-autarky decisions, but before trade starts in the market, the government approaches a fraction m of the population that entered the market and offers to exchange their corresponding endowment good for one unit of fiat money. 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 government's choice.

The market is organized as follows. We have 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.³ Since $K > 2$, there are no double coincidence of wants meetings. An agent, however, can trade his endowment for money

³We adopt this particular market structure for simplicity. As we show in the next section, it delivers a simple expression for the value functions within a period. The results do not change under a more general specification for the meeting technology.

and use money to buy the good he likes. More precisely, if an agent wants money, he goes to the sector that trades his endowment and searches for an agent with money. If he has money, he goes to the sector that trades the good he likes and searches for an agent with it. As soon as an agent obtains one unit of the good he likes, he consumes it. After that, he receives one more unit of his endowment good, that can be used for further trading.⁴ In every period, anyone going to the market faces n rounds of meetings, where $n > 1$ is fixed. If at the end of the last round of market meetings an agent has money with him, he has two options. He can either attempt to obtain his endowment good back from the government, or he can keep his money until the next period. If he obtains his endowment good from the government, he can use it to get utility $a < u$. Finally, we assume that agents don't discount within a trading period.

To summarize, the sequence of actions in this economy in any given period is as follows: (i) At the beginning of every period each agent decides between staying in autarky or going to the market, and the government chooses the amount $m \in \{m_L, m_H\}$ of money it wants to put in circulation; (ii) If an agent stays in autarky, he obtains utility \bar{a} ; (iii) If instead the agent chooses the market, he first has the chance (in case he is approached by the government) of exchanging his endowment for money; (iv) Following that, all agents face n rounds of market meetings; (v) After the last of these market meetings is over, agents with money can go back to the government and attempt to exchange their unit of money for their endowment. In the next period, the same process is repeated.

One important assumption that we make about the government's behavior is that once it chooses the value of m in the first period, it cannot change it afterwards. In the last section of the paper we discuss the role of this assumption, and how it affects our results. We say the monetary regime is soft if $m = m_H$ and tight if $m = m_L$.

Notice that in the description of the market, we took the behavior of the agents and the government as given. It is possible, in a natural way, to model the market environment itself as a game involving the agents and the government. As the next subsection makes clear, this game has an equilibrium where: (i) The agents always exchange their endowment for one unit of money if approached by the government; (ii) The agents' behavior in the market (sectors to visit and trading decisions) is as above; (iii) The government always returns the endowment goods if the agents that can claim them at the end of the trading period do so. Since in this paper we are interested in the agents' market/autarky decisions, and in how these interact with the choice of monetary regime by the government, we omit the above details for the sake of brevity.

⁴A detailed specification of the production process is not important for our analysis. Hence, we proceed as in KW (1993) and assume that consumption is the only input necessary for the production of a new endowment.

2.1 Agents' Payoffs

We now specify what are the agents' payoffs in this environment. We begin by describing how their flow payoffs from entering the market depend on m , the amount of money in circulation. An implicit assumption in all that follows is that in every period, no matter the nature of the monetary regime, a positive fraction of the agents enters the market. Otherwise the random matching process described above does not make sense. In Section 6 we show that there are conditions on the model parameters that justify this assumption.

Suppose the fraction of money in the economy is m . Let w_j^i indicate the current period expected payoff of an agent with j units of money right before his i^{th} meeting, where $j \in \{0, 1\}$ and $i \in \{1, \dots, n\}$. Let w_j^{n+1} indicate the current period expected payoff of an agent that has j units of money at the end of the market meetings. Then

$$\begin{aligned} w_0^i(m) &= mw_1^{i+1}(m) + (1-m)w_0^{i+1}(m), \\ w_1^i(m) &= mw_1^{i+1}(m) + (1-m)(u + w_0^{i+1}(m)). \end{aligned}$$

Consider the second equation. An agent with money right before his i^{th} trade opportunity has a probability m of meeting another agent with money. In this case no trade occurs and he moves to the next round of meetings with money. With probability $(1-m)$ he meets an agent without money, in which case he obtains utility u and moves to the next round of meetings without money. A similar interpretation holds for the first equation.

Below we are going to determine under what conditions an agent holding money at the end of the n^{th} market meeting exchanges it for his endowment good. If that is the case, then

$$w_0^{n+1}(m) = w_1^{n+1}(m) = a.$$

We can solve this problem recursively and obtain the current period expected payoff of an agent right before his first meeting in the market. If we let $w_0(m) = w_0^1(m)$ and $w_1(m) = w_1^1(m)$, then

$$\begin{aligned} w_0(m) &= a + (n-1)m(1-m)u, \\ w_1(m) &= a + (n-1)m(1-m)u + (1-m)u. \end{aligned}$$

Notice that $w_1(m) > w_0(m)$. An agent's current period expected payoff $w(m)$ from going to the market as a function of the amount of money in circulation is then equal to

$$w(m) = (1-m)w_0(m) + mw_1(m) = a + nm(1-m)u.$$

We stated before that the overall gain in autarky in every period is equal to \bar{a} . We now describe how this value is obtained. When in autarky, any agent can use his endowment good as an input

to a production technology. Each unit of endowment good produces one unit of a consumption good that yields utility a .⁵ If the agent has more production possibilities, he receives one more unit of his endowment good. We assume that each round of meetings in the market corresponds to a production possibility in autarky. Moreover, since an agent that goes to the market has an additional consumption opportunity at the end of his market meetings, we also give this opportunity to an agent that stays in autarky. Hence $\bar{a} = (n + 1)a$. Consequently, the payoff gain or loss from going to the market is equal to

$$w(m) - (n + 1)a = nm(1 - m)u - na.$$

Instead of looking at this difference, we look at the payoff gain or loss per unit of market meetings,

$$\frac{1}{n} [w(m) - (n + 1)a] = m(1 - m)u - a.$$

Denote $w(m)/n$ by $v(m)$. This normalization is useful when we make comparisons between the market and autarky for distinct choices of n , and when $n \rightarrow \infty$. In particular, this difference is independent of n , the number of market meetings. Therefore, when we study the effect of changes in n on an agent's optimal decision, we look at the informational effects only, not to the real effects that this change has over the flow payoffs.

In this paper we are interested in a situation where an agent's expected flow payoff from choosing the market decreases when the fraction of people with money increases. Otherwise, there are no trade-offs involved when the government decides how much money to issue. Therefore, we restrict attention to the region of parameters where $v(m)$ decreases with m . This leads to our first assumption:

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

Our next assumption states that the values of m_H and m_L correspond, respectively, to non-existence and existence of monetary equilibrium under full information. In other words, if agents know that the government issues m_H , they prefer autarky, while if they know the government issues m_L they prefer the market:

$$\textbf{Assumption 2} \quad m_H(1 - m_H) < \frac{a}{u} < m_L(1 - m_L).$$

Finally, we impose conditions such that an agent ending up with a note at the end of his n^{th} market meeting wants to exchange it for his endowment.⁶ In this way we ensure that the fraction

⁵We can think that this produced good is the type of good that the agent likes, but there is a disutility $c = u - a$ in producing each unit of it.

⁶The government is indifferent between keeping the goods or giving them back (as they are perishable). Since the agent has a gain in receiving the good back, the analysis that follows looks at an efficient equilibrium where the government always returns the goods.

of agents in the market that carry money does not change over time. Because an agent can carry at most one unit of money, we need:

Assumption 3 $\frac{a}{u} > \bar{\beta}(1 - m_L)^2$.

This ensures that even the most patient agent possible would rather recover his endowment than increase the probability that he begins the following period with money. Since $m_L \geq \frac{1}{2}$, we know that for all $\beta \in [0, 1)$, $\beta(1 - m_L)^2 < m_L(1 - m_L)$, and so assumptions 2 and 3 can be both satisfied.

Now that we know the full-information flow payoffs, we can describe payoffs in the incomplete information case, the case of interest. It is convenient to regard the choice of m by the government at the start of every period as a meeting for the agents. If m is the choice of money by the government, then with probability m an agent that enters the market meets the government and trades his endowment for one unit of fiat money. Therefore, if n is the number of market meetings in a given period, all agents that enter the market face $n + 1$ meetings. We refer to these $n + 1$ meetings as trade meetings.

We begin by introducing some notation. Let θ_0 be the common prior belief among all agents that the monetary regime is soft. A mixed strategy for the government is a function $M : [0, 1] \rightarrow \mathcal{P}\{m_L, m_H\}$ that maps discount factors into probability distributions over the set of possible monetary regimes. If we let $M_H(\delta)$ be the probability that a government with discount factor δ chooses a soft monetary regime, then θ_0 is given by

$$\theta_0 = \int_{[0,1]} M_H(\delta) f(\delta) d\delta.$$

Let Ω^t be the set of all possible histories up to period t that an agent can face, and Ω^∞ be the set of all possible infinite histories in this environment. Loosely speaking, an element h^t of Ω^t is a list made up of: (i) The agent's prior belief that $m = m_H$; (ii) All his previous action choices; (iii) All his previous good and money holdings; (iv) The good and money holdings of all his previous trade partners; (v) The outcome of his trade meetings.

Since money and goods are indivisible, they can only be exchanged on a one-to-one basis. Moreover, in our environment agents hold either one unit of money or one unit of their endowment good at any point in time. Therefore, the only relevant piece of information for an agent at the beginning of a period is the record of money holdings of his partners in all his previous trade meetings. This includes the meetings with the government, where being offered to exchange one's endowment for one unit of money is interpreted as the government having one unit of money. Hence $\Omega = \{0, 1\}^{n+1} \cup \{A\}$ summarizes an agent's experience in any given period. If he chooses autarky (A), he observes nothing. If he chooses to enter the market (M), he faces a succession of $n + 1$

meetings, his government meeting and the n market meetings. A zero indicates a meeting with an agent without money, a one a meeting with an agent with money. The set of histories up to period t is then $\Omega^t = \times_{\tau=1}^{t-1} \Omega$ if $t > 1$ and $\Omega^t = \{\theta_0\}$ if $t = 1$. The set of all infinite histories an agent can face in this environment is just $\Omega^\infty = \times_{\tau=1}^\infty \Omega$.

A behavioral strategy for an agent is a list $\{s_t\}_{t=1}^\infty$ of functions such that $s_t : \Omega^t \rightarrow \mathcal{P}\{A, M\}$ is the map describing the agent's (possibly random) choice of action in period t as a function of his private history. Here M stands for market and A for autarky.

Let $p(M|h^t) = s_t(h^t)(\{M\})$ and $p(A|h^t) = s_t(h^t)(\{A\})$ denote, respectively, the probabilities that the market and autarky are chosen after history h^t when the behavioral strategy $s = \{s_t\}_{t=1}^\infty$ is being followed. Let $v^E(h^t)$ denote the expected flow payoff from entering the market given a history h^t . Then, given the normalization in payoffs introduced above,

$$u_t(h^t|s) = p(M|h^t)v^E(h^t) + p(A|h^t)\frac{(n+1)a}{n}$$

is the expected flow utility in period t , given $h^t \in \Omega^t$, from playing s . To finish, observe that a behavioral strategy s , together with an initial prior belief θ_0 and a monetary regime m , determine a probability distribution over Ω^∞ . Denote this probability distribution by $\sigma(s, m, \theta_0)$. Therefore, if an agent's discount factor is β , his expected payoff from following s , given θ_0 , is

$$V(s|\theta_0) = (1 - \beta) \left\{ \theta_0 E_{\sigma(s, m_H, \theta_0)} \left[\sum_{t=1}^\infty \beta^{t-1} u_t(h^t|s) \right] + (1 - \theta_0) E_{\sigma(s, m_L, \theta_0)} \left[\sum_{t=1}^\infty \beta^{t-1} u_t(h^t|s) \right] \right\}.$$

2.2 The Government's Payoffs

Let $\mu_t(m)$ be the measure of the population that enters the economy in period t when m is the amount of money in circulation. This measure is determined by the agents' behavior and the value of m . As a consequence of Assumption 3, $\mu_t(m)m$ is the amount of goods that the government stores during this period. The government's utility, as a function of its discount factor δ and its choice of money supply m in the first period, is then given by

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

if $\delta < 1$. When $\delta = 1$, we take the government's utility to be given by

$$U(m, 1) = \liminf_T \frac{1}{T} \sum_{t=1}^T \mu_t(m)m.$$

A convenient feature of this utility specification for an infinitely patient government is that if $\mu_t(m)$ converges, then the following two facts hold. First, $U(m, 1) = \mu_\infty(m)m$ where $\mu_\infty(m)$ is the limit of $\mu_t(m)$. Second, $U(m, \delta)$ converges to $U(m, 1)$ as δ goes to 1.

2.3 Equilibrium Notion

In this paper we assume that all agents use Bayes rule to update their beliefs that the monetary regime is soft.⁷ A behavioral strategy $s = \{s_t\}$ for the agents together with a choice m of monetary regime by the government determines the measures $\mu_t(m)$ of agents that enter the market in each period t . An equilibrium of this game is then a pair (s^*, M^*) such that:

- (i) $E_m[U(m, \delta)|M^*(\delta)] = \sup\{E_m[U(m, \delta)|M(\delta)] \mid M : [0, 1] \rightarrow \Delta_2 \text{ measurable}\}$ for almost all $\delta \in [0, 1]$, where E_m denotes expectation with respect to m and Δ_2 is the unit simplex in \mathbb{R}_2 ;
- (ii) $V(s^*, \theta) = \sup\{V(s, \theta) \mid s \text{ is a behavioral strategy}\}$ for all $\theta \in [0, 1]$;
- (iii) Agents behave in a sequentially rational way.

We restrict attention to equilibria where a positive measure of agents enters the market in every period. This condition rules out the uninteresting non-monetary equilibrium where no agent ever enters the market.

3 Agent's Behavior

In this section we take the behavior of the government as given and study the agents' behavior when the number of market meetings is n . The agent's problem is an example of a two-armed bandit with one known arm.⁸ It is natural to look at such a problem as a Markovian decision problem where the state, that we denote by θ , is the belief that $m = m_H$ and it changes through Bayesian updating.

The way the θ changes is as follows. If at the beginning of a period t an agent chooses autarky, θ does not change, as he receives no new information about the value of m . If, instead, he goes to the market, his new updated belief is

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

where $c \in \{0, \dots, n+1\}$ is the number of meetings with money he faces. Note that since the decision of the government does not change over time, it does not matter the order in which agents with money are met, only the total number of those agents.

Let \mathcal{P} denote the set of all probability measures on the (Borel sets of) the unit interval and let N be the random variable denoting the possible number of meetings with money in one period of

⁷In the next section we see in detail how belief updating takes place.

⁸See Banks and Sundaram (1992) for a detailed discussion of multi-armed bandits.

market trading. The probability an agent with belief θ assigns to the event $N = c$ is

$$\theta \Pr_H(N = c) + (1 - \theta) \Pr_L(N = c), \quad (1)$$

where N has a multinomial distribution with parameters m_H or m_L , depending on whether the regime is soft or tight, and $n + 1$. Let $\pi(\theta)$ be the probability distribution over $\{0, \dots, n + 1\}$ induced by N and θ ; that is, let $\pi(c|\theta)$ be given by (1). From $\pi(\theta)$ and B , we can then construct the transition probability $q^n : [0, 1] \rightarrow \mathcal{P}$ such that

$$q^n(D|\theta) = \int I_D(B(\theta, c)) \pi(dc),$$

where D is a measurable subset of $[0, 1]$ and I_D is the indicator function of this set. This transition probability describes, for an agent with belief θ that goes to the market on a given period, the distribution of his possible updated beliefs after facing the government and his n market meetings.

Let W_β be the value function for an agent with discount factor β . Its corresponding Bellman equation is then given by

$$W_\beta(\theta) = \max \left\{ (1 - \beta) \left(a + \frac{a}{n} \right) + \beta W_\beta(\theta), (1 - \beta) \left(v(\theta) + \frac{a}{n} \right) + \beta \int W_\beta(s) q^n(ds|\theta) \right\}, \quad (2)$$

where $v(\theta) = \theta v(m_H) + (1 - \theta) v(m_L) - \frac{a}{n}$. If in a given period an agent with belief θ chooses autarky, he gets a flow payoff of $a + \frac{a}{n}$ and his belief stays the same. If instead he chooses the market, his expected flow payoff is $v(\theta) + \frac{a}{n}$ and he updates his belief based on his market experience. Recall, once more, the normalization of payoffs introduced in Section 2. In both cases he faces the same market/autarky decision in the following period. Let $V_\beta(\theta) = W_\beta(\theta) - a/n$. Then (2) can be rewritten as

$$V_\beta(\theta) = \max \left\{ (1 - \beta)a + \beta V_\beta(\theta), (1 - \beta)v(\theta) + \beta \int V_\beta(s) q^n(ds|\theta) \right\}.$$

From now we use the above equation as the Bellman equation for the agents' problem and V_β for his value function. We omit the dependence of V_β on n until Section 5, where we study what happens when we let n vary. The following proposition is a well-known result from the bandit literature. Its proof is in Appendix A.

Proposition 1. *There exists, for each $\beta \in (0, 1)$, a unique $\theta^G(\beta) \in (0, 1)$ such that an agent with discount factor β always goes to the market if $\theta < \theta^G(\beta)$ and always stays in autarky if $\theta > \theta^G(\beta)$.*

An interesting implication of this result is that even when the monetary regime is tight, and so the market is objectively better than autarky, an optimizing agent may stay in the market only a finite number of periods and already feel informed enough to drop out of the market for good.

It is important to note that an optimal decision rule involves not only the comparison between the agent's expected flow payoff from entering the market and the flow payoff from staying in autarky. It also takes into consideration the fact that by entering the market, the agent obtains additional information about the monetary regime. To make this point more clear, consider the optimal decision rule of a myopic agent (one that has $\beta = 0$). In this case, his decision to enter or not in the market does not take into account any gains from experimentation and depends solely on the comparison between flow payoffs. He enters as long as $v(\theta) \geq a$. Let θ_m be the unique value of θ for which we have equality. Notice that θ_m is independent of n . It is possible to show that $\theta^G(\beta) > \theta_m$ for all $\beta \in (0, 1)$ and all $n > 1$. Therefore, even when the flow expected payoff in the market is smaller than in autarky, the benefits of obtaining additional information induce the agent to choose the market. In other words, the option value of market experimentation is positive. This also implies that $\theta^G(\beta)$ is increasing in β , as the more patient an agent is, the more he values experimentation. The following result is proved in Appendix A.

Proposition 2. *$\theta^G(\beta) > \theta_m$ for all $n > 1$ and all $\beta \in (0, 1)$. Moreover, $\theta^G(\beta)$ is increasing in β . In particular, $\theta^G(\beta) \leq \theta^G(\bar{\beta}) < 1$ for all $\beta \in [\underline{\beta}, \bar{\beta}]$.*

Throughout the paper we assume, without loss of generality, that when indifferent between the market and autarky, an agent always chooses the market. Since we restrict attention to equilibria where a positive fraction of the population enters the market in every period no matter the monetary regime, the only feasible equilibria have an agent with discount factor β entering the market if, and only if, his belief is less or equal than $\theta^G(\beta)$. Therefore, in order for such equilibria to exist, it is necessary to assume that $\theta_0 < \theta^G(\bar{\beta})$, so that at least in the first period a positive fraction of the population enters the market.⁹ Moreover, we assume that $\theta_0 \geq \underline{\theta}_0$, where $\underline{\theta}_0$ is a fixed positive number, and so all agents put some probability on the event that the monetary regime is soft.¹⁰ In Appendix B we see how such a lower bound for θ_0 is determined.

4 Aggregate Behavior and Government's Behavior

This section is divided in two parts. First we determine how the aggregate behavior of all agents depends on the nature of the monetary regime. We show that if the monetary regime is soft, then over time the fraction of agents accepting money and entering the market converges to zero. If, instead, the monetary regime is tight, there is a positive measure of agents that always accepts

⁹This follows from the assumption that the distribution G of discount factors across the agents has a density.

¹⁰There is no equilibrium in which $\theta_0 = 0$. If this were the case, then the government would always choose $m = m_H$, which contradicts the fact that $\theta_0 = 0$.

fiat money and enters the market. It is this characterization result that allows us to determine the government's choice of monetary regime as a function of its discount factor. The study of the government's behavior is done in the second part.

4.1 Aggregate Behavior

Let $\mu_t(m)$ be the fraction of agents entering in the market in period t given that m is the amount of money in circulation.¹¹ We want to show that if $\theta_0 \in [\underline{\theta}_0, \theta^G(\bar{\beta})]$, then

$$\lim_{t \rightarrow \infty} \mu_t(m_H) = 0 \quad \text{and} \quad \lim_{t \rightarrow \infty} \mu_t(m_L) > 0.$$

For each $h^t \in \Omega^t$, an agent has a belief $\theta(h^t)$ about the value of m . It is possible to construct a random variable $\theta_t : \Omega^\infty \rightarrow [0, 1]$ describing the distribution of period t beliefs for an arbitrary agent.¹² Therefore

$$\eta_t(m, \beta) = \Pr\{\theta_t \leq \theta^G(\beta) \mid m\}$$

is the probability that any agent with discount factor β enters the market in period t when the monetary regime is m .¹³ Since there is no aggregate uncertainty, the probability that an agent faces a private history h^t is equal to the measure of agents experiencing the history h^t . This implies that

$$\mu_t(m) = \int \eta_t(m, \beta) g(\beta) d\beta.$$

As autarky is an absorbing state, it is easy to see that both $\{\eta_t(m_H, \beta)\}$ and $\{\eta_t(m_L, \beta)\}$ are non-increasing sequences. Hence the same is true of $\{\mu_t(m_L)\}$ and $\{\mu_t(m_H)\}$, and so these fractions must converge to some numbers in $[0, 1]$, as they are bounded.

Let us first establish that for almost all $\beta \in [\underline{\beta}, \bar{\beta}]$, $\lim_{t \rightarrow \infty} \eta_t(m_H, \beta) = 0$. Suppose, by contradiction, that a positive measure of agents enters the market an infinite number of periods. If that is the case, they learn the true value of m by the consistency of Bayes estimates for the multinomial distribution. See De Groot (1970). In particular, when $m = m_H$, the belief of these agents converges to one with probability one. Because an agent with discount factor β should always choose autarky when his belief gets above $\theta^G(\beta)$, and $\theta^G(\beta) < 1$ for $\beta \in [\underline{\beta}, \bar{\beta}]$, we cannot have both $m = m_H$ and a positive measure of agents entering the market an infinite number of periods. Therefore, $m = m_H$ implies that with probability one any agent enters the market only a finite

¹¹The fractions $\mu_t(m)$ also depend on θ_0 and n . Since both are fixed for now, we omit this dependence.

¹²This construction is identical to the one from Easley and Kiefer (1988).

¹³Let $\gamma_t(m, \theta) = \Pr\{\theta_t \leq \theta \mid m\}$. Then γ_t , as a function of θ , has a finite number of discontinuities. Hence it is a Borel-measurable function of θ , as any function with a finite number of discontinuities can be written as the pointwise limit of a sequence of continuous functions. Because $\theta^G(\beta)$ is monotonic in β , it is a Borel-measurable function of β . Consequently $\eta_t(m, \beta) = \gamma_t(m, \theta^G(\beta))$ is a Borel-measurable function of β .

number of periods. Therefore it must be that $\lim_{t \rightarrow \infty} \eta_t(m_H, \beta) = 0$ for almost all $\beta \in [\underline{\beta}, \bar{\beta}]$, as we wanted to show.

Now we make use of Theorem 5.1 in Banks and Sundaram (1992). It says that for any multi-armed bandit problem with independent arms and finite type spaces, the following holds: If at any point in time an arm is selected by an optimal strategy, then there exists at least one type of this arm with the property that conditional on the arm's type being this particular one, this arm remains, with non-zero probability, an optimal choice forever after. In our setting, where the type of the market arm is m , this implies that if an agent with discount factor $\beta \in [\underline{\beta}, \bar{\beta}]$ enters the market at some point in time, then there exists $\alpha(\beta) > 0$ and $m \in \{m_L, m_H\}$ such $\eta_t(m, \beta) > \alpha(\beta)$ for all t . By hypothesis, there exists $\beta' < \bar{\beta}$ such that if $\beta \geq \beta'$, then $\theta^G(\beta) \geq \theta_0$. Moreover, we know, from the previous paragraph, that $\eta_t(m_H, \beta) \rightarrow 0$ for almost all $\beta \in [\underline{\beta}, \bar{\beta}]$. Hence it must be that for almost all $\beta \in [\beta', \bar{\beta}]$, $\eta_t(m_L, \beta) > \alpha(\beta) > 0$ for all $t \in \mathbb{N}$. We can then state the following result, which is an immediate consequence of the Lebesgue bounded convergence theorem, see Kolmogorov and Fomin (1970).

Proposition 3. *Suppose $\underline{\theta}_0 \leq \theta_0 < \theta^G(\bar{\beta})$. Then $\lim_{t \rightarrow \infty} \mu_t(m_H) = 0$ and $\lim_{t \rightarrow \infty} \mu_t(m_L) = \mu_L > 0$.*

Therefore, even though it is true that overissued money can circulate in the economy in the short-run, it gradually stops being used over time. Note that this process of overissue and abandonment of paper money may take a long time if the private experience of an agent in the market (measured by the number n of market meetings) is small. If, however, there is no overissue, a positive measure of agents always accepts fiat money and step in the market.

4.2 Government's Behavior

In the previous subsection we determined how the measures $\mu_t(m)$ evolve over time as a function of the monetary regime when all agents in the economy follow the decision rule described by Proposition 1. We now determine how the government's decision is affected by these dynamics.

It is obvious that an infinitely patient government prefers a tight monetary regime to a soft one. Suppose then that $\delta < 1$. The government's utility is

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

By choosing a soft monetary regime, the government enjoys a higher flow of utility in the first several periods, when the agents are still relatively uninformed. The government knows, however, that $\mu_t(m_H) \rightarrow 0$, and so its revenue from seignorage converges to zero. The alternative is to choose $m = m_L$. In this case, the revenue from seignorage, though smaller at the beginning, is

bounded away from zero. Hence, a sufficiently patient government should choose $m = m_L$, while an impatient one should go with $m = m_H$ despite the fact that monetary trade collapses. We have the following result. Its proof can be found in Appendix A.

Proposition 4. *Suppose $\theta_0 \in [\theta_0, \theta^G(\bar{\beta})]$. There exists a unique $\delta^* \in (0, 1)$ such that $U(m_L, \delta^*) = U(m_H, \delta^*)$. Moreover, $U(m_L, \delta) > U(m_H, \delta)$ if $\delta > \delta^*$ and $U(m_L, \delta) < U(m_H, \delta)$ if $\delta < \delta^*$.*

In other words, if the agents follow the cutoff belief rule described by Proposition 1, the government should follow a cutoff strategy as well: Choose m_H if $\delta < \delta^*$ and m_L if $\delta > \delta^*$, where $\delta^* \in (0, 1)$. From now on we adopt the convention that when indifferent between m_L and m_H , that is, when $\delta = \delta^*$, the government chooses m_L . Nothing that follows depends on this assumption. To finish, notice that δ^* depends on θ_0 and n .

5 Information Transmission

The purpose of this section is to study how n , the number of market meetings, affects the behavior of the government. This requires us to change slightly our notation, as we now need to make explicit the dependence of both the agents' and the government's problem on n . As such, we denote the individual agent value functions of Section 3 by $V_{\beta, n}$ and the corresponding cutoff beliefs by $\theta^G(\beta, n)$. We also need to consider how the payoffs to the agents and the government depend on the (common) prior θ_0 . For this reason, the fractions of agents entering the market are now $\mu_t(m_L, n; \theta_0)$ and $\mu_t(m_H, n; \theta_0)$, and the government's payoff is $U(m, \delta, n; \theta_0)$.

The normalization of the agents' payoffs introduced in Section 2 implies that they are in utils per number of market meetings. Therefore changing n changes only the degree of information transmission in the market, not the market's gain or loss relative to autarky. Specifically, as n increases, going to the market becomes more informative, and so the agents learn faster about the nature of the monetary regime. Our intuition would then suggest that a soft monetary regime breaks down in a shorter amount of time, thus reducing the incentives to overissue for any government.

It turns out that the above intuition is broadly correct, but not entirely accurate, and the reason is that changing n also affects the behavior of the agents. If the informational content of the market increases, the option value of trying it increases as well. In other words, the cutoff beliefs $\theta^G(\beta, n)$ should be increasing in n for all $\beta \in (0, 1)$. The first result that we establish is that $\theta^G(\beta, n)$ is indeed an increasing function of n , but it is bounded away from one.

To see why $\theta^G(\beta, n)$ must be bounded away from one, consider the hypothetical case where the number of market meetings is infinite. If an agent with discount factor $\beta \in (0, 1)$ and belief θ enters

the economy, his flow payoff is

$$v_\infty(\theta) = (1 - \beta)[\theta m_H(1 - m_H) + (1 - \theta)m_L(1 - m_L)]u.$$

Moreover, he learns the true amount of money in circulation. Given that before entering the market he expects the monetary regime to be soft with probability θ , his overall payoff from doing so is then

$$(1 - \beta)v_\infty(\theta) + \beta[a + (1 - \theta)m_L(1 - m_L)u].$$

If we let $V_{\beta,\infty}$ be the value function for this problem, then the corresponding Bellman equation is

$$V_{\beta,\infty}(\theta) = \max \{(1 - \beta)a + \beta V_{\beta,\infty}(\theta), (1 - \beta)v_\infty(\theta) + \beta[a + (1 - \theta)m_L(1 - m_L)u]\}. \quad (3)$$

The optimal decision rule for an agent is still a cutoff strategy, where the cutoff belief $\theta^G(\beta, \infty)$ is the unique solution to

$$a = (1 - \beta)v_\infty(\theta^G(\beta, \infty)) + \beta[\theta^G(\beta, \infty)a + (1 - \theta^G(\beta, \infty))m_L(1 - m_L)u].$$

In particular, we cannot have $\theta^G(\beta, \infty) = 1$, given that $v_\infty(1) = m_H(1 - m_H)u < a$, and so $\theta^G(\beta, \infty) < 1$ for all $\beta \in (0, 1)$. Since the option value of entering the market is highest when $n = \infty$, as the informational gain is the biggest possible, we must then have $\theta^G(\beta, n) < \theta^G(\beta, \infty)$ for all $\beta \in (0, 1)$ and all $n \in \mathbb{N}$, and this implies the desired result. It is possible to give a more rigorous justification for the above argument. It is a corollary to the following lemma.

Proposition 5. *For all $\beta \in (0, 1)$, $\{V_{\beta,n}\}$ is a non-decreasing sequence that converges uniformly to $V_{\beta,\infty}$ as $n \rightarrow \infty$.*

Corollary 1. *For all $\beta \in (0, 1)$, $\theta^G(\beta, n)$ is non-decreasing in n and $\sup_{n \geq 2} \theta^G(\beta, n) < 1$.*

Corollary 1 implies the following result: As long as $\theta_0 < \sup_{n \geq 2} \theta^G(\bar{\beta}, n)$ then, as n increases to infinity, $\mu_t(m_H, n)$ converges to zero and $\mu_t(m_L, n)$ converges to one for all $t \geq 2$. From this we get our next result. Its proof is in Appendix A.

Proposition 6. *Let $\underline{\delta} = (m_H - m_L)/m_H$ and fix $\bar{\theta}_0$ with $\underline{\theta}_0 < \bar{\theta}_0 < \sup_{n \geq 2} \theta^G(\bar{\beta}, n)$. Then, for all $\delta > \underline{\delta}$, there exists $n(\delta)$ such that $U(m_L, \delta, n; \theta_0) \geq U(m_H, \delta, n; \theta_0)$ if $\theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]$ and $n \geq n(\delta)$.*

Notice that $(1 - \delta)m_H$ is a lower bound on the payoff a government with discount factor δ obtains when it chooses a soft monetary regime. If the same government chooses a tight monetary regime, an upper bound for its payoff is m_L . Therefore, no matter the number n of market meetings, a government with discount factor $\delta < \underline{\delta}$ always chooses $m = m_H$.

6 Equilibrium

Here we bring together the results from the previous sections and establish the two main results of the paper. First, that for all $n \geq 2$ there are equilibria where a positive measure of agents is in the market in all periods and agents and the government follow the cutoff strategies of Propositions 1 and 4, respectively. In fact, these are the only equilibria where a positive measure of agents is in the market in all periods. Second, that as n increases, the cutoff discount factor of the government in all these equilibria approaches $\underline{\delta}$, the lowest cutoff discount factor possible.

Fix the number of market meetings, suppose the government's strategy is given by a function $M : [0, 1] \rightarrow \Delta_2$, and let θ_0 be the ex-ante probability of a tight monetary regime. The cutoff belief strategy described in Proposition 1 is, for any given agent, the unique best response to any profile of behavioral strategies by the other agents with the property that a positive measure of agents enters the market in all periods. It is, moreover, sequentially rational as a consequence of the principle of optimality for dynamic programming. Section 4 shows that if all agents follow the cutoff belief strategy described by Proposition 1, then, as long as $\theta_0 < \theta^G(\bar{\beta}, n)$, their aggregate behavior implies that the unique best response for the government is a cutoff discount factor strategy. Let $\delta^*(\theta_0, n)$ denote this cutoff discount factor. Unfortunately this is not enough to guarantee the existence of an equilibrium with the desired properties. We need to deal with the following consistency problem:

Consistency Problem: If the government follows a cutoff discount factor strategy with cutoff $\delta^*(\theta_0, n)$, the agent's common prior that the monetary regime is soft is given by $F(\delta^*(\theta_0, n))$. How do we know that $\theta_0 = F(\delta^*(\theta_0, n))$? In other words, how do we know that the agent's common prior belief about the nature of the monetary regime is justified. We need to determine if the maps Θ_n that take θ_0 into $F(\delta^*(\theta_0, n))$ have, for all $n \geq 2$, a fixed point. Moreover, we need to ensure that for each n , Θ_n has at least one fixed point lying in the interval $[\underline{\theta}_0, \theta^G(\bar{\beta}, n))$, for otherwise the analysis conducted so far breaks down.

In Appendix B we show that there are conditions on F ensuring that the maps Θ_n constructed above have, for each $n \geq 2$, a fixed point θ_0 in the desired interval. This establishes the first main result of this paper. Since the argument of Appendix B relies on Brouwer's fixed point theorem, there is no reason to expect uniqueness for fixed F and $n \geq 2$. An interesting possibility in case there is multiplicity of equilibria is that there may be a range of discount factors for the government where its choice of monetary regime differs across equilibria. The verification of this fact, however, is beyond the scope of this paper.

Proposition 7. *For all $n \geq 2$, there exist equilibria where a positive measure of agents enters the market in all periods. Moreover, all such equilibria have the government following the cutoff strategy described by Proposition 4 and the agents following the cutoff strategy described by Proposition 1. The self-fulfilling priors $\theta_0(n)$ corresponding to these equilibria lie in $[\underline{\theta}_0, \theta^G(\bar{\beta}, n))$.*

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

A consequence of this result is that incomplete information can bring instability to trade involving money. For the monetary system to be stable, the government must have incentives to keep money valuable over a long period of time. These incentives depend on the government's discount factor δ and the number n of per-period market meetings. Since the economy has no mechanism that fully reveals the government's policy, society has to rely on the possibility that the government is sufficiently patient. The second main result of this paper shows that the government's incentives to overissue decrease when n increases. It is an immediate consequence of Propositions 6 and 7.

Proposition 8. *For all $\delta > \underline{\delta}$ there exists $n(\delta)$ such that if $n \geq n(\delta)$, the equilibria of Proposition 7 are such that the cutoff discount factor of the government is less than δ .*

It seems reasonable to suggest that in modern economies the dissemination of information about decisions made by the government is much faster than in the past. We model this change as an increment in the number of opportunities an agent has to gather information about the state of the economy. This increment is interpreted here as an increase in n . Therefore, as a result of the above proposition, we expect that in modern societies the incentives of any government to overissue are less pronounced than in the past.

7 Comments

We make a number of assumptions that deserve some comment. We assume that money and goods are indivisible and there is an upper bound on money holdings. Moreover, we assume that the government can only choose two levels of m . Finally, once the government chooses m , it is not allowed to change it. We begin with this last issue.

7.1 Time Consistency

Consider an environment where the government makes an once and for all injection of money at the beginning of time. In this economy, the decision of whether to stay in the market or go to autarky depends on more than just the belief about the government's behavior. It also depends on an agent's money holdings. Agents with the same belief may behave differently depending on whether they carry money or not. As a result, conditional on the government's choice, the distribution of money holdings across the population changes over time, and so the choice between market and autarky cannot be time-independent. It turns out that working with such an environment is considerably more difficult, and it brings no new insights as far as the government's behavior is concerned. We would still have the same trade-off between overissuing and the eventual breakdown of trade.

We avoid the above problem by structuring the market and the government in such a way that the fraction of agents with money, conditional on the choice of monetary regime, does not change over time. This makes our model tractable. This modelling choice, however, brings out the issue of whether the government's behavior is time consistent. In other words, would the equilibria described in Proposition 7 still be equilibria of the game where the government can now choose at the beginning of every period the amount $m \in \{m_L, m_H\}$ of money to put in circulation? Unfortunately no, for reasons brought forth by the literature on Reputation.¹⁴

Here is a very brief and informal explanation of why the type of equilibrium we consider does not survive. Suppose it does, and take a government that chooses a tight monetary regime. Eventually, a period t is reached where almost all agents entering the market are nearly convinced that $m = m_L$. What happens in t if instead of choosing m_L , this government chooses m_H ? Any agent in the market that is nearly convinced that the monetary regime is tight attributes any negative experience (meeting someone else with money) to bad luck. Therefore virtually all of the agents that entered the market in period t reenter the market in the following period. This means that the government under consideration has an incentive to deviate, a contradiction. In a nutshell, in the environment considered in this paper no government has an incentive to maintain a reputation of being patient.

In a related paper, Araujo and Camargo (2005), we show how to modify our environment in a way that preserves the equilibria of Proposition 7 and the analysis that follows from it. This is done by introducing some exogenously driven uncertainty: In every period a fraction λ of the population dies and is replaced by newly born agents that are uncertain about the nature of the monetary regime. If λ is neither too high nor too low, this perpetual renewal of uncertainty induces on patient governments an incentive to build and maintain a reputation of being patient, so that they always choose a tight monetary regime.

¹⁴See Mailath and Samuelson (1998), for example.

7.2 Indivisibility

In our economy, an agent's experience in a given meeting is important as it allows him to make inferences about the government's behavior. This experience has two components: The money and good holdings of his partner and the outcome of the trade process between them. When we assume that both money and goods are indivisible, and there is an upper bound on money holdings, the money holdings of the agent's partner become the only relevant piece of information for him. This implies a very simple exchange process. More importantly for our purposes, it greatly simplifies the updating of beliefs while still capturing in a natural way the idea that agents learn from their private histories. If goods and/or money were divisible, we would have to specify how trade between agents depends on private histories. The analysis would become more complex and, given our goals, would not necessarily generate additional insights.¹⁵

7.3 The Choice of m

It is possible to work with a model where the government can choose among a finite number of different monetary regimes. In the end, however, what matters for the agent is not the *exact* amount of money in circulation, but whether it is best for him to enter the market or stay in autarky. In this sense, a model where the government is constrained to choose among two different levels of money supply, one where the market is better than autarky and another where it is worse, captures the relevant idea. Hence there seems little gain from studying the more complex case.

8 Conclusion

This paper addresses in a formal way the determinants of monetary stability in a decentralized economy where fiat money is endogenously created, information about its value is imperfect, and learning is decentralized. To the best of our knowledge, ours is the first paper to put together these features. We believe they constitute a reasonable model if one wants to analyze the instability of fiat money as observed throughout history. In particular, by assuming that learning only takes place over time, we are able to induce dynamics on the acceptability of fiat money that resembles the historical accounts on the rise and eventual demise of overissued paper money.

We show that the government's temptation to overissue is limited in two different ways. First, it depends on its commitment in maintaining the long-run value of money, here given by its patience. This result goes along with Ritter (1995), who emphasizes that patience is a key condition in

¹⁵Examples of papers that address learning from private histories when goods are divisible are Katzman, Keenan and Wallace (2003) and Araujo and Shevchenko (2004).

the transition from a barter to a fiat money economy. Second, it depends on society's ability to monitor the government's behavior, modelled by the number of transactions an agent faces in a given period. The ability to collect information, and the consequent stability it induces, has changed over time. More precisely, in modern economies, information flows much faster than in the past. This reasoning offers an explanation for the late widespread use of fiat money, despite its clear advantages.

We restricted attention to the government as the sole provider of money. However, our analysis can also offer insights into scenarios where a private agent issues notes that circulate in the economy. Consider, for example the U.S. banking experience in the 19th century. King (1983) explains the overissue of circulating notes during this period by arguing that "(...) holders of circulating notes are unlikely to closely monitor the activities of a note issuer, because notes represent a small fraction of an individual's wealth and are held only for a brief period." (p.136). Our paper offers a potential environment where the likelihood of an overissue of circulating notes can be formally addressed. For instance, assume that the reasons underlying the demand for information (the ratio of money holdings to wealth, or the length of time an individual holds a note, for example) are relatively invariant, but that the technology that supplies information improves over time. Our analysis then suggests that the probability of an overissue, and the consequent necessity of imposing stringent controls over the creation of notes, reduces as the economy evolves.

A Sections 3, 4, and 5

Proposition 1 *There exists, for each $\beta \in (0, 1)$, a unique $\theta^G(\beta) \in (0, 1)$ such that an agent with discount factor β always goes to the market if $\theta < \theta^G(\beta)$ and always stays in autarky if $\theta > \theta^G(\beta)$.*

Proof: Let $I = [0, 1]$ and $B(I)$ be the set of all real-valued, bounded, and measurable functions defined on I endowed with the sup-norm. Observe that $B(I)$ is a complete metric space. Consider the map T defined on $B(I)$ such that

$$Tf(\theta) = \max \left\{ (1 - \beta)a + \beta f(\theta), (1 - \beta)v(\theta) + \beta \int f(s)q^n(ds|\theta) \right\}$$

for all $f \in B(I)$. It is easy to see that if f is measurable, then

$$\int f(s)q^n(ds|\theta) = \sum_{c=0}^{n+1} f(B(c, \theta))[\theta \Pr_H(N = c) + (1 - \theta)\Pr_L(N = c)]$$

is also measurable. In fact, if $f(\theta)$ is continuous, then $\int f(s)q^n(ds|\theta)$ is also continuous; that is, the transition probabilities q^n have the Feller property. Since v is continuous, T maps $B(I)$ into itself.

A straightforward argument shows that T is a contraction. Therefore it has a unique fixed point V^* in $B(I)$ by the Banach fixed point theorem.

We now establish, in two steps, a property of V^* that leads to the desired result. First, suppose that $f \in B(I)$ is non-increasing and let $\theta_1 > \theta_2$. Because $B(c, \theta)$ is increasing in θ for all $c \in \{0, \dots, n+1\}$,

$$\int f(s)q^n(ds|\theta_1) - \int f(s)q^n(ds|\theta_2) \leq (\theta_1 - \theta_2) \sum_{c=0}^{n+1} f(B(c, \theta_2)) [P_H(N=c) - P_L(N=c)]. \quad (\text{A.1})$$

Let π_k , with $k \in \{L, H\}$, be the probability distribution on $\{0, \dots, n+1\}$ induced by N and the choice k of monetary regime. Since $m_H > m_L$, it is easy to see that

$$\frac{d}{dc} \left(\frac{P_H(N=c)}{P_L(N=c)} \right) = \frac{d}{dc} \left(\frac{m_H}{m_L} \right)^c \left(\frac{1-m_H}{1-m_L} \right)^{n+1-c} > 0.$$

Therefore, as a consequence of MLRP, π_H first-order stochastically dominates π_L . We can then conclude, as $B(c, \theta)$ is decreasing in c for all $\theta \in [0, 1]$, that the term on the right-hand side of (A.1) is non-positive. In other words,

$$\int f(s)q^n(ds|\theta)$$

is non-decreasing in θ if f is. A useful corollary of this result is that V^* is non-increasing in θ .

The fact established in the previous paragraph together with the same argument used in the proof of Lemma 1 in Kakigi (1983) lead to the following result: T has the property that if $f \in B(I)$ is such that

$$(1 - \beta)a + \beta f(\theta) - (1 - \beta)v(\theta) - \beta \int f(s)q^n(ds|\theta) \quad (\text{A.2})$$

is non-decreasing in θ , then Tf is such that

$$(1 - \beta)a + \beta Tf(\theta) - (1 - \beta)v(\theta) - \beta \int Tf(s)q^n(ds|\theta)$$

is strictly increasing in θ . Therefore, once more by standard arguments, V^* is such that (A.2) is strictly increasing in θ . Since

$$(1 - \beta)a + \beta V^*(0) < (1 - \beta)v(0) + \beta \int V^*(s)q^n(ds|0)$$

and

$$(1 - \beta)a + \beta V^*(1) > (1 - \beta)v(1) + \beta \int V^*(s)q^n(ds|1),$$

we can then conclude that there exists a unique $\theta \in [0, 1]$, that we denote by $\theta^G(\beta)$, such that

$$(1 - \beta)a + \beta V^*(\theta) = (1 - \beta)v(\theta) + \beta \int V^*(s)q^n(ds|\theta).$$

Moreover, $\theta^G(\beta)$ is interior and it satisfies $V^*(\theta^G(\beta)) = a$.

Because the action space is finite, the problem of the agents has an optimal decision rule. The value function V_β is the conditional expected reward from one such decision rule when the agent's discount factor is β . From Blackwell (1965), we know that V_β is bounded and measurable, and that it satisfies (2). Hence $V_\beta = V^*$. From the previous paragraph and the principle of optimality, we then have the desired result: For all $\theta < \theta^G(\beta)$, the market is chosen, and for all $\theta > \theta^G(\beta)$, autarky is chosen. \square

Proposition 2 $\theta^G(\beta) > \theta_m$ for all $n > 1$ and all $\beta \in (0, 1)$. Moreover, $\theta^G(\beta)$ is increasing in β . In particular, $\theta^G(\beta) \leq \theta^G(\bar{\beta}) < 1$ for all $\beta \in [\underline{\beta}, \bar{\beta}]$.

Proof: Let us first establish that $\theta^G(\beta) > \theta_m$ for all $n > 1$. For this fix $n > 1$ and let f_0 be given by $f_0(\theta) = v(\theta)$. Since f_0 is linear,

$$Tf_0(\theta) = (1 - \beta) \max\{a, v(\theta)\} + \beta f_0(\theta) \geq f_0(\theta)$$

for all $\theta \in [0, 1]$, where T is the operator introduced in the proof of the previous proposition. Let $f_1 = Tf_0$. Since $v(\theta_m) = a$ by definition,

$$\int f_1(s) q^n(ds|\theta_m) > (1 - \beta) \max\{a, v(\theta_m)\} + \beta f_0(\theta_m) = f_1(\theta_m).$$

Therefore $Tf_1(\theta_m) > f_1(\theta_m)$. Now let $f_k = T^k f_0$, with $k \in \mathbb{N}$. Since T is monotonic, $f_k \geq f_{k-1}$ for all k . Moreover, $T^k f_0$ converges uniformly to V_β . Because uniform convergence implies pointwise convergence, we can then conclude that $V_\beta(\theta) \geq f_0(\theta)$ for all $\theta \in [0, 1]$ and $V_\beta(\theta_m) > f_0(\theta_m)$. From the previous proposition we know that $V_\beta(\theta^G(\beta)) = a$. Consequently $V_\beta(\theta^G(\beta)) = f_0(\theta_m) < V_\beta(\theta_m)$, and so $\theta^G(\beta) > \theta_m$, as V_β is non-decreasing.

To prove the second part of the theorem, let $S = [0, 1] \times [\underline{\beta}, \bar{\beta}]$ and $B(S)$ be the set of all real-valued, bounded, and measurable functions defined on S endowed with the sup-norm. Consider now the map $Q : B(S) \rightarrow B(S)$ such that

$$Qf(\theta, \beta) = \max \left\{ (1 - \beta)a + \beta f(\theta, \beta), (1 - \beta)v(\theta) + \beta \int f(\theta, \beta) q^n(ds|\theta) \right\}.$$

It is straightforward to show that Q is a contraction, so that it has a unique fixed point in $B(S)$, given that this set is a complete metric space. Denote it by V^{**} . Now let \bar{B} be the subset of $B(S)$ such that if $f \in \bar{B}$, then f is non-decreasing in β and $f(\theta, \beta) \geq \max\{a, v(\theta)\}$ for all $(\theta, \beta) \in S$. Then \bar{B} is a closed subset of B that is mapped into itself by Q . Standard arguments show that $V^{**} \in \bar{B}$. Since $V^{**}(\theta, \beta) = V_\beta(\theta)$ for all $\beta \in [\underline{\beta}, \bar{\beta}]$, we can then conclude that V_β is non-decreasing

in β . To finish, note that since $V_\beta(\theta^G(\beta)) = a$, then $\theta^G(\beta)$ is the unique solution to

$$a = (1 - \beta)v(\theta) + \beta \int V_\beta(s)q^n(ds|\theta).$$

Because the right-hand side of the above equation is increasing in β and decreasing in θ , we have the desired result. \square

Proposition 4 *Suppose $\theta_0 \in [\underline{\theta}_0, \theta^G(\bar{\beta})]$. There exists a unique $\delta^* \in (0, 1)$ such that $U(m_L, \delta^*) = U(m_H, \delta^*)$. Moreover, $U(m_L, \delta) > U(m_H, \delta)$ if $\delta > \delta^*$ and $U(m_L, \delta) < U(m_H, \delta)$ if $\delta < \delta^*$.*

Proof: Let $G(\delta)$ be given by

$$G(\delta) = U(m_L, \delta) - U(m_H, \delta) = \sum_{t=1}^{\infty} \delta^{t-1} d_t,$$

where $d_t = \mu_t(m_L)m_L - \mu_t(m_H)m_H$. We want to show that there exists a $\delta^* \in (0, 1)$ such that $G(\delta) < 0$ if $\delta < \delta^*$ and $G(\delta) > 0$ if $\delta > \delta^*$. First observe that

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

By Proposition 3, $\mu_t(m_L)$ converges to some $\mu_L > 0$ while $\mu_t(m_H)$ converges to zero. Hence there exists $t' \geq k+1$ such that if $t \geq t'$, then $d_t \geq \frac{1}{4}\mu_L m_L$, so that

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

Since the first term after the above inequality is finite for all δ , we can conclude then that $\lim_{\delta \rightarrow 1^-} G^{(k)}(\delta) = +\infty$ for all $k \geq 0$. Let now \bar{t} be the smallest integer such that $d_t \geq 0$ for all $t \geq \bar{t}$. Such a t exists because of the asymptotic behavior of the population measures in the two monetary regimes. Since

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

we have that $G^{(\bar{t}-1)}(\delta) > 0$ for all δ . Notice that we cannot have $\bar{t} = 1$, since this would imply that $U(m_L, \delta) > U(m_H, \delta)$ for all $\delta \geq 0$, which is not true for δ sufficiently close to zero.

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

$$G^{(\bar{t}-2)}(\delta) = \sum_{t=\bar{t}-1}^{\infty} (t-1)\dots(t-\bar{t}+2)\delta^{t-\bar{t}+1} d_t = (\bar{t}-2)\dots 2d_{\bar{t}-1} + \sum_{t=\bar{t}}^{\infty} (t-1)\dots(t-\bar{t}+2)d_t,$$

and so $G^{(\bar{t}-2)}(0) < 0$. Since $G^{(\bar{t}-2)}$ is strictly increasing and $G^{(\bar{t}-2)}(1^-) = +\infty$, we have that there exists a unique δ_1 such that $G^{(\bar{t}-2)}(\delta) < 0$ if, and only if, $\delta < \delta_1$. If $\bar{t} = 2$, we are done, just set $\underline{\delta} = \delta_1$. So, suppose now that $\bar{t} > 2$. The same reasoning as above shows that $G^{(\bar{t}-3)}(0) < 0$. Since $G^{(\bar{t}-3)}(\delta)$ decreases until δ_1 , and after this point it increases strictly to $+\infty$, we have that there is a unique $\delta_2 > \delta_1$ such that $G^{(\bar{t}-3)}(\delta) < 0$ if, and only if $\delta < \delta_2$. If $\bar{t} = 3$, we are again done. Otherwise, we continue with this process. Since \bar{t} is finite, it eventually reaches an end. \square

Proposition 5. *For all $\beta \in (0, 1)$, $\{V_{\beta,n}\}$ is a non-decreasing sequence that converges uniformly to $V_{\beta,\infty}$ as $n \rightarrow \infty$.*

Proof: (1) Fix $\beta \in (0, 1)$, and let T_n be, for each $n > 1$, the map introduced in the proof of Proposition 1. In other words, $T_n : B[0, 1] \rightarrow B(I)$ is the map such that

$$T_n f = \max \left\{ (1 - \beta)a + \beta f(\theta), (1 - \beta)v(\theta) + \beta \int f(s)q^n(ds|\theta) \right\}.$$

Then $V_{\beta,n}$ is, for each $n > 1$, the unique fixed point of T_n in $B(I)$. A standard argument shows that the functions $V_{\beta,n}$ are in fact continuous, as the transition probabilities q^n satisfy the Feller property. By Lemma 3.1 in Banks and Sundaram (1992), we also have that they are convex.

(2) Let $C^\infty[0, 1] = \times_{i=1}^\infty C[0, 1]$ and endow this space with the product topology. Then $\{V_{\beta,n}\}$ is a fixed point of the map $\Gamma : C^\infty[0, 1] \rightarrow C^\infty[0, 1]$ such that $(\Gamma f)_n(\theta) = T_n f_n(\theta)$ for all $n > 1$.

Let $d : C^\infty[0, 1] \times C^\infty[0, 1] \rightarrow \mathbb{R}_+$ be such that

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

where $p_n(f) = \|f_n\|_{\sup}$ and $\{c_n\}$ is some strictly positive sequence of real numbers such that c_n converges to zero. For example, we can take $c_n = 2^{-n}$. One can show that d is a metric on $C^\infty[0, 1]$ and that it metrizes the product topology on $C^\infty[0, 1]$. Since the product of a countable number of complete metric spaces is also complete, we have that $(C^\infty[0, 1], d)$ is a complete metric space.

(3) We now prove that Γ is a contraction with respect to d , so that $\{V_{\beta,n}\}$ is its unique fixed point in $C^\infty[0, 1]$. For this suppose, by contradiction, that there is no $\delta < 1$ such that $d(\Gamma f, \Gamma g) \leq \delta d(f, g)$. So, for all $n \in \mathbb{N}$, there exists $j(n) \in \mathbb{N}$ such that

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

This implies that

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

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

$$\|Tf_{j(n)} - Tg_{j(n)}\|_{\sup} > \beta\|f_{j(n)} - g_{j(n)}\|_{\sup},$$

which contradicts the fact that $\forall j \in \mathbb{N}$ the map $f_j \mapsto (\Gamma f)_j = T_j f_j$ is a contraction of modulus β . Consequently Γ is a contraction in $C^\infty[0, 1]$, as we wanted to prove.

(4) Let $X = \{f \in C^\infty[0, 1] : \{f_n\} \text{ is uniformly convergent}\}$ and suppose that $\{f^m\}$ is a sequence in X such that converges to f . We want to show that $f \in X$, that is, that $f = \{f_n\}$ is a uniformly convergent sequence. From the definition of d we know that if $\{f^m\}$ is convergent in $C^\infty[0, 1]$, then $\{f_n^m\}$ is uniformly convergent in $C[0, 1]$ for all n , and it must converge to f_n . Now observe that for all $\theta \in [0, 1]$,

$$f_{n_1}(\theta) - f_{n_2}(\theta) = f_{n_1}(\theta) - f_{n_1}^m(\theta) + f_{n_1}^m(\theta) - f_{n_2}^m(\theta) + f_{n_2}^m(\theta) - f_{n_2}(\theta),$$

where the choice of m is arbitrary, so that

$$|f_{n_1}(\theta) - f_{n_2}(\theta)| \leq \|f_{n_1} - f_{n_1}^m\|_{\sup} + \|f_{n_1}^m - f_{n_2}^m\|_{\sup} + \|f_{n_2}^m - f_{n_2}\|_{\sup}.$$

Take $\epsilon > 0$. Since, by hypothesis, $\{f_n^m\}$ is uniformly convergent for all m , we know that there exists N such that if $n_1, n_2 \geq N$, then $\|f_{n_1}^m - f_{n_2}^m\|_{\sup} < \epsilon/3$. Take then n_1, n_2 greater than N . Because $\{f_{n_1}^m\}$ and $\{f_{n_2}^m\}$ converge uniformly to f_{n_1} and f_{n_2} , respectively, there is $m_0(n_1, n_2) \in \mathbb{N}$ such that if $m \geq m_0$, then $\|f_{n_1} - f_{n_1}^m\|_{\sup} < \epsilon/3$ and $\|f_{n_2}^m - f_{n_2}\|_{\sup} < \epsilon/3$. If we now take $m \geq m_0$, we can conclude that $|f_{n_1}(\theta) - f_{n_2}(\theta)| < \epsilon$ for all $\theta \in [0, 1]$; that is, $\|f_n - f_{n'}\|_{\sup} < \epsilon$. Consequently $\{f_n\}$ is Cauchy, and so uniformly convergent. This proves that X is indeed a closed subset of $C^\infty[0, 1]$.

Now let $Y = \{f \in C^\infty[0, 1] : f_n \text{ is convex for all } n\} \cap X$. As with X , we need to show that Y is a closed subset of $C^\infty[0, 1]$. For this suppose that $\{f^m\}$ is a convergent sequence in Y with limit $f \in C^\infty[0, 1]$ that is not convex. This means that there exist $q \in \mathbb{N}$, $\theta, \theta' \in [0, 1]$, and $\lambda \in (0, 1)$

such that $f_q(\lambda\theta + (1-\lambda)\theta') > \lambda f_q(\theta) + (1-\lambda)f_q(\theta')$. Now observe that

$$\begin{aligned} \lambda f_q(\theta) + (1-\lambda)f_q(\theta') - f_q(\lambda\theta + (1-\lambda)\theta') &= \underbrace{\lambda f_q(\theta) + (1-\lambda)f_q(\theta') - \lambda f_q^m(\theta) - (1-\lambda)f_q^m(\theta')}_{a_m} \\ &+ \underbrace{\lambda f_q^m(\theta) + (1-\lambda)f_q^m(\theta') - f_q^m(\lambda\theta + (1-\lambda)\theta')}_{b_m} \\ &+ \underbrace{f_q^m(\lambda\theta + (1-\lambda)\theta') - f_q(\lambda\theta + (1-\lambda)\theta')}_{c_m}. \end{aligned}$$

Since $f^m \rightarrow f$, we know (from the previous paragraph) that $f_q^m \rightarrow f_q$ uniformly for all $q \in \mathbb{N}$, and so it must be that $a_m, c_m \rightarrow 0$. By hypothesis, we know that $b_m \geq 0$ for all m . Moreover, the sequence $\{b_m\}$ is bounded above, since $\{f^m\}$ is convergent. Hence $\{b_m\}$ has a convergent subsequence. Without loss of generality, assume that $\{b_m\}$ itself is convergent and denote its limit by α . We know that $\alpha \geq 0$, and so the right-hand side of the above equality converges to a non-negative number, a contradiction. Hence Y is also closed, since X is.

Finally, let $W = \{f \in C^\infty[0, 1] : f_n \leq f_{n+1} \text{ for all } n\} \cap Y$. As above, we want to show that this set is closed. For this, suppose that $\{f^m\}$ is a sequence in W that converges to $f \in C^\infty[0, 1]$, but there are $q \in \mathbb{N}$ and $\theta \in [0, 1]$ for which $f_{q+1}(\theta) < f_q(\theta)$. Since

$$f_{q+1}(\theta) - f_q(\theta) = f_{q+1}(\theta) - f_{q+1}^m(\theta) + f_{q+1}^m(\theta) - f_q^m(\theta) + f_q^m(\theta) - f_q(\theta),$$

an argument similar to the one from the previous paragraph shows that the right-hand side of the above inequality has a subsequence converging to a non-negative real number, which is a contradiction. Therefore W is closed, since Y is.

(5) We now show that Γ maps W into itself, so that standard arguments allow us to conclude that $\{V_{\beta,n}\} \in W$. Suppose that $f \in W$. From Lemma 3.1 in Banks and Sundaram (1992), we know that if f_n is convex, then Tf_n also is. Since $q^{n+1}(\theta)$ is a mean-preserving spread of $q^n(\theta)$, we have that $q^{n+1}(\theta)$ second-order stochastically dominates $q^n(\theta)$. Therefore, since f_{n+1} is convex by assumption,

$$\int f_{n+1}(s)q^{n+1}(ds|\theta) \geq \int f_{n+1}(s)q^n(ds|\theta).$$

But $f_{n+1} \geq f_n$ by assumption as well, and so we also have that $(\Gamma f)_{n+1} \geq (\Gamma f)_n$. To finish we prove that $\{(\Gamma f)_n\}$ is uniformly convergent. For this, let f_∞ be the uniform limit of $\{f_n\}$, and let

$$Tf_\infty(\theta) = \max \{(1-\beta)a + \beta f_\infty(\theta), (1-\beta)v(\theta) + \beta[\theta f_\infty(1) + (1-\theta)f_\infty(0)]\}.$$

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

$$|Tf_n(\theta) - Tf_\infty(\theta)| \leq \max \left\{ 0, \beta \left[\int f_n(s)q^n(ds|\theta) - \theta f_\infty(1) - (1-\theta)f_\infty(0) \right] \right\}$$

It is possible to show that the consistency of the Bayes estimator for an i.i.d. sequence of Bernoulli trials implies that for all $\theta \in [0, 1]$, $q^n(\theta)$ converges in the topology of weak convergence of probability measures to $\theta\delta(1) + (1 - \theta)\delta(0)$. Here $\delta(x)$ is the Dirac measure with mass on x . Hence, by the same argument used by Easley and Kiefer (1988) in the proof of their Lemma 1, we know that $\int f_n(s)q^n(ds|\theta) \rightarrow \theta f_\infty(1) + (1 - \theta)f_\infty(0)$ for all $\theta \in [0, 1]$. Consequently $\{Tf_n\}$ is a sequence of continuous functions defined over a compact set that converges pointwise and monotonically to Tf . By Dini's theorem, see Rudin (1976), $\{Tf_n\}$ is uniformly convergent. We can then conclude that Γ maps W into itself.

From (5) we have the desired result: $\{V_{\beta,n}\}$ is a non-decreasing sequence of functions that converges uniformly to $V_{\beta,\infty}$. \square

Corollary 1 For all $\beta \in (0, 1)$, $\theta^G(\beta, n)$ is non-decreasing in n and $\sup_{n \geq 2} \theta^G(\beta, n) < 1$.

Proof: We first prove that $\theta^G(\beta, n)$ is increasing in n . From Lemma 5 we know that

$$\int V_{\beta,n+1}(s)q^{n+1}(ds|\theta) \geq \int V_{\beta,n}(s)q^n(ds|\theta)$$

for all $n > 1$. Since $\theta^G(\beta, n)$ is the unique solution to

$$(1 - \beta)v(\theta) + \beta \int V_{\beta,n}(s)q^n(ds|\theta) = a,$$

it must be that $\theta^G(\beta, n)$ cannot decrease with n , as the left-hand side of the above equation is non-decreasing in n and non-increasing in θ . This follows from the proof of Proposition 1.

Let us now prove that $\sup_{n \geq 1} \theta^G(\beta, n) < 1$ for all $\beta \in (0, 1)$. Fix $\beta \in (0, 1)$ and suppose not. Since $\{\theta^G(\beta, n)\}$ is a bounded and non-decreasing sequence, we know that it must converge to $\theta^G(\beta, \infty) = \sup_{n \geq 1} \theta^G(\beta, n)$. By assumption, we have that $\theta^G(\beta, \infty) = 1$. Now observe that

$$V_{\beta,n}(\theta^G(\beta, n)) = (1 - \beta)v(\theta^G(\beta, n)) + \beta \int V_{\beta,n}(s)q^n(ds|\theta^G(\beta, n)) = a.$$

Since $\{V_{\beta,n}\}$ converges uniformly to $\{V_{\beta,\infty}\}$, we have that $\lim_{n \rightarrow \infty} V_{\beta,n}(\theta^G(\beta, n)) = V_{\beta,\infty}(1)$. In fact, the sequence $\{V_{\beta,n}\}$ is an equicontinuous set, see Rudin (1976). An $\epsilon/2$ argument then implies that $V_{\beta,n}(\theta^G(\beta, n))$ converges to $V(1)$ as $n \rightarrow \infty$. We also know, from the previous lemma, that $\int V_{\beta,n}(s)q^n(ds|\theta)$ converges to $\theta V_{\beta,\infty}(1) + (1 - \theta)V_{\beta,\infty}(0)$ for all θ . Moreover, from the previous paragraph, this convergence is monotonic. Hence, by Dini's theorem, we have that if $h_n(\theta) = \int V_{\beta,n}(s)q^n(ds|\theta)$, then h_n converges uniformly to $h(\theta) = \theta V_{\beta,\infty}(1) + (1 - \theta)V_{\beta,\infty}(0)$. Consequently, see the last footnote, $h_n(\theta^G(\beta, n)) \rightarrow h(\theta^G(\beta, \infty)) = h(1)$, and so

$$a = V_{\beta,n}(\theta^G(\beta, n)) \rightarrow (1 - \beta)v(\theta^G(\beta, \infty)) + \beta a < a,$$

a contradiction. It must then be that $\theta^G(\beta, \infty) < 1$, the desired result. \square

Proposition 6 *Let $\underline{\delta} = (m_H - m_L)/m_H$ and fix $\bar{\theta}_0$ with $\underline{\theta}_0 < \bar{\theta}_0 < \sup_{n \geq 2} \theta^G(\bar{\beta}, n)$. Then, for all $\delta > \underline{\delta}$, there exists $n(\delta)$ such that $U(m_L, \delta, n; \theta_0) \geq U(m_H, \delta, n; \theta_0)$ if $\theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]$ and $n \geq n(\delta)$*

Proof: Suppose that $\delta > \underline{\delta} = (m_H - m_L)/m_H$. First notice that for all $t \geq 2$, $\mu_t(m_L, n; \bar{\theta}_0) \rightarrow 1$ and $\mu_t(m_H, n; \underline{\theta}_0) \rightarrow 0$ as $n \rightarrow \infty$. Now let $\epsilon = m_L - (1 - \delta)m_H > 0$. Since $\mu_2(m_H, n, \underline{\theta}_0) \rightarrow 0$, there exists $n_1(\delta)$ such that if $n \geq n_1(\delta)$, then

$$\mu_2(m_H, n, \underline{\theta}_0) \leq \frac{\epsilon(1 - \delta)}{4m_H\delta}.$$

Moreover, since $\mu_{t+1}(m_H, n; \underline{\theta}_0) \leq \mu_t(m_H, n, \underline{\theta}_0)$ for all t and n , we have, in fact, that the above inequality holds for all $t \geq 2$. Therefore,

$$U(m_H, \delta, n; \theta_0) \leq U(m_H, \delta, n; \underline{\theta}_0) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \mu_t(m_H, n; \underline{\theta}_0) m_H < (1 - \delta)m_H + \frac{\epsilon}{4}$$

whenever $n \geq n_1(\delta)$. We now need to find an appropriate lower bound for $U(m_L, \delta, n; \theta_0)$. For this, let $n_2(\delta, t)$ be such that if $n \geq n_2(\delta, t)$, then

$$\mu_t(m_L, n; \bar{\theta}_0) > 1 - \frac{\epsilon}{4Nm_L\delta^{t-1}},$$

where N is such that $(1 - \delta^{N+1})m_L - (1 - \delta)m_H + \frac{\epsilon}{2} > 0$. We know that such an N exists by hypothesis. Now let $n_2(\delta) = \max\{n_2(\delta, t) \mid t = 2, \dots, N\}$. If $n \geq n_2(\delta)$,

$$U(m_L, \delta, n; \theta_0) \geq U(m_L, \delta, n; \bar{\theta}_0) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \mu_t(m_L, n; \bar{\theta}_0) m_L > (1 - \delta^{N+1})m_L - (1 - \delta)\frac{\epsilon}{4}.$$

Hence, if $n \geq n(\delta)$, where $n(\delta) = \max\{n_1(\delta), n_2(\delta)\}$, then,

$$U(m_L, \delta, n; \theta_0) - U(m_H, \delta, n; \theta_0) > (1 - \delta^{N+1})m_L - (1 - \delta)m_H - \frac{\epsilon}{2} > 0,$$

for all $\theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]$, and so our claim is proved. \square

B Equilibrium

Here we show that there are conditions on the distribution function. F that ensure that the maps Θ_n defined in Section 6 have, for all $n \geq 2$, a fixed point in the interval $(0, \theta^G(\bar{\beta}, n))$. The key part of the argument is to find uniform (in n) bounds on the possible cutoff discount factors for the government. We begin with a lower bound.

We know, from Section 5, that if $\delta \leq \underline{\delta}$, where $\underline{\delta} = (m_H - m_L)/m_H$, then the government always chooses m_H no matter n . In other words, no cutoff discount factor δ^* can be smaller than $\underline{\delta}$. Therefore, a lower bound for the common prior is $F(\underline{\delta})$, and for this reason we set $\underline{\theta}_0 = F(\underline{\delta})$. Our first assumption about F is that $\theta^G(\bar{\beta}, n) > F(\underline{\delta}) > 0$ for all $n \geq 2$. It implies that all agents put positive probability on the monetary regime being soft. At the same time this probability is not too high to discourage the most optimistic of them from choosing the market in the first period. Since, by Corollary 1, $\theta^G(\bar{\beta}, n)$ is increasing in n , the above condition is satisfied for all $n > 1$ as long as it is satisfied for $n = 2$.

We now find an upper bound $\bar{\delta} < 1$ for the possible cutoff discount factors that is independent of n . For this, suppose that $\theta_0 \leq \bar{\theta}_0 < \theta^G(\bar{\beta}, 2)$. For example, we can set $\bar{\theta}_0 = \theta_m$, the myopic cutoff belief. We need the following auxiliary results.

Proposition 9. *Let $\{x_{nt}\}$ be an infinite double sequence such that: (i) For each $n \in \mathbb{N}$, there exists $x_\infty(n)$ such that $\{x_{nt}\}_{t=1}^\infty$ converges monotonically to $x_\infty(n)$; (ii) There exists x_∞ such that $\lim_{n \rightarrow \infty} x_{nt} = x_\infty$ for all $t \in \mathbb{N}$; (iii) $x_\infty(n) \rightarrow x_\infty$ as $n \rightarrow \infty$. Then $\lim_{t \rightarrow \infty} \sup_{n \in \mathbb{N}} |x_{nt} - x_\infty(n)| = 0$.*

Proof: Let $\bar{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$ and consider the function $d : \bar{\mathbb{N}} \times \bar{\mathbb{N}} \rightarrow \mathbb{R}$ such that

$$d(m, n) = \left| \frac{1}{m} - \frac{1}{n} \right|,$$

where, by definition, $1/n = 0$ when $n = \infty$. Then d is a metric on $\bar{\mathbb{N}}$. Let us show that $\bar{\mathbb{N}}$, when endowed with d , is a compact metric space. Suppose $\{A_n\}_{n=0}^\infty$ is an open cover of $\bar{\mathbb{N}}$. Then, for all $i \in \bar{\mathbb{N}}$ there exists $n_i \in \mathbb{N} \cup \{0\}$ such that $\infty \in A_{n_0}$ and $i \in A_{n_i}$ for $i \in \mathbb{N}$. Since $d(m, \infty) \rightarrow 0$ as $m \rightarrow \infty$, we know there exists $m_0 \in \mathbb{N}$ such that $m \in A_{m_0}$ if $m \geq m_0$. Therefore $A_{n_0}, A_{n_1}, \dots, A_{n_{m_0}}$ is a finite subcover of $\{A_n\}$, and so $(\bar{\mathbb{N}}, d)$ is indeed a compact metric space.

Now let, for each $t \in \bar{\mathbb{N}}$, $f_t : \bar{\mathbb{N}} \rightarrow \mathbb{R}$ be such that: (i) For $t \in \mathbb{N}$, $f_t(n) = x_{nt}$ if $n \in \mathbb{N}$ and $f_t(\infty) = x_\infty$; (ii) $f_\infty(n) = x_\infty(n)$ for all $n \in \bar{\mathbb{N}}$. Since for each $t \in \bar{\mathbb{N}}$, $f_t(n)$ converges to x_∞ as $n \rightarrow \infty$, we have that the functions f_t are continuous for all $t \in \bar{\mathbb{N}}$.¹⁶ Moreover $f_t(n) \downarrow f_\infty(n)$ for each $n \in \bar{\mathbb{N}}$ as $t \rightarrow \infty$. Therefore, $\{f_t\}_{t=1}^\infty$ is a sequence of continuous functions that converges pointwise on a compact metric space to a continuous function f_∞ . By Dini's theorem, see Rudin (1976), we then have that $\{f_t\}$ converges uniformly to f_∞ . In other words, $\sup_{n \in \bar{\mathbb{N}}} |f_t(n) - f_\infty(n)| = \sup_{n \in \bar{\mathbb{N}}} |x_{nt} - x_\infty(n)| \rightarrow 0$ as $t \rightarrow \infty$, which implies the desired result. \square

¹⁶The continuity of the functions f_t at the other points of $\bar{\mathbb{N}}$ is immediate.

Proposition 10. *The following two facts hold:*

- (i) *There exists $\gamma > 0$ such that $U(m_L, \delta, n; \theta_0) \geq \gamma$ for all $\theta \in [\underline{\theta}_0, \bar{\theta}_0]$ and all $n > 1$;*
- (ii) *$\lim_{\delta \rightarrow 1} \sup \{U(m_H, \delta, n; \theta_0) \mid n > 1 \text{ and } \theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]\} = 0$.*

Proof: Write $\eta_t(m, \beta, n; \theta_0)$ to denote the dependence of the probabilities η_t of Section 3 on both n and the prior θ_0 . Because η_t is non-increasing in θ_0 and non-decreasing in β , we have that

$$\mu_t(m_L, n; \theta_0) \geq \int \eta_t(m_L, \beta, n; \bar{\theta}_0) g(\beta) d\beta \geq \int \eta_t(m_L, \beta', n; \bar{\theta}_0) g(\beta) d\beta = \alpha \eta_t(m_L, \beta', n; \bar{\theta}_0),$$

where $\beta' = \inf\{\beta : \theta^G(\beta, 2) \geq \bar{\theta}_0\} < \bar{\beta}$ and $\alpha = \int_{[\beta', \bar{\beta}]} g(\beta) d\beta > 0$. Since, from Section 3, $\inf_t \eta_t(m_L, \beta', n, \bar{\theta}_0) = \eta_\infty(n) > 0$ for all $n \geq 2$, we then have that $\mu_t(m_L, n; \theta_0) \geq \alpha \eta_\infty(n) > 0$ for all $t \in \mathbb{N}$ and all $n \geq 2$. Notice that this lower bound depends on our choice of $\bar{\theta}_0$, but the latter can be made (and indeed is) independent of n . Moreover, $\eta = \inf_{n \geq 2} \eta_\infty(n) > 0$.¹⁷ Therefore, for all $\theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]$ and all $\delta \in [0, 1)$, we have that

$$U(m_L, \delta, n; \theta_0) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \mu_t(m_L, n, \theta_0) \geq \gamma > 0,$$

where $\gamma = \alpha \eta$ is independent of n . The same lower bound holds for $U(m_L, 1)$ for the reasons pointed at the end of Section 2.

Now observe that $\mu_t(m_H, n; \theta_0) \leq \eta_t(m_H, \bar{\beta}, n; F(\underline{\delta}))$ for all $t \in \mathbb{N}$, $n \geq 2$, and $\theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]$; that is, the fraction of the population entering the market in period t when $m = m_H$ is not bigger than the corresponding probability for the most initially optimistic and patient agents. Consequently,

$$\begin{aligned} U(m_H, n, \delta; \theta_0) &\leq \sup_{n \geq 2} \left\{ (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \eta_t(m_H, \bar{\beta}, n; F(\underline{\delta})) m_H \right\} \\ &\leq (1 - \delta) \left\{ \sum_{t=1}^{\infty} \delta^{t-1} \sup_{n \geq 2} \eta_t(m_H, \bar{\beta}, n; F(\underline{\delta})) m_H \right\} \end{aligned}$$

for all $\theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]$ and all $n \geq 2$. By Corollary 1, we know that $\sup_{n \geq 2} \theta^G(\bar{\beta}, n) < 1$, and so, since $F(\underline{\delta}) > 0$, $\lim_n \eta_t(m_H, \bar{\beta}, n; F(\underline{\delta})) = 0$ for all $t \in \mathbb{N}$. Therefore, as a consequence of Lemma 9, we have that $\sup_{n \geq 2} \eta_t(m_H, \bar{\beta}, n; F(\underline{\delta}))$ converges to zero as $t \rightarrow \infty$. Hence, the right-hand side of the above inequality converges to zero as δ goes to one. We can then conclude that $U(m_H, n, \delta; \theta_0)$ converges to zero as $\delta \rightarrow 1$ uniformly in n and θ_0 , the desired result. \square

This last lemma implies that there exists $1 > \bar{\delta} > \underline{\delta}$ such that if $\delta \geq \bar{\delta}$, then $U(m_L, n, \delta; \theta_0) \geq U(m_H, n, \delta; \theta_0)$ for all $\theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]$ and all $n > 1$. In other words, $\bar{\delta}$ is the desired upper bound for

¹⁷This is a consequence of the fact that there exists $n_0 \in \mathbb{N}$ such that $\eta_\infty(n) > \eta_\infty(2)$ for all $n \geq n_0$.

the possible cutoff discount factors for the government. Notice that $\bar{\delta}$ is independent of n and the c.d.f. F . The latter is true since the cutoff beliefs $\theta^G(\beta, n)$ are independent of F . This brings us to our second assumption about F , namely, that $F(\bar{\delta}) \leq \bar{\theta}_0$.

Summarizing, the only possible cutoff discount factors for the government lie in the interval $[\underline{\delta}, \bar{\delta}]$, no matter the number $n > 1$ of market meetings. The two assumptions made above ensure that if Θ_n has a fixed point, then it must lie in $[\underline{\theta}_0, \bar{\theta}_0] \subset (0, \theta^G(\bar{\beta}, 2))$. The final result in this appendix shows that this is indeed the case.

Proposition 11. *For all $n \geq 2$, the maps Θ_n have a fixed point.*

We need the following result before we can prove the above lemma.

Proposition 12. *For each $n \geq 2$, $U(m, n, \delta; \theta_0)$ is jointly continuous in δ and θ_0 if for all $t \in \mathbb{N}$, $\mu_t(m, n; \theta_0)$ is a continuous function of θ_0 .*

Proof: Suppose that for all $t \in \mathbb{N}$, $\mu_t(m, n; \theta_0)$ is a continuous function of θ_0 . Then $U(m, n, \delta; \theta_0)$ is, for all $\delta \in [0, 1]$, a continuous function of θ_0 . Since each μ_t is monotonic in θ_0 , U is monotonic in θ_0 as well. Moreover, for all $\theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]$, U is continuous in δ . We can then apply Lemma 2 in Dutta et. al. (1994) to conclude that U is indeed jointly continuous in θ_0 and δ . \square

Proof of Lemma 11: We know that for each $\theta_0 \in [\underline{\theta}_0, \bar{\theta}_0]$ and $n \geq 2$, $\delta^*(\theta_0, n)$ is the unique solution to the equation $U(m_L, \delta, n; \theta_0) = U(m_H, \delta, n; \theta_0)$. Hence, if we show that both $U(m_L, n, \delta; \theta_0)$ and $U(m_H, n, \delta; \theta_0)$ are jointly continuous in δ and θ_0 for all $n \geq 2$, the same is true of $\delta^*(\theta_0, n)$. Because the c.d.f. F associated with f is a continuous function as well, this implies that the maps Θ_n are all continuous. The existence of a fixed point for each one of them then follows from Brouwer's fixed point theorem together with the fact that they map the interval $[\underline{\theta}_0, \bar{\theta}_0]$ into itself.

By Lemma 12, the only thing we need to do is show that for each $t \in \mathbb{N}$ and $n \geq 2$, both $\mu_t(m_L, n; \theta_0)$ and $\mu_t(m_H, n; \theta_0)$ are continuous functions of θ_0 . For this, suppose $\{\theta_0^k\}$ is an infinite sequence such that $\theta_0^k \rightarrow \theta_0$. This implies that for each t and $n \geq 2$, $\eta_t(m, \beta, n; \theta_0^k) \rightarrow \eta_t(m, \beta, n; \theta_0)$ for almost all $\beta \in [\underline{\beta}, \bar{\beta}]$.¹⁸ By Egorov's theorem, see Kolmogorov and Fomin (1970), almost everywhere pointwise convergence implies convergence in L_1 . Therefore

$$\mu_t(m, n; \theta_0^k) = \int \eta_t(m, \beta, n; \theta_0^k) g(\beta) d\beta \rightarrow \mu_t(m, n; \theta_0) = \int \eta_t(m, \beta, n; \theta_0) g(\beta) d\beta$$

for all t and $n \geq 2$, and so we have the desired continuity result. \square

¹⁸We know the fractions $\eta_t(m, \beta, n; \theta_0)$ are not necessarily continuous in θ_0 , as there may be one or more private histories that, given θ_0 , leave an agent indifferent between the market and autarky for at least one period $t' \leq t$. If that is the case, perturbing θ_0 induces a discontinuous change in this fraction. However, for each t and n , the number of priors θ_0 at which this can happen is finite, and so of Lebesgue measure zero.

References

- [1] L. Araujo and B. Camargo, “Private Supply of Money and Time Consistency”, Unpublished Manuscript, 2005.
- [2] L. Araujo and A. Shevchenko, “Price Dispersion, Information and Learning”, Unpublished Manuscript, 2004.
- [3] D. Blackwell, “Discounted Dynamic Programming”, *Annals of Mathematical Statistics*, 36 (1965), 226-235.
- [4] J. Banks and R. Sundaram, “Denumerable-Armed Bandits”, *Econometrica* 60 (1992), 1071-1096.
- [5] P. Dutta, M. Majumdar, and R. Sundaram, “Parametric Continuity in Dynamic Programming Problems”, *Journal of Economic Dynamics and Control*, 18 (1994), 1069-1092.
- [6] D. Easley and N. Kiefer, “Controlling a Stochastic Process with Unknown Parameters”, *Econometrica*, 56 (1988), 1045-1064.
- [7] M. Friedman and A. Schwartz, “Has Government Any Role in Money?”, *Journal of Monetary Economics*, 17 (1986), 37-62.
- [8] J. K. Galbraith, *Money, Whence It Came, Where It Went*, Houghton Mifflin Company, Boston, 1975.
- [9] M. DeGroot, *Optimal Statistical Decisions*, McGraw-Hill, New York, 1970.
- [10] R. Kakigi, “A Note on Discounted Future Two-Armed Bandits”, *The Annals of Statistics*, 11 (1983), 707-711.
- [11] 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.
- [12] R. King, “On the Economics of Private Money”, *Journal of Monetary Economics*, 12 (1983), 127-158.
- [13] N. Kiyotaki and R. Wright, “On Money as a Medium of Exchange”, *Journal of Political Economy* 97 (1989), 927-954.
- [14] N. Kiyotaki and R. Wright, “A Search-Theoretic Approach to Monetary Economics”, *American Economic Review* 83 (1993), 63-77.

- [15] A.N. Kolmogorov and S.V. Fomin, *Introductory Real Analysis*, Dover, New York, 1970.
- [16] 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.
- [17] B. Norwood, “On the Emergence of Money as a Medium of Exchange”, Unpublished Manuscript, 2003.
- [18] J. Ritter, “The Transition from Barter to Fiat-Money”, *American Economic Review* 85 (1995), 134-149.
- [19] W. Rudin, *Principles of Mathematical Analysis*, Third Edition, McGraw-Hill, Singapore, 1976.
- [20] Y. Sik-Kim, “Liquidity, Prices, Seigniorage, and the Transition from Barter to Fiat Money”, *Macroeconomic Dynamics*, 5 (2001).
- [21] G. Tullock, “Paper Money-A Cycle in Cathay”, *Economic History Review* IX (1957), 393-407.
- [22] N. Wallace, “Short-run and Long-run Effects of Changes in Money in a Random Matching Model”, *Journal of Political Economy* 105 (1997), 1293-1307.
- [23] A. Wolinsky, “Information Revelation in a Market with Pairwise Meetings”, *Econometrica* 58 (1990), 1-23.
- [24] L-S Yang, *Money and Credit in China: A Short History*, Harvard University Press, Cambridge, MA, 1952.