

The Simple Analytics of Money and Credit in a Quasi-linear Environment*

David Andolfatto[†]

April 2008

Abstract

Lagos and Wright (2005) demonstrate how the essential properties of a money-search model are preserved in an environment that is rendered highly tractable with the use of quasi-linear preferences. In this paper, I show that this same innovation can be applied to closely related environments used elsewhere in the literature that study insurance and credit markets under limited commitment and private information. The analysis demonstrates clearly how insurance, credit, and money are interrelated in terms of their basic functions. The analysis also leads to a heretofore neglected result pertaining to the Friedman rule. In particular, I find that the same frictions that render money essential may at the same time operate to render the Friedman rule infeasible. Thus, even if the Friedman rule is a desirable policy, an incentive-induced lower bound on the rate of deflation may nevertheless entail a strictly positive rate of inflation.

1 Introduction

Microfounded models of money are appealing because they explain, rather than assume, the societal benefits of monetary exchange. For the most part, however, these models are difficult to study analytically; at least, in versions that do not impose restrictive assumptions on individual money holdings. An innovation in this regard is Lagos and Wright (2005), who exploit a quasi-linear preference structure to eliminate the distribution of money holdings as a state variable. The analytical tractability of their framework has led to a rapid proliferation of extensions and applications in the literature.

In this paper, I take one step back; rather than extending the original Lagos-Wright (LW) model, I choose to simplify it further still. The primary simplification is to dispense with the search-theoretic aspects of their environment.

*I thank Mei Dong, Janet Hua Jiang, Narayana Kocherlakota, Francesco Lippi, Ricardo Lagos, Randall Wright, and seminar participants at the Einaudi Center for Economics and Finance (Rome) for their many useful comments.

[†]Simon Fraser University and the Rimini Center for Economic Analysis. Email: dan-dolfa@sfu.ca

The idiosyncratic uncertainty and lack of double coincidence induced by random pairwise meetings in their “decentralized” market are replaced by random shocks to individual preferences and abilities operating on agents who meet at centralized locations. I also dispense with a needless complication in their “centralized” market and simply assume that utility is transferable; which is, after all, the gist of their quasi-linear preference structure.

One benefit of this exercise is that it renders clear the essential properties of the LW model. Fundamentally, it is an environment where agents value insurance. Moreover, it is an environment where agents can self-insure by trading through a sequence of debt markets. When private credit is unavailable, a monetary instrument serves as a substitute for the missing debt instrument.

Another benefit of this exercise is that it allows one to see clearly how their framework maps into an existing literature that studies insurance and credit arrangements under private information and limited commitment. For example, Atkeson and Lucas (1992) study an optimal insurance arrangement in a dynamic model with private information and full commitment. In the context of the LW environment, this problem is rendered highly tractable; and I demonstrate how a quasi-linear mechanism can implement the first-best allocation. Moreover, under full commitment, I show that the first-best allocation can also be implemented as the equilibrium of a sequence of competitive debt markets.

Next, I dispense with the assumption of full commitment. In the spirit of Kehoe and Levine (1993), I assume that agents can be banished from an existing trading arrangement; the threat of which serves to motivate voluntary debt repayment. This type of penal code relies on the availability of what Kocherlakota (1998) has labeled “societal memory;” i.e., a public-access database containing the records of all individual trading histories. Under quasi-linear preferences, the amount of societal memory needed to implement a constrained-efficient allocation is significantly reduced; once again, the analysis is rendered highly tractable. Here, I demonstrate that the first-best allocation is implementable when agents are sufficiently patient; and if agents are sufficiently impatient, some agents are effectively debt-constrained. I also demonstrate the conditions under which a constrained-efficient allocation can be implemented as the equilibrium of a sequence of competitive debt markets.

Finally, I demonstrate what is by now a well-known result that absent societal memory a fiat money instrument is essential. Here, I show that with quasi-linear preferences, a monetary mechanism can implement any incentive-feasible allocation available under memory. That is, the first-best allocation is implementable without intervention (a constant money supply); which is a result similar to that derived by Hu, Kennan, and Wallace (2007). When I further restrict trades to occur on a sequence of competitive money-goods markets, implementation of a constrained-efficient allocation requires intervention (either a contracting or expanding money supply). The efficient monetary equilibria here correspond to the allocations achievable under a system of competitive debt markets when memory is available.

While much of the paper’s contribution can be viewed as pedagogical in nature, this last result provides—as far as I can tell— a novel insight pertaining to the Friedman rule. By construction, the environment I consider here is one in which the Friedman rule is a desirable policy.¹ Typically, the way this policy is implemented is with lump-sum taxes. The problem with this policy prescription is that it assumes that agents can, in effect, commit to paying a tax obligation; and it is precisely the absence of this form of commitment that renders money essential in the first place. I resolve this inconsistency by treating all debt obligations symmetrically; i.e., agents must be induced (rather than coerced) to comply with their promises. I then demonstrate that while the Friedman rule is be a desirable policy, it may not be an incentive-feasible policy. To put things another way, there may be an incentive-induced lower bound on the rate at which a monetary authority may deflate. In fact, the only constrained-efficient policy may entail a strictly positive inflation rate. The result bears a direct relation to why consumers may face a binding debt constraint in a market for private debt.

2 The Physical Environment

The economy is populated by a continuum of *ex ante* identical agents, distributed uniformly on the unit interval. Each period $t = 0, 1, 2, \dots, \infty$ is divided into two subperiods, labeled *day* and *night*. Agents meet at a central location in both subperiods; in particular, I abstract from the commonly employed assumption of random pairwise meetings in one of the subperiods.

All agents have common preferences and abilities during the day. Let $x_t(i) \in \mathbb{R}$ denote the consumption (production, if negative) of output in the day by agent i at date t . The key simplifying assumption is that preferences are linear in this term. The possibility of exchange then implies transferable utility. Output produced in the day is nonstorable, so an aggregate resource constraint implies:

$$\int x_t(i) di \leq 0; \tag{1}$$

for all $t \geq 0$.

At night, agents realize a shock that determines their type for the night. In particular, agents either have a desire to consume, an ability to produce, or neither. Refer to these types as *consumers*, *producers*, and *nonparticipants*, respectively. Types are determined randomly by an exogenous stochastic process. This process is *i.i.d.* across agents and time; there is no aggregate uncertainty. Let $\pi \in (0, 1/2)$ denote the measure of agents who become either consumers or producers; so that $(1 - 2\pi)$ denotes the measure of nonparticipants.²

¹This is in contrast to environments in which it is not; for example, Levine (1991), Deviatov and Wallace (2001), Molico (2006) and Williamson (2006).

²At the individual level, these measures represent probabilities.

A consumer has utility $u(c)$ and a producer has utility $-g(y)$; where $c \in \mathbb{R}_+$ and $y \in \mathbb{R}_+$ denote consumption and production of the night good, respectively. Assume that $u'' < 0 < u'$, $u(0) = 0$ and $g', g'' \geq 0$ with $g(0) = g'(0) = 0$. Nonparticipants neither value consume nor are they able to produce it; their utility is normalized to zero. As the night good is also nonstorable, there is another aggregate resource constraint given by:

$$\int c_t(i) di \leq \int y_t(i) di; \quad (2)$$

for all $t \geq 0$.

As agents are *ex ante* identical, their preferences can be represented as:

$$E_0 \sum_{t=0}^{\infty} \beta^t \{x_t(i) + \pi [u(c_t(i)) - g(y_t(i))]\}. \quad (3)$$

where $0 < \beta < 1$. Note that there is no discounting across subperiods.

Weighting all agents equally, a planner maximizes (3) subject to the resource constraints (1) and (2). As utility is linear in x , feasibility implies that agents are indifferent across any lottery over $\{x_t(i) : t \geq 0\}$ that delivers $E_0 x_t(i) = 0$. Anticipating what is to follow later, one such lottery is given by:

$$x_t(i) = \begin{cases} +x & \text{w.p. } \pi; \\ 0 & \text{w.p. } 1 - 2\pi; \\ -x & \text{w.p. } \pi; \end{cases} \quad (4)$$

for any $x \geq 0$.

Consider next how output is allocated at night. If g is strictly convex, all producers will be required to produce a common level of output $y \geq 0$.³ Given the strict concavity of u , all consumers will be allocated a common level of consumption $c \geq 0$. Given that the active population is divided equally among producers and consumers at night, the resource constraint (2) implies $c = y$. Hence, conditional on a given level of y (and invoking the fact that $E x_t(i) = 0$), *ex ante* welfare is represented by:

$$W(y) = \left(\frac{\pi}{1 - \beta} \right) [u(y) - g(y)]. \quad (5)$$

Clearly, $W(0) = W(\bar{y}) = 0$ for some unique $0 < \bar{y} < \infty$. Moreover, there exists a unique maximizer $y^* \in (0, \bar{y})$ characterized by:

$$u'(y^*) = g'(y^*). \quad (6)$$

In what follows, I refer to (x, y^*) as the *first-best* allocation; where x should be understood to satisfy a lottery in the form of (4).

³If g is linear, then y can be interpreted as an expected level of output.

As far as a social planner is concerned, the day subperiod is irrelevant; one may without loss restrict attention to the first-best allocation $(0, y^*)$. The pattern of trades that supports this allocation entails some form of social insurance. Specifically, agents face the risk of wanting consumption with no means of producing it. The solution entails having those agents with a contemporaneous ability to produce to satisfy those members of society with a contemporaneous want. Alternatively, the planner's solution may also be interpreted as a type of social credit system; with agents borrowing resources from society when they have a desire to consume and promising to discharge their debt to society when they have an ability to produce.

Of course, the first-best allocation can be decentralized as a competitive equilibrium. One obvious market structure that achieves this is a simple (static) contingent-claims market that opens at the beginning of every night (prior to the realization of types). Agents purchase claims to consumption redeemable in the event of desire to consumption; and finance the purchase of this insurance policy by issuing claims against output redeemable in the event that production is possible.

3 Private Information

Assume now that individual types are private information; but that agents retain access to a commitment technology. In a static version of this model, this informational asymmetry is sufficiently severe to render autarky the only implementable allocation. In a dynamic setting, however, more desirable allocations can be implemented by relying on the information contained in personal trading histories. In general, it is desirable to make use of an agent's entire trading history (or history of reports, in the case of direct mechanisms). It is a convenient property of this quasi-linear environment that one may without loss truncate histories to include actions observed (or reported) in the most recent past. Hence, in what follows, I restrict attention to (stationary) allocations (x, y) of the form:

$$x_t(i) = \begin{cases} +x & \text{if } (c_{t-1}(i), y_{t-1}(i)) = (0, y); \\ 0 & \text{if } (c_{t-1}(i), y_{t-1}(i)) = (0, 0); \\ -x & \text{if } (c_{t-1}(i), y_{t-1}(i)) = (y, 0); \end{cases} \quad (7)$$

with $x_0(i) = 0$ for all i . If agents reveal their types truthfully (either directly or indirectly), then any allocation (x, y) given by (7) is feasible and generates the *ex ante* welfare function (5); in particular, note that $E_t x_t(i) = 0$ for all i and all t .

In the mechanism considered here, agents report their types indirectly by submitting either a claim for consumption or by displaying a level of production (for deposit or sale); if they do neither, they indirectly represent themselves as nonparticipants. As consumers and nonparticipants are technologi-

cally restricted from producing output, they have no capacity to misrepresent themselves as producers. On the other hand, producers have an ability to misrepresent themselves as either consumers or nonparticipants; and, of course, consumers and nonparticipants may attempt to mimic each other.

For a producer, an allocation (x, y) is incentive-compatible (IC) if it satisfies:

$$-g(y) + \beta [x + W(y)] \geq \max \{ \beta [-x + W(y)], \beta W(y) \};$$

or, given that $x \geq 0$,

$$x \geq \beta^{-1}g(y). \tag{8}$$

That is, the promised future punishment/reward (x) must be sufficiently large to compensate the producer for the utility expense associated with producing output in the amount of y .

For a consumer, an allocation (x, y) is IC if it satisfies:

$$u(y) + \beta [-x + W(y)] \geq \beta W(y);$$

or,

$$\beta^{-1}u(y) \geq x. \tag{9}$$

That is, the consumer's future debt obligation cannot be so high as to discourage him from revealing his true type at night (he could, as an alternative, misrepresent himself as a nonparticipant).

It should be apparent that a nonparticipant will not want to misreport himself as a consumer. Hence, the following result is immediately apparent.

Result 1 *The allocation (x, y^*) with any $\beta^{-1}u(y^*) \geq x \geq \beta^{-1}g(y^*)$ is implementable when types are private information and when agents can commit.*

That the private information friction can be circumvented entirely relies heavily on the quasi-linear structure of preferences. In general, the first-best allocation is not implementable and the distribution of promised utilities widens over time; see Atkeson and Lucas (1992). That is, when utility is nontransferable, it is optimal to smooth punishments and rewards over time; whereas here one may, without any *ex ante* utility expense, discharge punishments and rewards fully on a period-by-period basis.

3.1 A Market Mechanism

There is also a decentralized solution to the resource allocation problem when types are private information. Of course, the simple contingent-claims market structure described earlier will not work here (producers would strongly prefer to misrepresent themselves as nonparticipants). As with the "centralized"

mechanism described above, efficiency can be enhanced by introducing a form of history-dependence in allocations. The way this can be done here is with the following market structure. Assume that a competitive debt market is available every night (opening subsequent to the realization of types). Debt issued at night is redeemed the next day at a competitively-determined (gross) real interest rate $R > 0$.

The quasi-linear preference structure allows one to cast decision-making in terms of a sequence of two-period (from one night to the next day) problems. For a consumer, the choice problem can be formulated as:

$$\max \{u(c) - \beta x : c \leq R^{-1}x\}.$$

Consumer demand is characterized by:

$$u'(c^D) = R\beta;$$

with the supply of future (day) output given by $x^S = Rc^D$. For a producer, the choice problem can be formulated as:

$$\max \{-g(y) + \beta x : x \leq Ry\}.$$

The supply of output is characterized by:

$$g'(y^S) = R\beta;$$

with the demand for future (day) output given by $x^D = Ry^S$. The equilibrium allocation (x^*, y^*) and price-system R^* satisfies:

$$u'(y^*) = g'(y^*); \quad R^* = \beta^{-1}g'(y^*); \quad x^* = R^*y^*. \quad (10)$$

Result 2 *The competitive debt market equilibrium (10) implements the first-best allocation when types are private information and when agents can commit.*

Despite the absence of an insurance market, agents are nevertheless able to insure themselves fully through borrowing and lending. The equilibrium distribution of net asset positions is endogenous, but can be derived recursively from equilibrium behavior (this is to say that behavior and prices do not depend on the distribution of assets). At the beginning of the day, there is a three-point distribution of assets (corresponding to the three possible types realized at night); and at the end of the day, the distribution of assets collapses over zero. The simplicity afforded by this recursive structure and the fact that the first-best is implementable are both results that rely heavily on the quasi-linear structure of preferences.

It is also interesting to note how this competitive market for “short-term” debt evidently replicates the mechanism’s ability to elicit (indirectly) truthful

reports. This must be true of any *ex post* competitive spot market; buyers and sellers that gather together for the purpose of exchange have little incentive or ability to misrepresent themselves on the spot (and at terms of trade that they view beyond their control). Producers eager to make a sale will reveal themselves by displaying their wares; and consumers eager to make a purchase can do no better than demonstrate their wants at the going price. This is in contrast to the *ex ante* contingent-claims market described above where, *ex post*, producers would find it optimal to mimic nonparticipants.

Finally, note that both the centralized and decentralized mechanisms described above make use of individual trading histories (the commitment to honor a debt presumably relies on, among other things, the ability to demonstrate evidence of its past occurrence). As such, one might say that circumventing the problem of private information requires the use of memory; although the way this term is employed here differs from the concept of *societal memory* emphasized by Kocherlakota (1998). That is, given the commitment power assumed here, societal memory (a public database rendering all private trading histories observable to society) is not necessary. All that is required is *private memory* in the sense that a creditor need only remember his own debtor(s).

4 Limited Commitment

Assume now that agents lack commitment so that individual behavior is restricted to be *sequential rational* (SR). I assume that society has the power to banish individuals to a state of perpetual autarky.

For an indirect mechanism recommending allocation (x, y) , the producer is obliged to deliver y units of output at night. This obligation will be honored if it is sequential rational to do so; i.e.,

$$-g(y) + \beta [x + W(y)] \geq 0. \quad (11)$$

Likewise, a consumer is obliged to deliver x units of output in the day. It will be sequentially rational to honor this obligation if:

$$-x + W(y) \geq 0. \quad (12)$$

Observation A If an allocation (x, y) satisfies (12) and (8), then (11) is necessarily satisfied.

By this observation, any allocation that is IC for the producer is also SR for the producer. In what follows then, we can ignore this latter constraint and focus on the consumer's IC and SR constraints and the producer's IC constraint (IC and SR are trivially satisfied for the nonparticipant). The set of incentive-feasible allocations is therefore described by:

$$F \equiv \{(x, y) \in \mathbb{R} \times \mathbb{R}_+ : \min\{W(y), \beta^{-1}u(y)\} \geq x \geq \beta^{-1}g(y)\}. \quad (13)$$

The set F is clearly non-empty, convex, and compact. Moreover, with an additional mild restriction on preferences, the set will include a strictly positive y .

It follows as a corollary that the problem of choosing $(x, y) \in F$ to maximize $W(y)$ is well-defined. Moreover, as $W(y)$ is strictly concave, there is a unique solution $y > 0$. Associated with this solution is an x (not necessarily unique) satisfying (13). The exact nature of the solution depends on parameters; and in particular, on the discount factor β .

Define β^* as the solution to:

$$\left(\frac{\pi}{1-\beta^*}\right) [u(y^*) - g(y^*)] = \left(\frac{1}{\beta^*}\right) g(y^*); \quad (14)$$

where clearly, $\beta^* \in (0, 1)$. That is, for $\beta = \beta^*$, the SR and IC constraints for the consumer and producer (respectively) just bind at the first-best allocation. As well, the consumer's IC constraint is clearly satisfied, since $u(y^*) > g(y^*)$.

As the LHS (RHS) of (14) is increasing (decreasing) in β , it follows that the first-best continues to satisfy consumer SR and producer IC for all $\beta \in [\beta^*, 1)$. Moreover, as $u(y^*) > g(y^*)$ is independent of β , it follows that consumer IC remains satisfied as well.

Result 3 *If $\beta \in [\beta^*, 1)$, then the allocation (x, y^*) with any x satisfying (13) is implementable when types are private information and agents lack commitment.*

Result 3 is stated as a sufficient condition, but it is clearly necessary as well. It follows as a corollary that when agents are sufficiently impatient, the first-best is not implementable. To characterize the constrained-efficient allocation, we need to know which constraints—consumer SR or consumer IC—bind as $\beta \rightarrow 0$. As it turns out, the relevant constraint for the consumer is SR and not consumer IC.

Result 4 *If $\beta \in (0, \beta^*)$, then the constrained-efficient allocation (x_0, y_0) is characterized by $0 < y_0 < y^*$ satisfying $W(y_0) = \beta^{-1}g(y_0) = x_0$.*

In any constrained-efficient allocation then, the allocation will be determined by consumer SR and producer IC; that is, consumer IC will remain slack. To prove this result, note any constrained-efficient allocation $0 < y_0 < y^*$ must satisfy $\min\{W(y_0), \beta^{-1}u(y_0)\} = \beta^{-1}g(y_0)$. The claim is that $W(y_0) = \beta^{-1}g(y_0)$. Suppose to the contrary that $\beta^{-1}u(y_0) = \beta^{-1}g(y_0)$. By the properties of u and g , this must then imply that $0 < y^* < y_0$; which is a contradiction. In what follows, I will refer to the constrained-efficient allocation (x_0, y_0) as *second-best*.

4.1 A Market Mechanism

I demonstrated earlier that under private information over types and full commitment, restricting the mechanism to deliver allocations according to a linear price-system is in no way constraining. In general, this will no longer be the case under limited commitment.

Assume now that debt contracts must be self-enforcing. By Observation A, we already know that IC is stronger than SR for the producer; and as IC is satisfied in a competitive spot market, we can ignore SR for the producer. Hence, the analysis rests on whether SR is satisfied for the consumer; recall condition (12). As with the non-market mechanism, this constraint will remain slack if agents are sufficiently patient; i.e., if $\beta \geq \beta^0$ for some $\beta^0 \in (0, 1)$. For $\beta < \beta^0$, consumers will be debt-constrained. A characterization of β^0 now follows.

To begin, note that any competitive equilibrium allocation (x, y) must lie on the locus:

$$x = \beta^{-1} g'(y)y; \quad (15)$$

this will be true whether or not consumers are debt-constrained.⁴ Next, consumer SR requires that $W(y) \geq x$; which when combined with (15) implies:

$$\left(\frac{\pi}{1-\beta}\right) [u(y) - g(y)] \geq \left(\frac{1}{\beta}\right) g'(y)y. \quad (16)$$

Clearly, there is a unique value $\beta^0 \in (0, 1)$ satisfying:

$$\left(\frac{\pi}{1-\beta^0}\right) [u(y^*) - g(y^*)] = \left(\frac{1}{\beta^0}\right) g'(y^*)y^*. \quad (17)$$

Together, conditions (16) and (17) imply the following result:

Result 5 *If $\beta \in [\beta^0, 1)$, then the competitive equilibrium (x^*, y^*, R^*) is sequentially rational (consumer debt-constraint remains slack) and incentive-compatible.*

On the other hand, if agents are sufficiently impatient, then the first-best is not implementable.

Result 6 *If $\beta \in (0, \beta^0)$, then the consumer debt-constraint binds and the competitive equilibrium allocation (x_1, y_1) and price-system R_1 is characterized by:*

$$x_1 = \left(\frac{\pi}{1-\beta}\right) [u(y_1) - g(y_1)] = \left(\frac{1}{\beta}\right) g'(y_1)y_1; \quad (18)$$

$$R_1 = \beta^{-1} g'(y_1); \quad (19)$$

⁴In case the debt-constraint binds, producers will be supplying all they desire at the going interest rate, but consumers will not be borrowing as much as they desire.

where $0 < y_1 < y^*$, $0 < x_1 < x^*$, $R_1 < R^*$, and $u'(y_1) > g'(y_1)$.

It remains to ascertain the relationship between β^* and β^0 . As it turns out, one may easily establish the following result:

Result 7 *If $g'' > 0$, then $\beta^* < \beta^0$; and if $g'' = 0$, then $\beta^* = \beta^0$.*

To see this, simply compare (17) with (14) while noting that the convexity of g implies $g'(y)y \geq g(y)$. Hence, the restriction to a linear price-system is not constraining only in the special case for which g is linear. If g is strictly convex, then the non-market mechanism strictly dominates the market mechanism for all $\beta \in (0, \beta^0)$; and indeed, the non-market mechanism can implement the first-best allocation for all $\beta \in [\beta^*, \beta^0)$ while a market mechanism cannot. This result should not be surprising, as the non-market mechanism is free to adopt a non-linear pricing scheme.⁵ In what follows, I refer to the market allocation when $\beta \in (0, \beta^0)$ as *third-best*.

To conclude this section, it is worth emphasizing that implementation here relies on *societal* memory; as defined in Kocherlakota (1998). That is, it assumes that each individual's trading history is observed by society; or, at least, by whatever mechanism that is acting on its behalf. The necessity of societal memory stems from the lack of commitment (and not the private information) together with a penal code that relies on the ability to ostracize individual members from a societal trading arrangement in the event of noncompliance.

5 Money

The term *money* should be understood here to mean a fiat token that circulates as a means of payment. The issue of how an intrinsically useless object might come to possess exchange value is a central question in the theory of money. A related—and somewhat deeper—question pertains to identifying the circumstances under which an intrinsically useless token might be *necessary* for promoting an efficient trading arrangement.

An answer to the first question—which is a matter of sufficiency—is relatively straightforward. Consider, for example, a *monetary mechanism* that creates and issues one indivisible, durable, and non-counterfeitable token to each member of society. Let $z_t(i)$ denote the number of tokens held by an agent i at the beginning of the day; and let $m_t(i)$ denote the number of tokens held entering the night. While it is conceivable that agents might accumulate any number of tokens, an educated guess suggests that in equilibrium, $m_t(i) = 1$ and $z_t(i) \in \{0, 1, 2\}$ for all i and all t .

⁵Whether an allocation supported by non-linear prices survives *ex post* renegotiation here is an interesting question but one that is not pursued here.

Now, consider some stationary allocation (x, y) . The monetary mechanism issues the following instructions. At night, subsequent to the realization of types, but prior to trade, the mechanism asks each agent to display their money holdings. If $m_t(i) = 0$ (or if money is not displayed) then the agent is prohibited from trading in the contemporaneous subperiod. Each producer is asked to deposit y units of output and each consumer is asked to deposit one unit of money; consumers receive the output and producers receive the money (nonparticipants remain idle). Hence, the relevant choice for an agent who enters the night with $m_t(i)$ units of money is over the following three options:

$$(c_t(i), y_t(i), z_{t+1}(i)) \in \{(0, y, m_t(i) + 1), (0, 0, m_t(i)), (y, 0, m_t(i) - 1)\}. \quad (20)$$

At the beginning of each day, the monetary mechanism intermediates an exchange of output for tokens. In particular, x units of output can be purchased with one token; with the supplier of output receiving the token. Hence,

$$x_t(i) = \begin{cases} +x & \text{if } m_t(i) = z_t(i) - 1; \\ 0 & \text{if } m_t(i) = z_t(i); \\ -x & \text{if } m_t(i) = z_t(i) + 1. \end{cases} \quad (21)$$

Note that the monetary mechanism described above makes no direct reference to any individual's personal trading history; although an indirect reference is made to the extent that contemporaneous money holdings reveal something about past behavior. It should be clear enough that the following result is true.

Result 8 *The monetary mechanism defined by (20) and (21) can implement any incentive-feasible allocation $(x, y) \in F$.*

To see this, note that all agents enter the initial night with one unit of money. Nonparticipants have no desire to consume and no ability to produce, so they leave the night with one unit of money. Consumers want output and, by incentive-feasibility, are willing to acquire it in exchange for an implicit obligation to deliver x the next day. Hence, consumers purchase output by surrendering their token; they leave the night with zero units of money. The next day, consumers are willing (again, by incentive-feasibility) to produce x in exchange for a token; recall that the failure to acquire this token prevents them from participating in the night market.⁶ Finally, since the allocation is incentive-feasible, producers are willing to deliver output at night in exchange for an implicit promise to x the next day; the latter which may be purchased with the additional money token acquired at night. At the end of the day, all agents are once again reduced to holding one unit of money.

⁶Implicitly, I assume here that if it is rational to default in a single day, it remains rational to default forever. Hence, while society cannot banish any individual in perpetuity (it cannot keep records), banishment in each contemporaneous night period forever amounts to the same thing.

The example above demonstrates how a fiat money object can be valued in exchange. In particular, it can be valued to the extent that individual money balances reveal something important about a person’s trading history. In the example above, entering the day with ‘large’ money balances constitutes evidence of a past sacrifice, worthy of current redemption. Leaving the day with ‘low’ money balances constitutes evidence of free-riding; an act deserving of punishment (the threat of which is used to discourage the practice). In short, money—like societal memory—is a record-keeping device. Societal memory—itsself an intrinsically useless object—is valued to the extent it helps mitigate the problems associated with limited commitment. Fiat money is just a physical manifestation of some relevant aspect of an individual’s past trading behavior. In the context of this particular environment, money turns out to be a perfect substitute for memory.

What the example above does not demonstrate, however, is why money might be *essential* in the sense of being the *only* way in which to implement a constrained-efficient allocation. Clearly, if societal memory is available, then money is not necessary. Whether money is essential or not hinges on the following question: Is money a technologically superior method of encoding societal memory? Not surprisingly, the answer to this question depends on what further assumptions are made with respect to the environment.

One might, for example, assume that a public access database recording the private histories of all individuals is absent; perhaps because such a record-keeping technology would entail an unbearable resource cost. Alternatively, one might assume that private trading histories constitute private information; and that simple ‘intangible’ reports of histories can be costlessly fabricated. Or, if private histories are somehow recorded in a central data bank, one might imagine that these data are subject to (identity) theft. Under any of these circumstances, if society can create tokens that are less easily counterfeited and more difficult to steal (relative to identity fabrication or theft), then a ‘tangible’ money token can serve as the device by which individuals can credibly record their past behavior and make this record accessible (through a spot trade of money for output) to other agents.⁷

5.1 A Market Mechanism

Following the discussion above, imagine that memory is unavailable, so that money is essential. Assume that a mechanism (suitably interpreted here as a monetary/fiscal authority) initially creates and distributes money evenly to all members of society. I assume here that money is perfectly divisible.

Agents are assumed to trade on a sequence of competitive spot markets (with money being exchanged for goods in both the day and night).⁸ Other

⁷A topic not pursued here regards the coexistence of money and credit; see, for example, Kocherlakota and Wallace (1998) and Sanches, Williamson and Wright (2007).

⁸The model in this section is a variant the competitive market model presented in Ro-

than this restriction, the mechanism considered here has all the powers of the monetary mechanism described above; in particular, it has the power to exclude individuals from trading in the night market. Keep in mind that agents are free to skip either the day or night market of their own volition.

There are three things that are by now well-known about this environment. First, in the absence of money, the only equilibrium is autarky. Second, if the money supply is held constant, the monetary equilibrium improves upon autarky, but falls short of the first-best. Moreover, with a constant money supply, the monetary equilibrium is incentive-feasible even absent any ability on the part of the government to prohibit trades. Third, if the money supply is made to grow (by way of monetary injections of any form), the resulting inflation reduces *ex ante* welfare.

It is also well-known that the desired policy in the present context is to contract the money supply in accordance with the Friedman rule. The requisite destruction of money is usually assumed to be achieved by way of lump-sum taxes.⁹ The problem with this policy prescription, at least, within the context of the environment laid out here, is that coercive taxation is ruled out by assumption. To put things another way, it assumes that agents can, in effect, commit to paying a tax obligation; and it is precisely the absence of this form of commitment that renders money essential in the first place. There is an obvious inconsistency here that needs to be resolved.

The issue to be resolved here is whether there is any way in which the mechanism can contract the money supply in a manner that is not inconsistent with the assumed structure of the environment. In what follows, I demonstrate how this question can be answered in the affirmative, but subject to a natural restriction that may render the first-best allocation outside of the set of incentive-feasible allocations (so that the Friedman rule is not implementable).

To begin, assume that the government asks agents to pay a lump-sum tax of money at some point in the day. As with the private debt market described earlier, this public debt obligation must be self-enforcing; i.e., agents must somehow be induced to pay the tax voluntarily. The question is how. When memory is available, noncompliance can be punished by the threat of exclusion. But there is no memory here. Moreover, simply displaying one's money balances provides no evidence that the tax has been paid.¹⁰

The problem here is that there are two distinct forms of credit that need to be recorded. The first is the standard one; i.e., whether or not an agent made or received a gift in the night-market. As described above, money is the device

cheteau and Wright (2005).

⁹Fabiano Schivardi has pointed out to me that it might also be achieved by embedding within monetary objects a form of planned obsolescence; where (say) a fixed fraction of money exogenously evaporates. In what follows, I rule out this technology.

¹⁰For example, one might anticipate that all agents will have the same money balances at the end of the day; but there is no way of knowing whether this dollar amount was acquired by working and paying taxes, or avoiding the tax and shirking.

that records these activities. The second is whether the tax has been paid or not. We cannot ask that a single record-keeping device keep track of these two distinct actions. The solution, therefore, is to introduce a second record-keeping device.

This second record-keeping device can be in the form of a dated tax receipt. In keeping with our earlier assumptions, this government-issued tax receipt can be issued in a physical and non-counterfeitable form. It is an object that can be carried into the night-market and displayed to the government as a sort of license to trade in the night-market. In a manner consistent with what we have assumed earlier, the government has the power to exclude agents from trading in the night-market; a power that the government can and should exercise absent evidence of the tax receipt. Finally, as the tax receipt is dated, it expires at the end of the night; it must be renewed at the beginning of each day (and will not circulate as a different form of money).

With this mechanism in place, it is conceivable that agents may voluntarily pay their taxes if it grants them access to the night-market. This is exactly analogous to the private debt market studied earlier; where agents were seen to voluntarily make good on their private debt obligations in exchange for the prospect of future trading opportunities.

Assume then that the government expands the money supply M at (gross) rate $\mu \geq \beta$. New money $(\mu - 1)M$ is injected (withdrawn) by way of a lump-sum transfer (tax) every day; let $\tau = (\mu - 1)M$ denote the transfer (tax, if negative). Moreover, assume for the moment that agents are willing to pay this tax. Then it is an easy matter to establish that the competitive monetary equilibrium allocation (x^e, y^e) satisfies:

$$u'(y^e) = \pi^{-1} \left[\left(\frac{\mu}{\beta} \right) - 1 + \pi \right] g'(y^e); \quad (22)$$

$$x^e = \beta^{-1} g'(y^e) y^e. \quad (23)$$

Observation B $y^e = y^*$ when $\mu = \beta$; i.e., the Friedman rule implements the first-best allocation.

Observation C Condition (23) is identical to condition (15).

That is, the competitive equilibrium allocation in a sequence of money spot markets is required to satisfy the same restriction as the competitive equilibrium allocation in a sequence of private debt markets (where memory was employed). One implication of this is that the term $\beta^{-1} g'(y^e)$ represents the (gross) real rate of return on money (rather than private debt) from the night to the next day.

The final thing to check is whether the first-best allocation, achieved under the Friedman rule, is in the set of incentive-feasible allocations. As in the case of the private debt market, the answer is clearly no if $W(y^*) < x^*$. By Observation C, the following result is readily apparent:

Result 9 *Results 5-7 apply to the competitive monetary equilibrium. In particular, if $\beta \in [\beta^0, 1)$, then the Friedman rule implements the first-best allocation. If $\beta \in (0, \beta^0)$, the Friedman rule is not incentive-feasible; the constrained-efficient allocation in this case is second-best if $g'' = 0$ and third-best if $g'' > 0$.*

Result 9 implies that if agents are sufficiently impatient, then there is an incentive-induced lower bound on the rate at which a monetary authority can deflate. That is, when $\beta \in (0, \beta^0)$, the best incentive-feasible allocation (x_1, y_1) satisfies (18). But as y_1 must also satisfy (22), this condition identifies the lower bound on deflation; i.e.,

$$\mu \geq \mu_1 = \beta + \pi\beta \left[\left(\frac{u'(y_1)}{g'(y_1)} \right) - 1 \right].$$

Clearly, $\mu_1 > \beta$ as $u'(y_1) > g'(y_1)$.

In fact, it is evident that if SR binds sufficiently tightly for consumers, then it is possible that the only incentive-feasible monetary policy requires a strictly positive inflation. There is a clear and direct analog here to the private debt market considered earlier. That is, if consumers have a strong incentive to default on their private debt obligations, producers are compelled to restrict the amount of output produced severely. The equilibrium real interest rate in this case is very low (consumers would like to borrow more at this low interest rate, but cannot). The low real interest rate in this case corresponds to the high inflation rate in the monetary economy.

6 Conclusion

Lagos and Wright (2005) demonstrate how the essential properties of a money-search model can be preserved in an environment that is rendered highly tractable with the use of quasi-linear preferences. My own contribution has been to point out that this same innovation can be applied to closely related environments studied elsewhere in the literature. In the context of dynamic models with private information and limited commitment, the effect of quasi-linearity is reduce the amount of memory required to implement a constrained-efficient allocation. In the context of models where agents can self-insure through competitive asset markets, the effect of quasi-linearity is to remove the distribution of assets as a state variable. As has been highlighted elsewhere in the literature, the role of fiat money is to allow people to credibly record some aspect of their past transactions and make that record accessible to other people when societal memory is too costly (effectively rendering agents anonymous). This basic insight holds true whether or not trade is imagined to occur in centralized or decentralized locations.

Relating the LW model to the existing literature also led to an interesting discovery; namely, that the same frictions that render money (memory) essential

may at the same time render the Friedman rule infeasible. When this is so, there is an incentive-induced lower bound to the rate of deflation away from the Friedman rule. In some circumstances, the best incentive-feasible monetary policy may entail a strictly positive rate of inflation.

References

1. Deviatov, Alexei and Neil Wallace (2001). “Another Example in Which Lump-sum Money Creation is Beneficial,” *Advances in Macroeconomics* 1, Article 1.
2. Hu, Tai-wei, John Kennan and Neil Wallace (2007). “Pairwise-Core Monetary Trade in the Lagos-Wright Model,” Manuscript.
3. Kocherlakota, Narayana (1998). “Money is Memory,” *Journal of Economic Theory*, 81: 232–251.
4. Kocherlakota, Narayana and Neil Wallace (1998). “Incomplete Record-Keeping and Optimal Payment Arrangements,” *Journal of Economic Theory*, 81(2): 272–289.
5. Levine, David (1991). “Asset Trading Mechanisms and Expansionary Monetary Policy,” *Journal of Economic Theory*, 54: 148–164.
6. Lagos, Ricardo and Randall Wright (2005). “A Unified Framework for Monetary Theory and Policy Analysis,” *Journal of Political Economy*, 113: 463–484.
7. Molico, Miguel (2006). “The Distribution of Money and Prices in Search Equilibrium,” *International Economic Review*, 47(3): 701–722.
8. Rocheteau, Guillaume and Randall Wright (2005). “Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium,” *Econometrica*, 73: 175–202.
9. Sanches, Daniel, Steve Williamson, Randall Wright (2007). “Money and Credit with Limited Commitment,” Manuscript.
10. Williamson, Stephen (2006). “Search, Limited Participation, and Monetary Policy,” *International Economic Review*, 47: 107–128.