

On the Benefits of Currency Reform*

R. VIJAY KRISHNA† OKSANA LEUKHINA‡

1ST APRIL 2019

Abstract

Money allows agents to achieve allocations that are not possible without it. However, currency in most economies is a uniform object, and there may be incentive compatible allocations that cannot be implemented with a uniform currency. We show that *currency reform*, ie, changing the monetary base by replacing one currency with another, is a powerful tool that can enable a monetary authority to achieve a desired allocation. Our monetary mechanism with currency reform is *anonymous* and features a *nonlinear* exchange rate between currencies and a monotone value of money. These results help interpret the characteristics of currency reforms observed in practice.

Keywords: currency reform, monetary models

1. Introduction

Money has been a fact of economic life for thousands of years because it helps economic agents attain more desirable allocations. As has been well understood, at least since Ostroy (1973), an important property of money is that it provides a signal about private histories, and is therefore a proxy for private histories. There

(*) We thank David Andolfatto, Narayana Kocherlakota, Fernando Martin, and Neil Wallace for helpful comments. The editor and an anonymous referee's comments have helped greatly improve the exposition. Krishna is grateful to the NSF for financial support via grant SES 1461502.

(†) Florida State University <rvk3570@gmail.com>

(‡) Federal Reserve Bank of St. Louis <oksana.m.leukhina@gmail.com>

is now an extensive literature where money plays this role more-or-less explicitly, and is part of what Wallace (2010) refers to as the ‘mechanism design approach to monetary theory’. However, the models in this literature in particular, and in monetary theory more generally, are typically concerned with changing the *amount* of money in circulation and the effect this has on allocations. Less attention has been paid to currency reform, ie replacing one currency with another.¹

Currency reform is employed more often than one might suspect and it typically features nonlinear pricing — ie, the rate at which current currency holdings can be exchanged for new currency may depend on the amount of money exchanged and/or the agents’ observable characteristics. One recent example is the 2016 currency reform in India which was used as a way to devalue undeclared wealth and savings. In particular, ₹500 and ₹1,000 banknotes of the Mahatma Gandhi Series were declared as invalid and could be exchanged for ₹500 and ₹2,000 banknotes of the Mahatma Gandhi New Series up to a limit which changed over time and only by individuals who could provide documentation of past income. These last requirements effectively imply nonlinearity of the exchange rate.

In this paper, we show that *currency reform with nonlinear pricing* expands the set of attainable outcomes and increases the monetary authority’s welfare measure, whatever it may be. In particular, we demonstrate that currency reform is a tool that allows the monetary authority to incentivise agents to reveal their unobservable savings which reduces their future informational rents,² thereby allowing the monetary authority to achieve outcomes unattainable under uniform currency.³ Our monetary mechanisms have the familiar feature that the value of money is monotonic. We detail several notable examples of currency reform in Section 6. The monetary authorities’ objectives may vary across these examples and they are not always clear. Whatever these objectives may be, however, these examples illustrate the power of currency reform.

(1) Kocherlakota (2002) and Dong and Jiang (2010), discussed further in Section 2, consider the contemporaneous use of multiple currencies, but not currency reform.

(2) Observable savings are typically not an issue.

(3) We emphasise that the monetary authority need not maximise social welfare. We only require that the authority be like a Central Banker who exchanges one currency for another at a certain rate. This Central Banker is more like a price discriminating monopolist, in that it sets nonlinear exchange prices between currencies, than a traditional planner.

More specifically, we consider two well-studied economies — the environment described in Kehoe, Levine and Woodford (1992) (henceforth KLV) and the one described in Atkeson and Lucas (1992) (henceforth, AL)⁴ — and show that currency reform with a nonlinear exchange rate can help the monetary authority (subject to informational constraints) achieve all second best outcomes.

We emphasise that the nonlinearity of exchange rate is necessary for the currency reform to be effective. Currency reform at a constant exchange rate is entirely ineffective in the environments we consider: Under such reform, the monetary authority effectively incentivises agents to reveal their private savings histories, but never acts upon that information. This type of currency reform is equivalent to the currency reenumeration that is often observed after periods of hyperinflation. It is the nonlinearity of exchange rate that allows the monetary authority to condition future outcomes upon revealed information.

To obtain intuition for how currency reform with nonlinear pricing can improve on allocations achieved with a uniform currency, consider an economy with infinitely lived agents. Suppose that histories are unobservable,⁵ and so the monetary authority cannot distinguish between agents, and that the monetary authority uses a uniform currency to implement certain allocations. Notice that the monetary authority would like to achieve two things at any instant. It would like agents to truthfully reveal their current private information, and it would also like for them to reveal their past history. The first can be achieved by nonlinear pricing. However, the second is much harder to achieve if agents can save currency over time and if these savings are not observable.⁶

Currency reform in every period solves these problems by allowing for a nonlinear rate of exchange between currency at one date and the next. This is akin to nonlinear pricing and works in exactly the same fashion. But crucially, currency reform also kills any incentive to save currency for future use. This is because with currency reform, the introduction of a new currency means that a currency that is

(4) In both of these environments, money is essential in the sense that it allows the agents to achieve allocations that are unachievable without it.

(5) If histories are observable, the threat of autarchy can implement the monetary authority's desired outcome.

(6) Indeed, principal-agent problems where the agent has private savings are notoriously difficult. We provide below a discussion of how the literature addresses this problem.

saved will never be used again, and so becomes worthless. Now that savings are no longer an issue, the monetary authority can keep track of histories by summarising them in money holdings and having nonlinear pricing for the consumption good.

For nonlinear pricing to be incentive compatible, agents must not trade multiple times or refrain from trading. Currency reform can be carried out so that (i) trade in a period results in cash holdings that are only useful in the subsequent period, and (ii) all the current cash holdings are used up in the trade, so that there are no repeat purchases. Of course, abstention means that an agent no longer has the currency necessary for trade in subsequent periods, and so is banished to autarchy.

In a sense, we confirm and extend the general intuition in the literature — see, for instance, Kocherlakota (1998) — which says that for money to be *essential* (in the sense that without money, some allocations would be rendered unattainable), there must be some imperfection in the monitoring of private histories. We show that currency reform is sufficient to incentivise truthful revelation of histories.

Our monetary mechanism with currency reform has the feature that it is *predictable* in the sense that it only depends on information available to the monetary authority at the start of the period, and it is *anonymous* in that agents' private histories remain private and are not known to the monetary authority. As noted above, it also features *nonlinear* pricing of consumption goods and future assets, as observed in practice (see Section 6 for some examples),⁷ and the value of money is increasing in the quantity held (in contrast to Kocherlakota (2002), where the value of money is not monotone in holdings.)

We now put our results in context. We first consider the K_LW economy, where agents are anonymous, and their endowments are random and unobservable. Here, trade is feasible only in the presence of an asset, called money.⁸ In such a setting, K_LW analyse the impact of inflationary monetary policy using a uniform currency. They show that under some circumstances, inflation can be welfare enhancing. They do not consider the question of whether monetary policy can achieve the second best. Indeed, they restrict attention to a special class of equilibria that

(7) Nonlinear pricing also plays an essential role in other works that we review below.

(8) Money is intrinsically valueless, and derives its value in equilibrium. Indeed, there is an equilibrium in the K_LW model where the price (and hence value) of money is always 0.

are essentially history independent, which precludes attainment of second best outcomes.

However, *if* the monetary authority could identify agents (though not necessarily their endowments), then it *could* incentivise them to make suitable and welfare improving transfers. We show in Theorem 1 that by using *currency reform* (whereby the currency in use is changed) in every period, the monetary authority can implement such second best outcomes, even if it cannot identify agents.

The environment in K LW is special in that an agent's type is persistent over time. Thus, once an agent reveals his type, it is known forever. This persistence (which manifests itself as the monetary authority knowing the currency holdings of a type) is at the heart of the equilibria that K LW consider. But persistence also means that if the monetary authority can identify agents, then after the first period, when types are revealed, there are no further incentive constraints to consider. This raises the question of whether our approach is pertinent when the private information is not persistent.

To illustrate that it is, we consider an environment where agents face iid taste shocks, as in the AL economy. As before, if the monetary authority can identify agents, it can record histories of reported taste shocks. We demonstrate in Theorem 2 that even if the monetary authority cannot identify agents, it can nonetheless incentivise them to reveal their private histories of shocks by changing the currency in every period via a nonlinear exchange rate.

Even though the monetary authority does not know any individual agent's history, we assume that the monetary authority knows the *distribution* of money and asset holdings in the economy. (Such an assumption is also made in K LW and Levine (1991).) Theorems 1 and 2 show that as long as the monetary authority can issue (and trade) currency and have centralised trading mechanisms with nonlinear prices, it can implement any allocation that is incentive compatible, even ones that rely on observable histories.

We note that Levine (1991) considers an environment similar to K LW, except that it is the monetary authority who has to make allocations, and agents themselves have no endowments. He shows that if agents are sufficiently patient, currency reform and nonlinear pricing can implement the first best, although he says nothing about the second best. We show, more generally, that even for an arbitrary discount

factor and for multiple states and types, the second best can be achieved by repeated currency reform.

This is not to say that we are advocating currency reform as a matter of practice. As Levine (1991) acknowledges, currency reform is costly in practice, and preventing forgery is also costly (though this is a cost that must also be borne in the case of a uniform currency). And although currency reform in every period may be impractical (even though it achieves the second best), our results suggest that periodic currency reform, which is more realistic, can nonetheless approximate the second best, and thereby improve welfare. It goes without saying that currency reform is a powerful tool, so powerful, in fact, that it could do a lot of damage in the hands of an ill-meaning monetary authority. In Section 6, we discuss instances of currency reform in practice, which seem to occur mainly in countries that are closer to the ‘totalitarian’ rather than ‘democratic’ end of the spectrum.

The remainder of the paper is organised as follows. In Section 2, we discuss related literature. Section 3 introduces the KLW environment and recursive revelation mechanisms that characterise second best allocations, while Section 4 describes a monetary mechanism featuring currency reform, ie, a monetary mechanism with a new currency in each period, that implements second best outcomes. Section 5 studies the AL environment and describes a monetary mechanism with currency reform that implements second best outcomes. In Section 6, we describe some events in the last century that have featured currency reform. This is not an exhaustive list, but is intended to provide a flavour of what currency reform entails in practice.

2. Related Literature

We consider two environments, first the model of KLW and second, the model of AL. In both models, we consider centralised trade. The monetary authority plays the role of market maker and executes trades and allocations. Trade in both economies has been analysed under a uniform currency (by KLW in the former, and by Lucas (1978) in the latter), and the outcome is always far from the second best.⁹

(9) Indeed, in the AL economy, the unique equilibrium allocation with a uniform currency is stationary (ie, constant across agents and over time), as is the distribution of money holdings.

We show that if the monetary authority is endowed with the tools of currency reform, then the second best allocation can be implemented.

Our use of money is inspired by the literature that takes the mechanism design approach to monetary theory; see, for instance, the survey by Wallace (2010). However, most of this literature seeks to use the theory of mechanism design, as Wallace (2010, p. 21) puts it, ‘... to explain as an optimum three features of most actual economies: currency is a uniform object; currency is (usually) dominated in rate of return; some transactions are accomplished using currency and others are accomplished in other ways.’ In contrast, we directly use money and the mechanism design approach to implement optimal mechanisms, while emphasizing that we explicitly relax the assumption of a uniform currency. Our finding that relaxing the assumption of uniform currency attains a much richer set of allocations reinforces Wallace’s call to explain the use of uniform currency. One possible explanation is the inability of the monetary authority to *commit* to using currency reform benevolently.¹⁰

Our findings are broadly in line with the findings in papers that accomplish improving the equilibrium allocation by endowing the monetary authority with additional tools. For instance, Kocherlakota (2003) shows how nonlinear pricing allows both the existence of money and an asset that dominates it, while also providing the monetary authority an additional instrument to screen agents with different taste shocks (though he does not find the optimal mechanism in his environment.) As noted above, the seminal paper Kocherlakota (1998) shows that money serves as an important record-keeping device in an economy. Most papers that follow in this tradition focus on a single money and the allocations that can be achieved with such a currency restriction. A notable exception is Kocherlakota (2002) who demonstrates that two monies are necessary and sufficient to replace any record-keeping device in an economy with limited commitment (though without private information). Also of note is Andolfatto (2010), who demonstrates that nonlinear pricing can lead to Pareto improvements in the model

See Lucas (1992) for a fuller discussion and comparison of Lucas (1978) and AL.

(10) Athey, Atkeson and Kehoe (2005) also consider an environment where the Central Bank has information (and preferences) different from Society at large. They note that putting restrictions on how the Central Bank can react to (publicly available) information is one way to provide the Central Bank with the appropriate incentives.

of Lagos and Wright (2005).

Another important paper that considers two monies is Dong and Jiang (2010). They study the environment in Lagos and Wright (2005), but without the search frictions. In their quasilinear environment with limited commitment and no private information, Dong and Jiang (2010) show that one money, in variable amounts, can achieve the second best. If, in addition, the agents also have private information about their marginal utilities, which is perfectly correlated with an observable (but randomly evolving, iid) state, they show that two monies can achieve the second best. Notice that in Dong and Jiang (2010), an agent's type is fixed over time because types are defined by the dependence of marginal utility on the observable state, as is the case in our version the KWL environment, discussed in Section 3. The major difference from the present paper is that because their environment is quasilinear, two monies will suffice, while in our nonlinear environment, a new currency is required in each period to replicate incentive compatible reporting and record keeping.¹¹ Importantly, they do not consider a counterpart of our model with iid private information, described in Section 5, where we show that nonlinear pricing along with period-wise currency implements the second best allocation.

Also closely related is Bajaj et al. (2017), who show that the constrained-efficient allocation in Lagos-Wright can be implemented either with indivisible money or an inflationary scheme. They show that when agents are impatient (ie, not sufficiently patient) and the first best cannot be implemented, the equilibrium allocation approaches the constrained-efficient allocation as the growth rate of money approaches infinity. In particular, the inflationary scheme now looks like the currency reform described in this paper,¹² because the high inflation means that agents get increased amounts of currency only if they have never deviated in the past. This is akin to replacing the old currency with a new one, but at a nonlinear exchange rate, and only if the agents have enough of the old currency; see Section 4.3 for a discussion along these lines.

The model of Lagos and Wright (2005) considers trade in pairwise matches and then trade in a centralised market. The trading protocol in the pairwise match

(11) Indeed, in the private information version of Dong and Jiang (2010), the second money is used only to screen types in the first period, because types do not change over time. This is reminiscent of the screening role that illiquid bonds play in Kocherlakota (2003).

(12) We thank an anonymous referee for this reference, and for noting this connection.

is fixed and time (and history) invariant. Hu, Kennan and Wallace (2009) take the mechanism design approach and allow for nonlinear pricing within trades. They show that, with sufficiently patient agents, the first best allocation is implementable and immune to coalitional deviations. In particular, they conclude that the Friedman rule is not necessary.

Hu and Rocheteau (2013) and Hu and Rocheteau (2015) are two other papers that build on Lagos and Wright (2005) and use nonlinear pricing in an essential way. The former considers fiat money and production and concludes that all optimal allocations must feature the coexistence of money and higher-return assets. The latter investigates the interaction between monetary policy and asset prices and shows how asset price ‘bubbles’ are related to liquidity in the economy.

Green (1987) considers an economy with private taste shocks, which also results in uninsurable risk, and then computes the efficient allocation from risk sharing with an outside lender for a specific class of preferences. His environment is a precursor to the one considered by AL. Lucas (1992) provides an exhaustive literature review of earlier work on risk sharing with private information.

Even though we work with the environment studied in AL, the construction of our mechanisms suggests that our results should extend to all environments with the one-dimensional summary statistic property. These are environments where past signals and/or reports can be summarised by a one-dimensional variable. Kocherlakota (2002) shows that there is a large class of environments with this property. The main difference between his environment and ours is that we allow for instantaneous private information that affects social surplus, while he doesn't.

Our results are robust to the possibility that agents can privately save money. The ability to store money privately¹³ means that preferences over continuation problems are no longer common knowledge. Private savings in an economy with money is a harder problem to handle, because the value of stored money is determined in equilibrium.¹⁴ It is precisely this hurdle that the creation of new money overcomes, by ensuring that the value of saved currency is exactly zero which

(13) For example, Cole and Kocherlakota (2001) consider a setting where agents have endowments that are iid over time and can also be stored, both publicly as well as privately, and both at the same rate of return.

(14) For a more complete discussion of dynamic moral hazard with and without saving, see Sannikov (2008).

forces agents to exchange all their money holdings in each period.

3. Allocations under Limited Commitment

3.1. Environment

We describe a special case of the environment considered in K LW.¹⁵ Time is discrete. There are two types of agents, $i = 1, 2$, and two states of the world $\eta = 1, 2$. Let $\omega_{i,\eta}$ denote type i 's endowment when the state is η . In state 1, agents of type 1 have a *high* endowment h , while type 2 agents have a *low* endowment ℓ , where $h > \ell > 0$. In state 2, agents of type 2 have the higher endowment, while type 1 agents have the lower endowment. (Thus, $\omega_{i,\eta} = h$ if $i = \eta$ and is ℓ otherwise.) The state η follow a Markov process: the probability of transitioning from state η to state $\eta' \neq \eta$ is $\pi \in (0, 1)$. All agents discount the future at the same rate $\beta \in (0, 1)$, and have the same utility function $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ that is concave and increasing.

The state η is observable. However, endowments are private and agents in this economy cannot identify each other.

3.2. Planner's Problem

Let us consider the planner's problem as a dynamic mechanism design problem. More specifically, the planner is a *mediator* in the sense of Myerson (1991, Chapter 6). That is, in the first period, agent's report their type to the planner. He then recommends transfers. The planner doesn't need to seek further reports of types, because types are fixed through time and knowledge of η (which is observable) tells him each type's endowment. For the problem to be well defined, we assume that the planner can identify agents, ie, he can condition his recommendations on the initial reports of types by agents.

Let (η_t) be the stochastic process of observable states and η^t denote the history of shocks up to (and including) time t . (We will sometimes write η_t as $\eta(t)$.) Let H^t

(15) K LW allow for time varying utilities and endowments. We keep utility functions fixed over time, and only let endowments vary.

denote the space of t -period histories, and $\mathcal{H} := \bigcup_{t \geq 0} H^t$ the space of all histories. A *sequential mechanism* is a pair of functions $c^1, c^2 : \mathcal{H} \rightarrow \mathbb{R}_+$. The function c^i induces a process (c_t^i) that represents the consumption at time t as a function of η^t of an agent who declares himself to be type i at date 1. The *planner's problem* is to maximise ex ante welfare subject to certain constraints that we now describe.

Notice first that the *value of autarchy* to an agent of type i given history η^t is¹⁶

$$V_{\text{aut}}^i(\eta^t) := (1 - \beta) \mathbf{E}_0 \left[\sum_{s=t}^{\infty} \beta^{s-t} u(\omega_{i, \eta(s)}) \mid \eta^t \right]$$

The mechanism is *individually rational* if

$$[\text{IR}_t] \quad (1 - \beta) \mathbf{E} \left[\sum_{s=t}^{\infty} \beta^{s-t} u(c_s^i) \mid \eta^t \right] \geq V_{\text{aut}}^i(\eta^t)$$

holds for all η^t . If an agent of type i initially declares that he is of type j , then his utility from the consumption stream (c_t^i) is given by

$$[\text{3.1}] \quad W^i(\eta_t; \eta^{t-1}, j) = \max \left[(1 - \beta)u(c_t^j) + \beta \mathbf{E} W^i(\eta_{t+1}; \eta^t, j), V_{\text{aut}}^i(\eta^t) \right]$$

Notice that $W^i(\eta_t; \eta^{t-1}, j)$ in [3.1] reflects the possibility that the agent can enter autarchy at any point in time, but that this decision is irreversible. Finally, this allows us to define an agent's *interim incentive constraint* at time $t = 1$.

$$[\text{IC}_{ij}] \quad W^i(\cdot; \eta^0, i) \geq W^i(\cdot; \eta^0, j)$$

The mechanism should also be *consistent*, which amounts to requiring that the initial declarations of types be consistent with those present in the environment, ie,

$$[\text{Cons}] \quad \text{Leb}\{\text{report} = i\} \neq 1 \text{ implies } c_t^i = \omega_{i, \eta(t)} \text{ for all } i, t$$

where $\text{Leb}(\cdot)$ is Lebesgue measure. Of course, the mechanism should also be feasible over all mechanisms that satisfy these constraints, and can be written as

$$[\text{Feas}] \quad c_t^1 + c_t^2 \leq h + \ell \quad \text{for all } t$$

(16) The term $(1 - \beta)$ ensures that expected lifetime utility is measured in the same units as instantaneous consumption.

The *planner's problem* then is to maximise ex ante expected utility

$$\begin{aligned}
 \text{[P]} \quad & \max_{(c_t^i)} \mathbf{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \left(\frac{1}{2}u(c_t^1) + \frac{1}{2}u(c_t^2) \right) \right] \\
 & \text{subject to [IC}_{ij}], [\mathbf{IR}_t], [\mathbf{Cons}], [\mathbf{Feas}]
 \end{aligned}$$

The solution to [P] is the *second best* outcome. The *first best* outcome is when the planner maximises the objective in [P] subject to the feasibility constraint, but without considering the incentive and individual rationality constraints. In particular, the first best outcome is consuming $(h + \ell)/2$ in each period.

The mechanism in the planner's problem is *anonymous* in the sense that all agents who report a particular initial type are treated identically. Also, the allocation only depends on the agent's *initial* report at date $t = 1$ (after η_1 has been realised), and the subsequent, publicly observable, realisations of η_t for $t \geq 2$.

Notice also that an agent's type is *persistent* in that it does not change over time. So all the information rent that accrues to him does so initially, at date $t = 1$. In particular, if the agent has revealed his type truthfully at date $t = 1$, then the planner knows the agent's endowment at each subsequent date. Of course, in principle, a positive measure of agents of type 2 could pretend to be type 1, while all type 1 agents report truthfully. The constraint [Cons] ensures that in this case, the allocation is the autarchic one. It is possible that there is a group deviation where a positive measure $\varepsilon \in (0, 1)$ of each type pretends to be of the other type. Such a deviation will not be detected by the planner. Nevertheless, the incentive constraints [IC_{ij}] ensure that this is never optimal for any agent.

3.3. Solution to Planner's Problem

As noted above, *after* the first period, there is no private information in this economy. At this point, this becomes a standard risk-sharing game without commitment. As such, it is isomorphic to the game considered in Kocherlakota (1996) (or at least isomorphic to the version without aggregate uncertainty) with the following caveats. The game in Kocherlakota (1996) has only two agents who are identical and have stochastic endowments. They cannot commit to future allocations, but instead have to rely on self-enforcing agreements. Our model has a continuum of

agents who are anonymous and *cannot* interact repeatedly with anyone else (for technological reasons). Nonetheless, the maximisation problem for the planner is the same as in Kocherlakota (1996), and our constraints are the same as his. Therefore, the set of optimal allocations in Kocherlakota (1996) is also optimal here.

Furthermore, for each $\eta = 1, 2$, the optimal allocation in Kocherlakota (1996) can be summarised recursively, by a variable $v_1 \in [V_{\text{aut}}^1(\eta), V_{\text{max}}^1(\eta)]$ that denotes the utility promised to an agent of type 1, and a function $f_\eta : [V_{\text{aut}}^1(\eta), V_{\text{max}}^1(\eta)] \rightarrow \mathbb{R}$ so that $f_\eta(v_1)$ denotes the corresponding utility promised to an agent of type 2. It is easy to show (see Kocherlakota 1996 for details) that f_η is decreasing. The range of f_η is denoted by $[V_{\text{aut}}^2(\eta), V_{\text{max}}^2(\eta)]$.¹⁷

In particular, the *Pareto frontier* (or more precisely, the *utility possibility frontier*) in state η is $\mathcal{U}_\eta := \{\mathbf{v} = (v_1, v_2) : v_1 \in [V_{\text{aut}}^1(\eta), V_{\text{max}}^1(\eta)], v_2 = f_\eta(v_1)\}$, as shown in Figure 1. The set \mathcal{U}_η denotes all Pareto optimal payoffs that are attainable in some equilibrium, beginning in the state η . It is also, clearly, the graph of the function f_η .

The optimal allocation can be summarised by a pair $\mathbf{v} = (v_1, v_2) \in \mathcal{U}_\eta$. In this case, consumption in state η is given by $c_\eta^i : \mathcal{U}_\eta \rightarrow \mathbb{R}$ and the next period's utility (one for each possible state) is determined by functions $\mathbf{w}_\eta^i : \mathcal{U}_\eta \rightarrow \mathcal{U}_1 \times \mathcal{U}_2$.

Finally, define the set $\mathcal{U}_\eta^\dagger := \{\mathbf{v} \in \mathbb{R}^2 : v_i \geq V_{\text{aut}}^i(\eta) \text{ and there exists } \hat{\mathbf{v}} \in \mathcal{U}_\eta \text{ s.t. } \mathbf{v} \leq \hat{\mathbf{v}}\}$. (In Figure 1, this is the convex hull of \mathcal{U}_η and the origin, and is shaded.) Clearly, $\mathcal{U}_\eta \subset \mathcal{U}_\eta^\dagger$ and \mathcal{U}_η^\dagger is convex. It is easy to show that the set \mathcal{U}_η^\dagger is *self generating* and is the largest set of payoffs from subgame perfect equilibria.

An equilibrium then generates a Markov process (\mathbf{v}_t) with state space $\mathcal{U}_1 \cup \mathcal{U}_2$. In each state, there is consumption, and then a transition to a new state in the subsequent period. The function \mathbf{w}_η^i determines next period's η -contingent promised utility as a function of the current period's promised utility. In particular, $\mathbf{w}_{\eta, \eta'}^i$ is the promised utility for an agent of type i if the current state is η and the following period's state is η' .

We shall use these observations to implement the planner's optimal allocation

(17) Here, $V_{\text{max}}^i(\eta)$ represents the maximum utility that an agent of type i can get in any subgame perfect equilibrium when the state is η .

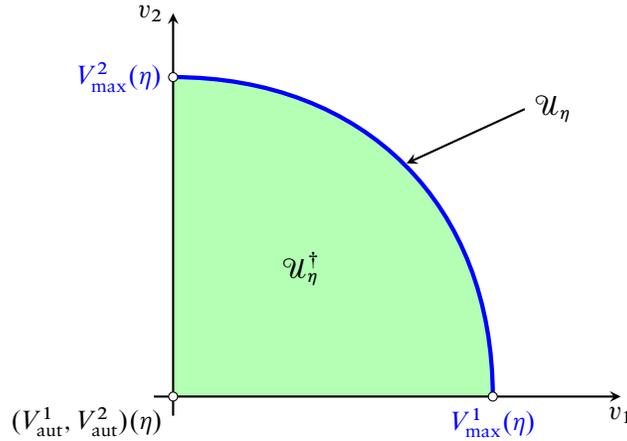


Figure 1: *Equilibrium Payoffs in state η*

using a monetary mechanism that features currency reform.

4. Monetary Mechanisms under Limited Commitment

A *monetary mechanism* is a mechanism where agents trade currency with the monetary authority in return for the consumption good and some more currency. A monetary mechanism of the kind described in Lucas (1978) or Kehoe, Levine and Woodford (1992) (but also see Lucas 1992) features a uniform currency, ie, where the same currency is used in every period. A monetary mechanism features *currency reform* if the currency is transformed into another currency at some point in time. (This is also referred to as changing the monetary base.) Note that the transformation need not be 1-to-1. A monetary mechanism features *radical* currency reform if the currency base is changed in every period.

Our monetary mechanism has two key differences from the mechanism considered in K LW. First, as noted above, they only consider one money, which the monetary authority can print and distribute uniformly. In contrast, we allow the monetary authority to print money of different colours in each period, and let him pull money out of circulation in exchange for new money. Secondly, trade in K LW occurs under linear pricing—this is a standard centralised trading mechanism. We consider centralised trading, with the monetary authority being the monopolistic

market-maker, but with nonlinear pricing.

It is easiest to see how our mechanism works in a simple example. We provide a general implementation result in the sequel.

4.1. Example: Monetary Mechanism with High Patience

For the allocational mechanism described in Section 3.3, there exists a β^* such that if $\beta \in (\beta^*, 1)$, then the first best allocation of $\omega^* := (h + \ell)/2$ is attainable. Set $\Delta := h - \omega^* = \omega^* - \ell$. We shall now construct a monetary mechanism that implements the first best allocation.

As in KLW, assume that the agents of type i initially have money holdings of m_0^i . The monetary authority is a monopolist and buys and sells the consumption good. Our mechanism requires that the agents place orders, the monetary authority evaluates all the orders, and if feasible, executes them, subject to the restriction that it must sell less than it buys of the good.

Date $t = 1$ prices. Let the observable state be $\eta_1 = 1$. Then, the monetary authority sets the following prices:

- Sell (ie, place orders with the monetary authority for) Δ units of the consumption good *and* m_0^1 units of existing currency, in exchange for 1 unit of *blue* period 1 currency.
- Buy (ie, place orders with the monetary authority for) Δ units of the consumption good *and* 1 unit of *red* period 1 currency, in exchange for m_0^2 units of existing currency.

If the state is $\eta_1 = 2$, the monetary authority sets the following terms of trade:

- Sell (ie, place orders with the monetary authority for) Δ units of the consumption good *and* m_0^2 units of existing currency, in exchange for 1 unit of *red* period 1 currency.
- Buy (ie, place orders with the monetary authority for) Δ units of the consumption good *and* 1 unit of *blue* period 1 currency, in exchange for m_0^1 units of existing currency.

Notice that type 1 agents always have *blue* currency while type 2 agents always

have *red* currency.

Issuing Coloured Currency: The monetary authority can issue new currency of a different *colour* for use in the next period. The blue date t currency is distinct from red and blue currency issued at any other date, as well as from m_0^i for $i = 1, 2$. There is no counterfeiting.

Executing Orders: If the total demand for the consumption good is less than the total supply, the monetary authority executes the orders. If not, it doesn't execute the orders at this date or any future date.

The monetary authority then repeats this at every date. More specifically, we have the following.

Date t prices. Let the observable state be $\eta_t = 1$. The monetary authority sets the following prices:

- Sell (ie, place orders with the monetary authority for) Δ units of the consumption good *and* 1 unit of *blue* period $t - 1$ currency, in exchange for 1 unit of *blue* period t currency.
- Buy (ie, place orders with the monetary authority for) Δ units of the consumption good *and* 1 unit of *red* period t currency, in exchange for 1 unit of *red* period $t - 1$ currency.

If $\eta_t = 2$, then red and blue currencies are interchanged, as was discussed above for date $t = 1$.

Notice that if agents of type i in state $\eta = i$ choose to be sellers (and others are buyers), then the monetary mechanism gives every agent the allocation ω^* in every period and state. We now comment on two important and distinctive features of our monetary mechanism.

Nonlinear pricing: An important feature of the monetary mechanism is that it features *nonlinear pricing*, with the monetary authority playing the role of market-maker. In particular, nonlinear pricing ensures that agents can only trade multiples of Δ units of consumption. Of course, these currency holdings ensure that, in fact, they can only trade Δ units of consumption.

Coloured money: A second important feature of the mechanism is that it features *coloured money*. Coloured money holdings that change in every period

and as a function of past holdings ensure that an agent of putative type i cannot assume the role of an agent of type $j \neq i$ in a subsequent period. If he doesn't continue playing as a type i agent, his only option is autarchy.

To see that the mechanism implements the allocation ω^* at each date and state, first consider date $t = 1$. Choosing to be a buyer or seller reveals an agent's type. Furthermore, once this type is revealed, the agent cannot assume the role of another type at any point in the future. Because the allocation ω^* is incentive compatible, it is optimal for each agent to become a seller in η if his type is $i = \eta$, and become a buyer otherwise.

After effectively reporting (or more precisely, choosing) a type in period 1, it is optimal for the agent to continue with the mechanism if he chose truthfully at $t = 1$. The mechanism then effectively requires transfers of Δ so that everyone consumes ω^* . This is optimal for an agent precisely because the allocation ω^* is incentive compatible and individually rational.

For sufficiently high β , the participation constraint does not bind. In such a case, our monetary mechanism is very similar to the dated asset mechanism in Proposition 3.1 of Levine (1991). For lower discount factors, the second best becomes more complicated and depends on the history η^t in a nontrivial way. We show next that a monetary mechanism with dated colours can implement the second best allocation for any discount factor $\beta \in (0, 1)$.

4.2. A General Monetary Mechanism

Recall that a solution to the monetary authority's problem is given recursively by functions $c_\eta^1, c_\eta^2 : \mathcal{U}_\eta \rightarrow \mathbb{R}$ for instantaneous allocations, and with continuation utilities given by $\mathbf{w}_\eta^1, \mathbf{w}_\eta^2 : \mathcal{U}_\eta \rightarrow \mathcal{U}_1 \times \mathcal{U}_2$. All the monetary authority need do is specify some initial utilities \mathbf{v}_0 to the agents. A *monetary mechanism* is a game form where the monetary authority issues currency of different colours in each period, the agents submit their orders (as buyers or sellers) to the monetary authority who is also a market maker, and the market maker then executes orders and issues new currency. We now state our main result.

Theorem 1. *Let $(c_\eta^i, \mathbf{w}_\eta^i, \mathbf{v}_0)$ be a recursive solution to the monetary authority's*

problem. Then, there is a mechanism that implements this allocation.

Proof. Let m_0^i be the initial, date $t = 0$ money holdings of agents of type i . As before, first consider date 1.

Date $t = 1$ prices. Let the observable state be $\eta_1 = 1$. Then, the monetary authority sets the following prices (ie, terms of trade):

- Sell (ie, place orders with the monetary authority for) $h - c^1(\mathbf{v}_0)$ units of the consumption good *and* m_0^1 units of existing currency, in exchange for $\mathbf{w}_{1,1}^1(\mathbf{v}_0)$ units of *blue* period 1 state $\eta_2 = 1$ currency, and $\mathbf{w}_{1,2}^1(\mathbf{v}_0)$ units of *blue* period 1 state $\eta_2 = 2$ currency.
- Buy (ie, place orders with the monetary authority for) $c^2(\mathbf{v}_0) - \ell$ units of the consumption good *and* $\mathbf{w}_{1,1}^2(\mathbf{v}_0)$ units of *red* period 1 state $\eta_2 = 1$ currency, and $\mathbf{w}_{1,2}^2(\mathbf{v}_0)$ units of *blue* period 1 state $\eta_2 = 2$ currency, in exchange for m_0^2 units of existing currency.

If the state is $\eta_1 = 2$, the monetary authority interchanges the colours and quantities in the above pricing scheme, as in Section 4.1.

Date t prices. Let the observable state be $\eta_t = 1$. Let observable money holdings be $\mathbf{v}_t = (v_t^1, v_t^2)$. (Recall that $v_t^i = \mathbf{w}_{\eta_{t-1},1}^i(\mathbf{v}_{t-1})$ for $i = 1, 2$.) Then, the monetary authority sets the following prices (ie, terms of trade):

- Sell (ie, place orders with the monetary authority for) $h - c^1(\mathbf{v}_t)$ units of the consumption good *and* v_t^1 units of *blue* period $t-1$ state 1 currency, in exchange for $\mathbf{w}_{1,1}^1(\mathbf{v}_t)$ units of *blue* period t state $\eta_2 = 1$ currency, and $\mathbf{w}_{1,2}^1(\mathbf{v}_t)$ units of *blue* period t state $\eta_2 = 2$ currency.
- Buy (ie, place orders with the monetary authority for) $c^2(\mathbf{v}_t) - \ell$ units of the consumption good *and* $\mathbf{w}_{1,1}^2(\mathbf{v}_t)$ units of *red* period 1 state $\eta_2 = 1$ currency, and $\mathbf{w}_{1,2}^2(\mathbf{v}_t)$ units of *blue* period 1 state $\eta_2 = 2$ currency, in exchange for v_t^2 units of *red* period $t-1$ state 1 currency.

If $\eta_t = 2$, then the colours are swapped and the agent's roles are swapped, so as to ensure that type 1 agents still trade with blue currency and type 2 agents trade with red currency.

Notice that at each step, the monetary authority's constraint requires that demand is less than supply, ie, $c^2(\mathbf{v}_t) - \ell \leq h - c^1(\mathbf{v}_t)$, which is equivalent to the

feasibility constraint [Feas].

In contrast to the monetary mechanism for a high discount factor discussed in Section 4.1, the monetary mechanism issues state contingent currency in every period for use in the next period. This makes the monetary mechanism sunspot-like in the sense that in the period following the issue of currency, only one of the currencies issued will be used, as a function of the realised state η .

It is easy to see that the mechanism implements the desired allocation. Incentive compatibility and participation constraints are all satisfied because the mechanism basically mimics the mechanism described as part of the solution to the monetary authority's problem, as described at the end of Section 3.2. \square

4.3. Discussion

Time-varying vs Stationary Money Holdings: Notice that the money holdings in the general mechanism vary over time. Intuitively, they correspond to a type's promised utility in the game. In this case, prices and currency issuance in any period are a function of current money holdings (of both colours) and the state. This makes our mechanism Markovian.

Alternatively, we could let the money holdings in each period be unity (as in the mechanism in Section 4.1), but instead vary the prices accordingly so as to attain the same allocations. The only drawback to this approach is that prices would now depend on the entire history η^t , because there isn't an obvious state variable to make the mechanism Markovian.

Monochromatic Currency: For high discount factors, the mechanism in Levine (1991) uses monochromatic currency, ie, currency of only one colour, in any period, and allocates it exclusively to type 2 agents, with agents of type 1 not using currency. Because his allocation is very simple (a consequence of discount factors being sufficiently high), this suffices for his purposes. (In particular, he doesn't consider general second best allocations.)

One might think that such monochromatic currency might also work for arbitrary discount factors, as long as we varied (say) type 1 money holdings appropriately. After all, knowing type 1's promised utility v^1 pins down type 2's

promised utility. Unfortunately, this intuition is not complete. To see this, suppose promised utility to type 1 agents is arbitrarily close to $V_{\text{aut}}^1(\eta)$ in state η (see Figure 1). Then, it is possible (and this will depend, in general, on the parameters of the model (π, h, ℓ, β) and the agents' risk aversion) that it is profitable for an agent of type 1 who holds currency to claim that he is of type 2 from then on (so that he never holds currency again). This is not a profitable action for sufficiently high discount factors (as in Levine 1991), but may be so for lower discount factors. Monochromatic currency cannot prevent such deviations.

However, such deviations are impossible with coloured currency. This is because an agent of type 1 who holds blue currency cannot pretend to be an agent of type 2 who holds red currency, precisely because he doesn't have red currency and it cannot be forged.

That being said, for the special case where discount factors are sufficiently high (as in the example in Section 4.1) where the first best allocation is implemented, we can make do with one currency. Moreover, as in Levine (1991), the allocation can be implemented with a single currency, ie, without currency reform, as long as there is sufficient inflation of the currency base via lump sum transfers.

The Role of Dated Currency: As in Levine (1991), we use dated currency. If we only had coloured currency that was in use in every period (ie, without currency reform), an agent could choose not to participate in some periods, say, when he is asked to make large transfers. Again, this is not an issue if the monetary authority can identify agents and threaten them with autarchy. When agents are anonymous, participation in the mechanism is indicated by dated currency holdings. Dated currency also ensures that agents have no incentive to save currency for use at a later stage. In a sense, these are the *only* roles that dated currency plays.

Nonlinear Pricing and Inflation: One of the principal rôles of inflation is to ensure continued participation. For instance, someone who left the US in the 1960s with all his currency holdings to live on an island would be able to buy a lot less with his currency now, were he to re-enter the economy again, and hence would be significantly worse off. In this sense, inflation can be welfare-enhancing, because it provides agents incentives to participate in all future periods in the mechanism.¹⁸

(18) Wallace (2014) conjectures that in a large class of pure-currency economies, transfer schemes financed by inflation can improve welfare. Our approach to inflation is more targeted, but

Of course, the state variables now are money holdings, inflation (relative to some base level), and the current observable state η . This is similar to the inflationary policy via lump sum transfers of a uniform currency in Levine (1991). The major difference is that Levine’s lump sum transfer is type independent, while our lump sum transfers depend on the type of the agent and vary by colour of money in circulation. Again, this is because Levine restricts attention to sufficiently patient agents. (It is easy to see that type-dependent lump sum transfers will be necessary if there are more types and/or states. See Section 4.4 for more on this.)

Welfare Effects of Money Creation: Wallace (2014) conjectures that in a large class of monetary economies (which includes ours), money creation increases ex ante welfare relative to what can be achieved by holding the stock of money fixed. In particular, he places the conjecture in terms of very simple rules for money creation — if an agent has x units of currency after trade in a period, he gets $x + \tau_t(x)$ after transfers, where $\tau_t(x) := \max[0, a_t + b_t x] \geq 0$, with $b_t \geq 0$. Thus, in any period, an agent’s post-trade money holdings can never go down, and can only go up. Wallace’s conjecture therefore amounts to saying that non-trivial transfers (ie, $a_t \neq 0$) can be welfare enhancing. We have seen a specific instance of this, when discussing the inflationary policy via lump sum transfers of a uniform currency in Levine (1991). Notice that the conjecture is only about welfare improvements relative to a constant stock of money, and not about implementing the second best allocation, which is our primary concern.

To better understand how his mechanism compares to ours, consider, for the moment, Wallace’s non-linear transfer scheme in the K LW economy. His transfers must serve two purposes: (i) to appropriately change the agents’ relative buying powers, and (ii) to induce the agents to truthfully submit demands, ie, to induce the different types to reveal themselves via their behaviour in the market. Notice that he only has one instrument by which to achieve both goals. Our monetary mechanism must also achieve these two goals, but it offers two instruments to do so — currency reform and a non-linear exchange rate.

The similarity between the two mechanisms is the availability of non-linear transfers as an instrument which, if types are perfectly known, can help adjust the relative purchasing power of agents in accordance with promised utility values.

also provides, in a special case, a verification of the conjecture.

The difference between the two mechanisms (and approaches) is that the currency reform featured in our mechanism effectively ensures type revelation — currency reform in each period indeed ensures agents cannot pretend to have the wrong colour of money. By contrast, Wallace must make sure that someone with a higher holding of money does not pretend to have a lower amount of money. Incentivising this may well be costly and will likely curtail the ability to change the relative buying powers of the agents.

There is yet another reason for why non-linearity in Wallace scheme may face an impediment to adjusting the relative buying power. This impediment is the assumption of *no-monitoring* in his environment, so that any taxes in his environment must be uniform across all agents. (More precisely, because there is no monitoring, and because money holdings are unobservable, he can only permit lump-sum taxes that don't distinguish between buyers and sellers.) In contrast, with currency reform, we are also effectively able to tax agents. So it is clear that with currency reform, we can do at least as well as with Wallace's scheme, and the argument above suggests that sometimes we should be able to do strictly better.

Group Defections: A notion of (stability against) group defection is *coalition-proofness*, which becomes the *pairwise core* in Ravikumar and Wallace (2001), Hu, Kennan and Wallace (2009), Hu and Rocheteau (2013, 2015), for instance. Our model allows for another kind of group defection, namely where the entire population replaces the monetary authority with another monetary authority, with possibly better currency strategies. Our mechanism is immune to such deviations precisely because it implements the second best solution.

In particular, the pairwise core in Ravikumar and Wallace (2001) requires that trades in any meetings are individually rational and are pairwise efficient, given the future values of (respective) monies. Our mechanism does not have pairwise meetings, but our mechanism is clearly immune to *pairwise coalitions* of the following sort. Suppose two agents meet and decide to go off together and leave the economy with their endowments. If they are of the same type, they are at essentially at autarchy because they cannot insure each other. If they are of different types, they can do no better than the second best, which is also our solution. Hence such deviations are not profitable. Of course, additional group defections must be accounted for if there are more types, as we discuss next. Thus, our mechanism is

stable against such pairwise deviations.

Yet another possibility is that a group of agents may pool their endowments and money together to trade with the planner. If all the agents are of the same type, then this does not improve on the planner's allocation. If, however, there are equal numbers of each type of agent, then the players don't need to trade with the planner at all; they can simply trade amongst themselves and simulate the planner's solution. The main concern is that for this to be feasible, the agents need to be able to identify each other. If this is possible, the agents don't need a planner to help them achieve the second best — they can simply play a repeated game amongst themselves. If they cannot identify each other, then such group defections are infeasible.

In sum, because the planner is implementing the second best solution which obtains under full observability of histories, no set of players can form a coalition and do strictly better. We focus on the equilibrium of the monetary mechanism that implements the planner's solution, and so if a coalitional deviation exists in the planner's problem, then there must also be one in our monetary mechanism, as long as agents can get together and form a new KLM-style subeconomy where they want to share risk. For more on this point, see Section 4.4 below.¹⁹

4.4. Extensions

Consider a generalisation of the environment presented in Section 3.1, with J types and N states, and where $J, N \geq 2$. Indeed, there is no reason to require aggregate certainty, as we have done above.

The monetary authority's problem from Section 3.2 easily generalises to this environment, as does the recursive, self generating solution. Once again, there are sets $\mathcal{U}_\eta \subset \mathbb{R}^J$ for each $\eta \in N$, where each type of agent (after reporting his type in period 1) is identified by his promised utility v^i , which in period t , lies in $\mathcal{U}_{\eta(t)}$.

The monetary mechanism described in Section 4.2 easily generalises to this setting. The monetary authority issues currencies with J different colours, and in

(19) There are, of course, many other equilibria of the monetary mechanism, most notably the one where all money holdings are ignored and all agents obtain the autarchic allocation.

each period, he simply issues N different state contingent coloured currencies for use in the next period. The currency that will actually be used by each type will depend on the state that realises in the subsequent period.

This expands on a conjecture of Levine (1991), who says on p148 (while referring to the KLV environment but without endowments), ‘With more states and/or types the first best may be unobtainable by any mechanism, and mechanisms with several assets may dominate a mechanism with a single asset.’

That multiple assets may be beneficial is already apparent even with 2 states and types for lower discount factors. It is also clear that they are necessary for attaining the second best for arbitrary discount factors with an arbitrary (but finite) number of states and types.

Such an environment must be cognisant of group defections. In particular, it is possible that entire types of agents form their own economy, with its own currency issuer. Such deviations must also be accounted for in the monetary authority’s problem. As long as they are, our monetary mechanism will still work as described in the proof of Theorem 1.

We emphasise that the solution to the planner’s problem with J types of agents is well defined, but it is *not* immune to group deviations. For instance, for low discount factors, there will be histories that occur with positive probability such that it is incentive compatible for types 1 and 2 (say) to stay in the mechanism, but it is a *strict improvement* for types 1 and 2 to take their endowments, form a coalition, and insure each other by essentially playing the two-person version of the insurance game described in Section 3. Because such coalitional deviations are possible in the planner’s problem, they must, perforce, also be possible in our monetary mechanism. However, if the planner were cognisant of such group defections (which would decrease the set of implementable allocations), then our monetary mechanism would also be immune to these group defections.

5. Model with IID Private Information

5.1. Environment

We now describe the environment of Atkeson and Lucas (1992). Time is discrete. There is a continuum \mathcal{H} of agents with a typical agent being denoted by h , and at each date $t = 0, 1, \dots$, there is an amount y of resources, that we shall call *cake*, that is to be divided among the agents. Each agent's utility from cake is $\theta u(c)$, where c is the agent's consumption of cake, $u : \mathbb{R}_+ \rightarrow \mathbb{R} \cup \{-\infty\}$ is an increasing, smooth, and strictly concave function, and $\theta \in \{\theta_1, \theta_2\}$ is a *taste shock* where $0 < \theta_1 < \theta_2$. The probability of θ_i occurring is μ_i , and taste shocks are iid across agents and over time. The shock at time t will be denoted by θ_t . We assume, without loss of generality, that $\theta_1\mu_1 + \theta_2\mu_2 = 1$. All agents discount the future at the common rate $\beta \in (0, 1)$.

Let $D := u(\mathbb{R}_+) \subset \mathbb{R} \cup \{\pm\infty\}$, where we allow $u(0) = -\infty$. Then, the utility of a (possibly random) consumption sequence (c_t) is $(1 - \beta) \mathbf{E} \sum_{t=0}^{\infty} \beta^t \theta_t u(c_t)$. Therefore, each agent's lifetime utility must lie in D . Let $C : D \rightarrow [0, \infty)$ be defined as $C(x) = u^{-1}(x)$, so $C(x)$ is the amount of cake consumed by an agent who gets x utiles of instantaneous consumption, modulo taste shocks.

There are two (equivalent) versions of the planner's problem. First, given a fixed amount of cake per period (say, unit cake per capita), the planner's problem is to maximise ex ante expected social welfare. A dual characterisation of the problem is to consider an initial level of utility to be promised to each agent, and then to find the least amount of cake required to make good on the promises.

We first consider the case where planner can observe past reports.

5.2. Allocation Mechanisms in Atkeson and Lucas (1992)

In this section, we assume that the planner can identify individuals via their history of reports. Of course, the planner cannot actually observe an agent's taste shocks.

An agent in this economy only cares about his expected utility. Thus, we may identify an agent with the utility promised to him; an agent's promised lifetime

utility is $w \in D$. Because taste shocks are iid, promised utility completely summarises an agent's past history of reports. Let \mathbf{M} denote the space of probability of measures on D . A probability measure $\psi \in \mathbf{M}$ denotes the distribution of promised utilities to agents in the economy.

A *recursive revelation mechanism* is a triple $\sigma = (f, g, \psi_0)$ where $f, g : \{\theta_1, \theta_2\} \times D \times \mathbf{M} \rightarrow D$ are functions that respectively denote the instantaneous and continuation utilities of an agent who has taste shock θ , promised utility w , and where $\psi_0 \in \mathbf{M}$ is an initial distribution of promised utilities.²⁰ As mentioned above, each agent with a promised utility level of w is clearly identifiable as such, ie, each agent is identified with the utility promised him. In each period, he makes a report $\hat{\theta} \in \Theta$ to the planner and then gets an allocation of cake which yields instantaneous utility $f(\hat{\theta}, w, \psi)$, and some continuation utility $g(\hat{\theta}, w, \psi)$ that will identify him tomorrow.

Given the function g and the distribution ψ , the following period's distribution of promises is denoted by $S_g \psi$, so $S_g : \mathbf{M} \rightarrow \mathbf{M}$ can be regarded as a Markov transition operator.²¹ Finally, ψ_0 denotes the initial distribution of promised utility. Therefore, given w and ψ , an agent's (ex ante) lifetime expected utility under the policy σ is $\sum_i \mu_i [(1 - \beta)\theta_i f(\theta_i, w, \psi) + \beta g(\theta_i, w, \psi)]$.

A recursive revelation mechanism σ

- *attains ψ with resources y* if the total amount of cake that is disbursed is feasible, ie, if

$$\sum_i \mu_i \int C(f(\theta_i, w, \psi)) d\psi(w) \leq y$$

- *is incentive compatible* if for all $\theta, \theta' \in \{\theta_1, \theta_2\}$

$$(1 - \beta)\theta f(\theta, w, \psi) + \beta g(\theta, w, \psi) \geq (1 - \beta)\theta f(\theta', w, \psi) + \beta g(\theta', w, \psi)$$

(20) For simplicity, we take ψ_0 to be given. In particular, the socially welfare maximising amount has ψ_0 as a Dirac point measure.

(21) Formally, $(S_g \psi)D_0 := \int_{B_g(D_0)} d\mu d\psi$, where $B_g(D_0) := \{(\theta, w) \in \Theta \times D : g(\theta, w, \psi) \in D_0\}$ for all Borel $D_0 \subset D$. Notice that every strategy by the agents induces a transition operator S_g . In what follows, we focus on the transitions induced by the truthtelling strategy. (We thank a referee for this observation.)

- and satisfies *promise keeping*²² if

$$\sum_i \mu_i [(1 - \beta)\theta_i f(\theta_i, w, \psi) + \beta g(\theta_i, w, \psi)] = w$$

We shall restrict attention to mechanisms with the following boundedness property. Given the mechanism σ and an initial level of utility w_0 , we can inductively define the agent's continuation utility in period $t + 1$ as a function of all the past reports as follows: $w_{t+1} := g(\hat{\theta}_t, w_t, \psi_t)$, where $w_t : D \times \Theta^{t-1} \rightarrow D$. Intuitively, this expression reflects the fact that the mechanism is Markovian in promised utility. We require that the mechanism satisfies the following version of the transversality condition: for all initial $w_0 \in D$ and $(\hat{\theta}_t) \in \Theta^\infty$,

$$\lim_{t \rightarrow \infty} \inf_{(\hat{\theta}_t) \in \Theta^\infty} \beta^t w_t(w_0, \hat{\theta}^{t-1}) = 0$$

where $\hat{\theta}^{t-1} \in \Theta^{t-1}$ for all t .

AL consider the dual version of the planner's problem, wherein it tries to make good on the promises ψ_0 with the least amount of resources. AL show that at the (second best) optimum, income distribution eventually becomes extremely skewed, with the fraction of agents who become immiserated (ie, whose utility goes off to minus infinity) eventually becoming 1. AL also try and decentralise this optimal allocation, but standard instruments like money are insufficient in that they cannot achieve the second best outcome. This is because standard instruments like money or additional assets are unable to screen agents across all histories. We address this issue next.

In trying to implement the optimal allocation, AL consider money or assets, and assume that there is a monetary authority who can issue these instruments (presumably on behalf of the planner). Following AL, we too consider a monetary authority that can issue currency and certain assets, but cannot identify agents because it cannot observe the currency and asset holdings of any individual agent (though it does know the distribution of both). The monetary authority is also a clearinghouse where assets are exchanged for other assets at some exchange rate that it sets.

(22) Because there is a continuum of agents, it suffices to consider the promise keeping condition in expectation

5.3. Monetary Mechanisms

We shall now assume that at each point in time, the monetary authority (or the Central Banker) can costlessly print paper currency of different amounts. In particular, we shall assume that the monetary authority can print (perfectly divisible) paper currency. We shall also allow the monetary authority to print paper currency in different *colours*. Instead of enumerating the colours as *red*, *blue*, *green*, \dots , and so on, we shall, for simplicity, let the set of colours available be $\mathbb{N} = \{0, 1, 2, \dots\}$. Let ξ_t denote a typical distribution of money holdings of colour t , and let \mathbf{M}_t denote the space of all distributions of money of colour t . We assume that the Central Banker

- Can observe the distribution of money holdings in each period (but cannot observe individual money holdings),
- Can create money of requisite amounts in any colour, and
- Is the only entity in the economy that can print money, so there is no counterfeiting and nor is there the threat of counterfeiting.

As before, a *monetary mechanism* is a mechanism where agents trade money with the monetary authority in return for cake and some more money. It features *currency reform* if the currency is transformed into another currency at some point in time, and has *radical* currency reform if the currency base is changed in every period.

It is useful to define our notion of implementation.

Definition 5.1. Let $\sigma = (f, g, \psi_0)$ be a recursive revelation mechanism. It is implementable by a monetary mechanism if the outcome induced by the monetary mechanism corresponds to the outcome of the recursive revelation mechanism.

We can now state our second implementation theorem.

Theorem 2. *Let σ be a recursive revelation mechanism. Then, there is a monetary mechanism featuring radical currency reform that implements it. In particular, second best (Pareto efficient) allocations are implementable by monetary mechanisms with radical currency reform.*

Theorem 2 is proved in Appendix A.

A typical problem with trying implement an allocation with a monetary mechanism is that agents can burn money or hide it for future use. As a result, the observed money holdings do not effectively reveal past histories. The monetary mechanism that we propose effectively eliminates the agents' incentives to burn or save money.

Radical currency reform, which changes the currency base in each period by withdrawing currency of a particular colour from distribution after one period of use, ensures that an agent does not have an incentive to save money from one period to another. (This is a fact well recognised by the North Korean regime, for instance; see Section 6 below.) By using an additional instrument that we introduce in the next section, the monetary authority can ensure that an agent also does not have an incentive to burn money. Put together, these features ensure that any incentive compatible recursive revelation mechanism can be realised as the outcome of a monetary mechanism with radical currency reform.

5.4. Some Aspects of our Monetary Mechanisms

Some comments on the monetary mechanism are in order.

- The monetary authority need not observe past trades made and option contracts held. It only observes the current distribution of money holdings, but not individual money holdings.
- This implies that the monetary mechanism is *predictable* in the sense that conditional on the information available to the monetary authority at time t , namely the observed distribution ξ_t of money holdings after options have been exercised, the price of the option contract A_t is deterministic. In contrast, the recursive revelation mechanism σ is not predictable, because the allocations and continuation utilities are random functions that depend on reported types. It is easy to see that this is just a version of the *taxation principle* from static mechanism design; see, for instance, Salanié (2005, p. 17).²³

(23) Intuitively, instead of asking the agent for his type, the taxation principle says you can give the agent a set of options, and ask him to choose his desired option. If this set is chosen so as to correspond to what the agent would receive in a direct revelation mechanism, then the final allocations will be the same.

- The monetary mechanism described above is *anonymous* in the sense that all that matters are money holdings. This is because the monetary mechanism maps promised utilities in the recursive mechanism σ to money holdings, and because agents are identical up to promised utilities in the recursive mechanism.
- The pricing functional \mathbf{p}_t is not required to be linear, and may well be nonlinear. Indeed, at the optimum, it is *necessarily* nonlinear.
- The choices from the option contract A_t^h determine new money holdings in the currency m_t .
- The monetary authority cannot observe individual holdings of m_t but knows the distribution ξ_t of m_t .
- Each colour of currency is used at exactly one date. In particular, money of colour t is used exactly at date t , and is used to convert wealth in period t into wealth in period $t + 1$. This conversion is typically nonlinear. In that sense, this represents a generalisation of an Arrow security, as it allows agents to transfer wealth from one date to the next, but in such a way so as to respect incentive compatibility.

6. Currency Reform and Nonlinear Pricing in Practice

Currency reforms characterized by nonlinear conversion rates are not uncommon. The replacement of the East German MARK (M) by the West German DEUTSCHE MARK (DM) implemented during the process of economic reunification of the East and West Germany in 1990 is one example. Savings accounts of the East Germans were converted at a 1M:1DM rate up to a limit of 4,000²⁴ for 15–59 year olds, a limit of 2,000 for the younger citizens and 6,000 for the older citizens. All other assets and liabilities were converted at the rate of 2M-1DM; see Bofinger (1990) for more details.

In most socialist countries, currency reforms were carried out for the purpose of undermining the wealthy and evening out the wealth distribution.²⁵ In the Soviet

(24) Approximately \$3,000 in U.S. dollars from the year 2000.

(25) As Keynes (1919, 1971) notes, ‘There is no subtler, no surer means of overturning the existing basis of society than to debauch the currency.’

Union, for example, the 1947 reform converted currency in circulation for the new currency at the rate of 10:1. Savings were converted at the rate of 1:1 on the first 3,000 RUBLES, the rate of 3:2 on the amounts between 3,000 and 10,000, and the rate of 2:1 on higher amounts; see Atlas and Drozdov (1978). Similar post-war currency reforms took place in Poland, Czechoslovakia, Romania, Bulgaria, North Korea and East Germany, all of which entailed relatively low ceilings on the amount that could be exchanged for the newly issued currency. Savings in excess of the established ceilings effectively disappeared.

In 2009, revaluation of the North Korean WON effectively wiped out the savings of many people. Only savings in the amount up to 150,000 WON in cash and 300,000 WON in bank accounts could be converted for the new currency. Excess savings could not be converted.²⁶ Of course, this being North Korea, precise details of the reform are not known. More to the point, the reform achieved its objective, namely, that of wiping out savings. This had the happy²⁷ consequence of bringing to nought the endeavours of a nascent entrepreneurial class. It goes without saying that there is a welfare loss with the North Korean reform, but there is no reason to think of the North Korean monetary authority as maximising social welfare.

More recently, in November of 2016, India decided to take some currency out of circulation in exchange for new currency partly of the same denomination. Again, as in our mechanisms, the purpose of currency reform is to ensure that agents who have not been compliant with existing laws and have unaccounted for currency, can no longer derive value from their illicit holdings. In particular, ₹500 and ₹1,000 currency notes of the Mahatma Gandhi Series were declared as invalid and could be exchanged for ₹500 and ₹2,000 banknotes of the Mahatma Gandhi New Series. The banknotes could also be exchanged over the counter of bank branches up to a limit that varied over the days: Initially, the limit was fixed at ₹4,000 per person from November 8 to 13, 2016. This limit was increased to ₹4,500 per person from November 14 to 17 of 2016. The limit was reduced to ₹2,000 per person from November 18, 2016. All exchange of banknotes was abruptly stopped from November 25, 2016.²⁸ The change had some effect on

(26) See, for instance, <http://news.bbc.co.uk/2/hi/asia-pacific/8394987.stm>.

(27) But happy only from the point of view of the rulers of North Korea.

(28) See the Wikipedia article <https://en.wikipedia.org/wiki/2016_Indian_banknote_demonetisation> for an excellent summary.

reducing illicit holdings, although the overall effect on the economy seems to have been negative. This is attributed, in part, to the haphazard administration of the currency reform.²⁹

Nonlinear pricing is widely observed in practice in other contexts. Firms offer quantity discounts and set both the quality and price of their wares. These are static policies. Firms also offer history-dependent nonlinear pricing. For instance, airlines offer frequent flier miles (that may actually depend on how much was paid for the ticket), and subsequent ticket prices can depend on the number of miles earned. Moreover, airlines also offer additional amenities as a function of the number of airline miles earned. Here, airline miles are a proxy for past purchases, and airlines incentivise consumers to fly with them again instead of a competitor by offering such pricing strategies. Another example is software licensing by Adobe and Mathematica. The user essentially buys the software for a year, after which the license expires. New users have to pay larger sums to get the product, but long-term users can renew the license at a reduced fee.

The construction of our mechanism illustrates why non-linear pricing is effective and helps interpret the characteristics of currency reforms observed in practice.

Appendix

A. Monetary Mechanism

We begin with some notation. A *period- t option contract* $A_t = \{(c_i, m_{i,t}) : c_i \geq 0 \text{ and } m_{i,t} \in \mathbb{R} \text{ for } i = 1, 2\}$, where c_i is an amount of cake for consumption and $m_{i,t}$ is an amount of currency of colour t . (Intuitively, we provide one pair $(c_i, m_{i,t})$ for each type of an agent.) The set of all period- t option contracts is denoted by \mathcal{A}_t . Before describing the mechanism, we shall sketch the timing of the interactions.

(29) See, 'India's Botched War on Cash', by Bhaskar Chakravorti, *Harvard Business Review*, December 2016, <<https://hbr.org/2016/12/indias-botched-war-on-cash#>>.

Consider period t . Every agent enters the period with an option contract, that is, a menu of choices. Each choice in this menu will give him some consumption in period t and some amount of money holdings m_t .

Each period has four stages, and the timing is as follows:

Time $t.0$: Agent $h \in \mathcal{H}$ enters the period with the option contract $A_t^h \in \mathcal{A}_t$.

Time $t.1$: Agent h realises taste shock θ_i , and makes a choice, ie, makes a choice from the set A_t^h .

Time $t.2$: The monetary authority executes the options chosen by each agent, so each agent gets some instantaneous consumption c_i^h and some amount of currency $m_{i,t}^h$.

Time $t.3$: Agent h takes his money holdings $m_{i,t}^h$ and goes to a market to buy an option contract for the next period, $A_{t+1}^h \in \mathcal{A}_{t+1}$, taking prices as given.

We are now ready to describe our monetary mechanism.

Definition A.1. A *radical monetary mechanism* consists of currencies with colours $t \in \mathbb{N}$ (and amounts in \mathbb{R}), option contracts $A_t \in \mathcal{A}_t$, a distribution of period t money holdings $\xi_t \in \mathbf{M}_t$, a period- t pricing function $\mathbf{p}_t : \mathcal{A}_{t+1} \times \mathbf{M}_t \rightarrow \mathbb{R}$ for period $t + 1$ -option contracts, in terms of currency of colour t , and an initial, period 0 allocation of option contracts $A_1^h \in \mathcal{A}_1$ for every agent $h \in \mathcal{H}$. The mechanism proceeds as follows:

Date 0 At date $t = 0$, agent h is given the option contract (ie, the menu) $A_1^h \in \mathcal{A}_1$.

Enter period $t = 1$.

Date 1.1 After agent h observes his private taste shock (at time 1.1), let him exercise his option, ie, let agent h choose $(c_i^h, m_{i,1}^h)$ from the menu A_1^h .

Date 1.2 Execute options (at time 1.2) throughout the economy, so agent h gets consumption c_i^h and money holdings $m_{i,1}^h \in \mathbb{R}$. Record resulting distribution of money holdings

ξ_1 .³⁰

Date 1.3 Fix a pricing function $\mathbf{p}_1 : \mathcal{A}_2 \times \mathbf{M}_1 \rightarrow \mathbb{R}$ for period-2 option contracts, where the price is in terms of currency of colour 1. Let agents exchange (at time 1.3) their money holdings of first period money for an option contract in \mathcal{A}_2 , with the understanding that there is no credit, ie, agent h can purchase $A_2 \in \mathcal{A}_2$ if, and only if, his money holdings m_1^h exceed $\mathbf{p}_1(A_2, \xi_1)$, the price of option contract A_2 .

End period $t = 1$.

Enter period $t = 2$

Date 2.1

Repeat as in period 1 ...

We are now in a position to prove Theorem 2.

Proof of Theorem 2. The proof consists of constructing the appropriate monetary mechanism. Fix a recursive revelation mechanism $\sigma = (f, g, \psi_0)$. As we will see below, all money holdings $m_{i,t}$ will be in the set D , and will correspond to promised utilities in an appropriate way.

At time 0, consider agent h who will receive expected utility w from the mechanism. Give agent h the option contract $A_1^h \in \mathcal{A}_1$, where

$$A_1^h := \{(c_i^h, m_{i,1}^h) : c_i^h := C(f(\theta_i, w, \psi_0)), \\ m_{i,1}^h := g(\theta_i, w, \psi_0), i = 1, 2\}$$

where ψ_0 be the original distribution of promised utilities. More generally, consider period t , and in particular, time $t.3$, after the monetary authority has executed the options. Suppose the distribution of money holdings of colour t is given by ξ_t . Construct option contracts $A_{t+1}(m_t, \xi_t) \in \mathcal{A}_{t+1}$ as follows:

$$A_{t+1}(m_t, \xi_t) := \{(c_i, m_{i,t+1}) : c_i := C(f(\theta_i, m_t, \xi_t)), \\ m_{i,t+1} := g(\theta_i, m_i, \xi_t), i = 1, 2\}$$

(30) Thus, it is because the monetary authority disburses money that it knows the distribution of money holdings at any point in time.

and let the price of such an option contract at time t .³ be $\mathbf{p}_t(A(m_t, \xi_t), \xi_t) = m_t$.

In words, under the revelation mechanism, agents enter a period with a promised lifetime utility of w . After their taste shock is realised, they choose (by virtue of their reported shock) one of the two bundle choices, where each bundle contains instantaneous consumption and continuation promised utility. Under the proposed monetary mechanism, agents enter a period with an option menu which corresponds to the two bundle choices under σ . The price for this contract at the end of the previous period corresponds to the promised utility w .

Consider, now, agent h with menu A_1^h entering period $t = 1$. If this agent has the taste shock θ_i , then his optimal (expected utility maximising) choice will be $(c_i, m_{i,1})$, for $i = 1, 2$. This follows immediately from the fact that σ is incentive compatible. Moreover, the consumption from the monetary mechanism in period 1 is exactly the same as under σ . This leads to a distribution ξ_1 of money holdings of color 1. Notice that by construction, $\xi_1 := S\psi_0 =: \psi_1$. Thus, the set of money holdings is D and the amount of issued currency is the sum of elements in D according to the distribution of money holding ξ .

In the monetary mechanism, consider the following strategy for agent h : After any history, when presented with the option contract A_t^h , choose the option $(c_i, m_{i,t})$ if the taste shock is θ_i , for $i = 1, 2$. Notice that this corresponds – in terms of outcomes – to reporting his type truthfully in the recursive revelation mechanism σ . Moreover, if there is any strategy in the monetary mechanism that can give agent h strictly greater utility than that afforded by truth-telling, such a strategy has a corresponding sequence of reports in σ . But this is not possible, because σ is incentive compatible. Therefore, in each period t , it is optimal for agent h (regardless of his choices in the past), to choose ‘truthfully’ (in the sense that it is incentive compatible for agent h with taste shock θ_i to choose $(c_i, m_{i,t})$ from his option contract), just as he would in the mechanism σ . Thus, the monetary mechanism presented implements σ . \square

References

- Andolfatto, David. 2010. 'Essential Interest-Bearing Money'. *Journal of Economic Theory* 145 (4): 1495–1507. (Cit. on p. 7).
- Athey, Susan, Andrew Atkeson and Patrick J. Kehoe. 2005. 'The Optimal Degree of Discretion in Monetary Policy'. *Econometrica* 73 (5): 1431–1475. (Cit. on p. 7).
- Atkeson, Andrew, and Robert E. Lucas Jr. 1992. 'On Efficient Distribution with Private Information'. *Review of Economic Studies* 59 (3): 427–453. (Cit. on pp. 3, 25).
- Atlas, Z. V., and V. A. Drozdov. 1978. 'Currency Reforms'. In *The Great Soviet Encyclopedia*, ed. by A. M. Prokhorov. Macmillan. (Cit. on p. 31).
- Bajaj, Ayushi, et al. 2017. 'Decentralizing Constrained-Efficient Allocations in the Lagos-Wright Pure Currency Economy'. *Journal of Economic Theory* 167:1–13. (Cit. on p. 8).
- Bofinger, Peter. 1990. 'The German Monetary Unification (GMU): Converting Marks to D-Marks'. *Federal Reserve Bank of St. Louis Review*. (Cit. on p. 30).
- Cole, Harold L., and Narayana R. Kocherlakota. 2001. 'Efficient Allocations with Hidden Income and Hidden Storage'. *Review of Economic Studies* 68 (3): 523–542. (Cit. on p. 9).
- Dong, Mei, and Janet Hua Jiang. 2010. 'One or Two Monies?' *Journal of Monetary Economics* 57:439–450. (Cit. on pp. 2, 8).
- Green, Edward J. 1987. 'Lending and the Smoothing of Uninsurable Income'. In *Contractual Agreements for Intertemporal Trade*, ed. by Edward C. Prescott and Neil Wallace. University of Minnesota Press. (Cit. on p. 9).
- Hu, Tai-Wei, John Kennan and Neil Wallace. 2009. 'Coalition-Proof Trade and the Friedman Rule in the Lagos-Wright Model'. *Journal of Political Economy* 117 (1): 116–137. (Cit. on pp. 9, 22).
- Hu, Tai-Wei, and Guillaume Rocheteau. 2015. 'Monetary Policy and Asset Prices: A Mechanism Design Approach'. *Journal of Money, Credit, and Banking* 47:39–76. (Cit. on pp. 9, 22).
- . 2013. 'On the Coexistence of Money and Higher Return Assets and its Social Role'. *Journal of Economic Theory* 148:2520–2560. (Cit. on pp. 9, 22).
- Kehoe, Timothy J., David K. Levine and Michael Woodford. 1992. 'The Optimum Quantity of Money Revisited'. In *Economic Analysis of Markets and Games*, ed. by Douglas Gale and Oliver D. Hart. Cambridge, Ma: MIT Press. (Cit. on pp. 3, 14).

- Keynes, John Maynard. 1919, 1971. *The Collected Writings of John Maynard Keynes. Volume 1: The Economic Consequences of the Peace*. Macmillan. (Cit. on p. 30).
- Kocherlakota, Narayana R. 1996. ‘Implications of Efficient Risk Sharing without Commitment’. *Review of Economic Studies* 63 (4): 595–609. (Cit. on pp. 12, 13).
- . 1998. ‘Money is Memory’. *Journal of Economic Theory* 81:232–251. (Cit. on pp. 4, 7).
- . 2003. ‘Societal Benefits of Illiquid Bonds’. *Journal of Economic Theory* 108:179–193. (Cit. on pp. 7, 8).
- . 2002. ‘The Two-Money Theorem’. *International Economic Review* 43 (2): 333–346. (Cit. on pp. 2, 4, 7, 9).
- Lagos, Ricardo, and Randall Wright. 2005. ‘A Unified Framework for Monetary Theory and Policy Analysis’. *Journal of Political Economy* 113 (3): 463–484. (Cit. on pp. 8, 9).
- Levine, David K. 1991. ‘Asset Trading Mechanisms and Expansionary Policy’. *Journal of Economic Theory* 54 (1): 148–164. (Cit. on pp. 5, 6, 17, 19–21, 24).
- Lucas, Robert E., Jr. 1978. ‘Equilibrium in a Pure Currency Economy’. *Economic Inquiry* 18:203–220. (Cit. on pp. 6, 7, 14).
- Lucas, Robert E., Jr. 1992. ‘On Efficiency and Distribution’. *Economic Journal* 102 (411): 233–247. (Cit. on pp. 7, 9, 14).
- Myerson, Roger B. 1991. *Game Theory: Analysis of Conflict*. Harvard University Press. (Cit. on p. 10).
- Ostroy, Joseph M. 1973. ‘The Informational Efficiency of Monetary Exchange’. *American Economic Review* 63 (4): 597–610. (Cit. on p. 1).
- Ravikumar, B., and Neil Wallace. 2001. *A Benefit of Uniform Currency*. Tech. rep. Pennsylvania State University. (Cit. on p. 22).
- Salanié, Bernard. 2005. *The Economics of Contracts: A Primer*. 2nd edition. MIT Press. (Cit. on p. 29).
- Sannikov, Yuliy. 2008. ‘A Continuous-Time Version of the Principal-Agent Problem’. *Review of Economic Studies* 75 (3): 957–984. (Cit. on p. 9).
- Wallace, Neil. 2014. ‘Optimal Money Creation in ‘Pure Currency’ Economies: A Conjecture’. *Quarterly Journal of Economics* 129 (1): 259–274. (Cit. on pp. 20, 21).
- . 2010. ‘The Mechanism Design Approach to Monetary Theory’. In *Handbook of Monetary Economics, Volume 3*, ed. by Benjamin M. Friedman and Michael Woodford, 3–23. Elsevier. (Cit. on pp. 2, 7).