Sixteenth-Century Replacement Costs of Coins: Implications for Optimal Financing and Divisibility*

Manjong Lee and Neil Wallace
Department of Economics, The Pennsylvania State University

October 30, 2004

Abstract

Replacement costs of money that approximate those in the sixteenth century are embedded in a parameterized matching model of money. Two issues are explored numerically: Financing of the costs by a user fee is compared with financing by a lump-sum tax; and the optimal size (denomination) of monetary items is determined. For some parameters, the user fee is better; while for others, the lump-sum tax is better. The optimal size of a monetary item is large and in order of magnitude is similar to the size of the single silver coin that was minted in France in the sixteenth century.

JEL classification #: E40, E42

1 Introduction

Currency—whether it consists of intrinsically valuable coins, as it did in most of the world prior to the nineteenth century, or of tokens, as it has in most of the world since then—wears out and, therefore, must be replaced. The older coinage systems differed from modern systems in two respects. First, replacement costs were higher—in part, because the currency was made of

*We are indebted to John Kareken and to Christopher Phelan for helpful discussions based on an earlier version of this paper.

stuff that was valuable. Second, the costs were financed by user fees; an owner of a worn coin bore the cost of replacing it with a new coin. In this paper, we insert into a parameterized matching model of money replacement costs of money that approximate what they were in the sixteenth century and explore two issues. We compare welfare under financing of the costs by a user fee with welfare under financing by a lump-sum tax. And we find the optimal degree of divisibility of currency.

Although neither the financing of currency replacement costs nor the degree of divisibility of currency arouses much attention in modern economies, both seemed to be important issues under the earlier coinage systems. Jevons devoted a substantial part of his well-known treatise, *Money and the Mechanism of Exchange,* to the financing question. In the Preface, he listed it among a short list of "questions that press for solution:"

> How long shall we in England allow our gold coinage to degenerate in weight? Shall we recoin it at the expense of the State or of the unlucky individuals who happen to hold light sovereigns? [2, page viii]

(As this suggests, Jevons came down on the side of re-coinage at the expense of the State, a policy we translate loosely as financing by lump-sum taxation.) The degree to which currency was indivisible has received much more attention; there is widespread agreement among historians that prior to the nineteenth century currency was not available in conveniently small denominations.[1]

In the matching model we study, the stock of money is exogenous and constant and is composed of a single denomination. The single denomination approximates the situation in France in the sixteenth and seventeenth centuries when, at any given time, just one silver coin was being minted (see Redish [5]).[2] Aside from the background matching model—which provides a simple representation of the benefits of greater divisibility of money—the critical assumptions are about the production and depreciation of money.

We make the somewhat extreme assumption that money "depreciates"

[1] See, for example, Sargent and Velde [6].

[2] In addition, one gold coin was being minted. We ignore the very valuable gold coin in our analysis. See Lee, Wallace, and Zhu [3] for a model with multiple denominations of currency. There, however, replacement costs are ignored.

only with use—when it is traded.[3] However, if depreciation actually occurs, then the money that a seller receives is more worn than the money that the buyer surrenders and the distribution of money by wear is a state variable of the model. In order to avoid that, we assume that money in the model does not depreciate. Instead, we assume that some perishable output is lost to nature when money passes from a buyer to a seller—the loss being such that the money is maintained in its pristine (undepreciated) form. In the user-fee version of the model, that loss is financed by having the seller produce more output than the buyer receives, the difference being the loss to nature implied by the transaction. In the lump-sum tax version, the loss to nature is financed by a per period lump-sum tax spread evenly over the population.

We assume that the loss to nature has two components: one component varies directly with the number of monetary items (the labor cost) and the other does not (the material cost). The first gives rise to a trade-off between the benefits of additional divisibility for trade and the costs of maintaining a more divisible stock of money.[4] We parameterize the two components using mainly two bits of historical (sixteenth century) information: the depreciation rate of money and the number of coins that a person working at a mint could produce.

The model gives ambiguous results regards financing. For high frequencies of trade, the user fee is better; for low frequencies of trade, the lump-sum tax is better. The results for optimal divisibility are more uniform. The optimal size of the monetary unit is large and in order of magnitude is similar to the size of the silver coin being minted in France in the sixteenth century. We measure divisibility by the ratio of the mean amount of money to the size of monetary item and find that the optimal ratio for different parameterizations ranges from 11 to 57. As a basis for comparison, if the mean (per household) amount of currency in the U.S. is \$100, then the current ratio—\$100 in pennies—is $10,000$. In other words, a ratio of 20 corresponds roughly to having a \$5 bill be the smallest U.S. denomination.

[3]An alternative extreme assumption is that money depreciates with age and not with use. A more general model would include both sources of depreciation.

[4]There are other ways to limit optimal divisibility. In [3], a carrying cost per monetary item is assumed.

2 The model

Aside from costs of maintaining the stock of money, the model is the version of the Shi [7] and Trejos and Wright [10] models studied by Zhu [11]: indivisible money with a sufficiently large upper bound on individual money holdings and take-it-or-leave-it offers by buyers.

Time is discrete. There is a unit measure of each of $N \geq 3$ types of infinitely lived agents and there are N distinct produced and perishable types of divisible goods at each date. A type n agent, $n \in \{1, 2, ..., N\}$, produces type n good and consumes only type $n+1$ good (modulo N). Each agent maximizes expected discounted utility with discount factor $\beta \in (0, 1)$. The realized utility in a period for a type n agent who consumes $y_{n+1} \in \mathbb{R}_+$ and who produces $y_n \in \mathbb{R}_+$, is $u(y_{n+1}) - y_n$ where $u : \mathbb{R}_+ \to \mathbb{R}$ is strictly concave, differentiable, increasing, satisfies $u(0) = 0$, and is such that there exists $y' > 0$ such that $u(y') = y'$.

There are three exogenous nominal quantities that describe the stock of money: $(s, \bar{s}, S)$, where s is the size of a unit of money, $\bar{s}$ is the average (per type) amount of money, and S is the bound on individual holdings, a bound that is needed to achieve compactness. We assume that S/s is an integer. The set of possible individual money holdings is $\mathbb{Z} = \{0, s, 2s, ..., S\}$ and money is the only asset.

At each date, each person meets one other person at random. In a single-coincidence meeting, the only relevant kind of meeting, the buyer makes a take-it-or-leave-it offer—either a deterministic offer (the non-lottery version), or a lottery offer (the lottery version).

If $x \in \mathbb{Z}$ is the amount of money traded in a meeting, then that meeting contributes a loss to nature of θx amount of goods, where $\theta \in \mathbb{R}_+$. This is the cost of maintaining the stock of money. In the user-fee version, the seller's output in the meeting must exceed the buyer's consumption by θx. Under financing by lump-sum taxes, the buyer consumes what the seller produces and each person in the model pays a lump-sum tax in goods to finance the maintenance cost. Although agents in the model treat θ as given, we assume that θ depends on a pair of parameters given by the technology for production of money, (θ_1, θ_2), and on $(s, \bar{s})$ according to

$$\theta(s, \bar{s}) = \frac{\theta_1}{s} + \frac{\theta_2}{\bar{s}}. \tag{1}$$

Equivalently, we have

$$\theta(s,\bar{s})\bar{x} = \frac{\theta_1 \bar{x}}{s} + \frac{\theta_2 \bar{x}}{\bar{s}}. \tag{2}$$

where $\bar{x} > 0$ is the average amount of money traded per period.

In (2), the first term on the right-hand side reflects the per monetary item cost, which we later associate with a labor cost of producing money. With $\theta_1 > 0$, a decrease in s holding $\bar{x}$ and $\bar{s}$ constant, implies an increase in the goods lost in trade. This reflects the assumption that there is a per item cost of production of money independent of the size of the monetary item. But (2) is also consistent with neutrality in the following sense. Let economy A have the vector of nominal quantities $(s, \bar{s}, S)$, let economy λA have the vector $\lambda(s, \bar{s}, S)$, where λ is a positive integer, and let A and λA be otherwise identical—and, in particular, have the same (θ_1, θ_2). According to (1), $\theta(\lambda s, \lambda \bar{s}) = \theta(s, \bar{s})/\lambda$. Therefore, if (y_c, y_p, x) is an equilibrium trade in economy A when the buyer has money m and the seller has m', where y_c is the buyer's consumption, y_p is the seller's production, and x is the amount of money, transferred from the buyer to the seller, then $(y_c, y_p, \lambda x)$—the same real trade—is a candidate equilibrium trade in economy λA when the buyer has λm and the seller has $\lambda m'$. That is, according to (2), $\theta(s, \bar{s})\bar{x} = \theta(\lambda s, \lambda \bar{s})\lambda \bar{x}$.

3 Steady states and welfare

A symmetric steady state, symmetric across the N specialization types, consists of a money distribution, $\pi : \mathbb{Z} \to [0, 1]$, and a value function, $w : \mathbb{Z} \to \mathbb{R}$, where $\pi(z)$ is the fraction of people with money z at the start of a period and $w(z)$ is the expected discounted value of beginning a period with money z. Because this model is essentially identical to Zhu [11] (for the non-lottery version) and to that in [3] (for the lottery version), we present only the part of the definition that is special to our version; namely, the buyer's problem in the user-fee version.

Problem 1 *Consider a single-coincidence meeting in which the buyer has money m and the seller has m' when the value function for the next period is w. Let $\Omega(m, m')$ be the set of lotteries over $\mathbb{Z}$ that places zero weight on infeasible offers of money; namely, offers that exceed $min\{m, S - m'\}$. The buyer chooses y and $\sigma \in \Omega(m, m')$ to maximize $u(y) + \beta E_{\sigma(x)} w(m - x)$ subject* *to*

$$-y - \theta E_{\sigma(x)} x + \beta E_{\sigma(x)} w(m' + x) \geq \beta w(m'). \quad (3)$$

This is the lottery choice problem. (A lottery over output received by the buyer is not considered because it is never an optimal choice for the buyer.) The non-lottery version is the special case in which σ is required to be degenerate. And for the lump-sum tax version, the term in θ in the constraint is dropped. Evidently, financing in the user-fee version is via a sales tax, or, equivalently in this model, a tax on production.

A solution to problem 1 exists and satisfies (3) with equality. Hence, the seller's payoff is $\beta w(m')$. If w is concave, then the buyer's payoff is concave in m. Consider the mapping from the Cartesian product of a set of (next period) values functions and a set of (current period) distributions into that set implied by the set of solutions to problem 1—the usual mapping studied in heterogeneous-agent models. A steady state is a (w, π) that is a fixed point of that mapping.[5]

The main analytical result for such models is the existence result due to Zhu [11] (for the non-lottery version), which is extended to the lottery version in [12] (see also [3]): If $u'(0)$, $\bar{s}/s$, and $S/\bar{s}$ are sufficiently large and if θ is sufficiently small, then there exists a steady state with w strictly increasing and strictly concave and with π having full support.[6] Because w is strictly increasing, this steady state has trade.

Our welfare comparisons are made across steady states, while taking into account, where appropriate, the one-time costs of adjusting the initial stock of money. The comparisons are comparisons of expected utility prior to the assignment of money to people, where the assignment is made according to the respective steady-state distribution of money. For the comparison across financing methods, no adjustment has to be made. There, we view the economy as beginning with a given stock of money, a given vector $(s, \bar{s}, S)$. Then we choose both the financing method and an initial distribution of money over people, subject to the restriction that the initial distribution

[5]With some standard adjustments, the numerical algorithm we use computes iterates of the mapping just described. Although convergence is not guaranteed and although uniqueness, even in the class described in the existence result, is not guaranteed, the iterates always converge and different initial conditions do not seem to produce different limits.

[6]The proof uses the fact that the mapping preserves concavity of value functions. That property of the mapping depends on take-it-or-leave-it offers by buyers.

be a steady state for the financing method. The comparisons for different degrees of divisibility require an adjustment. There, we view the economy as beginning only with $(\bar{s}, S)$—as if it begins with some stock of "raw silver" that will then be converted into coins. Then (for a given financing method) we choose s and an initial distribution of money over people, subject, again, to the restriction that the initial distribution be a steady state for that s. As explained in detail later, for such comparisons, we adjust welfare by the different one-time costs of producing a stock of money composed of units of size s.

Aside from the one-time adjustment, given a steady state, (w, π), we view its associated welfare to be the inner product πw. It is easy to show that

$$N(1-\beta)\pi w = \pi R \pi' - \theta \pi X \pi'. \tag{4}$$

Here, r_{ij}, the row i, column j element of R, is given by

$$r_{ij} = u(y_{ij}) - y_{ij} \tag{5}$$

where y_{ij} is the buyer's consumption when the buyer enters trade with $i \in \mathbb{Z}$ and the seller with $j \in \mathbb{Z}$ and x_{ij}, the row i, column j element of X, is the average amount of money transferred when the buyer enters trade with $i \in \mathbb{Z}$ and the seller with $j \in \mathbb{Z}$. This partitioning of welfare holds under both the user-fee and lump-sum-tax versions. An upper bound on $\pi R \pi'$ is $\max_y[u(y) - y]$.

Our welfare cost measures are computed in terms of consumption equivalents as follows. Suppose economy 2 has lower welfare than economy 1. Then we report $\Delta/[\pi^{(2)} Y^{(2)} \pi^{(2)\prime}]$, where $y_{ij}^{(2)}$, the row i, column j element of $Y^{(2)}$, is the buyer's consumption in economy 2 when the buyer enters trade with $i \in \mathbb{Z}$ and the seller with $j \in \mathbb{Z}$ and where Δ satisfies

$$N(1-\beta)\pi^{(1)} w^{(1)} = \pi^{(2)} R^{(c)} \pi^{(2)\prime} - \theta \pi^{(2)} X^{(2)} \pi^{(2)\prime}, \tag{6}$$

with $r_{ij}^{(c)}$, the row i, column j element of $R^{(c)}$, given by

$$r_{ij}^{(c)} = u(y_{ij}^{(2)} + \Delta) - y_{ij}^{(2)}. \tag{7}$$

Notice that the compensating amount of consumption, Δ, is additive rather than multiplicative. We make it additive because consumption is 0 in some meetings and those are the meetings where an addition to consumption is most valuable. Notice also that the compensating consumption is given only to buyers (consumers in single-coincidence meetings) and that its magnitude is reported relative to average buyer consumption in economy 2, not to per capita consumption.

4 The choice of parameters

First we describe the parameters for the background matching model; N, u, and β. We assume $N = 3$, the smallest magnitude consistent with our assumptions, and $u(y) = y^{1/2}$. This choice of u makes the relevant first-order condition in the lottery version a linear equation. (It implies that $\max_y[u(y) - y] = .25$ and is attained at $y = .25$.) We study three different values for β, which we regard as arising from an annual discount factor of 0.9 and different trading frequencies: bi-weekly meetings, monthly meetings, and bi-monthly meetings. Because all of the above selections are fairly arbitrary, our findings should be regarded as suggestive.

As regards $(s, \bar{s}, S)$, we set $S = 4\bar{s}$, which is large enough so that the results would not be affected if a higher S were imposed.[7] We take $\bar{s}/s$ from data for France in the second half of the sixteenth century. According to Redish [5], the French were minting an almost pure silver coin in the second half of the sixteenth century, a coin that was approximately 1/3 ounce of silver. We need to measure that size relative to an estimate of $\bar{s}$. Spooner [9] provides estimates of the total population for France and of the total amount of coin. Call the former P and the latter M in ounces of silver.[8] Because it seems reasonable to treat an agent in the model as a household, not as a person, we divide P by 5 to get a rough estimate of households. If we treat each household as a model agent, then $\bar{s} = 5M/P$. This implies $\bar{s}/s = 12$. However, treating every household in France as engaged in search for monetary trade seems extreme. Hence, we also use $\bar{s}/s = 24$, which assumes that half of the households held all the money and were engaged in such trade.

Before we describe our choices for (θ_1, θ_2), we report some steady states for the above choices and $\theta = 0$.[9] This description serves two purposes. First, we use these results as the benchmark relative to which we measure welfare

[7] We selectively checked that setting $S = 7\bar{s}$ does not affect the results.
[8] According to Spooner, the population of France was 16 million in 1600.
[9] A detailed description of the computational routine is given in the Appendix.

costs. Second, it helps us interpret the later results by providing information about how the above parameters—the different meeting frequencies and the two choices for $\bar{s}/s$—and the lottery and non-lottery versions affect the steady state.

We begin by displaying some distributions for one combination of parameters. Figure 1 presents distributions of consumption over single-coincidence meetings for monthly meetings and for $\bar{s}/s = 24$. Recall that the *first-best* distribution is degenerate at .25. The lottery distribution seems to be centered near the *first-best* distribution, while the non-lottery distribution is centered to the right of it. As is well-known from other examples for such models, take-it-or-leave-it-bargaining can produce outputs that are too large.[10] That sets the stage for the possibility that a user fee will improve welfare.

Figure 1: Consumption distributions over single-coincidence meetings: monthly meetings, $\bar{s}/s = 24$, and $\theta = 0$.

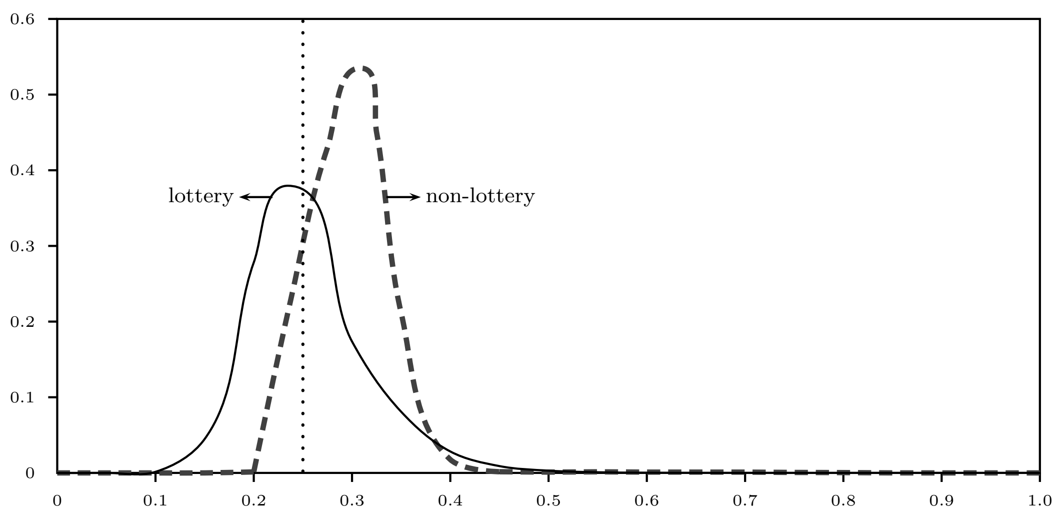

[10]Lotteries were introduced into matching models of money by Berentsen, Molico, and Wright [1]. In a model with a unit upper bound on money holdings, they showed that the output in pairwise-core lottery trades are bounded above by the *first-best* level. That result, however, does not generalize to more general money holdings. (This can be inferred from Molico [4]. He studies take-it-or-leave-it offers numerically under an approximation to divisible money and finds that output can be higher than the first best in some meetings.)

In Figure 2, we show the distributions of money for the steady states that correspond to the Figure 1 consumption distributions. Although the upper bound is, in fact, binding, there is a very little measure at or near the upper bound. Notice that the lottery version is more tightly bunched around the mean.

Figure 2: Money distributions: monthly meetings, $\bar{s}/s = 24$, $s = 1$, and $\theta = 0$.

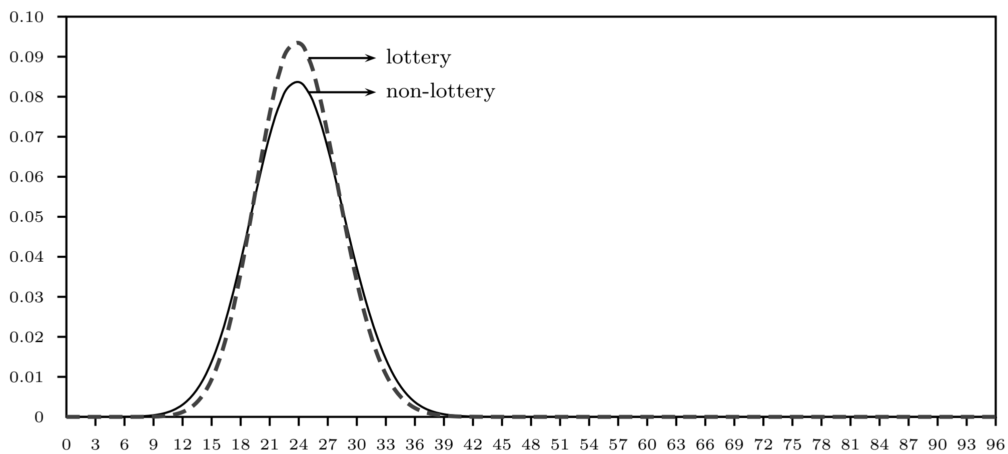

Table 1 contains summary statistics for all the parameter combinations. There, we report the average amount of money traded relative to the size of a unit, $\bar{x}/s$, and average consumption, $\bar{y}$.

Table 1: $(\bar{x}/s, \bar{y})$ for $\theta = 0$

	non-lottery			lottery		
$\bar{s}/s$	bi-week	month	bi-month	bi-week	month	bi-month
12	$(1.0, 0.73)$	$(1.0, 0.51)$	$(1.0, 0.27)$	$(0.3, 0.25)$	$(0.6, 0.25)$	$(1.0, 0.23)$
24	$(1.0, 0.55)$	$(1.0, 0.28)$	$(2.0, 0.22)$	$(0.5, 0.25)$	$(1.1, 0.24)$	$(2.2, 0.21)$

Notice that more frequent meetings (lower discounting of the future) produce larger average output (more pronounced in the non-lottery version) and (weakly) smaller average amounts of money traded. The lower discounting, other things the same, makes sellers more willing to incur current disutility in exchange for future consumption. As we will see, the patterns displayed here persist even when we introduce replacement costs financed by the user fee.

Now we turn to our choices for (θ_1, θ_2). Our plan is to first obtain θ and then to partition it between θ_1 and θ_2 in a way that is consistent with (2). For given parameters, including θ, a steady state for the user-fee version implies a flow of output devoted to maintenance of the stock of money; namely, $\theta\bar{x}$. For the purpose of choosing θ, we assume that the cost of producing money per unit is equal to its average value in exchange in the model; namely, $\bar{y}/\bar{x}$. This cost and the flow of output devoted to maintenance of the stock implies gross investment in money equal to $\theta(\bar{x})^2/\bar{y}$. Dividing that by the given stock, $\bar{s}$, implies a percentage rate of production of money. We choose θ to equate that percentage rate to the depreciation rate estimated by Spooner [9], denoted δ. (Spooner provides an estimate of the half-life of coins in France in the years $1500-1750$ that implies an annual depreciation rate of 2.3%.[11]) That is, we choose θ to satisfy

$$\theta = \frac{\delta \bar{s} \bar{y}(\theta)}{[\bar{x}(\theta)]^2} \tag{8}$$

for each combination of the other parameters, where, as indicated, we take into account the dependence of $\bar{x}$ and $\bar{y}$ on θ.[12]

Table 2 contains the estimates of θ. As noted above, $\bar{x}/s$ and $\bar{y}$ are not much different from what as they are in Table 1. Some perspective on these magnitudes of θ can be obtained by considering the maintenance cost as a percentage of output; namely, $\theta\bar{x}/\bar{y}\times 10^2$. This percentage and the associated θ vary considerably with the parameters. The lowest percentage cost is .60% and the highest is 2.45%.

Given θ, we obtain θ_1 as follows. The gross investment in money, $\theta(\bar{x})^2/\bar{y}$, implies that $\theta(\bar{x})^2/(\bar{y}s)$ monetary items are produced when each item is of size s. The labor cost of such production is given by

$$(\frac{\theta(\bar{x})^2}{s\bar{y}})(\frac{\bar{y}}{q}) = \frac{\theta(\bar{x})^2}{sq}, \tag{9}$$

[11]We convert an annual depreciation rate to a bi-weekly, monthly, and bi-monthly rate to get δ; namely, $\delta = 0.001$ for bi-weekly meetings, $\delta = 0.002$ for monthly meetings, and $\delta = 0.004$ for bi-monthly meetings.

[12]As described in the Appendix, the right-hand side of (8) is almost constant. That makes it easy to find θ that satisfies (8).

Table 2: Estimates of $\theta \times 10^2$, $(\bar{x}/s, \bar{y})$ and $[\theta\bar{x}/\bar{y} \times 10^2]$

	non-lottery			lottery		
$\bar{s}/s$	bi-week	month	bi-month	bi-week	month	bi-month
12	0.433	0.599	0.604	1.770	0.792	0.450
	$(1.0, 0.72)$	$(1.0, 0.50)$	$(1.0, 0.27)$	$(0.3, 0.24)$	$(0.6, 0.24)$	$(1.1, 0.22)$
	[0.60]	[1.20]	[2.34]	[2.09]	[1.97]	[2.26]
24	0.642	0.627	0.261	0.825	0.414	0.202
	$(1.0, 0.54)$	$(1.0, 0.27)$	$(2.0, 0.21)$	$(0.6, 0.25)$	$(1.1, 0.23)$	$(2.2, 0.21)$
	[1.20]	[2.37]	[2.45]	[1.99]	[2.11]	[2.17]

where $\bar{y}$ is the average wage in terms of output according to the model and q is the rate at which a person can produce coins. (According to Spulford [9], a person employed at a mint could produce approximately 10 coins per day. We, therefore, set q to be 50 per week.[13]) Then, we equate the expression in (9) with $\theta_1\bar{x}/s$, the first term on the right-hand side of (2), to get

$$\theta_1 = \frac{\theta\bar{x}}{q}. \tag{10}$$

Given θ_1 and θ, θ_2 is obtained from (1). Table 3 contains the associated estimates of θ_1 and θ_2.

Table 3: Estimates of $(\theta_1 \times 10^5, \theta_2 \times 10^2)$

	non-lottery			lottery		
$\bar{s}/s$	bi-week	month	bi-month	bi-week	month	bi-month
12	$(4.3, 2.6)$	$(3.0, 3.6)$	$(1.6, 3.6)$	$(1.5, 10.6)$	$(1.5, 4.8)$	$(1.3, 2.7)$
24	$(6.4, 7.7)$	$(3.2, 7.5)$	$(2.5, 3.1)$	$(3.0, 9.8)$	$(2.7, 5.0)$	$(2.5, 2.4)$

The estimates in Table 3 play a role when we study the model's implications for divisibility.

[13]All the computations were also done for $q = 75$ per week. They produced almost identical results.

5 User-fee versus lump-sum tax

Here we use only the estimates of θ. For each parameter combination and the associated θ, we compare welfare under user-fee and lump-sum-tax financing. Before presenting the welfare comparison in terms of consumption equivalents, we present results for the two terms on the right-hand side of (4), $\pi R\pi'$ and $\theta\pi X\pi'$. These are given in Table 4. In the table, $\theta\pi X\pi'$ is presented in parentheses below the corresponding $\pi R\pi'$.

Table 4: $\pi R\pi' \times 10^2$ and $\theta\pi X\pi' \times 10^2$ (in parenthesis)

		lump-sum			user-fee		
	$\bar{s}/s$	bi-week	month	bi-month	bi-week	month	bi-month
non-lottery	12	12.28	20.35	24.69	12.80	20.66	24.73
		(0.439)	(0.604)	(0.617)	(0.433)	(0.598)	(0.617)
	24	19.00	24.80	24.40	19.60	24.81	24.21
		(0.653)	(0.637)	(0.523)	(0.642)	(0.635)	(0.512)
lottery	12	24.05	24.10	24.20	23.99	24.09	24.16
		(0.456)	(0.449)	(0.472)	(0.509)	(0.482)	(0.478)
	24	24.57	24.59	24.50	24.57	24.56	24.44
		(0.440)	(0.470)	(0.447)	(0.492)	(0.472)	(0.446)

For most cases in the non-lottery version, the welfare comparison between user-fee and lump-sum-tax financing is dominated by the differences between the corresponding $\pi R\pi'$ terms. Moreover, except in one case, $\pi R\pi'$ is larger under user-fee financing. This happens because the non-lottery trades under the lump-sum tax produce outputs that are on average too high. The results in the lottery version are quite different. There the financing method matters less.

Underlying the results in Table 4 are distributions of consumption. In Figure 3, we present those for monthly meetings and for $\bar{s}/s = 24$. As expected, financing by the user fee gives a distribution that is shifted somewhat to the left. In the non-lottery version, it is clearly a better distribution, better in the sense of being closer to the *first-best* distribution.

Table 5 contains the welfare costs in consumption equivalents relative to the steady state with $\theta = 0$. The results summarize what we presented in

Figure 3: Consumption distributions over single-coincidence meetings: monthly meetings and $\bar{s}/s = 24$.

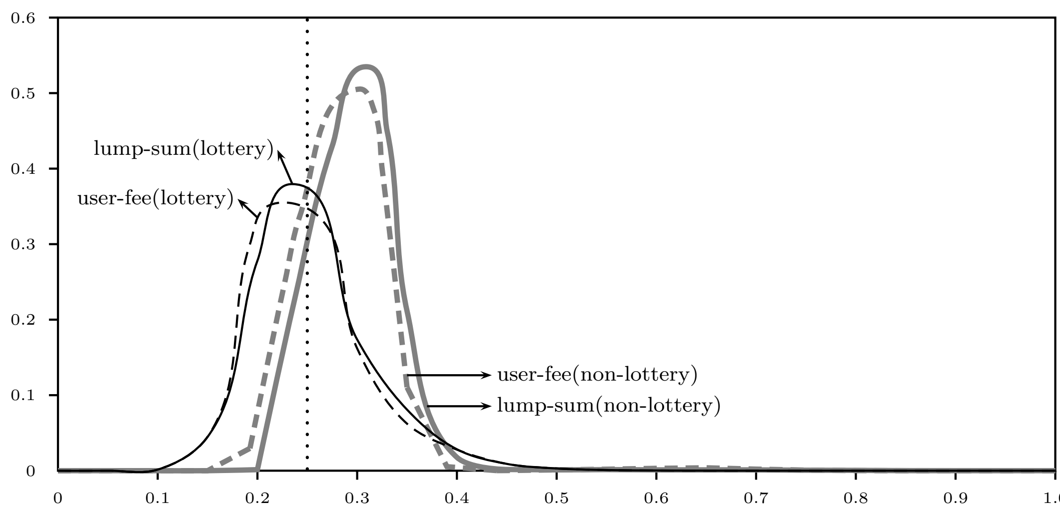

Table 4. Lump-sum taxation is better in the lottery version and is worse in the non-lottery version except in one case.

Table 5: Consumption equivalent welfare costs relative to $\theta = 0$

		lump-sum			user-fee		
	$\bar{s}/s$	bi-week	month	bi-month	bi-week	month	bi-month
non-lottery	12	7.02	5.80	4.13	1.32	2.50	3.83
	24	6.98	5.18	4.22	2.60	4.68	4.89
lottery	12	3.85	3.73	3.83	4.59	3.82	4.14
	24	3.64	3.83	3.58	3.77	4.13	4.11

6 The optimal size of coins

Here we use the Table 3 estimates of (θ_1, θ_2) to find the optimal s for each parameter combination and each financing scheme. That is, for each (θ_1, θ_2) pair in Table 3, we set $\bar{s} = 1$, $S = 4\bar{s}$ and find a steady state and its associated welfare for each $\bar{s}/s$ in the set $\{3, 4, 5, ...\}$. Then we find the s that gives the highest welfare. In doing this we take into account that as we vary $\bar{s}/s$ in the set $\{3, 4, 5, ...\}$, there is one-time cost of producing the additional divisibility. According to our description of the technology for coin production, the one-time cost depends on $\bar{s}/s$ according to $(\bar{s}/s)(\bar{y}/q)$, where $\bar{y}$ is determined in the model. Because $(\bar{s}/s)(\bar{y}/q)$ is a one-time cost per person, the welfare associated with a given $(\bar{s}/s)$ is the associated magnitude of πw as given in (4) minus $(\bar{s}/s)(\bar{y}/q)$.

Before we present the overall results, we present in Figure 4 a detailed picture of some of the underlying output. There we show for one set of parameters in the non-lottery, user-fee version three functions of $\bar{s}/s$: one function is $\pi R \pi'$ (the left-hand vertical scale), another is $\theta \pi X \pi' = (\frac{\theta_1}{s} + \frac{\theta_2}{\bar{s}})\bar{x}$ (the right-hand vertical scale), and a third is $N(1-\beta)(\bar{s}/s)(\bar{y}/q)$ (the right-hand vertical scale), where π, R, and X are all functions of s.

Figure 4: Welfare components as functions of $\bar{s}/s$: non-lottery, user-fee, monthly meetings, and θ for $\bar{s}/s = 24$.

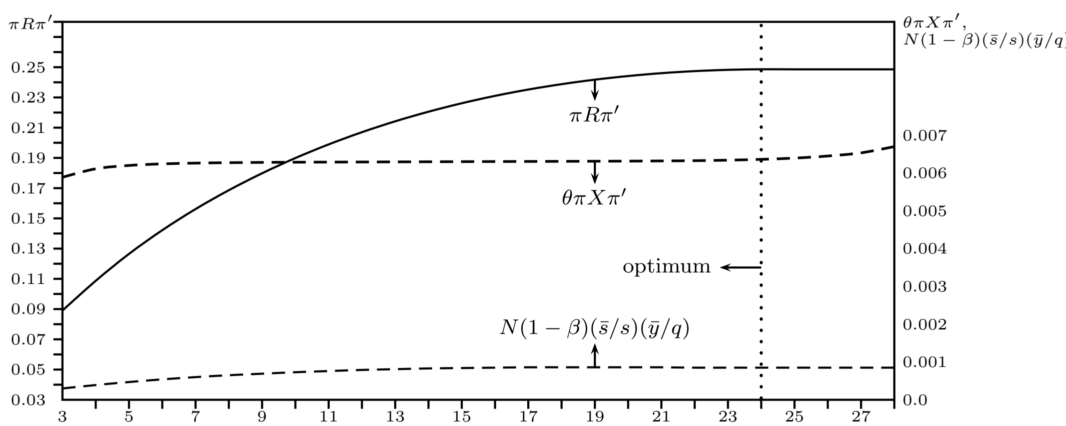

The optimum occurs where $\pi R \pi' - \theta \pi X \pi' - N(1-\beta)(\bar{s}/s)(\bar{y}/q)$ is a maximum, in this case at $\bar{s}/s = 24$.

Behind the shape of the $\theta\pi X\pi'$ function is the following behavior. At magnitudes of $\bar{s}/s$ between 3 and 5, $\bar{x}$ is increasing as more units of money enter trade. Then, for $\bar{s}/s$ between 5 and 25, $\bar{x}$ is almost constant and the slope of the function is θ_1, which is positive but very small. Then for higher $\bar{s}/s$, $\bar{x}$ again increases. However, by then, the benefits of additional divisibility as measured by the function $\pi R\pi'$ have largely disappeared.

The results for the optimal $\bar{s}/s$ are given in Table 6. (The column in Table 6 headed by ($\bar{s}/s$) is used only to indicate the (θ_1, θ_2) estimate from the corresponding cell in Table 3.)

Table 6: Optimal $\bar{s}/s$

		non-lottery			lottery		
	$\bar{s}/s$	bi-week	month	bi-month	bi-week	month	bi-month
lump-sum	12	57	25	12	17	18	11
	24	57	25	12	24	20	12
user-fee	12	56	24	12	17	18	11
	24	55	24	12	23	20	11

Although the optimal size varies with the parameters, the somewhat startling result is that the optimal size turns out to be not all that different from what we estimated the actual size to be in the second half of the sixteenth century in France. Thus, although by modern standards such sizes give rise to a very indivisible currency, the above model suggests that those sizes may have been close to optimal.

7 Concluding remarks

We have made many assumptions. The assumptions regarding the background matching model were chosen quite arbitrarily. We simply took an off-the-shelf model and, essentially, off-the-shelf parameters for it. We devoted more care to the assumptions about replacement costs.

Our numerical findings regarding whether it is better to finance replacement costs with a user fee or with a lump-sum tax are ambiguous. Nor, in a sense, do they deal with Jevons' concerns. Jevons thought that different degrees of wear for coins produced a recognizability problem under user-fee financing. Because no such recognizability problem appears in our model, we cannot claim to have addressed his concerns about user-fee financing of replacement costs.

Our findings regarding the optimal size of a money in a one-denomination system are more uniform. The optimal degree of divisibility is low and in order of magnitude is similar to what was in place in sixteenth-century France. In this regard, it is worth noting that from 1555 to 1700, the size of the single silver coin being minted in France rose from roughly one-third of an ounce to one ounce. Over this period, there is roughly a 12-fold increase in the stock of money and a 7-fold increase in the price level.[14] Before we formulated our model of replacement costs, we wondered why France did not keep the absolute size of the coin fixed and take advantage of the inflation to have a substantially more divisible currency. Our cost structure, which, as we noted above, is consistent with neutrality, goes some way toward explaining why that was not done.

So far as we know, this is the first study of currency replacement costs and their implications for how they should be financed and for the optimal degree of divisibility of money. The background matching model implies how trades depend on the financing method and on the degree of divisibility of money. The replacement costs provide the magnitude of what is to be financed and how that magnitude varies with additional divisibility. Together, they give conclusions about optima. The new ingredient relative to previous work on these issues is the use of a model that is able to depict the benefits of additional divisibility of money.

8 Appendix

The computational procedure is essentially to iterate on the mapping mentioned in section 3.[15] We start with an initial guess, denoted $g^{(0)} = (w^{(0)}, \pi^{(0)})$, where $w^{(0)}$ is interpreted as an end-of-period value function and $\pi^{(0)}$ as a beginning-of-period distribution. For the given $w^{(0)}$, we find the best deterministic offer for each single-coincidence meeting by a way of global search over all feasible offers of money in problem 1. (Although, in principle, the solution need not be unique, not surprisingly, we never find multiplicity.) Call

[14] Price level is estimated in terms of wheat price in Paris.

[15] A copy of the computer code is available on request.

this amount of money $x(m, m')$. For the non-lottery version, the function $x(m, m')$ can be used to generate an (end-of-period) distribution, $\pi^{(1)}$. In addition, the implied payoff for the buyer and $\pi^{(0)}$ give a (beginning-of-period) value function, $w^{(1)}$. In the lottery version, the best lottery offer is obtained by finding the best deterministic offer from the pseudo-value function determined by linear interpolation of $w^{(0)}$. Concavity of $w^{(0)}$ implies that the best lottery is either in $[x(m, m'), x(m, m')+1)$ or $(x(m, m')-1, x(m, m')]$, where $x(m, m')$ is the best deterministic offer. With our choice of $u(y) = y^{1/2}$, each of the two possibly relevant first-order conditions is a linear equation. Given the best lottery, the distribution $\pi^{(0)}$ can be updated to produce a $\pi^{(1)}$ and the value function $w^{(1)}$ can be computed.

The values for the $(i+1)$-th iteration, denoted $g^{(i+1)} = (w^{(i+1)}, \pi^{(i+1)})$, are formed by a weighted average of the elements $g^{(i)}$ and $g^{(i-1)}$, where different weights are used for the different components. This updating process is repeated until the components of g satisfy the convergence criterion: $\max_j \left|(w_j^{(i+1)} - w_j^{(i)})/w_j^{(i)}\right| < 10^{-4}$ for w and $(\sum_j (\pi_j^{(i+1)} - \pi_j^{(i)})^2)^{1/2} < 10^{-4}$ for π.

To obtain the estimates of θ, we looked for a θ that satisfies (8). We display the right-hand side function of θ for monthly meetings in Figure 5. These functions are such that a solution to (8) is easily obtained. Similar shapes turn up for the other meeting frequencies.

Figure 5: The right-hand side of (8) for monthly meetings.

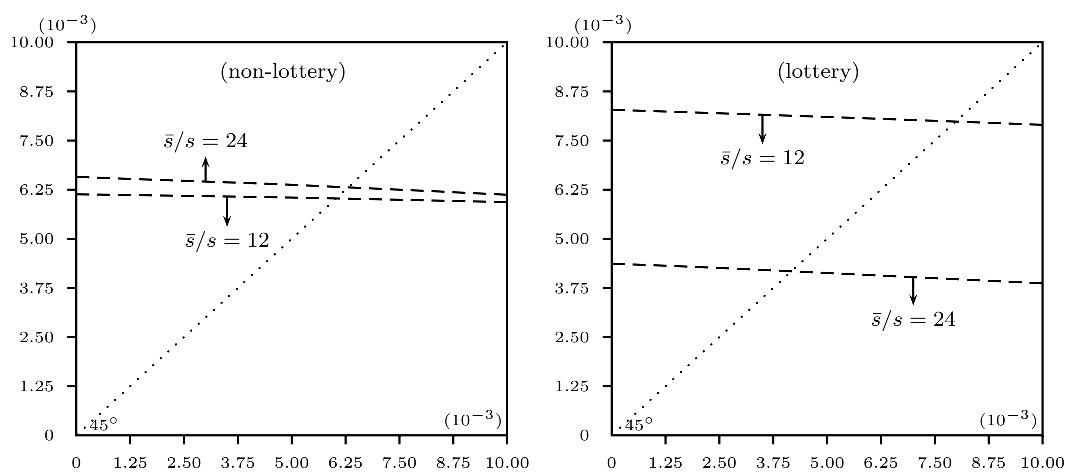

References

[1] Berentsen, A., M. Molico, and R. Wright, Indivisibilities, lotteries, and monetary exchange, *J. of Econ. Theory* 107 (2002) 70-94.

[2] Jevons, W., *Money and the Mechanism of Exchange*, London: D. Applegate and Company, 1899.

[3] Lee, M., N. Wallace and T. Zhu, Modeling Denomination Structures, manuscript, The Pennsylvania State University, 2004.

[4] Molico, M., *The Distribution of Money and Prices in Search Equilibrium*, Ph.D. Thesis, University of Pennsylvania, Philadelphia, 1997.

[5] Redish, A., *Bimetallism: An Economic and Historical Analysis*, Cambridge University Press, 2000.

[6] Sargent, T., and F. Velde, *The Big Problem of Small Change*, Princeton University Press, 2002.

[7] Shi, S., Money and prices: a model of search and bargaining, *J. of Econ. Theory* 67 (1995) 467-498.

[8] Spooner, F., *The International Economy and Monetary Movements in France, 1493-1725*, Harvard University Press, 1972.

[9] Spulford, P., Mint organisation in late medieval Europe, *Later Medieval Mints: Organisation, Administration and Techniques,* BAR International Series 389 (1988), 7-29.

[10] Trejos, A., and R. Wright, Search, bargaining, money and prices, *J. of Political Economy* 103 (1995) 118-141.

[11] Zhu, T., Existence of a monetary steady state in a matching model: indivisible money, *Journal of Economic Theory*, 112, 2003, 307-324.

[12] Zhu, T., and N. Wallace, Coexistence of Money and Higher-Return Assets as a Bargaining Norm, manuscript, The Pennsylvania State University, 2004.