Heterogeneity and the Welfare Cost of Inflation[1]

Paola Boel   Gabriele Camera

Bowdoin College   University of Iowa

July 25, 2008

Abstract

We quantify the welfare cost of inflation in a calibrated heterogeneous-agent model of the U.S. economy. There are three main findings. Inflation has a small social cost; for instance, on average agents would give up less than one percent consumption to avoid ten percent inflation. Second, the distribution across the population of the social cost of inflation depends on the source of heterogeneity. If agents differ in their labor productivity, then inflation does not redistribute monetary wealth, though it hurts the more productive and benefits the less productive. Instead, if agents differ in trading shocks, then there is equilibrium dispersion in monetary wealth and inflation has a redistributive effect. Third, the direction of wealth redistribution depends on whether money is the only asset in the model. If it is, then inflation benefits the poor—who hold less-than-average balances—and hurts the rich. The converse is true if agents can insure against consumption risk with a competing asset.

Keywords: Money, Heterogeneity, Friedman Rule, Welfare cost of inflation

JEL codes: E4, E5

[1]We thank Julián P. Díaz, Paul E. Hanouna, Ricardo Lagos, Randall Wright as well as seminar participants at the Cleveland Fed Workshop on Money, Banking and Payments, the Midwest Macroeconomic Meetings, and the SAET 2007 meetings.

1 Introduction

A considerable amount of theoretical work, based on disparate modeling approaches, supports the notion that efficiency in a monetary economy hinges on a deterministic deflationary policy known as the 'Friedman rule.' Yet, in practice deflationary policies are not implemented and low predictable inflation is widely tolerated (and sometimes advocated).

This discrepancy has motivated a literature aimed at quantifying the social cost of inflation. Recent papers on this subject can be sorted into one of two strands. Some studies have a representative-agent structure that display explicit micro foundations, the so-called matching model of money in which money has an explicit medium of exchange function and there is no role for private credit. These studies usually abstract away from wealth heterogeneity, or assume money is the only store of value, or must do some heavy lifting to compute analytically complex monetary distributions; for instance, see [11, 20, 23]. A second strand includes works based on models in which money has a more "descriptive" role, though it is not the only store of value. Examples include cash-in-advance models with costly credit, inventory-theoretic or precautionary balances models; e.g., see [1, 2, 14].

Our work ties these two strands of literature. We enrich the typical matching model introducing forms of heterogeneity that generate tractable equilibrium distributions of wealth and then accomplish a quantitative analysis. To sum up our findings, the analysis suggests that, once we account for wealth heterogeneity, inflation has quantitatively significant redistributive effects that depend on type of heterogeneity and financial structure. Initially we work under the (usual) assumption that money is the only asset. We quantify the impact of inflation on social welfare and trace inflation's redistributive effects for two typical sources of heterogeneity. The average welfare cost of inflation is not far from earlier studies but the redistributive implications of such a "money-only" model seem empirically implausible. Consequently, we introduce the option to insure against consumption risk by means other than simply holding money. The quantitative impact of inflation does not vary much in this augmented model, but the redistributive implications are more reasonable.

The model is based on the two sequential markets model in [20], with agents who trade in large anonymous competitive markets, as in [3, 9]. In the benchmark model agents can hold only cash to insure against consumption risk and are ex-ante heterogeneous in their trading risk. In stationary competitive equilibrium different agent types hold different money balances, since the ones who are more likely to trade hold more cash. This heterogeneity disappears if nominal interest rates converge to zero. We also consider a variant in which agents ex-ante differ in labor productivity. This case is interesting because there is no dispersion of money holdings but only in earning profiles. We calibrate the two variants of the benchmark model to the U.S. economy and quantify the welfare cost of inflation. For the representative agent, we find that ten percent inflation is worth less than one percent of consumption. This is in line with previous studies based on a variety of models such as in [12, 16, 20, 21, 25]. The burden of inflation, however, depends on the distribution of types and the kind of heterogeneity considered. If agents differ in labor productivity, then everyone suffers from inflation, though the burden differs across agent types. This is because agents can costlessly adjust their labor effort across the two markets, in each period. Instead, if agents differ in their trading risk, then the impact of inflation is qualitatively different. Positive inflation can in fact be welfare-increasing for a segment of the population, i.e., those who have less than average money balances. For these agents, inflation generates a beneficial redistribution of wealth because lump-sum money transfers more than offset their inflation-tax burden. This result is in line with the quantitative findings in [11] and the theoretical intuition developed in [5], in the same class of models. However, the redistributive implications of inflation appear to be empirically implausible; for instance, see the discussions in [2, 14].

We thus augment the model by introducing an additional nominal asset, as an alternative to cash. This asset can provide consumption insurance, much as money, but it can better shield agents from the inflation tax. In this setting we find a quantitatively similar impact of inflation on average welfare to the earlier analysis, but a redistributive effect that is more in line with empirical evidence.

The remainder of the paper is organized as follows. Section 2 presents the model economy. Section 3 introduces the definition of efficiency considered throughout the paper. Section 4 derives the stationary monetary equilibrium allocations. Section 5 discusses the quantitative analysis for the case of a representative agent and for the different types of heterogeneity studied. Section 6 concludes.

2 The model

The model is a variation of the one in [20]. Time is discrete, and the horizon is infinite. There is a large population of heterogeneous infinitely-lived agents who want to consume perishable goods and discount only even to odd dates. Thus, we work with trading cycles indexed by $t = 1, 2, ...$ each including an odd and an even date. There are infinitely many spatially separated trading groups, each of which defines a market as in [9]. Every market includes infinitely many anonymous agents who have never met before (a suitable matching process is described in [3]). Thus, in each trading cycle agents may visits two anonymous markets, denoted 'one' in the odd date and 'two' in the even.

On every date a single perishable consumption good can be generated by producers, i.e., agents who can supply labor to a technology that transforms each unit of labor into one good. Everyone can produce and consume on even dates. Instead, at the start of each odd date agents draw i.i.d. trading shocks determining whether the agent can produce, consume, or do neither (idle). Consuming or producing are assumed to be equally likely. Hence, on odd dates agents face idiosyncratic trading (consumption) risk, but not on even dates. In addition, we assume ex-ante heterogeneity, which can take one of two forms; agents can either differ in their odd-date trading shocks or productivity. For convenience we divide the population into two types $j = H, L$ in proportions ρ and $1 - \rho$, respectively.

Even-date preferences are assumed homogeneous and quasilinear. An agent of type j who consumes $q_j \geq 0$ goods and supplies $x_j \geq 0$ labor in market two (equivalently, produces x_j goods) has utility $U(q_j) - x_j$. On odd dates consumers of any type j derive utility $u(c_j)$ from $c_j \geq 0$ consumption. Producers of type j suffer $\phi_j(y)$ disutility from producing y goods, so type L agents must work longer than type H to produce the same amount of output. The functions u, ϕ_j and U are twice continuously differentiable, strictly increasing, with $u'' < 0$, $\phi'_j \geq 0$ and $U'' < 0$. Also, $u(0) = U(0) = \phi_j(0) = 0$ and we attach a star to the quantities that uniquely solve $u'(c) = \phi'_j(c)$ and $U'(q) = 1$. For heterogeneity in trading shocks we let $\phi_j(y) = \phi(y)$ for $j = H, L$, while α_j is the probability of trading on market one for any type j agent, with $0 < \alpha_L < \alpha_H \leq 1$. For heterogeneity in productivity we fix $\alpha_j = \alpha \in (0, 1)$ for $j = H, L$, and let $\phi'_H(y) < \phi'_L(y)$ for each $y \geq 0$. Agents are price takers.

We impose a (standard) assumption of limited enforcement and limited commitment. This means that agents have exclusive rights to their assets and endowments, so that trading plans must be compatible with individual incentives. This together with the frictions assumed above implies an essential role for money (see [3]) since on odd dates trade is *quid pro quo* but consumers cannot produce. Thus, a consumption shock on odd dates corresponds to a need for liquidity.

For the moment we assume there is only one asset, fiat money. A government exists that is the sole supplier of fiat currency, of which there is an initial stock $\bar{M} > 0$. The money stock evolves deterministically at gross rate π by means of lump-sum cash transfers at the beginning of even dates.

3 Stationary monetary allocations

Consider the allocation selected by a planner who maximizes the agents' lifetime utilities, treating them identically, and constrained by the same physical and informational restrictions faced by agents. Such allocation, called the efficient allocation, is unique and stationary across trading cycles.[2] Indeed, the planner would equate the marginal rates of substitutions of the different types of agents, on each date. In what follows we thus focus on stationary monetary outcomes. These are outcomes in which consumption is invariant across trading cycles and the sequence of nominal prices evolves so that the money stock has constant positive real value.

To simplify notation we omit t subscripts and use a prime to identify next-cycle vari-

[2] This is the same allocation that would arise if agents could coordinate and commit to a non-monetary trading plan on each odd date, before realizing their individual shocks. See the Appendix.

ables, when necessary. Accordingly, p_1 and p_2 denote the nominal price of goods on odd and even dates (markets one and two) of an arbitrary trading cycle t. We also normalize nominal variables by p_2, so that trades in market one occur at real price $p = \frac{p_1}{p_2}$. The timing of events during cycle t for the arbitrary agent of type j is as follows. He enters cycle t with real money holdings $m_j \geq 0$, saved in the preceding cycle. Subsequently, trade occurs and after market one closes the agent enters market two on the even date with $m_{j,k}$ real balances, where $k = n, s, b$ denotes the idiosyncratic trading shock experienced in market one (n if idle, b for buyer, s for producer).

Real money holdings evolve within the cycle according to

$$\begin{aligned} m_{j,b} &= m_j - pc_j \\ m_{j,s} &= m_j + py_j \\ m_{j,n} &= m_j \end{aligned} \tag{1}$$

In market one, a buyer spends pc_j and a producer earns py_j. Cash left over can be used in market two, when the real price is one, q_j is consumption bought and $x_{j,k}$ is production sold by an agent who experienced shock k (the notation q_j is without loss in generality, see [9]). In market two, agents also save money to self-insure against consumption shocks. Let $m'_j \geq 0$ denote real balances held at the start of the next trading cycle.

In a stationary monetary economy real balances must be positive and constant, i.e., $m'_j = m_j > 0$. If M is cash at the start of a cycle and $M' = \pi M$ is cash available in market two, then

$$\tfrac{p'_2}{p_2} = \tfrac{M'}{M} = \pi \ , \tag{2}$$

i.e., the inflation rate equals the rate of growth of money. This rate is controlled by means of per-capita lump-sum transfers τ in market two, so if all agents of type j hold the same amount of money m_j, then the government budget constraint can be written as

$$\tau = [\rho m_H + (1 - \rho) m_L](\pi - 1). \tag{3}$$

Stationarity and money market clearing imply that real balances available in market one must equal the real money stock, denoted $\bar{m}$, at each date, i.e.

$$\tfrac{M}{p_2} = \bar{m} = \rho m_H + (1 - \rho) m_L. \tag{4}$$

3.1 Trade on even dates

Given the recursive nature of the problem, we use a dynamic programming approach to describe the problem faced by the representative agent of type j on any date. We let $V_j(m_j)$ be the expected lifetime utility of this agent when he starts the trading cycle with m_j real balances before trading shocks are realized. We let $W_j(m_{j,k})$ be the expected lifetime utility from entering an even date with $m_{j,k}$ real balances.

The agent's budget constraint at the start of an even date is:

$$x_{j,k} = q_j + \pi m'_j - (m_{j,k} + \tau) \tag{5}$$

The resources available to the agent in market two partly depend on the realization of the trading shock k, as he carries $m_{j,k}$ real balances from market one. Other resources are $x_{j,k}$ receipts from current sales of goods and the lump-sum transfer τ.[3] These resources can be used to finance current consumption q_j, or simply to save $\pi m'_j$ real money balances. Notice that short-selling is not allowed, and agents can save only with money and cannot lend to each other. The factor $\pi = \frac{p'_2}{p_2}$ multiplies m'_j because the budget constraint lists current real values. Indeed, the real rate of return on monetary savings is $1/\pi$.

The agent's problem at the start of an even date can be represented as follows:

$$W_j(m_{j,k}) = \max_{q_j, m'_j \geq 0} \{U(q_j) - q_j - \pi m'_j + m_{j,k} + \tau + \beta V_j(m'_j)\} \tag{6}$$

It follows that in a stationary monetary economy

$$\frac{\partial W_j(\omega_{j,k})}{\partial m_{j,k}} = 1 \quad \text{for } j = H, L. \tag{7}$$

The result hinges on the linearity of production disutility and the use of competitive pricing, linear in the quantity sold. It follows that the marginal value of money reflects the price of one unit of real balances, which is of course one. The economic implication is the marginal valuation of real balances in market two neither hinge on the agent's type j, nor his (monetary) wealth, i.e., trading history.

So, we have

$$W_j(m_{j,k}) = \quad W_j(0) + m_{j,k}, \tag{8}$$

[3]Notice that $x_{j,k} \geq 0$ so we must verify that this is true for all k in equilibrium.

i.e., the agent's expected value from having portfolio $m_{j,k}$ at the start of an even date is the expected value from having no wealth $W_j(0)$, plus the current real value of $m_{j,k}$ real balances. This implies agents' savings choice m'_j is independent of trading histories as in [20]. However, it is important to realize that, unlike [20], different agent types might save different amounts because $W_j(0)$ may differ.

To see this, note that everyone consumes identically in market two since (6) implies

$$q_j = q^* \quad \text{for } j = H, L. \tag{9}$$

The reason is agents in market two can produce any amount at constant marginal cost. Goods market clearing on even dates requires

$$q^* = (1-\rho)[\tfrac{\alpha_L(x_{L,s}+x_{L,b})}{2} + (1-\alpha_L)x_{L,n}] + \rho[\tfrac{\alpha_H(x_{H,s}+x_{H,b})}{2} + (1-\alpha_H)x_{H,n}]. \tag{10}$$

Given (9) we write

$$W_j(m_{j,k}) = U(q^*) - q^* + m_{j,k} + \tau + \max_{m'_j \geq 0} [-\pi m'_j + \beta V_j(m'_j)]. \tag{11}$$

The central implication is that the choice of savings hinges on the expected marginal benefit of carrying real balances in market one. In turn, since market one trades depend on the availability of money balances, then efficiency will depend on the agents' choice of savings m'_j. This is studied next.

Given that we are focusing on monetary outcomes, i.e., $m'_j > 0$, we must have:

$$1 = \frac{\beta}{\pi} \times \frac{\partial V_j(m'_j)}{\partial m'_j} \tag{12}$$

Recalling that one unit of real balances buys one unit of consumption, the left hand side of the expression simply defines the marginal cost of real balances. The right hand side is the expected marginal benefit from holding money discounted according to time preferences and inflation. The expected benefit of holding money will generally differ across types j. To see how, we must study trades on odd dates.

3.2 Trade on odd dates

Consider an agent of type j with m_j real balances at the start of market one. His expected lifetime utility must satisfy

$$\begin{aligned} V_j(m_j) = \quad & \max \tfrac{\alpha_j}{2}[u(c_j) + W_j(m_{j,b})] + \max \tfrac{\alpha_j}{2}[-\phi_j(y_j) + W_j(m_{j,s})] \\ & +(1-\alpha_j)W_j(m_{j,n}), \end{aligned} \tag{13}$$

where the maximization is over $c_j \leq \frac{m_j}{p}$ as a buyer and $y_j \geq 0$ as a producer.

The optimal choice y_j of a producer must satisfy

$$-\phi_j'(y_j) + \frac{\partial W_j(m_{j,s})}{\partial m_{j,s}} \frac{\partial m_{j,s}}{\partial y_j} = 0. \tag{14}$$

Since $\frac{\partial W_j(m_{j,s})}{\partial m_{j,s}} = 1$ from (7), and $\frac{\partial m_{j,s}}{\partial y_j} = p$ from (1), then

$$p = \phi_j'(y_j), \quad \text{for} \quad j = H, L. \tag{15}$$

Hence, agents of different productivity produce different amounts.

Now consider a buyer. Recall from (1) that $m_{j,b}$ depends on c_j. The first order condition for the buyer's problem, omitting the multiplier on the inequality constraint, is

$$u'(c_j) + \tfrac{\partial W_j(m_{j,b})}{\partial m_{j,b}} \tfrac{\partial m_{j,b}}{\partial c_j} \geq 0.$$

Since $\frac{\partial W_j(m_{j,b})}{\partial m_{j,b}} = 1$ from (7), $\frac{\partial m_{j,b}}{\partial c_j} = -p$ from (1) and $p = \phi_j'(y_j)$ from (15), we get

$$u'(c_j) \geq p. \tag{16}$$

If the constraint is not binding, then $u'(c_j) = p$, solved uniquely by $c(p) > 0$, independent of j, so any unconstrained buyer spends $m^* = pc(p)$. If the constraint is binding, then $u'(c_j) > p$. The buyer of type j consumes $c_j < c(p)$ and spends $m_j < m^*$. Thus, c_j will never exceed $c(p)$. That is,

$$c_j = \min\{\tfrac{m_j}{p}, c(p)\}. \tag{17}$$

Efficiency of the allocation depends on whether some agents are cash constrained. Indeed, the planner's allocation satisfies $u'(c_j) = \phi_j'(y_j)$, which can be sustained only if $c_j = c(p)$, since $p = \phi_j'(y_j)$.

To find optimal savings of type j we use (1) and (8) in (13) to obtain

$$V_j(m_j) = m_j + \tfrac{\alpha_j}{2}[u(c_j) - \phi_j(y_j)] + \tfrac{\alpha_j}{2}p(y_j - c_j) + W_j(0) \tag{18}$$

where c_j satisfies (17). The expected lifetime utility $V_j(m_j)$ depends on the agent's real wealth m_j and two additional elements. First, the expected surplus from market one trades. With probability $\alpha_j/2$ the agent is a buyer and spends pc_j real balances to enjoy utility $u(c_j)$. With probability $\alpha_j/2$ the agent is a producer, earns py_j real balances but suffers disutility $\phi_j(y_j)$, where $p = \phi_j'(y_j)$. Second, there is the continuation payoff $W_j(0)$. We emphasize that the third term on the RHS of (18) does not appear in the representative-agent formulation in [20]. The change in wealth expected from market one trades, $p(y_j - c_j)$, is invariably zero in a representative agent model since $y = c$ by market clearing. Here, instead, agents may have unequal real balances or productivity. Hence, some may produce amounts that differ from what they would consume. Indeed, goods market clearing on odd dates implies

$$\alpha_H \rho y_H + \alpha_L(1-\rho)y_L = \alpha_H \rho c_H + \alpha_L(1-\rho)c_L. \tag{19}$$

From (18) we can calculate the expected marginal value of money:

$$\tfrac{\partial V_j(m_j)}{\partial m_j} = 1 + \tfrac{\alpha_j}{2}[u'(c_j) - p]\tfrac{\partial c_j}{\partial m_j} \tag{20}$$

where $\frac{\partial c_j}{\partial m_j} = \frac{1}{p}$ if the agent is liquidity constrained (zero otherwise). It follows that $V_j(m_j)$ is strictly concave in cash holdings if buyer j is liquidity constrained, and linear otherwise. For a cash-constrained buyer $m_j < m^*$ and the marginal value of money is

$$\tfrac{V_j(m_j)}{\partial m_j} = 1 + \frac{\alpha_j}{2}\left[\tfrac{u'(c_j)}{p} - 1\right], \tag{21}$$

which depends on the marginal utility of consumption and is decreasing in m_j since $u'' < 0$. If instead an agent is not cash constrained, the marginal value of money is constant and equal to one, the real value of one unit of money.

We can now provide a definition of equilibrium.

Definition 1 *Given an initial money stock $\bar{M} > 0$ and a government policy specified by (π, τ), a competitive stationary monetary equilibrium is a time-invariant list of real* *quantities* $(c_j, y_j, q, x_{jk}, m_j)$ *and cycle-dependent prices* $(p_{1,t}, p_{2,t})$ *that solve the agent's problems (6) and (13), satisfy (15), the government budget constraint (3), and market clearing (4), (10), and (19).*

Using the expressions (11) and (18) we define ex-ante welfare for a type j in stationary equilibrium:

$$\begin{aligned}(1-\beta)V_j(m_j) = \quad & \tfrac{\alpha_j}{2}[u(c_j) - \phi(y_j)] + \tfrac{\alpha_j}{2}p(y_j - c_j) + U(q^*) - q^* \\ & +(\pi - 1)(\bar{m} - m_j).\end{aligned} \tag{22}$$

Here, c_j, y_j and m_j are optimal consumption, production and money holdings, $p = \phi_j'(y_j)$ and $\bar{m} = \rho m_H + (1-\rho)m_L$, i.e., the real money supply equals average real balances. The expression in (22) is standard, except for two terms that do not appear in the representative agent model in [20]. First, there is an expected *change in wealth* from market one trades, $\frac{\alpha_j}{2}p(y_j - c_j)$. Low productivity agents might produce less than they consume, for example. Second, there is a *redistributive effect* of inflation $(\pi-1)(\bar{m}-m_j)$, i.e., the (lump-sum) transfer minus the inflation tax on money holdings. This term is nonzero if agents of different types hold unequal balances. Those with less than average balances $\bar{m}$ end up with a net transfer, and the others with a net tax. That is, inflation redistributes real balances from the top to the bottom of the distribution. Since money is usually the only asset in this class of models, it follows that inflation can only redistributes wealth from the "rich" to the "poor." Indeed, a similar result emerges in the random matching model in [11], which studies the redistributive impact of anticipated inflation in a model with richer distributions of money balances, and in [5], that studies the redistributive effect on unanticipated inflation.

3.3 Heterogeneity in trading shocks

In this section we let agents differ only in the probability of market-one trading, setting $0 < \alpha_L < \alpha_H \leq 1$ and $\phi_j(y) = \phi(y)$ for $j = H, L$. Using (15) we have $p = \phi'(y)$, so everyone produces the same amount y of output.

In a monetary economy $m_j > 0$. Using (12) and (21) we have the Euler equation

$$\tfrac{\pi}{\beta} = 1 + \tfrac{\alpha_j}{2}\left[\tfrac{u'(c_j)}{p} - 1\right] \text{ for } j = H, L. \tag{23}$$

The equation is akin to an arbitrage condition. On the left hand side we have the (gross) nominal interest rate on an illiquid bond (e.g., see [5, 9, 20]). On the right hand side we have the nominal yield on money, one, plus its expected liquidity premium. It is positive because $u'(c_j) \geq \phi^{'}(y)$ from (17) and it arises because money is needed to trade in market one. This premium grows with the severity of the cash constraint and the likelihood of consumption shocks.

Letting $i = \frac{\pi}{\beta} - 1$ denote the net nominal interest rate and $p = \phi^{'}(y)$, we write (23) as

$$i = \frac{\alpha_j}{2} \left[\frac{u'(c_j)}{\phi^{'}(y)} - 1 \right] \quad \text{for } j = H, L. \tag{24}$$

We will use (24) extensively in our quantitative analysis. It defines two equations in two unknowns, c_H and c_L, which can be uniquely determined as a function of the model's parameters and the interest rate i, which summarizes policy parameter in our model.

Policy affects the return of money the choice of monetary savings, and therefore consumption in market one. The next result immediately follows.

Lemma 1 *In any stationary equilibrium we must have $\pi \geq \beta$, i.e., $i \geq 0$. A unique stationary monetary equilibrium exists for $\pi > \beta$ and it is such that $m_L < m_H < m^*$, hence $c_L < c_H < c(p)$. As $\pi \to \beta$ we have $m_j \to m^*$ and $c_j \to c(p)$ for all j.*

Proof. By way of contradiction, suppose a monetary equilibrium exists with $\pi < \beta$. From (23) we need $\pi \geq \beta + \beta(\alpha_j/2)[u'(c_j)/\phi^{'}(y) - 1] \geq \beta$. This is in contradiction with $\pi < \beta$. So, let $\pi > \beta$. From (23), an outcome with $m_j > 0$ must satisfy

$$\pi = \beta \left\{ 1 + \frac{\alpha_j}{2} \left[\frac{u'(c_j)}{\phi^{'}(y)} - 1 \right] \right\} \quad \text{for } j = H, L.$$

As $\pi \to \beta$ then $u'(c_j) \to \phi^{'}(y)$ for $j = H, L$. Therefore, $m_H \to m^*$ and $m_L \to m^*$. By concavity of u, if $\pi > \beta$, then $u'(c_j) > \phi^{'}(y)$ for all j and so $c_L < c_H < c(p)$ and $m_L < m_H < m^*$. That is market one buyers are unconstrained only if $i = 0$. Existence follows from inspection of the individual optimality and market clearing conditions.■

Lemma 1 has several implications. First, only one steady state equilibrium exists for $\pi > \beta$. Intuitively, the rate of return on money $\frac{1}{\pi}$ cannot exceed the shadow interest rate $\frac{1}{\beta}$. If that were the case, then agents would want to keep accumulating money, which is not a stationary monetary equilibrium. Second, the equilibrium distribution of money has two mass points, with type H agents holding more money than type L. Intuitively, type H agents are more likely to buy (as well as to sell) in market one, so they self-insure more against consumption shocks by holding more money. However, notice that every agent is liquidity constrained as long as $i > 0$, in which case the allocation is inefficient. Third, as nominal interest rates approach zero, agents become indifferent between having a dollar today or one tomorrow. In this case, trade-frequency considerations do not enter saving decisions, hence all money holdings converge to the average value m^*. Clearly, this allocation is efficient as it satisfies $u'(c_j) = \phi'(y)$ for $j = H, L$.

3.4 Heterogeneity in productivity

In this section we assume agents differ only in market one productivity, i.e., we assume $\phi'_L(y) > \phi'_H(y)$ for each $y \geq 0$, and fix $\alpha_j = \alpha \in (0,1)$ for $j = H, L$. Clearly, $p = \phi'_j(y_j)$ from (15), which implies $y_L < y_H$ for all $p > 0$. Intuitively, agents are price takers so the least productive will sell less goods.

In a monetary outcome $m_j > 0$, and optimal savings must satisfy

$$\tfrac{\pi}{\beta} = 1 + \tfrac{\alpha}{2}\left[\tfrac{u'(c_j)}{p} - 1\right] \text{ for } j = H, L, \tag{25}$$

which is similar to (23). Since equation (25) has to hold for all j, we have $c_j = c$ for $j = H, L$ and therefore we have

$$i = \tfrac{\alpha}{2}\left[\tfrac{u'(c)}{\phi'_j(y_j)} - 1\right] \text{ for } j = H, L. \tag{26}$$

Since $c_H = c_L = c$ and $y_H > y_L$ from (15), equation (5) implies that $x_{Hs} < x_{Ls}$, i.e. type L agents must work more in market two than the more productive type H agents, to make up for lower sales in market one. As before, $\pi \geq \beta$ in any stationary monetary equilibrium. Therefore, we prove the following result.

Lemma 2 *Consider $\pi > \beta$. A unique stationary monetary equilibrium exists and it is such that $m_j = m < m^*$ and $c_j = c < c(p)$ for all j. As $\pi \to \beta$ we have $m \to m^*$ and $c \to c(p)$.*

Proof. Consider (26). Clearly $c_H = c_L = c$ for any $i \geq 0$. Hence, $m_H = m_L = m < m^*$. If $i > 0$ (i.e., $\pi > \beta$), then $c < c(p)$, so $m < m^*$. As $\pi \to \beta$ then $u'(c) \to p = \phi'(y_j)$, implying that $c \to c(p)$ and $m \to m^*$. Existence easily follows from inspection of the individual optimality and market clearing conditions.■

The key difference from the earlier case of heterogeneous trading shocks, is that this heterogeneous agent economy does not display equilibrium heterogeneity in money balances. This is because agents have identical preferences over consumption and face the same consumption shocks. Therefore they identically optimally self-insure holding identical money balances. The main consequence is that consumption is identical for all agent types. Once again, we find that for any $i > 0$ both agent types are liquidity constrained, i.e. $c_j = c < c(p)$ for $j = L, H$. As $i \to 0$ money holdings converge to m^* and the allocation is efficient. Notice that this means some agents produce more than others, depending on their productivity since $u'(c) = \phi_j'(y_j)$ for $j = H, L$.

4 Quantitative analysis of the benchmark model

In this section we calibrate the model and then proceed with the quantitative analysis. We start by considering a representative agent economy in order to determine the value of the preference parameters common across agents (the type-specific parameters will be considered average values). Then, we quantify the welfare cost of inflation for a representative agent. This is helpful as a benchmark and comparison to related models. Once this is done, we re-introduce heterogeneity and quantitatively study the redistributive impact of inflation. Throughout the analysis for simplicity we report the welfare cost of x percent inflation as a comparison to an economy with no inflation, unless otherwise specified.

Calibration of common parameters. In the representative agent model $\alpha_j = \alpha$ and $\phi_j(y) = \phi(y)$ for $j = H, L$. It is straightforward that in monetary equilibrium the relative price p satisfies $p = \phi'(y)$ and $pc = m$, and $c = y$ satisfies the agent's Euler equation

$$i = \tfrac{\alpha}{2}\left[\tfrac{u'(c)}{\phi'(y)} - 1\right]. \tag{27}$$

We consider standard functional forms (e.g., see [4, 20]): $u(c) = \frac{c^{1-a}}{1-a}$ with $a > 0$ and

$\phi(y) = \frac{y^\delta}{\delta}$ with $\delta \geq 1$; $U(q) = A\ln(q)$, which implies $q^* = A$. Now, using (27) and $c = y$ we find c as a function of the model's parameters and the nominal interest rate i

$$c = \left(\frac{\alpha}{2i+\alpha}\right)^{\frac{1}{\delta+a-1}}. \tag{28}$$

We consider a yearly model mainly to facilitate comparison with [20, 22].

The vector of parameters to identify is

$$\Theta = (\alpha, \beta, a, A, \delta).$$

We can easily assign numbers to two β and a. The annual rate of time preference is the standard value 0.04 (e.g., see [20]), so we fix $\beta = 0.96$. The parameter a is set to 0.71, following the recent empirical study on risk aversion in [24]. The remaining parameters require some more thought.

To pin down δ, notice it corresponds exactly to the elasticity of disutility of labor with respect to labor effort (derivation in the Appendix). The elasticity of labor supply with respect to p (the real wage in our model) is $\frac{1}{\delta-1}$. Therefore, we set δ to match average elasticity of labor supply with respect to own wage in the U.S.. However, estimates of the elasticity of labor supply vary according to the group considered (e.g., male versus female). From [15], estimates of labor supply elasticities are 0.0 for men and 0.80 for women. Consequently, we set δ to match the average of the two values with weights given by the proportion of men (0.55) and women (0.45) in the labor force for the period 1960-2006 as reported by the Bureau of Labor Statistics. We get $\delta = 3.78$.

The parameter α is set so that the theoretical interest elasticity of money demand, denoted by ε_m, matches -0.226, which is the estimated interest elasticity of $M1$ reported in [4]. In the Appendix we demonstrate that

$$\varepsilon_m = \frac{2i\phi'(y)}{\alpha c u''(c)},$$

so for the functional forms selected $\varepsilon_m = -\frac{2i}{(2i+\alpha)a}$. We measure i by the average nominal annualized yield on U.S. short-term commercial paper, which amounts to 0.044 for the period considered (1929-2006). Consequently, the calibrated value of α is 0.427. (Minimization of the distance between the data and the estimated elasticity yields similar calibrated values for α).

Last, we determine A to fit the real balances-income ratio $L = \frac{M}{PY}$, where P is the nominal price level, M is $M1$, and Y is real output. As suggested in [20, 22], the value L can be interpreted as money demand because real balances M/P are proportional to real spending Y with a factor of proportionality $L(i)$ that depends on the nominal interest rate i. For the empirical counterpart of L, we consider U.S. data for the sample period 1929-2006. We measure P by the GDP deflator and Y by real GDP.[4]

To construct the theoretical expression for L in the model we proceed as follows. Nominal output is $\frac{\alpha}{2}p_1c$ in the first market and $p_2q^* = p_2A$ in the second market, so aggregate nominal output is $PY = p_1\frac{\alpha}{2}c + p_2A$. From (4), we know that in equilibrium the nominal money stock is $M = p_2m$, so normalizing by p_2 we have $L = \frac{m}{\frac{\alpha}{2}pc+A}$. Given the functional forms selected the L associated to our model is

$$L = \frac{1}{\alpha/2+Ac^{-\delta}},$$

with c defined in (28) (for a derivation, see the Appendix). We calibrate A in order to minimize the distance between L in the data and in the model, given the calibrated parameters $(\alpha, \beta, \delta) = (0.427, 0.96, 3.78)$. This gives us $A = 3.052$.

In Figure 1 we show the quality of the fit of the model to the data, generated by the calibrated parameters. For each year in the period 1929-2006, we plot the observed real balances-income ratio M/PY against the nominal interest rate i. The continuous line represents $L = \frac{1}{\alpha/2+Ac^{-\delta}}$ given the calibrated parameters' values.

The welfare cost of inflation for a representative agent. Now that we have the parameter vector Θ we can quantify the welfare cost of inflation with a procedure analogous to the one in [20]. The welfare cost of inflation is defined as the percentage adjustment in consumption the representative agent would require to be indifferent between a steady state with inflation rate (money growth rate) π and a lower inflation rate $x \in [\beta, \pi)$.

[4] For 1929-75, the nominal yield on short-term commercial paper is from [18, Table 4.8, Column 6]. For 1976-1996, it is from [13, Table B-69]. For 1997-2006, it is the Financial Commercial Paper with maturity 3-month, from [19]. The money supply M1 is in billions of dollars, December of each year, not seasonally adjusted. For 1929-58, it is from [17, pp. 708-718, Column 7]. For 1959-2006, it is from the Federal Reserve Bank of St. Louis *FRED Database*. For the period 1929-2006, nominal GDP is from [28].

Let w_π denote expected (lifetime) utility for the representative agent in stationary equilibrium given inflation rate π. Using (22) we have:

$$(1-\beta)w_\pi = \frac{\alpha}{2}[u(c_\pi) - \phi(c_\pi)] + U(q^*) - q^*. \tag{29}$$

Here, c_π denotes equilibrium c given π. If we reduce π to x and adjust consumption in both markets by the proportion $\bar{\Delta}_x$, then we define adjusted expected utility by

$$(1-\beta)w_x = \frac{\alpha}{2}[u(\bar{\Delta}_x c_x) - \phi(c_x)] + U(\bar{\Delta}_x q^*) - q^*. \tag{30}$$

The welfare cost of having π instead of x inflation is the value $\Delta_x = 1 - \bar{\Delta}_x$ that satisfies $w_\pi = w_x$. If $\Delta_x > 0$, then agents are indifferent between π inflation, or alternatively, x inflation *and* consumption reduced by Δ_x percent.

Using $x = 1$ to denote zero inflation and $x = \beta$ to denote the Friedman rule, we have that the welfare cost of inflation for the representative agent is rather small. For example, ten percent inflation is worth less than 1% of consumption: $\Delta_1 = 0.18\%$ and $\Delta_\beta = 0.21\%$. These results are in line with previous findings based on various models; e.g., see [4, 12, 16, 20, 21, 25].[5] Now, we extend the quantitative analysis to the heterogenous-agents case.

4.1 Heterogeneous trading shocks

Given inflation π, (22) and Lemma 1 imply equilibrium expected utility for type j

$$\begin{aligned}(1-\beta)w_{j\pi} = \quad & \tfrac{\alpha_j}{2}[u(c_{j\pi}) - \phi(y_\pi)] + U(q^*) - q^* \\ & + \tfrac{\alpha_j}{2}\phi'(y_\pi)(y_\pi - c_{j\pi}) + (\pi - 1)(\bar{m}_\pi - m_{j\pi}).\end{aligned} \tag{31}$$

Here, $c_{j\pi}$, y_π and $m_{j\pi}$ are equilibrium quantities; average real balances equal the real money supply so from (4), (15) and (17) we have

$$\bar{m}_\pi = \rho m_{H\pi} + (1-\rho)m_{L\pi} = \phi'(y_\pi)[\rho c_{H\pi} + (1-\rho)c_{L\pi}].$$

The social cost of inflation is now unequally distributed across the population. Indeed, as compared to the representative agent case in (29), the expression in (31) contains

[5] To compare our results to those in [20], we have also considered linear disutility from labor effort, setting $\delta = 1$ (recalibrating A, which becomes 2.88) and we obtained $\Delta_1 = 0.83\%$ and $\Delta_\beta = 0.95\%$.

additional terms that capture the distributional impact of inflation. To see why, note that inflation affects an agent's welfare in three ways. As usual, it distorts consumption in market one and so it affects expected trade surplus on that market; this impact is captured by $\frac{\alpha_j}{2}[u(c_{j\pi}) - \phi(y_\pi)]$, which is type-dependent. Now, however, the distribution of average net earnings in market one and the inflation tax burden also are affected by inflation.

Inflation affects *average earnings* because expected income and expenditure are unequal in market one. Every type produces y_π, which corresponds to average output and consumption (by market clearing). Type L agents, however, consume less than average, so on average earn net income from market one trades. The opposite holds for type H. Given that the value of goods is $p = \phi'(y)$, average net earnings are $\frac{\alpha_j}{2}\phi'(y_\pi)(y_\pi - c_{j\pi})$, which varies with π (y_π and $c_{j\pi}$ fall with π). In addition, inflation *redistributes* monetary wealth due to equilibrium real-balance heterogeneity. Type L save less than the average $\bar{m}_\pi$ but receive the same lump-sum transfer $(\pi - 1)\bar{m}_\pi$ as anyone else. Their reduced exposure to the inflation tax results in a net transfer. The converse holds for type H. Consequently, inflation can only redistribute monetary wealth from the "rich" to the "poor".

If we eliminate inflation setting $\pi = 1$ and simultaneously adjust consumption by the proportion $\bar{\Delta}_{j1}$, then adjusted expected utility of an agent of type j is

$$(1-\beta)w_{j1} = \tfrac{\alpha_j}{2}[u(\bar{\Delta}_{j1}c_{j1}) - \phi(y_1)] + \tfrac{\alpha_j}{2}\phi'(y_1)(y_1 - c_{j1}) + U(\bar{\Delta}_{j1}q^*) - q^*. \tag{32}$$

The welfare cost of π inflation for an agent of type j is the value $\Delta_{j1} = 1 - \bar{\Delta}_{j1}$ that satisfies $w_{j\pi} = w_{j1}$. Notice that without inflation, there is no wealth redistribution. At the Friedman rule, instead, we have

$$(1-\beta)w_{j\beta} = \tfrac{\alpha_j}{2}[u(\bar{\Delta}_{j\beta}c_\beta) - \phi(y_\beta)] + U(\bar{\Delta}_{j\beta}q^*) - q^*. \tag{33}$$

This is identical to the representative agent case because if $\pi \to \beta$, then $m_j \to m^* = \bar{m}$ and $c_{j\beta} \to c_\beta = y_\beta$ for $j = H, L$ (Lemma 1). Hence, not only there is wealth redistribution but, since $c_{j\beta} = c_\beta = y_\beta$ for all j, there is no disparity in average net income.

Calibration and results To measure the welfare cost of inflation we proceed as follows. First, we fix the common preference parameters (β, a, A, δ) to the values calibrated in the representative agent model. Second, we fix the average trading friction to the value $\alpha = 0.427$ (also form the representative agent model) and then consider mean preserving spreads $\rho\alpha_H + (1-\rho)\alpha_L = 0.427$ for some given value ρ. Several possibilities exist.

The route we take is to associate the types $j = L, H$ to different segments of the U.S. population, as follows. Recall that $pc_j + A$ is consumption expenditure (in real terms) for an agent of type $j = H, L$. Thus, consider the ratio of consumption expenditures for agent of type j to average consumption expenditure, i.e.,

$$\frac{pc_j + A}{p[c_L + \rho(c_H - c_L)] + A}$$

We let the ratio for type L be associated to average consumption for the bottom three quintiles of consumption expenditure in the U.S., and the ratio for type H to the top two quintiles. This implies $\rho = 0.4$. Then, we use data on consumption expenditure by income quintiles from the *Consumer Expenditure Survey* for the period 1989-2006. We find that for type L the above ratio is 0.644 and for type H it is 1.533. We calibrate α_H and α_L to minimize the sum of the squared residuals between the theoretical and empirical ratios, given that $\rho\alpha_H + (1-\rho)\alpha_L = 0.427$. In this manner, we obtain $\alpha_H = 1$ and $\alpha_L = 0.045$ (approximately).

The average welfare cost of 10 percent inflation inflation is approximately 0.19% and it is distributed as follows, $\Delta_{L1} = -0.47\%$ and $\Delta_{H1} = 1.18\%$ (in the case of the Friedman rule the average welfare cost is 0.36%, while $\Delta_{L\beta} = -0.16\%$ and $\Delta_{H\beta} = 1.14\%$). In sum, inflation does generate a welfare cost on average, but it is low and it has a very different impact on different segments of the economy. Inflation has a beneficial welfare effect on low-consumption (and low wealth) agents, while it hurts high-consumption (low wealth) agents.

Is this result a construct of our way to calibrate α_j? The answer is negative. To demonstrate it, we ran the analysis fixing ρ to the arbitrary value 0.5 and quantify the welfare cost of inflation as α_L varies on $(0, 0.427)$ to satisfy $\rho\alpha_H + (1-\rho)\alpha_L = 0.427$. Figure 2 reports the results, plotting the welfare costs Δ_{j1} (ten percent inflation as opposed to no inflation) against α_L. The main findings confirm the results obtained earlier:

- First, inflation lowers welfare of the average agent. Intuitively, inflation distorts consumption of everyone[6] so it lowers average welfare. The average welfare cost, however, is relatively small, being less than 1% and similar to the case of a representative agent.

- Second, the burden of inflation is unequally distributed across types. Some inflation can be welfare-increasing for agents who hold less-than-average balances (type L) and welfare-decreasing those with more-than-average balances (type H). Intuitively, inflation has a redistributive effect because balance holdings are heterogeneous, so its burden is generally unequal across agent types. Nominal wealth is redistributed top to bottom, i.e., from type H agents (who pay a large inflation tax) to type L (low inflation tax). If the redistributive effect dominates the consumption distortion, then welfare for that agent type rises.

- Third, the welfare cost for an agent type increases with the share of monetary wealth held. Indeed, in the figure the welfare cost for agent j rises with α_j (m_j rises with α_j). This explains why, generally, the welfare cost is positive for type H agents, negative for type L, and such a disparity shrinks as α_L approaches the average value of α_j. This is clear in Figure 2. Intuitively, redistributive effects are strongest under great wealth disparities. Given a fixed average value of α_j, the lower is α_L the greater is the disparity in monetary wealth.[7]

[6]The Euler equation (24) can be rewritten as

$$F(c_j, \pi) = 1 + \frac{\alpha_j}{2}\left[\frac{u'(c_j)}{\phi'(y)} - 1\right] - \frac{\pi}{\beta} = 0 \text{ for } j = H, L.$$

Since goods market clearing on odd dates implies

$$y = \frac{\alpha_H \rho c_H + \alpha_L(1-\rho)c_L}{\alpha_H \rho + \alpha_L(1-\rho)},$$

using the implicit function theorem for type L agents we have

$$\frac{\partial c_L}{\partial \pi} = -\frac{\partial F/\partial \pi}{\partial F/\partial(c_L)} = -\frac{-1/\beta}{\frac{\alpha_j}{2}\frac{u''(c_L)\phi'(y) - u'(c_L)\phi''(y)\frac{\alpha_L(1-\rho)}{\alpha_H\rho + \alpha_L(1-\rho)}}{(\phi'(y))^2}} < 0$$

since $u' > 0$, $u'' < 0$, $\phi' > 0$ and $\phi'' > 0$.

[7]Such a beneficial effect of inflation for those with low money holdings is similar to the finding in [23]. That result emerges in a model where agents can only self-insure at random, and so can be short on cash simply because they did not have an opportunity to replenish their holdings. This is unlike our model,

- Fourth, welfare costs are not necessarily higher for types H. The reason lies in the average net earnings component $\frac{\alpha_j}{2}\phi'(y_\pi)(y_\pi - c_{j\pi})$, which is positive for L but negative for H. The size of this term changes nonlinearly with π and it may dominate the wealth redistribution effect. If agents are relatively similar (α_L is close to α_H) then type L may suffer more than type H. To give an example we calculated the welfare cost of 10 percent inflation as opposed to the Friedman rule. For low values of α_L the welfare cost is positive for H and negative for L. However, as α_L gets close to the average value 0.427 not only type L also suffer a cost, but it can dominate that for type H. For instance, for $\alpha_L = 0.4$ we have $\Delta_{L\beta} = 0.22\%$ and $\Delta_{H\beta} = 0.19\%$.

Discussion. The analysis above has generated two main results. The average welfare cost of anticipated inflation is small; 10 percent inflation is worth less than one percent consumption on average (for the different combinations of α_j, the highest value is 0.31%), which is not much different from the results of a representative agent model. Second, the average welfare cost of anticipated inflation is unevenly distributed across the population; the rich suffer more than the poor and the poor can even benefit from inflation. The first result in line with previous studies and is quite robust to variations in most calibrated parameters; a notable exception is of course A because this parameter determines the relative size of market two (where money is unimportant). The second result is also in line with other quantitative studies of the redistributive consequence of inflation in a similar framework (see [11, 23]). In particular, a beneficial effect of inflation for those with low money holdings is found in [23], a model where agents can only self-insure at random, and so can be short on cash simply because they did not have an opportunity to replenish their holdings. This is unlike our model, where self-insurance is not a problem (self-insurance opportunities arise periodically, on even dates), but agents display differences in the desire to self-insure.

Interesting, however, the redistribution result is at odds with at least some empirical

where self-insurance is not a problem (self-insurance opportunities arise periodically, on even dates), but agents display differences in the desire to self-insure.

evidence; for instance, richer agents are less concerned about inflation than the poor (see the discussion in [2]). In addition, it is at odds with the findings in [14] that suggest inflation redistributes wealth from the bottom to the top of the wealth distribution. That is, inflation acts as a regressive tax in [14] and not as a progressive tax in our model. A possible reason for such differences is the structure of financial portfolios in the two models. In [14] agents can self-insure with cash as well as assets that are less liquid but have a greater return. When the cost to liquidate these assets displays increasing returns to scale, wealthier agents (more productive types, who consume more than average) hold a smaller-than-average fraction of their wealth in cash. The opposite occurs for low-productive types, who consume less than average hence hold a greater fraction of wealth in cash. In our model, instead, agents can only hold cash. We will relax this assumption and explore its implications in a later section (Section 5).

4.2 Heterogeneity in productivity

Now consider economies where agents have identical self-insurance needs but differ in labor productivity as in Section 3.4. We calibrate the common preference parameters as for the representative agent, $(\alpha, \beta, A, \delta) = (0.427, 0.96, 3.05, 3.78)$. With regards to the heterogeneous productivity parameters we proceed as follows.

It is assumed that different types of agents can exploit different linear production technologies, one of which is more efficient than the other. Type L agents must supply $\theta - 1$ more hours than agents of type H to produce the same amount of output y. In this case

$$\phi_j(y) = \frac{(\theta_j y)^\delta}{\delta}$$

with $\theta_L = \theta > \theta_H = 1$. We interpret $\theta_j y_j$ as hours worked by type j to produce y_j output. It should be clear that with this formulation the elasticity of labor is δ, independent of j, and $\phi_L(y) > \phi_H(y)$ for all $y > 0$. Also, since choosing output or hours worked is equivalent, our analysis is in terms of y_j instead of hours. In the Appendix we show that

c, y_L and y_H can be defined as explicit functions of the parameters:

$$\begin{aligned} y_L = & \quad \left[\left(1+\tfrac{2i}{\alpha}\right)\left(\rho\theta^{\frac{\delta}{\delta-1}}+1-\rho\right)^a\theta^\delta\right]^{\frac{1}{1-a-\delta}} \\ y_H = & \quad y_L\theta^{\frac{\delta}{\delta-1}} \\ c = & \quad y_L\left(\rho\theta^{\frac{\delta}{\delta-1}}+1-\rho\right). \end{aligned}$$

Productivity is measured by average output per hour in nonfarm private industries using date from the Bureau of Labor Statistics for 1987-2006. To calibrate the relative productivity parameter θ we match the ratio of productivity in the service sector (very productive) to the goods sector (less productive); we obtain $\theta = 4.24$. Then, we set $\rho = 77\%$ to match the proportion of employment in the service sector.

Using (22), the expected lifetime utility for type j under inflation π is

$$(1-\beta)w_{j\pi} = \quad \tfrac{\alpha}{2}[u(c_\pi) - \phi_j(y_{j\pi})] + \tfrac{\alpha}{2}\phi_j'(y_{j\pi})(y_{j\pi} - c_\pi) + U(q^*) - q^*. \tag{34}$$

The expression differs from (31) because now agents with different productivity hold identical balances (Lemma 2). So, inflation cannot redistribute wealth. However, there are still disparities in average net eaernings because $c_{j\pi} = c_\pi$ for $j = H, L$ but $y_{H\pi} > y_{L\pi}$. Hence, the social burden of inflation will be unequally distributed, in general.

If we reduce π to x and simultaneously adjust consumption (in both markets) by the portion $\bar{\Delta}_{jx}$, then the adjusted expected utility of an agent of type j is

$$(1-\beta)w_{jx} = \quad \tfrac{\alpha}{2}[u(\bar{\Delta}_{jx}c_x) - \phi_j(y_{jx})] + \tfrac{\alpha}{2}\phi_j'(y_{jx})(y_{jx} - c_1) + U(\bar{\Delta}_{jx}q^*) - q^*. \tag{35}$$

The welfare cost of inflation for type j is thus $\Delta_{jx} = 1 - \bar{\Delta}_{jx}$.

The quantitative analysis generates the following results. The average welfare cost is rather modest, 0.17 percent, but it is unequally distributed. Once again, the welfare cost is higher for agents who are more productive, type H, and lower (or negative) for types L. For example, we find that the cost of 10 percent inflation is $\Delta_{H1} = 0.47\%$ and $\Delta_{L1} = -0.79\%$. This might seem surprising because there is no equilibrium dispersion in monetary wealth. However, inflation affects agents identically, when it comes to consumption distortions, but unequally when it comes to earnings. Intuitively, only the most productive agents earn net income from market one trades, and this income falls with inflation. So, the social burden of inflation lies mostly (or entirely) on their shoulders.

5 Money is not the only asset

Money in this class of models is usually assumed to be the only asset. However, the quantitative impact of inflation can obviously depend on whether alternative assets can provide (some) consumption insurance. Thus, in this section we quantify the social cost of inflation when heterogenous agents can hold more sophisticated financial portfolios.[8]

We augment the financial sophistication of the economy introducing a competitive financial sector that offers risk-pooling services. On even dates agents have the opportunity to buy consumption insurance as well as money because on market two an intermediary sells one-period nominal assets to the public at price θ. Assets can (only) be redeemed in market one for claims to money, which are enforceable in market two and financed with the revenue from asset sales. The intermediary earns zero profits.

Market one buyers can redeem the asset spending its claims to buy consumption, while sellers can redeem the asset to cash its claims in the next market. Idle agents cannot participate in market one trades, so can trade neither on goods nor financial markets, i.e., cannot redeem the asset. This feature is a form of limited participation in financial markets (as well as goods), which affects agent types differently. Consequently, the consumption insurance offered by the asset is less attractive to types L because they are less frequent traders on market one.

For an agent of type j who holds $b_j \geq 0$ assets and $m_j \geq 0$ money the expressions in (1) become:

$$\begin{aligned} m_{j,b} &= m_j + b_j - pc_j \\ m_{j,s} &= m_j + b_j + py_j \\ m_{j,n} &= m_j \end{aligned} \tag{36}$$

The agent's budget constraint at the start of an even date is:

$$x_{j,k} = q_j + \pi\theta b_j' + \pi m_j' - (m_{j,k} + \tau), \tag{37}$$

hence

$$V_j(m_j, b_j) = m_j + \alpha_j b_j + \tfrac{\alpha_j}{2}[u(c_j) - \phi_j(y_j)] + \tfrac{\alpha_j}{2}p(y_j - c_j) + W_j(0,0). \tag{38}$$

[8]Some related models in which money competes with another asset are [10, 26].

Recall that an agent's need for liquidity insurance depends on α_j. The constraint for a buyer with b_j assets is $pc_j \leq m_j + b_j$, so $c_j = \min\{\frac{m_j+b_j}{p}, c(p)\}$. For a constrained buyer we still need (23) for $m_j \geq 0$. Using (38) we have

$$\frac{\partial V_j(m_j,b_j)}{\partial b_j} = \alpha_j + \frac{\alpha_j}{2}[u'(c_j) - p]\frac{\partial c_j}{\partial b_j}$$

where $\frac{\partial c_j}{\partial b_j} = \frac{1}{p}$. Hence, we have $b_j \geq 0$ if

$$\theta\frac{\pi}{\beta} \geq \alpha_j + \frac{\alpha_j}{2}\left[\frac{u'(c_j)}{p} - 1\right] \text{ for } j = H, L. \tag{39}$$

As in our previous analysis, consider outcomes in which all market-one buyers are constrained. We wish to study an equilibrium in which those who consume less than average hold more money but less other assets than average, as in the U.S. data. We focus on the simplest possible scenario: $b_H > b_L = 0$ and $m_L > m_H = 0$. It is an equilibrium if (23) is an equality and (39) a strict inequality for $j = L$ (the converse must hold for $j = H$).

We claim that this is an equilibrium for some sufficiently small inflation rate bounded away from β. To demonstrate it observe that if only types H buy assets, then the repayment constraint faced by the intermediary is

$$\pi\theta b = \alpha_H b, \tag{40}$$

which pins down the price θ consistent with zero profits. Under the conjecture that only types H buy πb assets at price θ, the asset's (gross) average return $\frac{\alpha_H}{\theta}$ is simply the inflation rate π. Indeed, α_H is the portion of asset holders who redeem it.

In the conjectured equilibrium, if (39) is an equality then $b_H > 0$, which using (40) implies

$$\alpha_H(\tfrac{1}{\beta} - 1) = \frac{\alpha_H}{2}\left[\frac{u'(c_H)}{p} - 1\right]. \tag{41}$$

Notice that $\alpha_H(\frac{1}{\beta} - 1) < \frac{\pi}{\beta} - 1$ for all $\pi > \bar{\pi} = \beta + \alpha_H(1-\beta)$, with $\bar{\pi} \in (\beta, 1)$ since $\alpha_H < 1$. If (41) holds, then $\frac{\pi}{\beta} - 1 > \frac{\alpha_H}{2}\left[\frac{u'(c_H)}{p} - 1\right]$ for all $\pi > \bar{\pi}$ (so $m_H = 0$). As $\pi \to \beta$ types H hold only money and $u'(c_H) = p$ (efficiency). Intuitively, if $\pi \leq \bar{\pi}$, then inflation is small and assets offer consumption insurance that is 'too expensive' relative to the consumption insurance offered by money. Otherwise, type H prefers holding assets but not money, since they can consume more.

Now consider a type L. If $b_L = 0$, then (39) must hold as

$$\theta \tfrac{\pi}{\beta} > \alpha_L + \tfrac{\alpha_L}{2}\left[\tfrac{u'(c_L)}{p} - 1\right],$$

which requires $\pi < \tilde{\pi} = \beta + \alpha_H - \beta\alpha_L$, with $\tilde{\pi} > \bar{\pi}$ and $\tilde{\pi} > 1$ if $\beta > \frac{1-\alpha_H}{1-\alpha_L}$. To see this, note that $\frac{\alpha_L}{2}\left[\frac{u'(c_L)}{p} - 1\right] \leq \frac{\pi}{\beta} - 1$ from (23). So $b_L = 0$ whenever $\theta\frac{\pi}{\beta} > \alpha_L + \frac{\pi}{\beta} - 1$. Using (40) we get $\pi < \tilde{\pi}$. It follows that $b_L = 0$ and $m_L > 0$ with $\frac{\alpha_L}{2}\left[\frac{u'(c_L)}{p} - 1\right] = \frac{\pi}{\beta} - 1$. Intuitively, if $\pi < \tilde{\pi}$, then assets offer consumption insurance that is 'too expensive' for type L agents. These agents do not trade as frequently as type H, so the insurance offered by the asset is not as valuable. Type L agents buy financial asset only if inflation is sufficiently high, i.e., if money is a sufficiently poor store of value.

To sum up, if $\pi \in (\bar{\pi}, \tilde{\pi})$, then H agents only hold financial assets, while L agents only hold money. Hence, $c_H = b/p$ and $m_H = 0$, while $c_L = m_L/p$ and $b_L = 0$ The expression for c_L is obtained from (23) as before. The expression for c_H is obtained from (41).

Now we have

$$W_j(m_{j,k}) = U(q_j) - q_j - \pi\theta b'_j - \pi m'_j + m_{j,k} + \tau + \beta V_j(b'_j, m'_j) \tag{42}$$

that differs from (6) due to asset holdings. Using the expressions (42) and (38), ex-ante welfare for an agent of type j in stationary equilibrium is

$$\begin{aligned}(1-\beta)V_j(b_j, m_j) = \quad & \tfrac{\alpha_j}{2}[u(c_j) - \phi(y_j)] + \tfrac{\alpha_j}{2}p(y_j - c_j) + U(q^*) - q^* \\ & +(\pi - 1)(\bar{m} - m_j) + b_j(\alpha_j - \pi\theta).\end{aligned} \tag{43}$$

Here, c_j, y_j, and m_j are optimal choices, $b_j = pc_j - m_j$ are optimal asset holdings, $p = \phi'(y_j)$ and $\bar{m} = \rho m_H + (1-\rho)m_L$ from (4). The expression in (43) is standard, except for the new term $b_j(\alpha_j - \pi\theta)$ capturing the redistributive effect of inflation on asset holdings.

Given $m_H = b_L = 0$ we have $(1-\rho)m_L = \bar{m}$ and $b = pc_H$. For types L we have

$$(1-\beta)V_L(0, \bar{m}) = \tfrac{\alpha_L}{2}[u(c_L) - \phi(y_L)] + \tfrac{\alpha_L}{2}p(y_L - c_L) + U(q^*) - q^* - (\pi - 1)\tfrac{\rho}{1-\rho}\bar{m}.$$

Since $\bar{m} = (1-\rho)m_L = (1-\rho)pc_L$ then $(\pi - 1)\bar{m}\frac{\rho}{1-\rho} = (\pi - 1)\rho pc_L$.

For a type H instead we have

$$(1-\beta)V_H(b,0) = \frac{\alpha_H}{2}[u(c_H) - \phi(y_H)] + \frac{\alpha_H}{2}p(y_H - c_H) + U(q^*) - q^* + (\pi - 1)\bar{m}$$

because $\pi\theta = \alpha_H$, $m_H = 0$ and $b_H = b = pc_H$.

The inflation tax for type L is $-(\pi - 1)\frac{\rho}{1-\rho}\bar{m}$, and for a type H is $(\pi - 1)\bar{m}$. Since there is a proportion $1 - \rho$ of type L and ρ of types H, these two terms sum up to zero. That is, inflation generates a net wealth transfer from types L to types H. Assets holdings of types H are not subject to the inflation tax because the expected return on assets is $\alpha_H \frac{1}{\theta} = \pi$. Indeed, the price of nominal assets falls with inflation.

The quantitative analysis generates the following results. The average welfare cost is once again rather modest, 0.32 percent, but it is unequally distributed. The direction of the redistribution, however, is opposite to the one we find when money is the only asset. For example, we find that the cost of 10 percent inflation is $\Delta_{H1} = -1.18\%$ and $\Delta_{L1} = 1.32\%$. Therefore, in this case inflation benefits the rich - who hold financial assets - and hurts the poor.

We think this result should hold as long as the model has the feature that some agents choose to insure against consumption risk mostly (though not necessarily only) by using money and not alternative assets. In our simple model this occurs because frictions in asset market trading affect agents differently. Other formulations may include costs from accessing liquidity services that display increasing returns (such as the model in [14]) or agent specific preferences.

6 Final remarks

We quantify the welfare cost of inflation in a calibrated heterogeneous-agent model of the U.S. economy. We provide three main contributions. First, the social cost of inflation is small social cost, in that on average agents would give up less than one percent consumption to avoid ten percent inflation. Second, the distribution across the population of the social cost of inflation depends on the type of heterogeneity considered. If agents differ in their labor productivity, then inflation does not redistribute monetary wealth, though it hurts the more productive and benefits the less productive. This occurs because inflation affects agents identically when it comes to consumption distortions, but unequally when it comes to earnings. Instead, if agents differ in trading shocks, then there is equilibrium dispersion in monetary wealth and inflation has a redistributive effect. Third, the direction of wealth redistribution depends on whether money is the only asset in the model. If it is, then inflation benefits the poor—who hold less-than-average balances—and hurts the rich. The converse is true if agents can insure against consumption risk with a competing asset.

References

[1] Akyol, A., (2004). "Optimal monetary policy in an economy with incomplete markets and idiosyncratic risk," *Journal of Monetary Economics*, 51, 1245-1269.

[2] Albanesi, S. (2007). Inflation and Inequality. *Journal of Monetary Economics,* 54 (4), 1088-1114.

[3] Aliprantis, C.D., Camera, G., and D. Puzzello (2007). Anonymous markets and monetary trading. *Journal of Monetary Economics*, 54 (7), 1905-1928.

[4] Aruoba, S., Waller, C.J., and R. Wright (2007). Money and Capital. *Manuscript.*

[5] Berentsen, A., G. Camera, and C. Waller (2005). The distribution of money balances and the non-neutrality of money. International Economic Review 46(2), 465-487

[6] Bhattacharya J., J. H. Haslag, and A. Martin (2005). Heterogeneity, redistribution, and the Friedman rule. *International Economic Review* 46 (2), 437-454.

[7] Blau, F., and L. Kahn (2005). Changes in the labor supply behavior of married women: 1980-2000. *NBER Working Paper # 11230.*

[8] Blundell, R., and T. MaCurdy (1999). Labor supply: a review of alternative approaches. In *Handbook of Labor Economics*, vol. 3A, Ch. 27, Ashenfelter D. Card editors, Amsterdam, North Holland.

[9] Boel, P., and G. Camera (2006). Efficient monetary allocations and the illiquidity of bonds. *Journal of Monetary Economics*, 53 (7), 1693-1715.

[10] G. Camera and Ruscitti F. (2007). Inflation and Intermediation Frictions. Working paper, University of Iowa.

[11] Chiu, J., and Molico, M. (2007). Uncertainty, Inflation, and Welfare. Unpublished paper. Bank of Canada.

[12] Cooley, T.F. and G.D. Hansen (1989) The Inflation Tax in a Real Business Cycle Model. *American Economic Review* 79, 733-748.

[13] *Economic Report of the President* (1996). Washington D.C.: Govt. Printing Office.

[14] Erosa, A., and G. Ventura (2002). On inflation as a regressive consumption tax. *Journal of Monetary Economics*, 49 (4), 761-795.

[15] Filer, R. K., D. S. Hamermesh, and A. E. Rees (1996). *The Economics of Work and Pay*, Sixth Edition. New York: Harper-Collins.

[16] Fischer, Stanley, "Towards an Understanding of The Costs of Inflation," Carnegie-Rochester Conference on Public Policy, K. Brunner and A. Meltzer, eds. Autumn 1981, 15, 5-42.

[17] Friedman, M. (1963). *A Monetary History of the United States 1867-1960.* Princeton: Princeton University Press.

[18] Friedman, M., and A. J. Schwartz (1982). *Monetary Trends in the United States and the United Kingdom, 1867-1975.* Chicago: University of Chicago Press, for the National Bureau of Economic Research.

[19] H.15 Selected Interest Rates, *Federal Reserve Statistical Release.*

[20] Lagos, R., and R. Wright (2005). A unified framework for monetary theory and policy analysis. *Journal of Political Economy*, 113 (3), 463-484.

[21] Lucas, R. (1981). Discussion of Stanley Fischers towards an understanding of the costs of inflation: II. *Carnegie-Rochester Conference Series on Public Policy*, 15, 43-52.

[22] Lucas, R. (2000). Inflation and welfare, *Econometrica*, 68 (2), 247-274.

[23] Molico, M. (2006). The distribution of money and prices in search equilibrium. *International Economic Review*, 47 (3), 701-722.

[24] Raj, C. (2006). A New Method of Estimating Risk Aversion. *American Economic Review*, 96 (5), 1821-1834.

[25] Reed R., and C. J. Waller (2006). Money and Risk Sharing. *Journal of Money, Credit and Banking* 38(6), 1599-1618

[26] I. A. Telyukova, and Wright R. (2007). A Model of Money and Credit, with Application to the Credit Card Debt Puzzle. Working paper, University of Pennsylvania

[27] U.S. Bureau of the Census (1960). *Historical Statistics of the United States, Colonial Times to 1957*. Washington D.C.: Govt. Printing Office.

[28] U.S. Office of Business Economics. *The National Income and Product Accounts of the United States*. Washington D.C.: Govt. Printing Office.

Appendix

A.1 The constrained-efficient allocation

Consider the allocation selected by a planner who maximizes the agents' lifetime utilities and treats agents identically. The planner is subject to the same physical and informational constraints faced by the agents and therefore cannot observe identities. However, the planner observes types. Basically, the planner can propose a type-dependent consumption plan in each trading cycle, but does not have the ability to transfer resources across agents over time. Equivalently, the planner maximizes expected utility of the arbitrary agent on each date. The planning problem thus corresponds to a sequence of static maximization problems, i.e., to maximize ex-ante welfare of the representative agent, subject to technological feasibility.

Recall that on each date agents have identical preferences ex-ante and there is an identical proportion of buyers and sellers. Moreover, on each odd date agents that are active can produce or consume with equal probability.

Letting $\rho_j = \rho$ for $j = H$ and $1 - \rho$ for $j = L$, the planner problem is to choose $\{c_j, y_j\}_{j=H,L}$, q, and x to solve:

$$\begin{aligned} \max \quad & \textstyle\sum_{j=H,L} \frac{\alpha_j}{2} \rho_j [u(c_j) - \phi_j(y_j)] + U(q) - x \\ \text{s.t.} \quad & \textstyle\sum_{j=H,L} \rho_j c_j \leq \sum_{j=H,L} \rho_j y_j \ \text{ and } q \leq x \end{aligned}$$

It is clear that, but non-satiation, the feasibility constraints should hold with equality. Letting λ denote the Lagrange multiplier on the first feasibility constraint, the FOCs are thus

$$\begin{aligned} & \tfrac{\alpha_j}{2} \rho_j [u'(c_j) - \lambda] = 0 \\ & \tfrac{\alpha_j}{2} \rho_j [-\phi_j'(y_j) + \lambda] = 0 \\ & U'(q) - 1 = 0 \end{aligned}$$

That is agents are assigned consumption and produce up to the point where the marginal utility of their consumption or labor equal the marginal utility of income, λ.

Hence, the efficient allocation is stationary across trading cycles, and it can be characterized as follows. On odd dates $c_j = c^* = \rho y_H + (1 - \rho) y_L$ and $y_L = y_L^* < y_H = y_H^*$ where the starred output values are the unique positive solutions to the two equalities

$u'(y_L+y_H) = \phi_j'(y_j)$ for $j = H, L$. It should be clear that $c^* = y^*$ such that $u'(c^*) = \phi'(c^*)$ if there is no heterogeneity in productivity. On even dates $q_j = x_j = q^*$ for each type j in each trading cycle, where q^* is the unique positive solution to $U'(q) = 1$.

A.2 Elasticities and the money demand ratio L

Consider a representative agent economy and focus on odd dates.

Elasticity of disutility of labor: The disutility of labor is $\phi(y) = \frac{y^\delta}{\delta}$, where in our model y is production as well as labor effort of the agent. So, the elasticity of disutility of labor is

$$\varepsilon_y = \frac{d\phi(y)/\phi(y)}{dy/y} = \frac{d\ln\phi(y)}{d\ln y} = \frac{y^{\delta-1}y}{y^\delta}\delta = \delta,$$

since the differential

$$d\ln\phi(y) = d\ln(y^\delta/\delta) = d(\delta\ln y - \ln\delta) = \frac{\delta}{y}dy.$$

Since $\phi'(y) = p$, the labor supply $y(p)$ satisfies

$$y^{\delta-1} = p \Rightarrow y(p) = p^{\frac{1}{\delta-1}}.$$

Elasticity of labor supply: In our model the wage of a worker on odd dates is p. The elasticity of the labor supply with respect to the relative wage is

$$\varepsilon_p = \frac{dy(p)/y(p)}{dp/p} = \frac{d\ln y(p)}{d\ln p} = \frac{1}{\delta-1},$$

because the differential

$$d\ln y(p) = d(\ln p^{\frac{1}{\delta-1}}) = d\left(\frac{1}{\delta-1}\ln p\right) = \frac{1}{\delta-1}\times\frac{dp}{p}.$$

Elasticity of money demand: From (17) we have $pc = m$, so the Euler equation (27) for the representative agent is

$$F(m/p, i) = \tfrac{\alpha}{2}\left[\tfrac{u'(m/p)}{\phi'(y)} - 1\right] - i = 0.$$

Using the implicit function theorem we have

$$\frac{\partial m/p}{\partial i} = -\frac{\partial F/\partial i}{\partial F/\partial(m/p)} = \quad -\frac{-1}{\frac{\alpha}{2\phi'(y)}u''(m/p)} = \frac{2\phi'(y)}{\alpha u''(m/p)}$$

Given $c = m/p$ and market clearing $c = y$, the elasticity of money demand is

$$\varepsilon_m = \frac{\partial m/p}{\partial i} \times \frac{i}{m/p} = \quad \frac{2\phi'(y)}{\alpha u''(c)} \times \frac{i}{c} = \frac{2i\phi'(y)}{\alpha cu''(c)} \tag{44}$$

We have $\phi'(y) = y^{\delta-1}$ and $y = c$. So (44) is $\frac{2ic^{\delta-1}}{\alpha cu''(c)}$. Substituting c from (28) we get

$$\varepsilon_m = -\frac{2i}{a(2i+\alpha)}$$

The money demand ratio L**:** We have $L = \frac{m}{\frac{\alpha}{2}pc+A}$ and from (17) we have $pc = m$. Also, $p = \phi'(y)$. Since $\phi'(y) = y^{\delta-1}$ and $y = c$ from market clearing, we can write $L = \frac{1}{\alpha/2+Ac^{-\delta}}$, with c defined in (28) as a function of parameters and nominal interest rate.

A.3 Explicit solutions for consumption and output

Heterogeneity in trading risk: In this environment $y_H = y_L = y$. Given the assumed functional forms we have $\phi'(y) = y^{\delta-1}$ and $u'(c_j) = c_j^{-a}$ so we can rewrite the Euler equation (24) as

$$1 + \frac{2i}{\alpha_j} = \frac{c_j^{-a}}{y^{\delta-1}} \quad \text{for } j = H, L.$$

which implies that

$$c_L = \left(\frac{(2i+\alpha_L)\,\alpha_H}{\alpha_L(2i+\alpha_H)}\right)^{-\frac{1}{a}} c_H.$$

From market clearing (19) we have

$$y = \frac{\rho\alpha_H c_H + (1-\rho)\alpha_L c_L}{\rho\alpha_H + (1-\rho)\alpha_L}.$$

Substituting for y and c_L in the Euler equation above we obtain

$$c_H = \left(\frac{\alpha_H}{2i+\alpha_H}\left(\frac{\alpha_H\rho + \alpha_L(1-\rho)\left(\frac{(2i+\alpha_L)\alpha_H}{\alpha_L(2i+\alpha_H)}\right)^{-\frac{1}{a}}}{\alpha_H\rho + \alpha_L(1-\rho)}\right)^{1-\delta}\right)^{\frac{1}{a+\delta-1}}$$

Heterogeneity in productivity: From Lemma (2) we have $c_H = c_L = c$. Given the assumed functional forms we have $\phi'_j(y_j) = \theta_j^\delta y_j^{\delta-1}$ and $u'(c) = c^{-a}$ so can rewrite the Euler equation (26) as

$$1 + \frac{2i}{\alpha} = \frac{c^{-a}}{\theta_j^\delta y_j^{\delta-1}} \quad \text{for } j = H, L.$$

From market clearing (19) we have $c = \rho y_H + (1-\rho)y_L$; from (15) we have $p = \phi'_H(y_H) = \phi'_L(y_L)$, which is

$$y_H = y_L \left(\frac{\theta_L}{\theta_H}\right)^{\frac{\delta}{\delta-1}} = y_L \theta^{\frac{\delta}{\delta-1}}$$

since we have normalized $\theta_L = \theta > \theta_H = 1$. So, market clearing implies:

$$c = y_L \left(\rho\theta^{\frac{\delta}{\delta-1}} + 1 - \rho\right).$$

Substituting for c in the Euler equation above we obtain

$$y_L = \left[\left(1 + \tfrac{2i}{\alpha}\right)\left(\rho\theta^{\frac{\delta}{\delta-1}} + 1 - \rho\right)^a \theta^\delta\right]^{\frac{1}{1-a-\delta}}.$$

Money is not the only asset: In this environment $y_H = y_L = y$. The expression for c_L is obtained from (23), whereas the expression for c_H is obtained from (41). Given the assumed functional forms we have $\phi'(y) = y^{\delta-1}$ and $u'(c_j) = c_j^{-a}$ so we can rewrite the Euler equation (23) as

$$1 + \frac{2i}{\alpha_L} = \frac{c_L^{-a}}{y^{\delta-1}},$$

and from (41) we have that

$$\tfrac{2-\beta}{\beta} = \tfrac{c_H^{-a}}{y^{\delta-1}}$$

which implies that

$$c_L = \left(\frac{\alpha_L + 2i}{\alpha_L}\frac{\beta}{2-\beta}\right)^{-\frac{1}{a}} c_H.$$

From market clearing (19) we have

$$y = \frac{\rho\alpha_H c_H + (1-\rho)\alpha_L c_L}{\rho\alpha_H + (1-\rho)\alpha_L}.$$

Substituting for y and c_L in (41) we obtain

$$c_H = \left(\frac{\beta}{2-\beta} \left(\frac{\alpha_H \rho + \alpha_L (1-\rho) \left(\frac{\alpha_L + 2i}{\alpha_L} \frac{\beta}{2-\beta} \right)^{-\frac{1}{a}}}{\alpha_H \rho + \alpha_L (1-\rho)} \right)^{1-\delta} \right)^{\frac{1}{a+\delta-1}}$$

Appendix B

In this appendix we discuss the case in which agents are only *ex-post heterogeneous*. That is, they receive idiosyncratic random type shocks in each odd date. The shock is such that the type assignment is i.i.d. across agents and time, with ρ being the probability that the arbitrary agent is assigned type H. So, by the law of large numbers the population of traders is randomly partitioned into two types; a fraction ρ of traders is of type H, and the complementary fraction $1-\rho$ is of type L, on each odd date. The model is otherwise identical to the case in which agents' types are fixed ex-ante. We will demonstrate that in this case there is no heterogeneity in money balances.

Heterogeneity in trading shocks

Agents receive a type shock at the beginning of each odd date. Before knowing the type shock he'll receive, each agent has a probability ρ of being of type H and a probability $1-\rho$ of being of type L. As in the case of fixed types, type H and type L agents have respectively a probability α_H and α_L of meeting a counterpart, with $0 < \alpha_L < \alpha_H \leq 1$.

At the start of an even date, the agent's problem can be represented as follows:

$$W_j(m_{j,k}) = \max_{q_j, m' \geq 0} \{U(q_j) - q_j - \pi m' + m_{j,k} + \tau + \beta EV(m')\} \tag{45}$$

and (7), (8) and (9) still hold. Note that at the end of the second market agents do not know their type for the following cycle and therefore they all choose the same money holdings m' based on the expected continuation utility $EV(m')$:

$$EV(m) = \rho V_H(m) + (1-\rho) V_L(m) \tag{46}$$

Given that we are focusing on monetary outcomes, i.e. $m' > 0$, we must have

$$1 = \tfrac{\beta}{\pi} \times \tfrac{\partial EV(m')}{\partial m'} \tag{47}$$

The intuition for (47) is analogous to the one for (12).

After an agent realizes his type shock j, his expected lifetime utility of entering a period with m must satisfy:

$$\begin{aligned} V_j(m) = \quad & \max_{c_j \in [0, \frac{m}{p}]} \tfrac{\alpha_j}{2}[u(c_j) + W_j(m_{j,b})] \\ & + \max_{y_j} \tfrac{\alpha_j}{2}[-\phi_j(y_j) + W_j(m_{j,s})] + (1-\alpha_j)W_j(m_{j,n}) \end{aligned} \tag{48}$$

The seller's problem is analogous to the case of fixed types and (15) still holds. Goods market clearing for odd dates is identical to what seen earlier.

The buyer's problem is similar to the case of fixed types, except for the fact that money holdings m do not depend on the agent's type, i.e. $m_{j,b} = m - pc_j$, so that (16) and (17) still hold.

To find the optimal cash holdings of an agent j we must calculate the expected marginal value of holding money, $\frac{\partial EV(m)}{\partial m}$:

$$\frac{\partial EV(m)}{\partial m} = \quad \rho \frac{\partial EV_H(m)}{\partial m} + (1-\rho)\frac{\partial EV_L(m)}{\partial m} \tag{49}$$

where $\frac{\partial V_j(m)}{\partial m}$ satisfies (21) for $j = H, L$.

Now we can calculate the equilibrium marginal value of money. Specifically,

$$\frac{\partial EV(m)}{\partial m} = \quad \rho[1 + \tfrac{\alpha_H}{2}(u'(c_H) - 1)\tfrac{\partial c_H}{\partial m}] + (1-\rho)[1 + \tfrac{\alpha_L}{2}(u'(c_L) - 1)\tfrac{\partial c_L}{\partial m}]$$

where $\frac{\partial c_j}{\partial m} = 1$ for $j = H, L$ if the agent is liquidity constrained, and zero otherwise. It follows that $EV(m)$ is strictly concave in cash holdings if at least a buyer is liquidity constrained, and linear otherwise. So, if $m \in (0, m^*]$ then the agent's optimal savings choice must satisfy

$$1 = \tfrac{\beta}{\pi}\left\{\rho[1 + \tfrac{\alpha_H}{2}(u'(c_H) - 1)] + (1-\rho)[1 + \tfrac{\alpha_L}{2}(u'(c_L) - 1)]\right\} \tag{50}$$

the intuition for which is analogous to the one for (23).

Goods market clearing on even dates is identical to (10). Given that all agents have identical money holdings, (4) becomes

$$\frac{M}{p_2} = m. \tag{51}$$

The definition of equilibrium is analogous to the layout of Definition 1 with the obvious modification. That is, we must account for equations (45), (46), (48), and (51).

In a stationary monetary equilibrium we must have $\pi \geq \beta$. Therefore, we prove the following result.

Lemma 3 *If $\pi > \beta$, then $c_L < c^*$ and $c_H < c^*$. If $\pi \rightarrow \beta$, then $c_L \rightarrow c^*$ and $c_H \rightarrow c^*$.*

Proof. From (50) we know that if $m > 0$ then $\frac{\pi-\beta}{\beta} = \frac{\alpha_H}{2}\rho(u'(c_H)-1)+\frac{\alpha_L}{2}(1-\rho)(u'(c_L)-1)$. Note that $\pi \geq \beta$ is necessary. If $\pi > \beta$, then $\frac{\alpha_H}{2}\rho(u'(c_H)-1)+\frac{\alpha_L}{2}(1-\rho)(u'(c_L)-1) > 0$. Since $m_H = m_L = m$, then $u'(c_H) > 1$ and $u'(c_L) > 1$. This implies $c_H < c^*$ and $c_L < c^*$.

As $\pi \rightarrow \beta$ then $u'(c_H) \rightarrow 1$ and $u'(c_L) \rightarrow 1$. Therefore, $m \rightarrow m^*$, $c_H \rightarrow c^*$ and $c_L \rightarrow c^*$ and neither type of agent is cash constrained. Thus, the Friedman rule can achieve the efficient allocation. Existence easily follows from inspection of the individual optimality and market clearing conditions.■

Lemma 3 has several implications. First, away from the Friedman rule all agents are cash-constrained. Remember that in this environment all agents carry the same money holdings m since they don't know their type for the following periods. Second, as $\pi \rightarrow \beta$ money holdings m converge to m^* and therefore neither type L nor type H agents are cash constrained, i.e. $c_H \rightarrow c^*$ and $c_L \rightarrow c^*$.

Heterogeneity in disutility from production

Agents receive a type shock at the beginning of each odd date. Before knowing the type shock he'll receive, each agent has a probability ρ of being of type H and a probability $1-\rho$ of being of type L. As in the case of fixed types, production of y output generates disutility $\phi_j(y)$ with $j \in \{L, H\}$ and $\phi'_L(y) > \phi'_H(y)$. As heterogeneity in this economy only affects marginal production disutility, money holdings choices will not be affected by the uncertainty faced by agents. That is, all agents still hold the same m at the beginning of every period as in the case of fixed types. Therefore, Lemma 2 still holds in this environment, i.e., allocations are inefficient away from the Friedman rule and efficient when the Friedman rule is implemented.

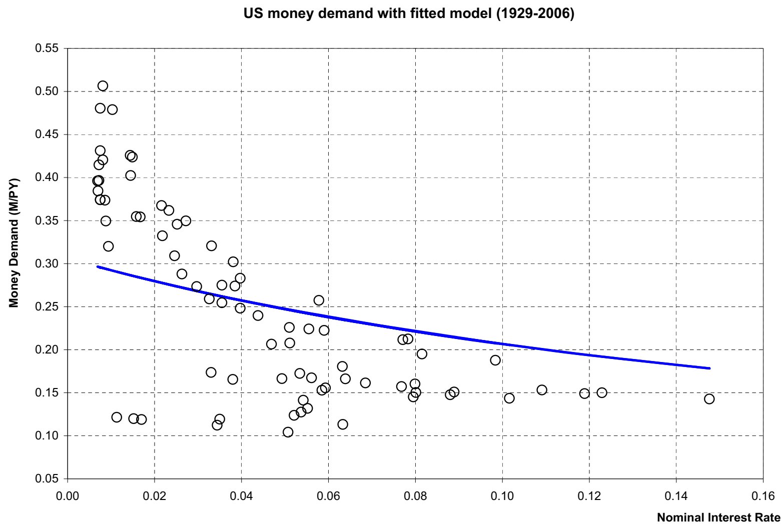

Figure 1: Quality of fit

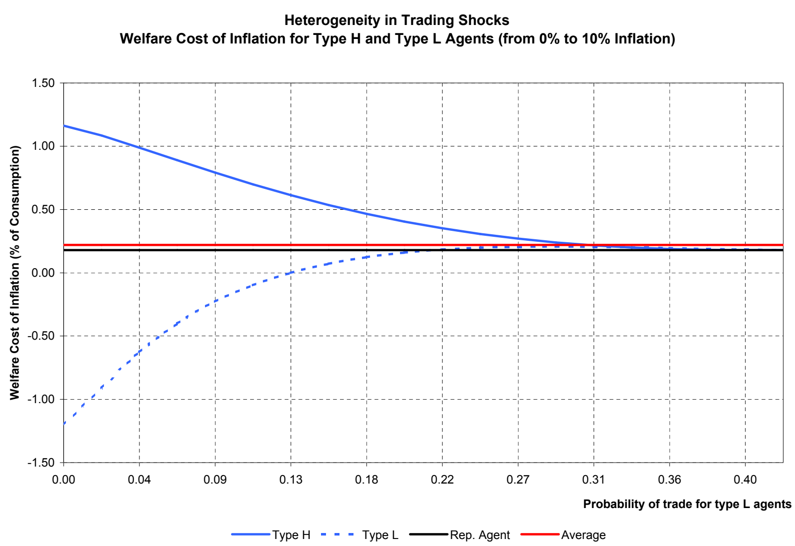

Figure 2: Welfare cost (zero inflation), heterogeneity in trading shocks