

Seigniorage and distortionary taxation in a model with heterogeneous agents and idiosyncratic uncertainty*

Sofia Bauducco[†]

December 2016

Abstract

In this paper we study the optimal monetary and fiscal policy mix in a model in which agents are subject to idiosyncratic uninsurable shocks to their labor productivity. We identify two main effects of anticipated inflation absent in representative agent frameworks. First, inflation stimulates savings for precautionary reasons. Hence, a higher level of anticipated inflation implies a higher capital stock in steady state, which translates into higher wages and lower taxes on labor income. This benefits poor, less productive agents. Second, inflation acts as a regressive consumption tax, which favors rich and productive agents. We calibrate our model economy to the U.S. economy and compute the optimal policy mix. We find that, for a utilitarian government, the Friedman rule is optimal even when we allow for the presence of heterogeneity and uninsurable idiosyncratic risk. Although the aggregate welfare costs of inflation are small, individual costs and benefits are large. Net winners from inflation are poor, less productive agents, while middle-class and rich households are always net losers.

JEL Classification: E52, E63, E21, D31.

Keywords: Seigniorage, Friedman rule, Heterogeneous Agents, Optimal Monetary and Fiscal Policy

*I would like to thank Albert Marcet and Josep Pijoan-Mas for fruitful discussions and solid guidance. I also want to thank Klaus Adam, Anton Braun, Andrea Caggese, Francesco Caprioli, Juan Carlos Conesa, Wouter den Haan, José Dorich, Martin Ellison, Andrés Erosa, Jordi Gali, Antonio Mele, Xavier Raurich, Michael Reiter, Thijs Van Rens, Hugo Rodríguez, and seminar participants at the UPF Macro Break, CEMFI Lunchtime Seminar, Drexel University, Temple University, University of Aarhus, Sveriges Riksbank, Norges Bank, Johns Hopkins SAIS, EPFL, Magyar Nemzeti Bank, Central Bank of Chile, 5th. European Workshop in Macroeconomics (Mannheim), XIV Workshop on Dynamic Macroeconomics (Vigo), CEA University of Chile, Royal Economic Society 2010 Conference, Encuentro Anual 2010 de la Sociedad de Economía de Chile and XXXV Simposio de Análisis Económico for useful comments and suggestions. All errors are my responsibility.

[†]Email: sbauducco@bcentral.cl

1 Introduction

The seminal papers by Friedman (1969) and Phelps (1973) opened a wide debate in the last decades over the issue of the optimal monetary and fiscal policy combination in representative agent frameworks. Friedman argued that optimality required setting the nominal interest rate to zero, so that the return on money holdings was equated to the return on any other interest-bearing nominal asset. This is known in the literature as the *Friedman rule*. Phelps, on the other hand, indicated that in economies in which lump-sum taxes are not available, the policy maker should tax all goods, including money. Moreover, since the money demand function is typically more inelastic than the demand for consumption goods, Phelps concluded that money should be taxed heavily. This apparent contradiction in the optimal policy prescription motivated some classical papers such as Chari et al. (1996) and Correia and Teles (1996) among many others. With some exceptions, the general consensus appears to be that Central Banks should follow the Friedman rule.

However, by construction all these early contributions overlook issues of heterogeneity and redistribution. By working with the representative agent assumption, their analysis of optimal policy focuses only on efficiency in distorting relative prices. Therefore, a crucial aspect of inflation is neglected, which is the fact that it does not affect all individuals in the same way.

In this paper we tackle this issue by building a heterogeneous agent model with uninsurable idiosyncratic shocks to labor productivity. We consider the problem of a benevolent government that has to finance an exogenous and constant stream of public expenditure. The available instruments are the inflation tax and a tax on labor income and, given that labor supply is endogenous, both instruments are distortionary. We look for the determination of the optimal monetary and fiscal policy mix in such an environment, assuming the government assigns equal weight to all agents in the economy¹.

Since we are interested in identifying and studying the effects of inflation over individuals with different income and wealth profiles, we need a model in which agents differ in these two dimensions. In order to assess the effect exerted by inflation on the incentives of households to consume, work and save, and consequently on aggregate long-run capital and output, we depart from the complete-markets assumption by assuming that agents cannot insure against their idiosyncratic shocks and are subject to a borrowing constraint. Finally, previous literature has pointed to the fact that inflation can be regarded as a regressive tax on consumption². We introduce this in the analysis by means of a transaction technology alternative to money in which richer, more productive agents have comparative advantages relative to poorer, less productive

¹Examples of other papers that use the same social welfare function are Domeij and Heathcote (2004), Floden (2001), Floden and Linde (2001) and Heathcote et al. (2008a).

²See Erosa and Ventura (2002).

ones. In addition, we exacerbate these advantages by assuming easier access to the transaction technology for more productive agents.

In the setup we propose there are two effects from inflation that are not present in representative-agent frameworks. On the one hand, more productive agents, by having easier access to alternative transaction technologies, can shelter better from inflation. This shifts the burden of taxation from richer, more productive agents to poorer households, thus benefiting the former group. If the planner cares sufficiently for poor agents, this effect goes in favor of the Friedman rule. On the other hand, in economies in which households cannot insure against their idiosyncratic shocks, inflation amplifies the motive for precautionary savings. When the bad shock hits and the individual is very poor (i.e., is close to hitting the borrowing constraint) the inflation tax reduces consumption and leisure and, consequently, utility, thus creating an incentive to save. The higher savings level translates into higher capital in steady state, higher wages and lower labor tax rates. By this means, a higher level of inflation increases welfare of poor, low productivity households that rely almost entirely on labor income. When the government cares about poor households, this effect goes against the Friedman rule. Considering these two effects jointly, we see that they operate in opposite directions. Deviating from the Friedman rule assures poor agents a higher labor income, while middle-class and rich agents have to endure lower levels of capital income. Nevertheless, they are benefited by a reduction in their tax burden associated to the increase in the inflation tax which, as explained before, is a regressive tax.

There is a distortion associated to inflation that is already present in environments with a representative agent, which is related to the uniform taxation argument from the public finance literature. When consumption goods can be bought either with cash or with an alternative transaction technology, deviating from the Friedman rule implies taxing the goods bought with cash more than the rest of the goods. If all goods enter in the utility function of the household in identical manner, this is not efficient³.

From the previous discussion it can be concluded that, in economies with uninsurable idiosyncratic uncertainty and heterogeneous agents, the determination of the optimal monetary and fiscal policy mix remains a quantitative question. In order to provide a suitable answer, we calibrate the model economy to match some selected statistics of the U.S. economy. Some key targets are the correlation between money and asset holdings, the fraction of consumption expenditures made with cash and the Gini coefficient of the asset distribution. We define the benchmark economy to be one that displays an annual rate of inflation of 2%⁴.

Given our parameterization we find that for a utilitarian planner⁵ the Friedman rule is

³To be more precise, for this argument to hold the utility function has to be separable in consumption and leisure and the subutility over consumption goods has to be homothetic.

⁴This is a reasonable annual inflation target for the Fed.

⁵A utilitarian planner is one that assigns equal weights to all households in the computation of the social welfare function.

optimal despite the introduction of heterogeneity and uninsurable idiosyncratic risk. The aforementioned beneficial effects of inflation on welfare are not sufficiently large, from a quantitative point of view, to offset its detrimental effects. Thus it is optimal to set the nominal interest rate equal to zero. Next, we perform a welfare analysis comparing the benchmark economy to an economy in which the Friedman rule is implemented. We find that the aggregate welfare gains of switching from the benchmark policy regime to the optimal one are rather small. The percent of life-time consumption that agents in the benchmark economy are prepared to give up to get the policy change is 0.51%. Following Floden (2001) we decompose these welfare gains into gains from increased levels of consumption, reduced uncertainty and reduced inequality. We find that most of the gains are due to increased levels of consumption and, to a lesser extent, to reduced uncertainty.

Despite these seemingly small aggregate welfare gains, the individual welfare gains and losses can be very large. A surprising finding is that poor, less productive agents are net losers from the policy change⁶. When inflation decreases, so does aggregate capital and, with it, wages. This effect can be very harmful for these agents: we find that, for the poorest individual, consumption should decrease permanently by about 4% in the benchmark economy for him to be indifferent between living in any of the two worlds proposed. On the contrary, middle-class and rich individuals are net winners from the change in regime. Rich, low-productivity households should obtain a permanent increase in consumption of around 4% to be indifferent between the two regimes. These large individual effects cancel out almost completely in the aggregate, thus yielding the small overall effects described before.

Although this paper is not the first one to look at inflation in heterogeneous-agent environments, it is, to our knowledge, the first one that identifies in a unified framework different effects of inflation that had been described separately and derives the optimal policy within such framework.

The contributions of Erosa and Ventura (2002) and Algan and Ragot (2006) study the redistributive aspects of inflation. Both papers point out different mechanisms through which inflation affects the agents' welfare depending on their level of wealth and labor productivity, namely, that inflation acts as a regressive tax on consumption and that it stimulates savings. They reach contradictory conclusions: while the former states that inflation is relatively more harmful to poorer households, the latter claims that the higher level of capital in steady state translates into higher wages and higher welfare for poor, labor-income dependent individuals. Neither of these papers addresses the issue of inflation from a normative point of view.

There are a number of papers that deal with the determination of the optimal inflation rate when taking into account issues of heterogeneity. Akyol (2004) studies an endowment economy in which only high productivity households hold money in equilibrium. Seigniorage revenues can

⁶See Algan and Ragot (2006) for a similar result.

be used to finance redistributive (anonymous) transfers and interest payments on government debt. Therefore, the main role of inflation is to redistribute resources from agents with high endowment shocks to those holding bonds, which improves risk-sharing. The author finds that the optimal inflation rate is of about 10%. Although in this paper not all high-productivity agents are also rich, the correlation between wealth and productivity is very high. Then, the idea that richer agents are the ones holding money in equilibrium is at odds with some stylized facts on transactions and cash holdings reported in Erosa and Ventura (2002), which indicate the opposite. In this paper, we construct the model such that its predictions in terms of cash holdings and proportion of purchases made with cash are in line with what the data suggests.

Albanesi (2005) and Bhattacharya et al. (2007) propose frameworks in which agents are ex-ante heterogeneous and there is no idiosyncratic uncertainty. In particular, they assume there are two types of households in the economy. While in Albanesi (2005), agents of different types have different labor productivities, Bhattacharya et al. (2007) assume that agents differ in the marginal utility they derive from real money balances. Albanesi (2005) finds that it may be optimal to depart from the Friedman rule, depending on the weight that the government assigns to each type of agent. On the other hand, Bhattacharya et al. (2007) conclude that, because inflation redistributes resources from one type of agent to the other, both types may benefit if the central bank deviates from the Friedman rule. Both papers abstract from capital, but given the lack of idiosyncratic uninsurable risk, monetary policy would not affect the long-run capital stock. This is a crucial element that we include in our analysis.

Some papers in the search literature study the implications of anticipated inflation, as for example Lagos and Rocheteau (2005) and Bhattacharya et al. (2005). These studies view monetary policy as a mechanism to induce agents to exert the correct amount of search effort. Search effort is related to the quantity of output produced and, consequently, welfare. The scope of this literature is substantially different from ours. Although we do regard money as a means of payment, our interest in inflation is related to the fact that it distorts the consumption, leisure and savings decisions. Thus, we do not address in detail how money facilitates transactions and to what extent monetary policy can enhance this role but, rather, focus on the effects of money growth on allocations.

Finally, da Costa and Werning (2008) propose a framework in which agents have private information on their labor productivity. The analysis abstracts from idiosyncratic risk by assuming that differences in productivity are permanent. Moreover, agents do not hold capital. The authors assume that money and work effort are complements, so that the demand for money, conditional on the expenditure of goods, weakly increases with the amount of work effort. They find that the Friedman rule is optimal if labor income is positively taxed. The reason for this result is that deviating from the Friedman rule does not aid the planner in designing a mechanism to ensure that individuals do not underreport their productivity, i.e., it does not help relaxing

the incentive-compatibility constraints in the planner’s problem. Although this study and ours reach similar conclusions in terms of policy prescription, the reasons behind this result are very different in the two setups. We assume the government cannot observe an agent’s productivity by restricting the set of fiscal instruments to an anonymous tax on labor, so in our framework agents do not have incentives to underreport by construction. Instead, we focus on the effects of inflation on households that differ in wealth as well as labor productivity.

The remainder of the paper is organized as follows. Section 2 describes the model. In section 3 we discuss the calibration strategy. Section 4 contains a description of the different effects operating in the model and shows the results in terms of optimal policy. In section 5 we perform the welfare analysis. Finally, section 6 concludes and discusses lines for future research.

2 The model

The model is close to Erosa and Ventura (2002) but with two important differences. First, in our model the labor supply is endogenous. Second, we introduce the idea that more productive agents have easier access to transaction technologies alternative to the use of cash, with respect to less productive agents. These two modifications to the basic setup alter completely the analysis of the effects of inflation over different population sectors. We defer the discussion of these effects until the next section.

Households face idiosyncratic labor productivity shocks, that we denote by ε_t . There is no aggregate uncertainty. Consequently, the economy is at its steady state and all aggregate real variables remain constant. For simplicity, we will omit the time subscript from aggregate real variables.

2.1 Households

The economy is populated by a continuum of mass 1 of ex-ante identical and infinitely lived households. Households are endowed with one unit of time each period and derive utility from consumption of final goods and leisure.

Markets are incomplete, in the sense that it is not possible to trade bonds which payoffs are contingent on the realization of the idiosyncratic shock. Moreover, we assume that agents cannot borrow. Consequently, agents can only save by holding one-period riskless assets and money. We denote by W_t the total nominal wealth an individual has in period t , where W_t is the sum of total money holdings M_t and assets A_t .

Agents consume a continuum of final goods indexed by $j \in [0, 1]$. We assume that the consumption aggregator, denoted by c , takes the form $c = \inf_j \{c(j)\}$ ⁷. The choice of the

⁷We choose this aggregator for simplicity reasons only. Working with a more general aggregator such as a Dixit-Stiglitz aggregator does not alter our results qualitatively.

aggregator implies that agents consume an equal amount of each good j , therefore

$$c = c(j) = c(m) \quad \forall j, m \in [0, 1]$$

Agents choose optimally whether to buy final goods with cash or with a costly transaction technology, which we will call credit. In order to buy an amount c of good j with credit, the consumer must purchase $\gamma(j)$ units of financial services.

Following Lucas and Stokey (1983), we assume that the financial market closes first and the goods market follows. More specifically, at the beginning of period t , after observing the current shock ε_t , agents adjust their portfolio compositions by trading money and assets in a centralized securities market and pay the credit obligations that they contracted in the previous period. In this sense, the transaction technology we consider does not allow households to transfer liabilities from one period to the other. Instead, it represents a way in which consumers can transform their interest-bearing assets into a means of payment that is not subject to the inflation tax.

The gross nominal return of a one-period bond A_t is the nominal interest rate R . Notice that the gross nominal return of money is 1, so money is (weakly) dominated in rate of return by assets. Nevertheless, households value money because it provides liquidity services to buy consumption goods.

After trade in the securities market has taken place, the goods market opens. At this stage, households buy consumption goods with money or credit, decide how much to work and save a fraction of their total income in the form of nominal wealth W_{t+1} that they will carry to the next period to transform it into assets and money in the securities market. Household i 's budget constraint at the time the goods market opens is

$$p_t c_t^i + q_t \int_0^{z_t^i} \gamma^i(j) dj + W_{t+1}^i = R A_t^i + M_t^i + (1 - \tau^l) \omega p_t (1 - l_t^i) \varepsilon_t^i \quad (1)$$

where p_t is the unitary price of the final good, q_t is the price of a unit of financial services, $\int_0^{z_t^i} \gamma^i(j) dj$ is the total amount of credit bought, ω is the real wage for one efficiency unit of labor, l_t^i is time devoted to leisure and τ^l is an anonymous linear tax on labor income.

As mentioned before, households need either cash or credit in order to buy consumption goods. Agents choose which goods they will buy with credit and which with cash, and $z_t^i \in [0, 1]$ stands for the fraction of goods bought with credit by household i ⁸. Since all goods that are not bought with credit need to be paid with money, household i faces the following cash-in-advance constraint in the goods market:

$$p_t c_t^i (1 - z_t^i) \leq M_t^i \quad (2)$$

⁸To be more precise, agents will buy goods indexed from 0 to z_t^i with credit and the rest of the goods (from z_t^i to index 1) with cash.

Notice that, if $R > 1$, money is strictly dominated in rate of return by assets and, consequently, constraint (2) will always be binding because agents can adjust their money holdings *after* the idiosyncratic shock is observed.

The problem that household i solves can be stated as

$$\max E_0 \sum_{t=0}^{\infty} \beta^t u(c_t^i, l_t^i)$$

subject to

$$\begin{aligned} p_t c_t^i + q_t \int_0^{z_t^i} \gamma^i(j) dj + W_{t+1}^i &= R A_t^i + M_t^i + (1 - \tau^l) \omega p_t (1 - l_t^i) \varepsilon_t^i \\ p_t c_t^i (1 - z_t^i) &\leq M_t^i \\ A_t^i &\geq 0 \end{aligned}$$

In the appendix we show the optimality conditions associated to this problem.

2.2 Firms

There are two types of competitive firms in this economy: firms that produce consumption goods and financial firms that produce transaction services. We assume that all markets are perfectly competitive and, in consequence, firms make zero profits in equilibrium.

2.2.1 Consumption goods sector

Let K denote the aggregate capital stock and L^g aggregate labor in efficiency units employed in the goods sector. Then the production technology in the consumption goods sector can be written as

$$y_t = F(K, L^g)$$

where $F(\cdot)$ is a neoclassical production function⁹. Optimality conditions of the firm are

$$r = F_K - \delta \tag{3}$$

$$\omega = F_L^g \tag{4}$$

where r is the before-tax real return on capital and δ is the depreciation rate.

⁹We assume that the production function is identical for any type of consumption good $i \in [0, 1]$ so the relative price of any two types of goods i and j is 1 $\forall i, j$.

Figure 1: Marginal cost of credit $\gamma(j)$ for $\varepsilon_1 < \varepsilon_2$

2.2.2 Financial services sector

We assume that a household with labor productivity ε_t^i can acquire a fraction of goods \tilde{z}_t^i with credit at zero marginal cost, and that \tilde{z}_t^i depends on the potential labor income of an agent. More specifically, $\tilde{z}_t^i = f(\varepsilon_t^i)$ with $\frac{\partial \tilde{z}_t^i}{\partial \varepsilon_t^i} > 0$. In words, this assumption means that the fraction of goods that can be purchased with credit at zero marginal cost increases with the productivity of an agent. Then, more productive agents have an advantage in the use of credit with respect to less productive ones. Figure 1 shows the marginal cost of credit for goods $i \in [0, 1]$ for agents n and m with $\varepsilon_t^n > \varepsilon_t^m$.

The assumption that \tilde{z}_t^i depends positively on the labor productivity of agent i is a shortcut to reflect the better access to commercial credit markets that high-income, rich households enjoy when compared to poor, low-income ones. Think of a similar scenario to the one proposed here, but now at the beginning of each period, after observing her individual shock, household i decides whether to repay her credit obligations contracted in the last period. If the household chooses not to repay she is excluded from the credit market for the next period, otherwise it can apply for a new credit line. The household will decide to pay the credit obligations contracted in period $t - 1$ only if the value of the credit in t is higher or equal than what she owes from $t - 1$. Obviously, the amount of credit in t is determined by the wealth and the labor productivity of the agent in t . In $t - 1$, when the agent applies for the loan, the financial services firm will charge an interest rate that reflects the risk that the agent defaults in the next period. This risk will be decreasing in wealth and productivity since richer, more productive agents are more likely to use credit more intensely in the next period.

With this argument in mind, it is natural to think that the cost of credit depends positively

on the earnings capacity of an individual. A natural way to capture this feature in the model is through \tilde{z}_t^i . Of course, given our previous discussion, one would naturally think that \tilde{z}_t^i should depend on W_t^i as well. Nevertheless, adding this dependence complicates the numerical solution of the model and, since ε_t^i and W_t^i are highly correlated anyway, it presumably does not change the results substantially.

For goods $j \in [\tilde{z}_t^i, 1]$ the nominal marginal cost of the use of credit is $q_t \gamma^i(j)$. The total cost of credit for agent i with productivity ε_t^i that buys a fraction $z_t^i > \tilde{z}_t^i$ of goods with credit is

$$q_t \int_0^{z_t^i} \gamma^i(j) dj = q_t \int_{\tilde{z}_t^i}^{z_t^i} \gamma^i(j) dj$$

Following Erosa and Ventura (2002), the function $\gamma(i)$ is strictly increasing in i for $i > \tilde{z}_t$ and satisfies $\lim_{i \rightarrow 1} \gamma(i) = \infty^{10}$. This assumption guarantees that some goods will be purchased with cash so that there is a well-defined demand for money.

The technology to produce transaction services requires one unit of labor (in efficiency units) per unit of service produced. We denote by L^c the amount of labor in efficiency units employed in the production of transaction services:

$$\int \int_{\tilde{z}_t^i}^{z_t^i} \gamma^i(j) dj d\lambda = L^c$$

where λ is the distribution of agents in the economy. Firms in the sector charge the price q_t per unit of credit sold. Competition ensures that firms make zero profits and set their prices such that $\omega p_t = q_t^{11}$.

2.3 Government

The government has to finance an exogenous stream of public spending through distortionary taxes on aggregate labor income and asset returns and through seigniorage revenues. The nominal budget constraint of the government is

$$p_t g + M_t = \tau^l \omega p_t L + \tau^k r p_t K + M_{t+1} \quad (5)$$

where L is total aggregate labor in efficiency units $L = L^g + L^c$ and τ^k is the tax rate on asset returns.

¹⁰The meaning and role of \tilde{z}_t will be analyzed in detail later.

¹¹The zero profit condition can be written as

$$q_t \int \int_0^{z_t^i} \gamma(j) dj d\lambda = p_t \omega L^c$$

2.4 Equilibrium

In our economy, each agent is characterized by the pair (w_t, ε_t) where w_t is wealth in real terms. Let $W \equiv [0, \bar{w}]$ be the compact set of all possible wealth holdings where \bar{w} is an upper bound on wealth and the lower bound is determined by the no borrowing condition¹². Let $E \equiv \{\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n\}$ be the set of all possible realizations of the labor productivity shock ε_t . The shock follows a Markov process with transition probabilities $\pi(\varepsilon', \varepsilon) = \Pr(\varepsilon_{t+1} = \varepsilon' | \varepsilon_t = \varepsilon)$. Define the state space S as the cartesian product $S = W \times E$ with Borel σ -algebra \mathcal{B} and typical subset $\mathcal{S} = (W \times E)$. The space (S, \mathcal{B}) is a measurable space, and for any set $\mathcal{S} \in \mathcal{B}$ $\lambda(\mathcal{S})$ is the measure of agents in the set \mathcal{S} . Denote Λ as the set of all probability measures over (S, \mathcal{B}) . Then,

$$\lambda_{t+1}(W \times E) = \int_S Q((w, \varepsilon), W \times E) d\lambda_t(w, \varepsilon)$$

where $Q((w, \varepsilon), W \times E)$ is the probability that an individual with current state (w, ε) is in the set $W \times E$ next period:

$$Q((w, \varepsilon), W \times E) = \sum_{\varepsilon' \in E} I\{w'(w, \varepsilon) \in W\} \pi(\varepsilon', \varepsilon)$$

Here $I(\cdot)$ is the indicator function and $w'(w, \varepsilon)$ is the optimal savings policy of an individual in state (w, ε) .

DEFINITION 1. *A stationary equilibrium is given by functions $\{g^c, g^l, g^z, g^{w'}, g^a, L^g, L^c\}$, a price system $\{p_t, \omega, q_t, r\}_{t=0}^\infty$ and government policies $\{\tau^l, \tau^k, R, M_t\}_{t=0}^\infty$ such that*

1. *Given prices and government policies, the allocations solve the household's problem.*
2. $r = F_K(K, L^g) - \delta, \omega = F_{L^g}(K, L^g)$
3. *Given the allocations and price system, the government budget constraint is satisfied.*

¹²Notice that the no borrowing condition means that

$$a_t \geq 0 \tag{6}$$

Nevertheless, since $w_t = a_t + m_t$ and a_t and m_t are decision variables of the agent, only by imposing $w_t \geq 0$ we make sure that condition (6) is satisfied always.

4. *Markets clear:*

$$\begin{aligned}\int g^c d\lambda + g + \delta K &= F(L^g, K) \\ \int \int_{\bar{z}_t}^{z_t} \gamma(j) dj d\lambda &= L^c \\ L^g + L^c &= \int \varepsilon(1 - g^l) d\lambda \\ K &= \int g^a d\lambda\end{aligned}$$

5. *The measure of households is stationary:*

$$\lambda^*(\mathcal{W} \times \mathcal{E}) = \int_S Q((w, \varepsilon), \mathcal{W} \times \mathcal{E}) d\lambda^*(w, \varepsilon)$$

2.5 Optimal Policy

The government is benevolent and seeks to set taxes and seigniorage so as to maximize a social welfare function. We define the objective function in the maximization problem of the government to be the *utilitarian social welfare function*, U :

$$U = \int E_t V(\{c_s, l_s\}_{s=t}^\infty) d\lambda \quad (7)$$

where $V(\{c_s, l_s\}_{s=t}^\infty) = \sum_{s=t}^\infty \beta^{s-t} u(c_s, l_s)$ is life-time utility at time t . As we can see in expression (7), all households receive an equal weight for the computation of social welfare. A standard interpretation for this criterion is that the planner maximizes welfare under the veil of ignorance; that is, ex-ante welfare for a hypothetical household before knowing in which point of the distribution she is.

The problem of the government can be stated as

$$\max_{\tau^l, R} \int E_0 \sum_{t=0}^\infty \beta^t u(c_t^i, l_t^i) d\lambda$$

subject to

Competitive equilibrium (equations (20) - (24))

$$p_t g + M_t = \tau^l \omega p_t L + \tau^k r p_t K + M_{t+1}$$

3 Effects of inflation

It is well known that in representative-agent frameworks the Friedman rule is the optimal policy recommendation for a wide variety of models (see Chari et al. (1996)). This result extends to our model economy if we shut down heterogeneity among households. In the case in which $\varepsilon_t^i = \varepsilon \forall i$ a benevolent planner sets $R = 1$. The intuition behind this result lies in the uniform commodity taxation argument from the public finance literature. Notice that in our framework goods bought with cash and with credit enter the utility function in identical manner¹³. Therefore, the tax on labor income implicitly taxes all goods, whether bought with cash or with an alternative transaction technology, at an identical rate. Setting $R > 1$ entails taxing more those goods bought with cash, which is not efficient. Given the representative-agent assumption, in the model we are describing there are no issues of redistribution or self-insurance. Moreover, the capital-labor ratio is pinned down by the intertemporal discount factor β , so inflation does not affect the return on capital or the wage rate. Thus, efficiency in the taxation of different goods is the only aspect that the planner should take into account when designing the optimal policy plan.

When we introduce idiosyncratic uninsurable risk the analysis changes substantially. Now inflation has effects over the level of capital in steady state. In addition, due to the presence of an alternative transaction technology like the one we have introduced, inflation acts as a regressive tax on consumption. We proceed to describe these effects in detail. We argue that, once these effects are taken into consideration, the determination of the optimal policy mix remains an open question that needs a quantitative answer.

3.1 Inflation as a regressive tax on consumption

As described by Erosa and Ventura (2002), inflation can act as a regressive consumption tax when, in a heterogeneous agent setup such as ours, we allow agents to substitute cash by an alternative transaction technology that displays economies of scale. For the moment we abstract from the presence of idiosyncratic risk, since all the analysis holds by allowing for heterogeneity in labor productivities only.

Without loss of generality, assume that an agent's productivity ε^i is constant $\forall t$, $\varepsilon^i \in E = [\varepsilon^1, \varepsilon^2]$ with $\varepsilon^1 < \varepsilon^2$ and there is an equal mass of each type of agent in the population¹⁴. Furthermore, assume for simplicity that initial wealth holdings w_0^1 and w_0^2 are such that the economy is in steady state from $t = 0$ onwards. Then, optimality for agent i requires that

¹³The argument that follows holds always that utility is separable in consumption and leisure, and the utility over consumption goods is homothetic, see Chari et al. (1996) for a formal proof in a similar model.

¹⁴The results of this section are robust to changes in the number of productivity states and in the composition of the population.

$$R = 1 + \frac{\omega\gamma(z^i)}{c^i} \quad (8)$$

The second term on the right hand side of the previous expression is the unitary cost of credit for the threshold good z^i . It is clear from expression (8) and from the functional specification of the transaction technology (11) that this unitary cost decreases when the volume transacted increases. Thus, the transaction technology displays economies of scale.

Assume for now that $\tilde{z}^i = 0$ for $i = 1, 2$. From (8) it is immediate to see that $z^1 = z^2 = 0$ when $R = 1$, i.e, when the planner follows the Friedman rule both types of agents buy all consumption goods using cash, since holding cash does not bring about any opportunity cost. On the contrary, if $R > 1$, $z^2 > z^1 > 0$, given that $c^2 > c^1$ because agent 2 enjoys a higher labor income and, therefore, a higher level of consumption¹⁵. Then it follows that the more productive agents use the credit technology more intensely. This feature of the model is consistent with the evidence on transaction patterns and portfolio holdings that Erosa and Ventura (2002) report in their paper, which can be summarized in three main facts: high income individuals buy a smaller fraction of their consumption with cash, the fraction of wealth in the form of liquid assets held by a household decreases with her wealth and income and, finally, a non-negligible fraction of households does not own a credit card.

Due to the presence of economies of scale in the transaction technology, buying goods with credit is relatively more expensive for less productive agents. Because these agents buy a larger fraction of goods with cash, they need to hold a relatively larger fraction of their income in liquid assets. It is in this sense that inflation acts as a regressive tax on consumption, since setting $R > 1$ corresponds to taxing low-income individuals more.

If $\tilde{z}^2 > \tilde{z}^1$ this asymmetric effect of the inflation tax is exacerbated. As we discussed in section 2.2.2, the introduction of \tilde{z}^i is a shortcut to model differences in the access to commercial credit markets that high-income, rich households enjoy when compared to poor, low-income ones. The regressive nature of the inflation tax implies that *for high productivity households it is optimal to set a gross nominal interest rate higher than 1*, since in this way the burden of taxation is shifted to poor, unproductive individuals. To see the intuition behind this statement, think of the limit case in which $\tilde{z}^2 = 1$. In this case the inflation tax does not affect agents of type 2 in any way, so they would want the government to set it as high as necessary to finance completely its expenditure from seigniorage revenue.

In the appendix we show (numerically) that, in the current setup, a benevolent government (a Ramsey planner) that assigns a sufficiently high Pareto weight on type 2 agents would find it optimal to deviate from the Friedman rule and set $R > 1$.

¹⁵Strictly speaking, this is only true for particular levels of initial wealth w_0^1 and w_0^2 . Here we are implicitly making the assumption that the more productive agent is at least as rich as the less productive one.

3.2 Inflation as a motive for precautionary savings

The effect described in the previous section is at work due to the assumption of heterogeneity and the transaction technology we have specified, but it does not depend on the presence of uninsurable idiosyncratic risk. In this section we argue that inflation accentuates such risk and that, consequently, households save more when inflation is high.

Consider first the case in which there is no uncertainty (aggregate or idiosyncratic). Then the capital/labor ratio is determined in steady state by the discount factor β and is completely independent of the inflation rate. In this sense *inflation is neutral* and it does not affect the wage rate or the real interest rate in steady state. This is also true if we allow for idiosyncratic uncertainty but assume that households can trade a complete set of Arrow securities contingent on the realization of the labor productivity shock. It is easy to show that in this case agents can do full risk sharing and, if there is no aggregate uncertainty and utility is separable in consumption and leisure, enjoy a constant level of consumption independently of their current labor productivity. As in the case with no uncertainty, β determines the aggregate capital-labor ratio, ω and r ¹⁶.

In models with incomplete markets and borrowing constraints, agents save not only to smooth consumption by transferring resources from one period to the other, but also to insure themselves against bad realizations of the shock that may push them close to the borrowing constraint and force them to consume very little. The absence of complete markets and the presence of borrowing constraints lead agents to save for precautionary reasons¹⁷. Moreover, the more uncertain future income (and consumption) becomes, the stronger the motive for precautionary savings. The increase in savings translates into an increase in the capital stock in steady state, with the consequent decrease in the real interest rate and increase in the wage rate.

In the economy we have described in previous sections, a higher level of steady state inflation implies that future consumption is more uncertain and, consequently, reinforces the incentives to save. To see this, consider the no-borrowing constraint of household i , which says that $A_t^i \geq 0$. As shown in the appendix, this constraint can be re-written as:

$$c_t^i(1 - z_t^i)(1 + \pi) \leq w_t^i \tag{9}$$

where π is the inflation rate. Consider a steady state with π^A , in which household i at time t hits the constraint:

¹⁶We have abstracted from the possibility that there is aggregate uncertainty. In this case, if there are incomplete markets with respect to the aggregate shock, inflation can have an active role as a mechanism to complete the markets. See Chari et al. (1991) for a discussion.

¹⁷When the marginal utility is convex, i.e., utility displays a positive third derivative, independently of the presence of borrowing constraints, agents save because of prudence.

$$c_t^{i,A}(1 - z_t^{i,A})(1 + \pi^A) = w_t^{i,A}$$

If inflation were higher, say $\pi^B > \pi^A$, to sustain the same level of consumption $c_t^{i,A}$ and the same fraction of goods bought with credit $z_t^{i,A}$, household i would need to have a higher level of wealth in order to satisfy constraint (9). Similarly, for a household i that has wealth holdings $w_t^{i,A}$ in an economy where the inflation rate is $\pi^B > \pi^A$, either $c_t^{i,B} < c_t^{i,A}$, $z_t^{i,B} > z_t^{i,A}$, or a combination of both. Raising z_t^i entails working more to be able to pay for the higher credit expenses; since the intratemporal optimality condition (23) has to be satisfied, consumption needs to decrease, so it has to be the case that $c_t^{i,B} < c_t^{i,A}$. This means that the higher level of inflation π^B renders consumption more uncertain.

The previous discussion points out to the fact that inflation raises the incentives to save and, as a consequence, the level of steady-state capital. Thus, an economy with higher inflation displays a lower real interest rate and higher wages. Also, because the budget constraint of the government (5) has to be satisfied, the higher seigniorage revenue calls for a decrease in τ^l . Poor agents, who rely almost entirely on their labor income and whose marginal utilities of consumption and leisure are very large, find this beneficial because a small increase in disposable income translates into a sizeable increase in utility. On the other hand, middle-class and rich households are harmed by the reduction in their capital income derived from the lower real interest rate.

4 Calibration and functional specification

The model described here cannot be solved analytically. Consequently, we need to resort to numerical techniques to obtain a solution. In what follows, we describe the calibration strategy and functional specification we work with.

The length of the model period is one year. We define the benchmark steady state as one in which the government sets its policy in accordance with what is observed for the U.S economy. In particular, we set inflation to be 2% annually in the benchmark steady state. Next, we select the model parameters so that in the benchmark steady state equilibrium the model economy matches some selected features of the U.S data. While some of the parameters can be set externally, others are estimated within the model and require solving for equilibrium allocations¹⁸. We summarize the values of externally set and internally calibrated parameters in Tables 4.1.2 and 4.2.1, as well as the targets they are related to and the values for the targets obtained from the model. Notice that, in the case of internally determined parameters, all parameters affect all

¹⁸The distinction between externally set and internally calibrated parameters corresponds to Heathcote et al. (2008b).

calibration targets. Nevertheless, since usually each parameter affects more directly only one target, we report in the table the target that it is more related to¹⁹.

4.1 Parameters set exogenously

4.1.1 Preferences

Following Domeij and Flodén (2006), the utility function we use is

$$u(c_t, l_t) = \frac{c_t^{1-\sigma} - 1}{1-\sigma} - \theta_0 \frac{(1-l_t)^{1+\frac{1}{\theta_1}}}{1+\frac{1}{\theta_1}} \quad (10)$$

This specification is convenient because the Frisch elasticity of labor supply is equal to θ_1 . We set $\theta_1 = 0.59$, which is in line with Domeij and Flodén (2006) estimates²⁰. It is common to find in the literature that $\sigma \in [1, 2]$, so we set $\sigma = 1.5$ ²¹.

4.1.2 Technology

Consumption goods sector: The technology for the production of consumption goods is a standard Cobb-Douglas function

$$y = K^\alpha L^{1-\alpha}$$

α is set such that the labor income share of GDP, $wL/Y = 1 - \alpha = 0.64$.

Transaction services sector: We adopt the following credit technology, which is a modified version of what Dotsey and Ireland (1996), Erosa and Ventura (2002) and Albanesi (2005) use:

$$\gamma^i(j) = \gamma_0 \left(\frac{j - \tilde{z}_t^i}{1 - j} \right)^{\gamma_1} \quad (11)$$

γ_0 and γ_1 are internally calibrated, as explained in section 4.2.1.

As was explained before, \tilde{z}_t^i depends on the labor productivity of an agent, ε_t^i . A proposed function for \tilde{z}_t^i is

$$\tilde{z}_t^i = \mu_0 + \mu_1 ((1 - \bar{l})\omega\varepsilon_t^i)$$

¹⁹As Pijoan-Mas (2006) explains, this calibration strategy can be seen as an exactly identified simulated method of moments estimation.

²⁰These authors claim that previous estimates of the labor-supply elasticity are inconsistent with incomplete market models because borrowing constraints are not considered in the analysis, in particular, they are downward-biased.

²¹Examples of papers that use these values are Erosa and Ventura (2002), Campanale (2007), Castañeda et al. (2003) and Domeij and Heathcote (2004), among many others.

Table 1: Parameters set exogenously

Parameter	Value
σ	1.5
θ_1	0.59
μ_0	-0.0117
μ_1	0.0835
G/Y	0.2003
τ^k	0.397
π	0.02

where $(1 - \bar{l})$ are average hours worked, which we set to be 1/3 of disposable time. Notice that we are making \tilde{z}_t^i depend on *potential* gross labor income rather than on actual labor income. This simplifies greatly the analysis. If \tilde{z}_t^i depended on current labor income and/or wealth w_t , then the agent would take into account that by changing her labor/leisure and consumption/savings decisions, she would be affecting the cost of credit she faces. This interaction complicates the task of obtaining the allocations c_t, l_t, z_t and w_{t+1} from the optimality conditions of the household²². We leave this exercise for future research.

μ_0 and μ_1 are determined so that agents with the lowest productivity level possible have $\tilde{z}_t^i = 0$ while agents with the highest productivity level have $\tilde{z}_t^i = 0.2$ in the benchmark parameterization.

4.1.3 Government

We adopt a very standard parameterization for fiscal variables. Usually it is accepted that the ratio of government expenditure over GDP lies in the interval $[0.19, 0.22]$ ²³. We set $G/Y = 0.2$. To make our study comparable to other papers in the optimal fiscal policy literature, we introduce a tax on asset returns τ^k . Nevertheless, we assume it is fixed at a value of 0.397, which is in line with Domeij and Heathcote (2004) and Floden and Linde (2001). The reason to introduce this tax is that, otherwise, almost all public revenues would have to come from labor income taxation, causing the distortion on the labor/leisure decision to be very large. Finally, we set inflation to be 0.02 in the benchmark parameterization.

²²In the Appendix we describe in detail the numerical algorithm used to solve the model.

²³Some studies that use values in this range are Erosa and Ventura (2002), Campanale (2007) and Floden and Linde (2001).

Table 2: Parameters set endogenously

Parameter	Value	Target (data)	Model
θ_0	25	$\int(1-l)d\lambda = 0.33$	0.34
β	0.952	$K/Y=3$	2.99
γ_0	0.035	$\text{corr}(M,K)=0.16$	0.175
γ_1	0.2	$\frac{\int(1-z_t^i)c_t^i d\lambda}{\int c_t^i d\lambda} = 0.8$	0.79
ε^1	0.3547	$\text{Gini } K = 0.78$	0.77
ε^2	0.9428	Fraction of wealth of Q1+Q2 = 0.0335	0.08
α	0.36	$\omega L/Y = 0.64$	0.64
δ	0.083	$I/Y=0.25$	0.25
p	0.921	$\rho=0.92$	0.92
q	0.9886	$\sigma_\varepsilon = 0.21$	0.21

4.2 Internally calibrated parameters

4.2.1 Preferences and technology

We need to pin down θ_0 from the utility function (10), the discount factor β , the depreciation rate δ and the two parameters γ_0 and γ_1 from the credit technology (11). The targets we choose are the following: the average fraction of disposable time devoted to work should be $1/3$ ²⁴, the capital to output ratio should be $K/Y = 3$ ²⁵, the investment to output ratio should be $I/Y = 0.25$, the correlation between money and asset holdings should be $\text{corr}(m, a) = 0.16$ ²⁶ and the fraction of consumption expenditures made with cash should be $\int(1 - z_t^i)c_t^i d\lambda = 0.8$ ²⁷.

²⁴See Pijoan-Mas (2006) and Castañeda et al. (2003) for a similar choice.

²⁵Pijoan-Mas (2006), Campanale (2007), Conesa and Krueger (1999) and Castañeda et al. (2003) among others.

²⁶The correlation between money and asset holdings is computed using the 2004 Survey of Consumer Finances. Since cash holdings are not reported in the survey, we use the amount of money held in checking accounts as a proxy for money holdings. All the remaining sources of net worth are considered to be assets a_t .

²⁷Erosa and Ventura (2002) report that 80% of transactions are made with cash (M1). Since we do not have a measure of the number of transactions in our model, we use as a proxy the value of transactions. Similarly, Algan and Ragot (2006) report that $M1/C \simeq 0.78$ for the 1960-2000 period. We use a value of 0.8 which is in the middle of these two.

4.2.2 Labor productivity process

Our main aim is to define what the optimal mix of fiscal and monetary policy should be in the presence of idiosyncratic uninsurable risk. As we will see in the following sections, individuals with different levels of wealth and labor income are affected differently by different policies. Moreover, the optimal policy prescription depends crucially on the presence and the extent in which agents are exposed to the idiosyncratic shock. Consequently, the definition of the process for the labor productivity shock ε_t is critical to the analysis.

We follow the approach of Domeij and Heathcote (2004) and set two goals that our specification of the shock should accomplish. The first goal is that the persistence and variance of earnings shocks in the model are consistent with empirical estimates from panel data. The second is that in equilibrium the model yields a distribution of households across wealth that resembles in some aspects the distribution observed in the US.

We assume that the labor productivity process can display only three values, i.e., $E = \{\varepsilon^1, \varepsilon^2, \varepsilon^3\}$ where $\varepsilon^1 < \varepsilon^2 < \varepsilon^3$. The transition probability matrix corresponding to the shock adds 6 free parameters²⁸ which, added to the three productivity levels, sum up to 9 parameters that need to be pinned down.

In order to restrict the number of free parameters available for calibration, we assume the following: households cannot jump directly from the lowest productivity state to the highest one and viceversa, and they face equal probability of going from the medium productivity state to the low one as to the high one. These restrictions yield the following transition probability matrix:

$$\Pi = \begin{pmatrix} p & 1-p & 0 \\ \frac{1-q}{2} & q & \frac{1-q}{2} \\ 0 & 1-p & p \end{pmatrix} \quad (12)$$

Finally, we impose average productivity in the economy to be equal to 1. This leaves us with 4 free parameters to pin down.

As mentioned before, our first goal is to have a process that replicates the persistence and variance of earnings shocks present in panel data. As described in Floden (2001) and Pijoan-Mas (2006), wages (in logs) can be decomposed into two components. The first component, which we call η , is constant for a given individual and represents ability, education and all other elements that influence wages and can be depicted as fixed idiosyncratic characteristics of an agent. The second component, ν_t , is a stochastic individual component meant to capture idiosyncratic uncertainty in the earnings process. It basically reflects changes in the employment

²⁸Since the shock is a 3-state Markov chain, the transition probability matrix is a 3-by-3 matrix so it has 9 values to be determined. Nevertheless, the rows of the matrix have to add up to one, therefore the number of free parameters is 6.

Table 3: Comparison of steady states

	Benchmark	Optimal
R	1.0426	1
r	0.0367	0.0376
K/Y	2.99	2.976
ω	1.1873	1.1823
τ^l	0.2308	0.259
Welfare	-39.9484	-39.7701

status of each agent, job changes to positions that match better or worse the individual's ability, health shocks that affect productivity, etc. This last component corresponds to the notion of (log of) ε in our model.

Floden and Linde (2001) have estimated the following process for $\nu_{i,t}$:

$$\nu_t^i = \rho^\nu \nu_{t-1}^i + \zeta_t^i \text{ with } \zeta_t^i \sim N(0, \sigma_\zeta^2) \quad (13)$$

and found $\rho^\nu = 0.92$ and $\sigma_\zeta = 0.21$. We use these as the targets for the persistence and variance of our productivity shock²⁹.

Our second goal is to have realistic heterogeneity in terms of wealth in equilibrium. As a consequence, we set as targets the Gini coefficient of the asset (total wealth minus money holdings) distribution, which is approximately 0.78 according to 2004 SCF data, and the fraction of total wealth in the hands of the two poorest quintiles of population, which is about 3.35% using the same data. The last target is important because, a priori, inflation is likely to affect more poorer agents that consume a relatively much larger fraction of their income than richer ones and thus need to have most (if not all) of their wealth in the form of money.

5 Results

5.1 Optimal policy

In the previous sections we have discussed the role that the inflation tax has as a regressive tax on consumption and as an incentive for capital accumulation. The two effects affect asymmetrically different sectors of the population: while the former benefits richer, more productive agents, the latter increases welfare of the poor, unproductive ones.

²⁹Different studies suggest that ρ^ν should belong to the [0.88, 0.96] interval, while σ_ζ should be between 0.12 and 0.25. See Domeij and Heathcote (2004) for references.

Figure 2: Consumption

Having exposed all the mechanisms by which inflation affects the agents in our economy, it should be clear by now that the determination of the optimal policy mix is a question that does not have an immediate answer. There are a variety of effects operating simultaneously, affecting different agents in contradictory ways. The only way to provide an answer is to find the optimal policy numerically, once we have a reasonably calibrated model economy.

The main result of our analysis is the following:

RESULT 1. *For a utilitarian social welfare function, that is, one that assigns equal weight to all individuals, the Friedman rule is optimal and the government sets $R = 1$.*

The previous result suggests that, despite the fact that some agents win with relatively high levels of inflation (compared to the one that arises when $R = 1$), the efficiency motive associated to uniform taxation dominates and the optimal policy prescription is the Friedman rule. Therefore, the optimality of the Friedman rule is not only robust to the introduction of distortionary taxation, as explained by Chari et al. (1996), but also to considerations of heterogeneity and uninsurable idiosyncratic risk, as we have shown here.

5.2 Steady state comparison

Table 3 shows some statistics for the benchmark and the optimal policy steady states, respectively. From the table we see that in the optimal policy steady state the capital to labor ratio is

Figure 3: Leisure

smaller than in the benchmark economy, thus the lower real interest rate and higher wage rate. This is a direct consequence of the fact that inflation reinforces the motive for precautionary savings, as discussed in the previous section. The lower wage rate, lower savings and lower seigniorage revenue force the government to increase the tax rate on labor income in order to balance its budget. Therefore τ^l is higher in the optimal policy economy.

Figures 2 and 3 show consumption and leisure policy functions in the benchmark economy and in the optimal policy one, for the three levels of labor productivity. It can be observed that, for ε^1 and ε^2 and very low levels of wealth, consumption and leisure are higher in the benchmark economy, the reason being the higher labor income that poor households enjoy. When wealth increases the return on capital holdings starts being a relevant source of income for the household. Since the real interest rate is lower in the benchmark economy, consumption and leisure decrease. For high productivity households the picture looks different. These households are always enjoying a high level of consumption, and even for very little levels of wealth their labor income is sufficiently high to finance high consumption and savings. The decrease in uncertainty as a consequence of lower inflation levels present in the optimal policy economy diminishes the incentives to save for precautionary reasons. Therefore, very productive agents can afford to work less and enjoy higher levels of leisure, even if this means giving up some of their consumption.

Figure 4 plots the difference in savings between the benchmark economy and the optimal

Figure 4: Savings

policy one. It is immediate to see that, for agents with ε^1 and ε^2 , this difference is always negative, thus savings are higher in the optimal policy steady state. For agents with ε^3 , however, the contrary statement is true. Again, this result hinges on the fact that high productivity agents need to save less for self-insurance reasons when $R = 1$. Agents with low and medium productivity and very little wealth need to save more because hitting the constraint is more harmful in this case. Notice that this is the reason why the curve of the difference in savings first rises and then goes down. As wealth increases, their total income goes up and they are able to better self-insure by saving more.

Finally, figure 5 plots the use of the transaction technology in the benchmark economy. Notice that, unless the borrowing constraint is binding, when $R = 1$ an agent will not use the transaction technology because the opportunity cost of holding money is zero³⁰, so the z^i policy function is trivial in the optimal policy economy. If the borrowing constraint is binding, from inspection of equation (9) it is clear that an agent will use credit, even if $R = 1$, because doing so allows her to relax such constraint.

Figure 5 shows that the use of the transaction technology becomes more intensive with higher wealth holdings. The reason for this lies in the increasing returns to scale nature of the transaction technology we have adopted. Since higher wealth implies higher consumption for all

³⁰To be precise, agent i will be indifferent between using money or buying up to good z_i^i with credit, since both entail zero costs.

Figure 5: Transaction technology use

levels of labor productivity, the unitary cost of credit goes down as w^i goes up, so z^i goes up as well. Increasing returns to scale are also responsible for the three lines, corresponding to the three labor productivity levels, becoming closer together as wealth increases, since for high w^i the differences in consumption become smaller.

6 Welfare analysis

We proceed to compare aggregate welfare in the benchmark economy (with a level of inflation of 2% annually) and in the economy in which the optimal policy is implemented. In this section we are comparing welfare in two different steady states, and we are not saying anything as to what happens during the transition from one to the other if there is a reform on policy. Obviously, studying the transition is a very interesting exercise, specially if we want to determine the “optimal transition”, i.e., the transition such that no agent loses from the policy reform. As shown by Greulich and Marcet (2008) the optimal transition can imply a policy during the transition very different from the long-run policy prescription. We leave the analysis of the transition for future research.

We need to start with some definitions. The overall utilitarian welfare gain of the policy reform, ϖ_U is such that

$$\int E_t V(\{(1 + \varpi_U)c_s^B, l_s^B\}_{s=t}^\infty) d\lambda^B = \int E_t V(\{c_s^O, l_s^O\}_{s=t}^\infty) d\lambda^O$$

where the superscript B stands for the benchmark economy and O for the economy with the optimal policy. ϖ_U can be thought of as the percent permanent change in consumption that

Table 4: Welfare gains

$\varpi_U = 0.0051$	$\varpi_{lev} = 0.0044$	$\varpi_{unc} = 0.0015$	$\varpi_{ine} = -0.00005$
---------------------	-------------------------	-------------------------	---------------------------

agents in economy B should receive to be indifferent between living in economy B or in economy O .

Notice that the utilitarian social welfare (eq. (7)) can increase for three reasons. The first is when consumption or leisure increase for all agents. This is called the *level effect*. The second, called *inequality effect*, is when inequality is reduced, since $u(\cdot)$ (and therefore $V(\cdot)$) is concave. Finally, since agents are risk-averse, if uncertainty is reduced U increases. This is the *uncertainty effect*. Following Floden (2001), we can approximately decompose the utilitarian welfare gain into the welfare gains associated to the three effects mentioned before. In order to do this, define the certainty-equivalent consumption bundle \bar{c} as:

$$V(\{\bar{c}, l_s\}_{s=t}^{\infty}) = E_t V(\{c_s, l_s\}_{s=t}^{\infty})$$

Call $C = \int c d\lambda$, $Leis = \int l d\lambda$ and $\bar{C} = \int \bar{c} d\lambda$ average consumption, leisure and certainty-equivalent consumption, respectively. Then the cost of uncertainty p_{unc} can be defined as

$$V(\{(1 - p_{unc})C, Leis\}_{s=t}^{\infty}) = V(\{\bar{C}, Leis\}_{s=t}^{\infty})$$

This is the fraction of average consumption that an individual with average consumption and leisure would be willing to give up to avoid all the risk from labor productivity fluctuations. When uncertainty increases, \bar{C} decreases and, since C and $Leis$ remain unchanged, p_{unc} necessarily goes up.

Define the cost of inequality p_{ine} as

$$V(\{(1 - p_{ine})\bar{C}, Leis\}_{s=t}^{\infty}) = \int V(\{\bar{c}, l_s\}_{s=t}^{\infty}) d\lambda$$

If we redistribute consumption from a rich household to a poor one, \bar{C} and $Leis$ remain unchanged. However, the right-hand side of the previous expression increases, so p_{ine} has to go down. Finally, define leisure-compensated consumption in economy O , \hat{C}^O as

$$V(\{\hat{C}^O, Leis^B\}_{s=t}^{\infty}) = V(\{C^O, Leis^O\}_{s=t}^{\infty})$$

which is the average consumption level that would make life-time utility in economy O equal to the one in an economy with the average leisure of economy B .

We are now ready to define the welfare gains associated to each one of the effects described before:

- The welfare gain of increased levels, ϖ_{lev} is

$$\varpi_{lev} = \frac{\hat{C}^O}{C^B} - 1$$

- The welfare gain of reduced uncertainty is

$$\varpi_{unc} = \frac{1 - p_{unc}^O}{1 - p_{unc}^B} - 1$$

- The welfare gain of reduced inequality is

$$\varpi_{ine} = \frac{1 - p_{ine}^O}{1 - p_{ine}^B} - 1$$

Table 4 shows the welfare gains in our setup. As we can see, the aggregate utilitarian welfare gains are very small, only 0.51% of life-time consumption. The majority of these gains are due to the change in consumption levels (0.44% of consumption), and the remaining is because of the decrease in uncertainty that is associated with the optimal policy (0.15% of consumption). The welfare gains associated to the decrease in inequality are, actually, welfare costs, and are negligible.

We proceed now to compute the individual welfare gains from the change in policy. We perform the following exercise: we calculate the percentage permanent increase in consumption ϖ^i that we should give to a household i with (w^i, ε^i) to be indifferent between living in the benchmark economy B or living in the optimal policy economy O with the same level of labor productivity and wealth. Figure 6 shows ϖ^i as a function of w^i and ε^i . We can see from the graph that, although the aggregate welfare gain is small, individual welfare gains and losses can be very large, depending on an agent's productivity and wealth holdings. It is by aggregation that the individual effects cancel out, thus yielding a mild aggregate effect.

From inspection of figure 6 we can determine who are the net winners and net losers from the reform. Because of the smaller level of steady state capital in the optimal policy economy, the wage rate is lower and labor taxes are higher. Then, very poor agents always lose with the change in policy, irrespective of their labor productivity, the reason being that poor agents rely almost entirely on their labor income to pay for consumption goods, so changes in labor income matter substantially for them. The welfare loss amounts to about a 4% of permanent consumption for low productivity agents, while it is less than 1% for high productivity households. The difference in these effects relies on the fact that utility is concave and agents with high ε are income-rich, so they can afford higher levels of consumption and leisure. On the contrary, net winners from the policy change are middle-class and rich households, again irrespective of their labor productivity, who see their returns on capital increased because of the higher real interest

Figure 6: Individual welfare gains

rate. Again, because of the concavity of the utility function, rich and low-productivity agents are the ones that benefit the more. Their welfare gains can reach a maximum of around 4% of permanent consumption, while for high-productivity households the maximum is about 2.5%.

7 Conclusions

The determination of the optimal monetary policy prescription in the long run is a crucial issue for policy makers as well as for academics. Arguably, central banks set their inflation targets according to some criteria related to the maximization of social welfare. The natural question that arises is what the long-run optimal inflation target should be.

The standard literature in optimal monetary and fiscal policy, by focusing on representative agent environments, looks into this problem in a partial way and only considers issues of efficiency in distributing the distortions associated to taxation. In this paper we have relaxed the representative-agent assumption by allowing for heterogeneity and uninsurable idiosyncratic risk. This allows us to include in the analysis issues of redistribution of the tax burden and long-run effects of different tax schemes over capital and output that cannot be addressed in the traditional framework.

We make the standard modeling assumption that agents demand cash because it provides liquidity services. Moreover, we allow them to use an alternative costly transaction technology by which they economize on their money holdings. This transaction technology reconciles the model with some stylized facts reported in the literature regarding transaction patterns for different sectors of population. We are able to identify the effects of inflation as a regressive tax

on consumption and as a motive to increase savings for precautionary reasons.

We calibrate the model to the U.S. economy and find that the optimal policy prescription that arises from the exercise is the Friedman rule. This result provides robustness to what is a classical result in representative-agent models. A surprising implication of the analysis is that, despite the fact that inflation taxes relatively more consumption of poor agents, these agents actually win with inflation, while middle-class and rich agents lose. Therefore, this analysis challenges the conventional wisdom that inflation hurts the poor and benefits the rich.

The analysis presented here opens many avenues for future research. Probably one of the most natural extensions is the study of the transition between the steady state with the benchmark policy and the one in which the optimal policy is implemented. Studying the transition allows to perform a more accurate analysis of the welfare gains from the change in policy for different individuals. Moreover, studying the *optimal* transition, i.e., the transition taking into account that all agents should benefit from the reform, can lead to policy plans very different than what is optimal in the long run, as is shown in Greulich and Marcet (2008).

On a related note, Doepke and Schneider (2006) have driven attention to the fact that unexpected inflation can have large redistributive effects for individuals with different portfolio holdings. As a further step, we would like to introduce aggregate fluctuations to a framework similar in spirit to the one we consider here, but taking into account this heterogeneity of portfolio holdings among different individuals. This type of environment is suitable for studying optimal monetary and fiscal policy as stabilizing mechanisms for macroeconomic fluctuations, an issue that has not been addressed in this paper.

References

- Akyol, A.: 2004, Optimal monetary policy in an economy with incomplete markets and idiosyncratic risk, *Journal of Monetary Economics* **51**, 1245–1269.
- Albanesi, S.: 2005, Optimal and time consistent monetary and fiscal policy with heterogeneous agents, *mimeo*.
- Algan, Y. and Ragot, X.: 2006, On the non-neutrality of money with borrowing constraints, *mimeo*.
- Bhattacharya, J., Haslag, J. and Martin, A.: 2005, Heterogeneity, redistribution and the Friedman rule, *International Economic Review* **46**, 437–454.
- Bhattacharya, J., Haslag, J., Martin, A. and Singh, R.: 2007, Who is afraid of the Friedman rule?, *Economic Inquiry* **46**, 113–130.
- Campanale, C.: 2007, Increasing return to savings and wealth inequality, *Review of Economic Dynamics* **10**, 646–675.
- Castañeda, A., Diaz-Gimenez, J. and Rios-Rull, J.: 2003, Accounting for the U.S. Earnings and Wealth Inequality, *Journal of Political Economy* **111**(4), 818–857.
- Chari, V. V., Christiano, J. and Kehoe, P. J.: 1991, Optimal fiscal and monetary policy: some recent results, *Journal of Money, Credit and Banking* **23**, 519–539.
- Chari, V. V., Christiano, J. and Kehoe, P. J.: 1996, Optimality of the Friedman rule in economies with distorting taxes, *Journal of Monetary Economics* **37**, 203–223.
- Conesa, J. C. and Krueger, D.: 1999, Social security reform with heterogeneous agents, *Review of Economic Dynamics* **2**, 757–795.
- Correia, I. and Teles, P.: 1996, Is the Friedman rule optimal when money is an intermediate good?, *Journal of Monetary Economics* **38**, 223–244.
- da Costa, C. and Werning, I.: 2008, On the optimality of the Friedman rule with heterogeneous agents and non-linear income taxation, *Journal of Political Economy* **116**, 82–112.
- Doepke, M. and Schneider, M.: 2006, Inflation and the redistribution of nominal wealth, *Journal of Political Economy* **114**, 1069–1097.
- Domeij, D. and Flodén, M.: 2006, The labor supply elasticity and borrowing constraints: why estimates are biased, *Review of Economic Dynamics* **9**, 242–262.

- Domeij, D. and Heathcote, J.: 2004, On the distributional effects of reducing capital taxes, *International Economic Review* **45**, 523–554.
- Dotsey, M. and Ireland, P.: 1996, The welfare cost of inflation in general equilibrium, *Journal of Monetary Economics* **37**, 29–47.
- Erosa, A. and Ventura, G.: 2002, On inflation as a regressive consumption tax, *Journal of Monetary Economics* **49**, 761–795.
- Floden, M.: 2001, The effectiveness of government debt and transfers as insurance, *Journal of Monetary Economics* **48**, 81–108.
- Floden, M. and Linde, J.: 2001, Idiosyncratic Risk in the U.S. and Sweden: Is there a Role for Government Insurance?, *Review of Economic Dynamics* **4**, 406–437.
- Friedman, M.: 1969, The optimum quantity of money, *The optimum quantity of money and other essays* pp. 1–50.
- Greulich, K. and Marcet, A.: 2008, Pareto-improving optimal capital and labor taxes, *mimeo*.
- Heathcote, J., Storesletten, K. and Violante, G.: 2008a, Insurance and opportunities: A welfare analysis of labor market risk, *Journal of Monetary Economics* **55**, 501–525.
- Heathcote, J., Storesletten, K. and Violante, G.: 2008b, The macroeconomic implications of rising wage inequality in the United States, *Working paper*, NBER.
- Lagos, R. and Rocheteau, G.: 2005, Inflation, output and welfare, *International Economic Review* **46**, 495–522.
- Lucas, R. E. and Stokey, N. L.: 1983, Optimal fiscal and monetary policy in an economy without capital, *Journal of Monetary Economics* **12**, 55–93.
- Phelps, E.: 1973, Inflation in the theory of public finance, *Swedish Journal of Economics* **75**, 67–82.
- Pijoan-Mas, J.: 2006, Precautionary savings or working longer hours?, *Review of Economic Dynamics* **9**, 326–352.

A Appendix

A.1 Optimality conditions of the household

The problem of household i is

$$\max E_0 \sum_{t=0}^{\infty} \beta^t u(c_t^i, l_t^i) \quad (14)$$

subject to

$$p_t c_t^i + q_t \int_0^{z_t^i} \gamma^i(j) dj + W_{t+1}^i = R A_t^i + M_t^i + (1 - \tau^l) \omega p_t (1 - l_t^i) \varepsilon_t^i \quad (15)$$

$$p_t c_t^i (1 - z_t^i) \leq M_t^i \quad (16)$$

$$A_t^i \geq 0 \quad (17)$$

We define

$$\begin{aligned} a_t &= \frac{A_t}{p_{t-1}} \\ w_t &= \frac{W_t}{p_{t-1}} \\ m_t &= \frac{M_t}{p_{t-1}} \end{aligned}$$

We can rewrite equations (15) and (16) in real terms by dividing both sides of the equations by p_t :

$$c_t^i + q + \omega \int_0^{z_t^i} \gamma^i(j) dj + w_{t+1}^i = (1 + \tilde{r}) a_t^i + \frac{m_t^i}{1 + \Pi} + (1 - \tau^l) \omega (1 - l_t^i) \varepsilon_t^i \quad (18)$$

$$c_t^i (1 - z_t^i) \leq \frac{m_t^i}{1 + \Pi} = \frac{m_t^i R}{1 + \tilde{r}} \quad (19)$$

where we have used the fact that $\frac{p_t}{p_{t-1}} = 1 + \Pi$, $q_t = \omega p_t$ and $R = (1 + \tilde{r})(1 + \Pi)$. $\tilde{r} = r(1 - \tau^k)$ is the after-tax real return on capital.

Plugging equation (19) into (18), using $w_t^i = a_t^i + m_t^i$ and rearranging, we obtain

$$c_t^i (1 - (1 - z_t^i)(1 - R)) + \omega \int_0^{z_t^i} \gamma^i(j) dj + w_{t+1}^i = (1 + \tilde{r}) w_t^i + (1 - \tau^l) \omega (1 - l_t^i) \varepsilon_t^i \quad (20)$$

Similarly, equation (17) can be rewritten as

$$c_t^i(1 - z_t^i) \frac{R}{1 + \tilde{r}} - w_t^i \leq 0 \quad (21)$$

The problem of the household becomes maximizing (14) subject to constraints (20) and (21). The optimality conditions of the household are:

$$R = 1 + \frac{\omega\gamma(z_t^i)}{c_t^i} \quad (22)$$

$$\frac{u_{l,t}}{\varepsilon_t^i u_{c,t}} \Gamma_t = (1 - \tau^l)\omega \quad (23)$$

where

$$\Gamma_t = \left(1 + (1 - z_t^i) \frac{\omega\gamma(z_t^i)}{c_t^i} \right)$$

along with the Euler equation. Call μ_t the multiplier associated to constraint (21). If the borrowing constraint in period $t + 1$ is not binding, i.e., if $\mu_{t+1} = 0$, then the corresponding Euler equation is

$$\frac{u_{c,t}}{\Gamma_t} = \beta(1 + \tilde{r})E_t \frac{u_{c,t+1}}{\Gamma_{t+1}} \quad (24)$$

If, on the contrary, $\mu_{t+1} > 0$ the Euler equation becomes

$$\frac{u_{c,t}}{\Gamma_t} = \beta(1 + \tilde{r})E_t \frac{u_{c,t+1}}{\Gamma_{t+1}} \left(1 + \frac{1}{R} \left((1 - R) + \frac{\omega\gamma(z_{t+1}^i)}{c_{t+1}^i} \right) \right) \quad (25)$$

Given that equation (21) can be binding in t and/or in $t + 1$, 4 possible cases need to be considered when solving for the allocations of agent i :

- $\mu_t = 0$ and $\mu_{t+1} = 0$. The relevant equations for obtaining the allocations are (22), (23), (24) and the budget constraint (20).
- $\mu_t > 0$ and $\mu_{t+1} = 0$. The relevant equations for obtaining the allocations are (21), (23), (24) and the budget constraint (20).
- $\mu_t = 0$ and $\mu_{t+1} > 0$. The relevant equations for obtaining the allocations are (22), (23), (25) and the budget constraint (20).
- $\mu_t > 0$ and $\mu_{t+1} > 0$. The relevant equations for obtaining the allocations are (21), (23), (25) and the budget constraint (20).

A.2 Optimal fiscal and monetary policy with constant heterogeneity and no idiosyncratic risk

As in section 3.1, assume that an agent's productivity ε^i is constant $\forall t$, $\varepsilon^i \in E = [\varepsilon^1, \varepsilon^2]$ with $\varepsilon^1 < \varepsilon^2$ and there is an equal mass of each type of agent in the population³¹. Furthermore, assume for simplicity that the initial wealth holdings w_0^1 and w_0^2 are such that the economy is in steady state from $t = 0$ onwards.

Consider the case of a benevolent government (a Ramsey planner) that has to decide on the level of R and τ^l in our economy, in order to maximize a social welfare function given by the weighted sum of the utilities of both types of agents. The problem of the government can be written as:

$$\max_{\{c^i, l^i, z^i\}} \sum_{t=0}^{\infty} \beta^t u(c^i, l^i)$$

s.t.

$$\frac{1}{1-\beta} \left[u_{c^i} c^i + \frac{u_{c^i} c^i}{c^i + (1-z^i)\gamma(z^i)F_{L^g}} F_{L^g} \int_{\bar{z}^i}^{z^i} \gamma(j) dj - u_{l^i} (1-l^i) \right] = \frac{u_{c^i} c^i (1+\tilde{r})}{c^i + (1-z^i)\gamma(z^i)F_{L^g}} W_{-1}^i \quad i = 1, 2 \quad (26)$$

$$c^1 + c^2 + g + \delta K = F(K, L^g) \quad (27)$$

$$\int_{\bar{z}^1}^{z^1} \gamma(j) dj + \int_{\bar{z}^2}^{z^2} \gamma(j) dj = (1-l^1)\varepsilon^1 + (1-l^2)\varepsilon^2 - L^g \quad (28)$$

$$R = 1 + \frac{F_{L^g} \gamma(z^1)}{c^1} = 1 + \frac{F_{L^g} \gamma(z^2)}{c^2} \quad (29)$$

$$\frac{u_{l^1}}{\varepsilon^1 u_{c^1}} \left(1 + \frac{(1-z^1)F_{L^g} \gamma(z^1)}{c^1} \right) = \frac{u_{l^2}}{\varepsilon^2 u_{c^2}} \left(1 + \frac{(1-z^2)F_{L^g} \gamma(z^2)}{c^2} \right) \quad (30)$$

where $K = K^1 + K^2$, $L^g = L^{g,1} + L^{g,2}$ and ψ^i is the Pareto weight that the Ramsey planner assigns an agent of type i .

Equations (26), (27) and (28) correspond to the implementability constraints and resource constraints respectively. Notice that, when there is more than one type of agent in the economy, we need to consider one implementability constraint *for each type of household*³².

³¹The results of this section are robust to changes in the number of productivity states and in the composition of the population.

³²By virtue of Walras' Law, the budget constraint of the government is automatically satisfied.

Figure 7:

We assume that the tax system is *anonymous*, in the sense that the tax rates on labor τ^l and τ^k are not agent-specific. τ^k is exogenously imposed at a certain level for all individuals. Equation (30) imposes the condition that τ_l is equal for both types of households³³. Finally, given that the gross nominal interest rate R has to be the the same across individuals, equation (29) needs to be imposed.

In this case, the solution to the Ramsey problem varies with the determination of \tilde{z}^i and with the Pareto weight ψ^i that corresponds to each type of agent. Consider first the case in which $\tilde{z}^i = \tilde{z}$ for $i = 1, 2$. Our numerical exercise yields the following result:

RESULT 2. *In the heterogeneous-agent case with no idiosyncratic uncertainty and $\tilde{z}^i = \tilde{z}$ for $i = 1, 2$, the Friedman rule is optimal for all possible Pareto weights $\psi^1, \psi^2 \in [0, 1]$.*

The intuition for the result lies in the uniform taxation argument explained in the main text of the paper. When we allow $\tilde{z}^i = f(\varepsilon^i)$ the optimal policy prescription varies depending on the Pareto weights we consider, as summarized below:

RESULT 3. *Assume $\tilde{z}^i = f(\varepsilon^i)$ and $\frac{\partial \tilde{z}^i}{\varepsilon^i} > 0$, so that $\tilde{z}^1 < \tilde{z}^2$. Then, for a high enough Pareto weight on the more productive agents, ψ^2 , the planner deviates from the Friedman rule and sets $R > 1$.*

Figures 7 and 8 illustrate what is stated in results 2 and 3. These figures show the utilities of the two types of agents for $R \in [1, 1.05]$ for the cases in which $\tilde{z}^i = 0$ for $i = 1, 2$ and

³³To see why this equation implies that τ_l should be equal across individuals, notice that the optimality conditions of household i require that

$$\frac{u_{li}}{u_{ci}\varepsilon^i} \left(1 + \frac{(1 - z^i)\gamma(z^i)}{c^i} \right) = \omega(1 - \tau^l)$$

Figure 8:

$\tilde{z}^2 = 0.2 > \tilde{z}^1 = 0$, respectively. As we can see, in the first case the utility of both agents is decreasing in R , while in the second case the utility of agents with ε^2 is humped-shaped, increasing first with R , reaching a maximum and then decreasing for higher values of the nominal interest rate. The behavior of this utility causes the utility of the planner to be humped-shaped as well, provided that the weight on agents with ε^2 is high enough.

To grasp the intuition behind result 3, rewrite equation (23) as

$$\frac{u_l^i}{\varepsilon^i u_c^i} = \frac{\omega(1 - \tau^l)}{1 + (1 - z^i)(R - 1)} \quad (31)$$

Now assume that \tilde{z}^2 is large, that $\psi^2 \rightarrow 1$ and $\psi^1 \rightarrow 0$. In this case the denominator of the right hand side of expression (31) will be close to 1, even if R is very large. Since the planner cares only about agent 2, it wants to distort as little as possible the intratemporal decision of agent 2, therefore it wants the wedge between the marginal utility of leisure and that of consumption to be as close to one as possible. From expression (31) it is clear that the way to achieve this is to set τ^l as low as possible while relying heavily on seigniorage revenues to finance its budget. Since agents with ε^2 can shelter from inflation by recurring to credit, this is clearly optimal.

A.3 Computational algorithm

As explained in the main text, the model depicted here cannot be solved analytically. Consequently, we need to use numerical tools in order to obtain a solution.

The problem of the household includes a borrowing constraint, which is occasionally binding constraint. This constraint translates into strong non-linearities in the policy functions of the

household. Since we are interested in computing welfare gains and loses at the individual level, we should compute these policy functions accurately. Thus, we choose not to use perturbation techniques but, instead, implement a collocation algorithm to obtain such functions.

The solution algorithm is as follows:

1. Define a grid for w_t^i .
2. Set an initial guess for the real interest rate r^g and an initial guess for the vector of parameters Ω^g that we want to calibrate.
3. For a given r^g and Ω^g , approximate the individual policy function:

$$w_{t+1}^i = f(w_t^i, \varepsilon_t^i) \simeq \tilde{f}(w_t^i, \varepsilon_t^i, \Lambda)$$

We approximate the policy function with linear splines and discretize the state space using 80 grid points. We solve for Λ by collocation, which implies that we need to find Λ such that the Euler equation (24)-(25) is satisfied at every (w_t^i, ε_t^i) in our grid. Therefore, we obtain a system of non-linear equations that we solve through a successive approximations strategy.

4. Simulate the economy for 5000 ex-ante identical individuals, for 700 periods in order to obtain an invariant distribution.
5. Obtain aggregate capital by integrating asset holdings across individuals:

$$K = \int a d\lambda$$

Compute r^f .

6. If $r^f = r^g$ stop, otherwise update r^g and iterate on 3-5 until convergence.
7. Compute the desired moments from the model, if they coincide with the targets from the data, stop, otherwise update Ω^g and iterate on 3-6 until convergence.
8. Repeat 2-7 for a fine grid of R , compute welfare and choose R such that social welfare is maximized.