

THE WELFARE COST OF INFLATION RISK UNDER IMPERFECT INSURANCE

OLIVIER ALLAIS^a, YANN ALGAN^b, EDOUARD CHALLE^c AND XAVIER RAGOT^d

What are the costs of inflation fluctuations and who bears those costs? In this paper, we investigate this question by means of a quantitative incomplete-market, heterogeneous-agent model wherein households hold real and nominal assets and are subject to both idiosyncratic labor income shocks and aggregate inflation risk. Inflation risk is found to generate significant welfare losses for most households, i.e., between 1 and 1.5 percent of permanent consumption. The loss is small or even negative for households at the very top of the productivity and/or wealth distribution. A key feature of our analysis is a nonhomothetic specification for households' preferences towards money and consumption goods. Unlike traditional specifications, ours allows the model to reproduce the broad features of the distribution of monetary assets (in addition to being consistent with the joint distribution of nonmonetary assets and consumption).

JEL Codes: E21, E32, E41.

Keywords: Money-in-the-Utility, Incomplete Markets, Inflation Risk, Welfare.

1. INTRODUCTION

What are the costs of inflation fluctuations and who bears those costs? In this paper, we investigate this question by means of a quantitative incomplete-market, heterogeneous-agent model wherein households hold real and nominal assets and are subject to both idiosyncratic labor income shocks and aggregate inflation risk. To be more specific, the model has the following three main features:

First, households solve an optimal portfolio choice between money and capital claims, subject to an occasionally binding borrowing constraint. Both assets can coexist in household portfolios because the lower real return paid by money balances is compensated by the liquidity services that it provides – as captured by the fact that real money balances affect current utility, in the money-in-the-utility, or “MIU” tradition. Pioneered by Patinkin (1965) and operationalized in dynamic general equilibrium by Sidrauski (1967), MIU models provide a convenient and flexible framework for generating a demand for money and studying a variety of monetary policy issues. In the present paper we exploit this flexibility to generate a full cross-sectional distribution of money holdings that can be matched with the data.¹ The coexistence of monetary and nonmonetary assets in our framework is in contrast with existing incomplete-market models with aggregate shocks, which either abstract from portfolio choice altogether (by considering a single asset or a menu of perfectly substitutable assets), or focus on the choice between two real assets (so that inflation

We are grateful to conference participants at SED and ESWC, as well as to seminar participants at various places, for their comments.

^aINRA. olivier.allais@inra.fr

^bSciences Po and CEPR. yann.algan@sciences-po.fr

^cCREST (CNRS) and Ecole Polytechnique. edouard.challe@gmail.com

^dSciences Po, OFCE and CNRS. xavier.ragot@gmail.com

¹See Walsh (2017) for a survey of the flexible-price MIU model under complete markets, and den Haan, Rendahl, and Riegler (2018) for a recent use within an incomplete-market framework.

risk does not affect the relative returns of the assets).²

Second, the model is designed to jointly match three cross-sectional distributions: consumption, money holdings and nonmonetary assets. Our investigation being primarily quantitative, matching those three distributions at once is a prerequisite for a meaningful analysis of the aggregate welfare impact of inflation risk: both monetary and nonmonetary assets affect households' ability to smooth consumption, while both consumption and monetary assets enter households' utility. Matching the distribution of money requires departing from the functional form commonly used to parameterize money demand in monetary models with a representative agent.³ As discussed in Ragot (2014), the empirical cross-sectional distribution of money holdings in the US economy is close to that of financial assets and hence very different from that of consumption.⁴ This property cannot follow from the usual MIU specification, which features a constant elasticity of substitution between money and goods and hence strict proportionality between the demands for money and consumption goods. If this were the case then inequalities in money holdings would simply mirror inequalities in consumption, which is strongly against the data.⁵ We overcome this limitation of the standard MIU specification by introducing a more general utility function that nests the constant-elasticity case but also accommodates a non-constant elasticity of substitution, thereby allowing individual money holdings to vary more than proportionally with individual consumption levels. This utility function allows us to reproduce the broad features of the distribution of monetary wealth in the US economy, whilst at the same time being consistent with the observation that, at the individual level, higher wealth is associated with greater absolute money holdings but lower money holdings relative to total wealth.

Third, the model incorporates as part of the solution households' rational portfolio response to the inflation risk that they face. As far as we are aware, existing work on the effect of inflation on welfare in heterogenous-agent models has focused on the welfare impact of either mean inflation or unexpected shocks to inflation, leaving aside (by construction) households' optimal portfolio response to their expectation of future inflation shocks.⁶ This dimension is crucial in the present study for at least two reasons. First, households are likely to respond *ex ante* to the inflation risk they are facing, i.e., we expect different levels of inflation risk to generate different mean holdings of nominal assets;

²Krusell and Smith (1998) and Heathcote (2005) consider environments where assets are perfect substitutes, so that there is no portfolio choices. Krusell and Smith (1997) consider the portfolio choice between a real riskless bond and capital claims. Krusell, Mukoyama, Şahin, and Smith Jr (2009) and Challe, Le Grand, and Ragot (2013) analyze the portfolio choice between real bonds (differing by their maturity). In Gornemann, Kuester, and Nakajima (2016), a mutual funds hold all assets and sell identical, perfectly substitutable claims to the households.

³See, e.g., Chari, Christiano, and Kehoe (1996); Chari, Kehoe, and McGrattan (2000).

⁴For example, Ragot (2014) finds that in the US Gini coefficient for individual consumption levels is 0.54, while that for money balances is 0.85. In Italy the corresponding figures are 0.30 and 0.68, respectively.

⁵Basic Cash-In-Advance ("CIA") economies, which by construction also feature a constant elasticity of substitution between money and goods, suffer from exactly the same limitation.

⁶İmrohoroğlu (1992) and Erosa and Ventura (2002) study the impact of mean inflation on welfare in incomplete-market environments. Doepke and Schneider (2006) attempt to evaluate the redistributive and welfare effects of a moderate but persistent inflation episodes in the US economy. They find these effects to be large, due to the significant degree of heterogeneity in nominal asset positions across US households and the large fraction of foreign holders of US assets. In the same spirit, Sterk and Tenreyro (2018) use an overlapping-generation model to evaluate the impact of open market operations on durables consumption.

and second, allowing significant redistributive effects of inflation but ignoring the ex ante portfolio response to inflation risk is likely to overestimate the welfare cost of inflation shocks.

The model that we consider features both aggregate (inflation) and idiosyncratic (labor productivity) shocks and is calibrated accordingly. Regarding the time-series dimension, we feed the model with an exogenous process for the money growth rate so that the equilibrium response of inflation match the volatility properties of actual inflation over the post-Volcker, pre-Great Recession period. Treating equilibrium inflation (rather than the money supply) as the forcing term in our analysis ensures that the inflation risk that households face in the model is consistent with the historical inflation risk. Concerning the cross-sectional dimension, we calibrate the wage income process as well as households' preferences towards money and consumption goods so as to match the joint cross-sectional distributions of monetary and nonmonetary assets. Our analysis can thus be interpreted as way to discipline preferences towards money and goods on the basis of the statistical moments of cross-sectional data on monetary positions.

To evaluate the cost of inflation fluctuations and the way it is distributed across households, we compute the ex ante welfare of an agent conditional on a particular idiosyncratic state and portfolio, in economies with and without inflation risk. To summarize, we find significant welfare costs of inflation fluctuations for all households except for those at the top end of the productivity scale and/or wealth distribution. To be more specific, we find the welfare gains from eliminating inflation risk to be equivalent to a 1 percent increase in permanent consumption for the average household, and even greater (about 1.5%) for a households at the lower end of the distribution of wealth and productivity. Such welfare effects are large: they are one order of magnitude larger than the cost of business cycles computed by Lucas Jr (1987, 2003) in a representative-agent context and about five times the welfare cost of the business cycle as computed by Imrohoroğlu (1989) within a (nonmonetary) incomplete-market, heterogenous-agent model.⁷ The magnitude of our welfare gains are not, however, unreasonable in the context of incomplete-insurance economies. For example, they are of comparable magnitude as the cost of the business cycle in economies wherein the extent of uninsured idiosyncratic risk is systematically related to the size of aggregate productivity shocks (a channel that is absent from our model).⁸

The next section presents the model. Section 3 discusses our parameterization. Section 4 presents our main results on the welfare costs of inflation fluctuations, and Section 5 concludes.

2. THE MODEL

The model is a MIU, heterogenous-agent model augmented with i) aggregate shocks to inflation (driven by underlying changes in money growth), and ii) a nonhomothetic utility functional designed to accommodate significant dispersion in holdings of monetary assets.

⁷Lucas' calculations only rely on the assumption of complete market and a specification for the representative agent's preferences. While maintaining complete markets and suitably choosing preferences can generate larger welfare costs of the business cycle, Otrok (2001) shows that this argument does not survive in general equilibrium wherein households optimally choose (and effectively smooth) consumption.

⁸See, e.g., Storesletten, Telmer, and Yaron (2001) and Krusell et al. (2009).

2.1. Preferences

Households are infinitely-lived and in constant mass equal to 1. They share identical and additively time-separable preferences over sequences of consumption, $c \equiv \{c_t\}_{t=0}^{\infty}$, and real money balances, $m \equiv \{m_t\}_{t=0}^{\infty}$. Thus, they maximize

$$(1) \quad U(c, m) = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, m_t),$$

where $\beta \in (0, 1)$ is the subjective discount factor, \mathbb{E}_t denotes expectations conditional on the information set at date t , and u is the instant utility function. This function is assumed to take the following parametric form:

$$(2) \quad u(c_t, m_t) = \frac{1}{1-\lambda} \left(\omega c_t^{1-\rho} + (1-\omega) m_t^{\theta(1-\rho)} \right)^{\frac{1-\lambda}{1-\rho}}, \rho, \theta, \lambda > 0.$$

When $\theta = 1$, equation (2) becomes a standard homothetic utility function similar to that used by Chari et al. (1996, 2000) and Algan and Ragot (2010), among others. In this case, the interest-rate elasticity of real money demand is $1/\rho$, while the intertemporal elasticity of substitution between consumption and real balances is $(1-\rho)/(\lambda-\rho)$. As discussed above, an important limitation of the homothetic specification is that it implies a strict proportionality between individual real money holdings and individual consumption levels (for any given values of the nominal interest rate and the inflation rate). This makes it impossible to reproduce the highly unequal distribution of money holdings that is observed in US data. Our baseline calibration will thus have $\theta \neq 1$.

2.2. Idiosyncratic uncertainty

In every period, households are subject to idiosyncratic labor income shocks. More specifically, every household is endowed with \bar{n} labor units at every point in time (supplied inelastically since leisure is not valued), but individual labor productivity e_t can take three different values: $e_t \in E$, $E = \{e^l, e^m, e^h\}$ with $e^l < e^m < e^h$, and where e^h stands for ‘high productivity’, e^m for ‘medium productivity’, and e^l for ‘low productivity’. e_t evolves according to a first-order Markov chain with the 3×3 transition matrix F . We denote by p^* the vector of stationary ergodic probabilities and normalize productivity levels so that the mean of the invariant distribution is one, i.e., $\sum p_i^* e_i = 1$. Given a population of measure one, we can interpret p^* as describing the distribution of the population across productivity states. It follows that, even though individuals transit across productivity states, the total number of effective labor units in the economy at every point in time is constant and equal to $\sum p_i^* e_i \bar{n} = \bar{n}$.

2.3. Production

Markets are competitive. In every period t , the representative firm uses aggregate capital $K_t \in \mathbb{R}_+$ and households’ labor to produce $Y_t \in \mathbb{R}_+$ units of a single good with the aggregate technology

$$Y_t = f(K_t, \bar{n}) = K_t^\alpha \bar{n}^{1-\alpha}.$$

Capital depreciates at the constant rate $\delta \in (0, 1)$ and accumulates according to the law

of motion:

$$(3) \quad K_{t+1} = I_t + (1 - \delta)K_t,$$

where I_t denotes aggregate investment. Perfect competition in the markets for the representative firm's inputs implies that the real interest rate, r_t , and the real wage, w_t , can be written as:

$$r_t = \alpha K_t^{\alpha-1} \bar{n}^{1-\alpha} - \delta, \quad w_t = (1 - \alpha) K_t^\alpha \bar{n}^{-\alpha}.$$

2.4. The household's problem

We assume that markets are incomplete, so that households cannot write insurance contracts contingent on their labor income. Moreover, they face borrowing constraints and are thus prevented from using private loans to fully smooth out individual income fluctuations. Each household i maximizes its expected lifetime utility (1) subject to the following constraints:

$$(4) \quad c_t^i + k_{t+1}^i + m_t^i = a_t^i + (1 - \tau_t^l) w_t e_t^i \bar{n},$$

$$(5) \quad k_{t+1}^i \geq 0, \quad m_t^i \geq 0, \quad \text{and} \quad c_t^i \geq 0,$$

where k_{t+1}^i and m_t^i denote the claims to the capital stock and the real money balances held by household i at the end of date t , and where

$$a_t^i = (1 + (1 - \tau_t^c) r_t) k_t^i + \frac{m_{t-1}^i}{1 + \pi_t}$$

is the household's financial wealth at the beginning of date t . In equations (4) and (5), τ_t^l and τ_t^c are proportional tax rates on labor and capital, respectively. In equation (5), the presence of the borrowing constraint is reflected in the fact that capital and money holdings must be nonnegative at all times, while no other assets (i.e., private bonds) can be issued by the households.

From the households' objective and constraints, we find that their optimal asset demands, m_t^i and k_{t+1}^i , must satisfy the following first-order conditions:

- Money:

$$(6) \quad u_c(c_t^i, m_t^i) - u_m(c_t^i, m_t^i) = \beta E_t \left[\frac{u_c(c_{t+1}^i, m_{t+1}^i)}{1 + \pi_{t+1}} \right].$$

- Capital:

$$(7) \text{ Either } u_c(c_t^i, m_t^i) = \beta E_t [(1 + (1 - \tau_t^c) r_{t+1}) u_c(c_{t+1}^i, m_{t+1}^i)] \text{ and } k_{t+1}^i > 0, \\ \text{ or } u_c(c_t^i, m_t^i) > \beta E_t [(1 + (1 - \tau_t^c) r_{t+1}) u_c(c_{t+1}^i, m_{t+1}^i)] \text{ and } k_{t+1}^i = 0.$$

The instant utility function (2) implies that $u_m(c_t^i, 0) = \infty$, so the demand for real balances is always interior. In contrast, the demand for capital may be corner (i.e., $k_{t+1}^i = 0$), in which case the household would like to raise current consumption by borrowing against future income (but is prevented from doing so by a binding borrowing constraint). The solution to the households' problem consists of sequences of policy functions $m_t(a, e)$, $k_t(a, e)$ and $c_t(a, e)$, $(a, e) \in \mathbb{R}_+ \times \{e_1, e_2, e_3\}$, where a and e denote

individual beginning-of-period asset wealth and productivity, respectively.

To better understand the implications of our assumed period utility function with non-constant elasticity of substitution (i.e., (2)), consider the optimal trade-off between consumption and real money holdings by an unconstrained household (so that $k_{t+1}^i > 0$ in (7)) and abstract from aggregate shocks momentarily. From (6)–(7) and the functional form (2), we find the relation between money holdings and consumption to be:

$$(8) \quad m_t^i = A((1 - \tau_t^c) r_{t+1}, \pi_{t+1}) (c_t^i)^{\frac{\rho}{1-\theta(1-\rho)}},$$

where $A((1 - \tau_t^c) r_{t+1}, \pi_{t+1})$ is a coefficient whose value depends on the returns on the two assets and the deep parameters of the utility function. In the constant-elasticity case (i.e., $\theta = 1$), we have $m_t^i = A((1 - \tau_t^c) r_{t+1}, \pi_{t+1}) c_t^i$: real money demand is exactly proportional to consumption, so that cross-sectional inequalities in these two variables exactly mirror each other. For individual money holdings to increase more than proportionally following an increase in individual consumption, so that money be more unequally distributed than consumption (as is observed in the data), one needs $\rho / (1 - \theta(1 - \rho)) > 1$ (whether θ must lie above or below 1 for this inequality to hold depends on the value of ρ .)

2.5. Market clearing

Define $\mu_t : \mathbb{R}_+ \times \{e^h, e^m, e^l\} \rightarrow \mathbb{R}_+$ as the joint cross-sectional distribution of wealth and individual productivity at the beginning of period t . The market-clearing conditions for real balances and capital claims are given by:

$$(9) \quad \int \int m_t^i(a_t, e_t) d\mu(a_t, e_t) = \Omega_t^s,$$

and

$$(10) \quad \int \int k_t(a_t, e_t) d\mu(a_t, e_t) = K_{t+1},$$

where Ω_t^s is the total quantity of real balances at date t . By Walras law, the goods market clears when both the money and capital markets clear.

2.6. Fiscal and monetary policy

Fiscal and monetary policies interact here because monetary policy generates senioriage revenues that determine government income jointly with tax collections. We specify both as follows.

First, we assume that inflation follows a two-state Markov chain: the inflation rate transits between π^L and $\pi^H > \pi^L$ according to the transition matrix T .⁹ We treat this inflation process as a forcing term and infer from the equilibrium the underlying process for money growth that has generated the corresponding inflation rates (that is, money growth at every

⁹Just as in Krusell and Smith (1998), households in our model have one forecasting rule for the stock of capital per value of the exogenous aggregate state (the inflation rate here), and each rule solves a fixed point problem between expectations and outcomes. In this regard the assumption that the inflation process has only two states eases the computation of the equilibrium, but it could be relaxed somehow.

point in time can be reverse-engineered once the equilibrium has been solved under the assumed inflation process). The two inflation levels as well as the transition probabilities between the two can then be calibrated according to the actual inflation process estimated on historical data. There are at least two reasons for adopting this approach rather than doing the opposite, i.e., calibrating a two-state Markov chain for money growth and let inflation adjust endogenously. First, directly calibrating inflation ensures that the households in the model face exactly the same amount of inflation risk as the historical inflation risk. With an exogenous process for money growth the model generates excess inflation volatility in equilibrium, which would overestimate the welfare cost of inflation volatility.¹⁰ Second, treating inflation as exogenous is computationally faster, because the return on money holdings (the inverse of the inflation rate) is no longer an endogenous variable – only the return on capital claims is.

For a given inflation process, the growth rate of the quantity of money γ_t is determined as follows. First, let Δ_t denote the nominal quantity of newly issued money at date t (relative to the stock of money at the end of date $t - 1$) and by $\Pi_t = P_t/P_{t-1} = 1 + \pi_t \in \{\Pi^L, \Pi^H\}$ the gross inflation rate between date $t - 1$ and date t . In real term, the quantity of newly issued money can be expressed as:

$$(11) \quad \Delta_t/P_t = \gamma_t P_{t-1} \Omega_{t-1}^s / P_t = \gamma_t \Omega_{t-1}^s / \Pi_t,$$

and the dynamics of total real money balances by:

$$(12) \quad \Omega_t^s = (1 + \gamma_t) \Omega_{t-1}^s / \Pi_t.$$

Second, given Ω_{t-1}^s and current gross inflation $\Pi_t (\in \{\Pi^L, \Pi^H\})$ we determine γ_t so that Ω_t^s clears the money market (i.e., so that (9) holds), given aggregate money demand (which is a function of the aggregate state, including the inflation factor Π_t).

It is assumed that the newly issued money is given to the government as part of its resources. This assumption captures the fact that in practice money is created by open market operations which, once the budgets of the government and the Central Bank are consolidated, are effectively transfers to the government (see, e.g., chapter 5 of Walsh (2017)). Our specification allows us to incorporate spending and taxes without introducing public debt as a third asset. Under this monetary/fiscal arrangement the budget of the government is balanced in every period and such that:

$$G_t = \gamma_t \Omega_{t-1}^s / \Pi_t + \tau_t^l w_t \bar{n} + \tau_t^c r_t K_t,$$

We assume that $G_t = G > 0$ is constant over time and then let the tax rates on capital and labor adjust endogenously to satisfy the government budget constraint in every period. The constancy of G ensures that there are no government spending-related aggregate wealth effects that could artificially amplify the welfare cost of inflation fluctuations. For simplicity we assume that the ratio of the two taxes $\psi = \tau_t^c / \tau_t^l$ is constant over time, so

¹⁰While in principle the money growth rates in each state could be adjusted to produce the same variance of inflation as in the data, this would imply that in at least one of the states money growth would no longer be aligned with inflation.

that the taxes on labor and capital are given by, respectively:

$$\tau_t^l = \frac{G - \gamma_t \Omega_{t-1}^s / \Pi_t}{w_t L + \psi r_t K_t}, \tau_t^c = \psi \tau_t^l.$$

2.7. Recursive problem and equilibrium

We are considering a recursive equilibrium in which the aggregate state, which includes the aggregate stock of capital, changes over time. Note that the way we have specified monetary policy implies that households need not keep track of the aggregate money supply, because inflation (the inverse of the return on money) enters as an exogenous forcing term here (with the rate of money growth adjusting to make a particular inflation path happen). It follows that the recursive problem of a household can be written as:

$$(13) \quad v(a_t, e_t; \pi_t, K_t) = \max_{m_t, c_t, k_{t+1}} u(c_t, m_t) + \beta E_t [v(a_{t+1}, e_{t+1}; \pi_{t+1}, K_{t+1}) | e_t, \pi_t],$$

subject to (4)–(5). Following Krusell and Smith (1998) and much of the subsequent literature, we posit that households are able to successfully forecast the dynamics of the endogenous aggregate state K_t by means of (log-)linear laws of motion (one per value of the exogenous aggregate state π_t) involving K_t and Ω_t^s only:

$$(14) \quad \ln(K_{t+1}) = b_1(\pi_t) + b_2(\pi_t) \ln(K_t) + b_3(\pi_t) \ln(\Omega_t^s).$$

The solution to (13) produces individual decision rules for consumption as well as holdings of real balances and claims to the capital stock, which we denote by $g_c(a_t, e_t; \pi_t, K_t)$, $g_m(a_t, e_t; \pi_t, K_t)$ and $g_k(a_t, e_t; \pi_t, K_t)$, respectively. The law of motion of the distribution of beginning-of period total wealth μ_t is denoted by H . For a given set of individual policy rules, this law of motion can be written as

$$\mu_t = H(\mu_{t-1}, \pi_t),$$

i.e., the cross-sectional distribution of wealth μ_t depends on its value in the previous period μ_{t-1} as well as the current realization of the exogenous aggregate state π_t .

Definition of the recursive equilibrium. A recursive equilibrium is a law of motion H , a set of optimal individual policy functions and value function $\{g_c, g_m, g_k, v\}$, a set of price functions $\{\pi, r, w\}$, and a law of motion for K such that:

1. given $\{\pi, r, w\}$, the law of motion for K_t and the transition matrices F and T , the policy functions $\{g_c, g_m, g_k\}$ solve the household's problem;
2. the money and capital markets clear;
3. the law of motion H is generated by the optimal decisions $\{g_c, g_m, g_k\}$, and the transition matrices F and T .

We solve for the recursive equilibrium using the same approach as in Krusell and Smith (1997, 1998). However, rather than using Monte Carlo simulations to generate an updated cross-sectional distribution, we use the grid-based simulation procedure proposed by Young (2010), which keeps track of the mass of households at a fine grid of wealth levels. This allows us to get rid of the cross-sectional sampling variations in the Monte Carlo simulation procedure. A detailed description of the algorithm is provided in the Appendix.

2.8. Discussion of modeling choices

It may be useful at this stage to discuss our choice of the MIU specification and compare it with alternative ways of generating a demand for money. A popular microfoundation for the demand for money, based on its transaction role, is the Cash-In-Advance (CIA) constraint initially introduced by Clower (1967). The cross-sectional implication of the baseline CIA model is that, by construction, individual money holdings and consumption levels are proportional to each other, which as we show below is strongly against the data (since money is much more unequally distributed than consumption). Ragot (2014) shows that a realistic distribution of money can be obtained with a participation constraint in the Baumol-Tobin tradition. Unfortunately, this approach cannot be handled in a model with aggregate shocks while maintaining computational accuracy.¹¹ The MIU framework with non-homothetic utility function offers a convenient alternative that makes it possible to solve the model with aggregate shocks and to match the joint distribution of consumption, money and financial assets.¹² Finally, time-series investigations of the money demand function also call for a non-homothetic specification when money is introduced as an argument in the utility function (Aruoba and Schorfheide, 2011).

3. PARAMETERIZATION

TABLE I
PARAMETER VALUES

<i>Preferences</i>	λ	1.00	<i>Production</i>	α	0.360
	β	0.99		δ	0.025
	ρ	20.00		\bar{n}	0.300
	ω	0.06			
	θ	0.21			
<i>Productivity</i>	e^l	3.940			
	e^m	0.849			
	e^h	0.213			
	$\pi_{l,l}$	0.975			
	$\pi_{m,m}$	0.992			

Table I reports the parameters of the model. The time period is a quarter. Following Chari, Kehoe, and McGrattan (2000), our benchmark value for the utility parameter λ is 1. The capital share is set to $\alpha = 0.36$, and the depreciation rate to 0.025. Finally, labor supply is constant for all households and set to 0.3 as in Heathcote (2005). The individual productivity states and the transition probabilities across states are calibrated as follows. As surveyed by Domeij and Heathcote (2004), various authors have estimated an AR(1) process for log labor earnings using PSID data, and found that the serial autocorrelation coefficient is in the range of 0.88 to 0.96, whereas the standard deviation of the innovation

¹¹The portfolio choice problem a la Baumol-Tobin requires solving binary participation decisions and hence computing and comparing intertemporal welfare levels at every point in time. This procedure is computationally very difficult to handle for economies with aggregate shocks.

¹²Erosa and Ventura (2002) introduce a more general CIA constraint with increasing returns to scale in the transaction technology and a limited participation to financial markets. Again, this framework is very difficult to solve with aggregate shocks.

term is in the range 0.12 to 0.25. Such a process can be parsimoniously represented by a three-state Markov chain with a transition matrix F , with the following restriction: 1) there is a zero probability to transit between extreme states (i.e., $F_{h,l} = F_{l,h} = 0$) and 2) there is an equal probability to reach any of the extreme states when in the intermediate state (i.e., $F_{m,h} = F_{m,l}$).¹³ The transition matrix is then fully identified once $F_{l,l}$, $F_{m,m}$ and $F_{h,h}$ are set. Following Castaneda, Diaz-Gimenez, and Rios-Rull (2003), the ratios of productivity levels are set to $e^h/e^m = 4.64$ and $e^m/e^l = 3.99$. The level of e^l , e^m , e^h are such that the average labor productivity is 1: $\sum_{k=l,m,h} p_k^* e^k = 1$. The transition probabilities are $F_{l,l} = F_{h,h} = 0.9750$, $F_{m,m} = 0.9925$. This process yields an autocorrelation of the real wage equal to 0.91 and a standard deviation of the innovation term equal to 0.22 at annual frequency, in line with the data.

Table I also provides the preference parameters of the utility function. There are five such parameters: λ , β , ρ , ω and θ . Following Krusell and Smith (1998) and Domeij and Heathcote (2004), our benchmark value for the utility parameter λ is 1. This value allows reproducing a realistic distribution of financial wealth. The discount factor β is set to $\beta = 0.99$ (as in, e.g., Krusell and Smith (1998)). The three remaining parameters ρ , ω and θ are set to match three moments. The first moment is the ratio of the quantity of money to GDP. The money aggregate that we target is deliberately narrow, in order to avoid overestimating the impact of inflation shocks: it is defined as the checking accounts in the Survey of Consumer Finances (SCF). The quarterly value of money over GDP over the 1982-2005 period is 0.32. The second moment is the Gini of money holdings, which is as high as 0.85 in the data (again using the SCF). Broader monetary aggregates have similar dispersion (Ragot, 2014). Third, in addition to the Gini coefficient, we target the share of money holdings held by the top 20% of the money distribution. This share is 88.20% in the data, and 87.69% according to the model. As discussed further below, a key advantage of our specification and associated calibration is that it also generates realistic consumption inequalities: the model-generated distribution of consumption has a Gini of 0.30, while that in the data is 0.27 for nondurables and services consumption Krueger and Perri (2006).

We model the dynamics of monetary conditions between 1982Q1 and 2005Q4 as a two-state, first-order Markov chain.¹⁴ More specifically, we estimate this chain using CPI inflation and extract the inflation levels that prevail in the ‘high-inflation’ versus ‘low-inflation’ regimes, as well as the probabilities to transit between those regimes. This gives the quarterly value $\pi^L = 0.64\%$, $\pi^H = 1.17\%$, and probabilities to stay in the same state $\Pr(\pi^L|\pi^L) = 0.944$ and $\Pr(\pi^H|\pi^H) = 0.889$.

Finally, we specify the fiscal policy parameters G and ψ as follows. According to Domeij and Heathcote (2004), for the period 1990-1996 the capital income tax rate averaged 39.7 percent, while the labor income tax rate averaged 26.9 percent. We thus set $\psi = 39.7/26.9$. Moreover, applying those rates to a version of our model without aggregate risk (and with inflation equal to the unconditional mean of inflation in the baseline model) gives $G = 0.25$ as a residual; we thus calibrate G to this value in the baseline economy.

¹³Recall that l , m and h refer to the low, medium and high individual productivity states, respectively.

¹⁴A two-state Markov chain conveniently captures the dynamics of inflation here whilst keeping the size of the state space limited (see also Krusell and Smith (1998)).

4. RESULTS

4.1. *Equilibrium distributions and laws of motion*

We first check that our calibration allows us to match the broad features of the three cross-sectional distributions that we focus on: monetary wealth, nonmonetary wealth and consumption. The distinguishing feature of our approach, relative to earlier studies, is the breakdown of the wealth distribution into monetary and nonmonetary wealth. Their empirical counterparts in the SCF are the following. First, we report the empirical distribution of money across U.S. households, as measured in households' checking accounts. Second, we compute the distribution of nonmonetary wealth by removing monetary assets from the financial assets held by the households in the SCF. Nonmonetary wealth refers to bonds, stocks, life insurance, retirement plans and other managed financial assets, and other liquid assets. Table II compares the properties of those two distributions with those generated by the model, under the parameter configuration specified in the previous section. We also report the model-generated wealth distributions under homothetic period utility case, just to illustrate its failure at getting anywhere close to the empirical distribution of monetary wealth.

Given our parametric utility function, all households hold some monetary assets in our model, even though the amount held may be very small. However, many households are not wealthy enough to hold both money and nonmonetary assets: they are borrowing-constrained and endogenously choose not to hold capital claims.

The benchmark model predicts a fairly high Gini index for the distribution of nonmonetary assets (0.81), as is consistent with the data (0.82). Moreover, the model does a reasonable job in matching the lower tail of the distribution of nonmonetary assets. Perhaps unsurprisingly, the model underestimates the nonmonetary wealth share of the top 1%, which is predicted to be 9.03% while it is 34.30% in the data. This flaw is common to many models that only use idiosyncratic income risk to generate wealth dispersion and ignore, for example, entrepreneurship –see, e.g., Quadrini (2000).

Our model predicts that 6.6% of the households on average face a binding borrowing limit. The range of available estimates for this share is notoriously large. Using information on the number of borrowing requests which were rejected in the SCF, Jappelli (1990) argued that up to 19% of families are liquidity-constrained. However, using updated SCF data, Rodriguez, Díaz-Giménez, Quadrini, Ríos-Rull, et al. (2002) reported that only 2.5% of the households have zero net worth. We regard this number as a lower bound for the number of borrowing-constrained households in the US economy. In particular, since the concept of “nonmonetary assets” that we use excludes net homeownership, and since the latter is more equality distributed than other nonmonetary assets, the number of households with no such assets is necessarily greater than 2.5%.¹⁵

The model with non-homothetic utility function allows matching the cross-sectional distribution of money holdings while being also consistent with the cross-sectional distribution of consumption. This is in contrast with the homothetic specification, which counterfactually generates roughly similar degrees of dispersion for consumption and money holdings (the model implied Gini coefficient for money is then 0.31, instead of 0.85 in the

¹⁵We exclude net homeownership from nonmonetary wealth here because housing wealth is typically not used for the smoothing of nondurables consumption (see Challe and Ragot (2016) for a discussion of this point).

data). Again, this property is a direct implication of the proportionality between consumption and money that is implied by the homothetic specification, which our non-homothetic specification is precisely designed to break down.

Table II: Wealth distribution

	Data	Homothetic utility*	Nonhomothetic utility*
Distribution of nonmonetary assets			
Gini	0.82	0.74	0.81
Share of constrained households	[2.5%, 20%]	8.76%	6.6%
Fraction of total assets held by			
Bottom 20%	0.00	0.03	0.12
Bottom 40%	0.20	0.08	0.26
Top 20%	84.70	76.85	78.22
Top 10%	71.20	51.68	52.89
Distribution of money holdings			
Gini	0.85	0.31	0.81
Fraction of total money held by			
Bottom 20%	0.001	12.50	0.84
Bottom 40%	0.90	23.50	1.66
Top 20%	88.20	41.33	87.69
Top 10%	78.10	24.62	67.34
Distribution of consumption			
Gini	0.27	0.31	0.30
Fraction of total consumption by			
Bottom 20%	10.10	12.68	12.57
Bottom 40%	23.84	23.45	24.20
Top 20%	35.98	41.11	39.41
Top 10%	21.21	24.43	22.50
Capital/GDP	12.00	10.26	10.37
Money/GDP	0.32	0.33	0.32

Notes: * The model properties are averages over a 10,000 periods simulation.

Finally, we find the following laws of motion for the capital stock:

$$\begin{aligned} \text{low inflation} & : \ln K_{t+1} = 0.4723 + 0.7870 \ln K_t + 0.1140 \ln \Omega_t^s, R^2 = 0.9982. \\ \text{high inflation} & : \ln K_{t+1} = 0.4421 + 0.8085 \ln K_t + 0.1155 \ln \Omega_t^s, R^2 = 0.9981. \end{aligned}$$

Thus, just as in Krusell and Smith (1998), we find the first-order moments of the distributions of capital and real balances to yield an almost perfect prediction of future capital (hence of the return on capital claims).

4.2. Individual policy rules

Figure 1 displays the individual policy rules when the inflation rate is π^L (the policy rules when $\pi = \pi^H$ have a very similar shape) and for the average holdings of capital claims. The policy rules for total wealth, nonmonetary assets, money and consumption are decomposed for each level of productivity e^h , e^m and e^l . For example, the third panel of Figure 1 reports the individual policy rules for nonmonetary assets. The policy rule lies above the 45-degree line for the most productive household (with productivity e^h) and below the 45-degree line for the other two types (with productivity e^m or e^l). This implies that the former accumulate nonmonetary assets for self-insurance purposes whereas the l-

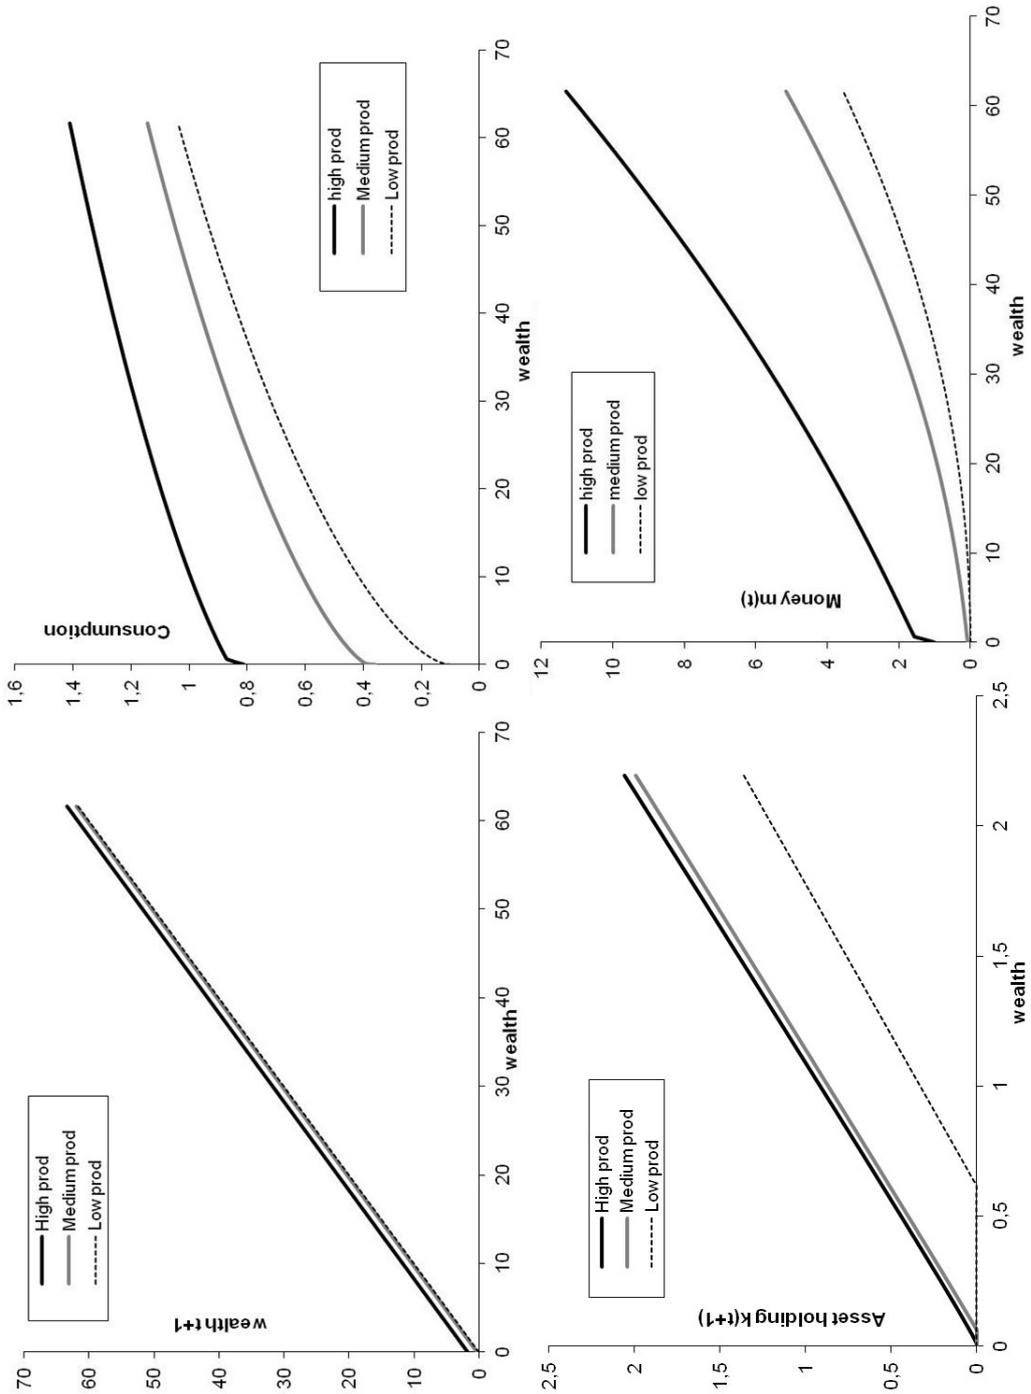


Figure 1: Individual policy rules as a function of total wealth for $\gamma = \gamma_1$.

after dis-save to smooth individual consumption. For medium- and low-productivity households, the policy rule displays a kink at low levels of wealth; this is because at such wealth levels these households hold money but no nonmonetary assets. The second and fourth panels show the policy rules for money holdings and consumption, which roughly display the same pattern as the policy rule for nonmonetary assets. The more productive the household (holding wealth constant), or the wealthier the household (holding productivity constant), the higher are individual consumption and money holdings. The close connection between the policy rules for consumption and money holdings stems from the complementarity between the two, a direct implication of our assumed nonhomothetic utility function.

4.3. Welfare

We now evaluate the welfare cost of inflation fluctuations within our incomplete-market economy. We do so by performing ex ante welfare comparisons conditional on a particular value of the individual state vector (productivity and wealth). The benchmark that we use for comparison is an economy without inflation risk, where the inflation rate is constant and equal to the unconditional mean of inflation (as computed from the baseline model with inflation risk). We wish to understand who are the losers (or winners) from fluctuating inflation and thus proceed as follows.

First, we compute the ergodic distribution of households over productivity e and beginning-of-period wealth a generated by our baseline model, and extract from the ergodic distribution the average beginning-of-period wealth over all households (denoted by \bar{a}) as well as the corresponding averages for low-, medium- and high-productivity households (denoted by a^l , a^m and a^h , respectively).

Second, we compute the ex ante welfare of typical households in the baseline economy. For example, to compute the ex ante welfare of a low-productivity (i.e., e^l) household, we first compute his ex post intertemporal welfare ($\sum_{t=0}^{\infty} \beta^t u(c_t, m_t)$) under a particular history of idiosyncratic and aggregate shocks after date 0, letting him start his life with the average portfolio of households with similar productivity (i.e., a^l). We then repeat the same computation for ten thousands histories of individual productivity, but under the same initial conditions (a and e) and history of aggregate shocks. Averaging over all idiosyncratic histories then gives us the ex ante welfare of this typical household (i.e., $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, m_t)$). We perform similar conditional welfare computations for the typical medium- (e^m, a^m) and high-productivity (e^h, a^h) households. To broaden the picture, we also compute the ex ante welfare of medium-productivity households (e^m) starting his life with the typical portfolio of a high- or a low-productivity household (ie. a^l or a^h). This would be the ex ante welfare level of a household who would have just transited into the medium-productivity class, after having stayed for a long time either in the high or in the low individual productivity state. We also compute the ex ante welfare levels by initial productivity types for households starting their life with the economy-wide average beginning-of-period wealth \bar{a} .

The welfare cost of inflation fluctuations, expressed in terms of equivalent life-time proportional consumption loss, is then calculated by finding the value of κ such that the following equality is satisfied:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(\kappa g_c(a_t, e_t; \bar{\pi}, \bar{K}), g_m(a_t, e_t; \bar{\pi}, \bar{K})) = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(g_c(a_t, e_t; \pi_t, K_t), g_m(a_t, e_t; \pi_t, K_t))$$

where the left hand side of the equation is the ex ante welfare in the benchmark economy without aggregate inflation shocks, i.e., wherein inflation and capital stay constant at the unconditional mean inflation rate $\bar{\pi}$ and capital stock \bar{K} of the baseline model (using the same individual histories as in the economy with inflation risk). In Table III we report the proportional welfare loss generated by inflation risk, i.e. $1 - \kappa$ (%). For example, for a currently mid-productivity household (e^m) holding the typical portfolio of this category (a^m), switching from the constant-inflation benchmark to an economy with inflation risk is as costly as staying in the constant inflation economy but experiencing a permanent 1.23 percent drop in individual consumption.

Table III: The welfare cost of inflation risk (%)

Productivity	Initial wealth	Welfare cost*
e^h	a^h	-0.36
e^m	a^h	-0.80
e^m	a^m	1.23
e^m	a^l	1.77
e^l	a^l	1.24
e^h	\bar{a}	0.87
e^m	\bar{a}	0.89
e^l	\bar{a}	0.93

Notes: * The welfare costs are averages over a 10,000 productivity histories, but under the same initial conditions (e^x, a^x) and history of aggregate shocks.

Differentiating households by productivity types only – that is, considering the welfare loss of individuals with different beginning-of-life productivity – but holding the same initial wealth – the economywide average \bar{a} – illustrates the fact that on average individuals tend to suffer from inflation fluctuations, with an equivalent permanent consumption loss close to 1%. As emphasised in the introduction, such welfare effects are one order of magnitude larger than the typical welfare cost of aggregate fluctuations in perfect-insurance economies, but in line with the welfare gains from eliminating productivity-driven fluctuations in incomplete-insurance economies with correlated unemployment risk. Interestingly, our welfare gains from eliminating inflation risk is large when compared to what existing work has found as to the benefits from lowering mean inflation.¹⁶

Importantly, the welfare cost of inflation fluctuations that we find are very unevenly distributed across households. In particular, while most households incur large welfare losses (up to 1.77% of permanent consumption), the most productive and/or wealthy households incur no loss or even a gain. Such is the case both for the typical high-productivity individual (starting with the average wealth of individuals with same productivity, a^h) and for a medium-productivity household having been lucky enough in the past to have accumulated a^h . To understand the welfare effects at work in our model economy, one must bear in mind how inflation risk affects households’ utility under our preference specification. More specifically, the impact of inflation risk on individual welfare can be traced back to three effects: a direct effect due to fluctuations in the marginal utility of money holdings, an indirect effect working through the distribution of aggregate capital income, and an

¹⁶See, e.g., İmrohoroğlu (1992) and Erosa and Ventura (2002).

indirect effect working through the equilibrium real wage. The first effect implies that inflation risk, which mechanically generates volatility in real money holdings, hurts all households holding cash (since utility is concave in real balances). The second effect is related to the complementarity between money and consumption. Because the marginal utility of real balances is concave and consumption and real balances are complement, an increase in the volatility of future real money balances (as implied by greater inflation risk) lowers the average marginal utility of future consumption and hence discourages current savings.¹⁷ This is true for all types of savings, including holdings of capital claims, and for all households, i.e., at all productivity levels. Hence, inflation risk lowers the capital stock and raises the return on capital claims, which tilts the distribution of total capital income towards high-productivity households (who holds greater capital holdings on average). This effect manifests itself by higher capital returns in the baseline model than in the constant-inflation model: the quarterly after-tax rates of return on capital claims are on average 0.60% and 0.59% in the high and the low inflation states of the baseline economy, respectively, while the same return is time-invariant at 0.55% in the economy without inflation risk. The third effect follows from the effect of inflation risk on the capital stock; namely, since the average capital stock is lower in the economy with inflation risk, so is the equilibrium price of an effective labor unit and hence wages at every individual productivity level. This effect of inflation risk is especially harmful for households who primarily rely on their wage income to provide for consumption and real balances, i.e., households located in the lower part of the wealth distribution.

To summarize, the first effect of inflation risk mechanically hurts all households (since all of them hold at least some cash), the second effect raises the capital income of wealthy households, and the third effect lowers the wage income of all households. This explains why initially wealth-poor households are hit the hardest by inflation risk. In contrast, the second effect turns out to (mildly) dominate the other two for wealthy households, so that they ultimately benefit from inflation risk in equilibrium.

5. CONCLUSION

Existing analyses of the welfare cost of fluctuations in general equilibrium (i.e., with endogenous consumption) have focused on the impact of business cycles driven by real factors, namely changes in total factor productivity and/or the extent of idiosyncratic income risk. As far as we are aware, our analysis is the first that explicitly examines the welfare impact of another key source of aggregate fluctuations, namely, inflation risk. In so doing, we have used the simplest monetary framework (i.e., Sidrauski's MIU, flexible price model), combined with the assumption that households face uninsured idiosyncratic income risk (in the tradition of Bewley, Hugget, Aiyagari and Krusell and Smith). We have then used the moments of the empirical cross-sectional distribution of monetary assets to discipline preferences towards money and consumption goods, which led us to adopt the nonhomothetic preference structure. Our main result is that a moderate amount of infla-

¹⁷Note that the lower future average marginal utility of consumption is accompanied by a greater *volatility* of future marginal utility of consumption. Since the marginal utility of consumption is convex at all levels of real balances, a "prudence" effect will limit, but not overturns, the impact of inflation risk on the average marginal utility of consumption (formally, the impact of prudence on current consumption-saving choices operates at one order of magnitude below the effect due to the complementarity between consumption and money).

tion risk (i.e., of magnitude equal to its empirical measure over the Volcker to pre-Great Recession period) generates substantial welfare losses for the average households (about one percent of permanent consumption), and is most harmful for households towards to the lower end of the distribution of income and wealth. Finally, note that by construction our welfare analysis ignores the potential benefits from stabilizing inflation that naturally arise under nominal rigidities. There are two well identified benefits from stable (zero) inflation in such economies—see Woodford (2003). First, stable inflation eliminates price dispersion, in a context where price dispersion is costly because it pushes households towards consuming inefficiently asymmetric consumption baskets. Second, when demand shocks are dominant relative to supply shocks then stabilizing inflation also stabilizes output around its efficient level. Such welfare effects would further raise the welfare benefits from stable inflation that we have documented here. We leave their study for future work.

APPENDIX A: NUMERICAL ALGORITHM

A.1. Overview of the Algorithm

The algorithm used to obtain the solution of the model is as follows.

1. Given the law of motion for capital, defined by (14), solve the individual problem given by equations (4), (5), (13), with $k_{t+1} > 0$.¹⁸
2. Simulate the economy to approximate the equilibrium law of motion for K . We use the grid-based simulation procedure proposed by Young (2010).

- (a) Set an initial wealth/employment-efficiency distribution $\mu_0(a, e)$ that provides $p_0^{i,e}$, i.e. the mass of agents of employment-efficiency type e with wealth a_i at the i th wealth grid point for, $i = 1, \dots, N_{grid}$.
- (b) Given individual decision rules $g_k(a_t, e_t; \pi_t, K_t)$ and $g_m(a_t, e_t; \pi_t, K_t)$ found in step 1, the wealth/employment-efficiency distribution $\mu(a_t, e_t)$, and a draw for π_t , compute

$$\int \int g_m(a_t, e_t; \pi_t, K_t) d\mu(a_t, e_t) = M_t$$

and

$$\int \int g_k(a_t, e_t; \pi_t, K_t) d\mu(a_t, e_t) = K_{t+1}$$

- c. Repeat steps (b) and (c) to get a long time series for K and M , of which the first part is discarded.
3. Use the time series obtained in step (2) to get the new equilibrium law of motion for K .
4. Compare the new equilibrium law of motion for K with that used in step (1). If they are similar, stop. Otherwise, update the coefficients of the laws of motion, and go to step (1).

A.2. Details on the Resolution of the Individual Problem

We solve the individual problem defined by the following FOC:

$$(15) \quad u_c(c_t, m_t) - u_m(c_t, m_t) = \beta E_t \left[\frac{v_a(a_{t+1}, e_{t+1}; \pi_{t+1}, K_{t+1})}{1 + \pi_{t+1}} \right]$$

$$(16) \quad u_c(c_t, m_t) = \beta E_t \left[(1 + (1 - \tau_{t+1}^c) r_{t+1}) v_a(a_{t+1}, e_{t+1}; \pi_{t+1}, K_{t+1}) \right] \text{ if } k_{t+1} > 0$$

¹⁸The next sub-section describes the algorithm.

and the budget constraint (4) by iterating on the derivative of the value function with respect to a , $v_a(\cdot)$. We stop the iteration process if the new derivative of the value function is sufficiently close to that in the previous step. Otherwise, we update $v_a(\cdot)$ with the new derivative of the value function.

We have two distinct cases, whether the demand for asset is constrained or not:

1) If $k_{t+1} = 0$, we solve equation (15) to get m_t with $c_t = a_t + (1 - \tau_t^l) w_t e_t \bar{n} - m_t$, knowing $v_a(a_{t+1}, e_{t+1}; \pi_{t+1}, K_{t+1})$ and K_{t+1} from the aggregate law of motion.

2) If $k_{t+1} > 0$, we find the solution for g_c, g_m, g_k , and v_a using nested bisection methods. First, we solve for m and k given a certain level of consumption using equation (15) and the budget constraint. Second, we solve equation (16) for c , where the m and k are given by the previous step, knowing $v_a(a_{t+1}, e_{t+1}; \pi_{t+1}, K_{t+1})$ and K_{t+1} from the aggregate law of motion.

REFERENCES

- ARUOBA, S. B., AND F. SCHORFHEIDE (2011): “Sticky prices versus monetary frictions: An estimation of policy trade-offs,” *American Economic Journal: Macroeconomics*, 3(1), 60–90.
- CASTANEDA, A., J. DIAZ-GIMENEZ, AND J.-V. RIOS-RULL (2003): “Accounting for the US earnings and wealth inequality,” *Journal of political economy*, 111(4), 818–857.
- CHALLE, E., F. LE GRAND, AND X. RAGOT (2013): “Incomplete markets, liquidation risk, and the term structure of interest rates,” *Journal of Economic Theory*, 148(6), 2483–2519.
- CHALLE, E., AND X. RAGOT (2016): “Precautionary saving over the business cycle,” *The Economic Journal*, 126(590), 135–164.
- CHARI, V. V., L. J. CHRISTIANO, AND P. J. KEHOE (1996): “Optimality of the Friedman rule in economies with distorting taxes,” *Journal of Monetary Economics*, 37(2), 203–223.
- CHARI, V. V., P. J. KEHOE, AND E. R. MCGRATTAN (2000): “Sticky price models of the business cycle: can the contract multiplier solve the persistence problem?,” *Econometrica*, 68(5), 1151–1179.
- CLOWER, R. (1967): “A reconsideration of the microfoundations of monetary theory,” *Economic Inquiry*, 6(1), 1–8.
- DEN HAAN, W., P. RENDAHL, AND M. RIEGLER (2018): “Unemployment (Fears) and Deflationary Spirals,” *Journal of the European Economic Association*, 16(5), 1281 – 1349.
- DOEPKE, M., AND M. SCHNEIDER (2006): “Inflation and the redistribution of nominal wealth,” *Journal of Political Economy*, 114(6), 1069–1097.
- DOMEIJ, D., AND J. HEATHCOTE (2004): “On the distributional effects of reducing capital taxes,” *International economic review*, 45(2), 523–554.
- EROSA, A., AND G. VENTURA (2002): “On inflation as a regressive consumption tax,” *Journal of Monetary Economics*, 49(4), 761–795.
- GORNEMANN, N., K. KUESTER, AND M. NAKAJIMA (2016): “Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy,” International Finance Discussion Papers 1167, Board of Governors of the Federal Reserve System (U.S.).
- HEATHCOTE, J. (2005): “Fiscal policy with heterogeneous agents and incomplete markets,” *The Review of Economic Studies*, 72(1), 161–188.
- IMROHOROĞLU, A. (1989): “Cost of business cycles with indivisibilities and liquidity constraints,” *Journal of Political Economy*, 97(6), 1364–1383.
- İMROHOROĞLU, A. (1992): “The welfare cost of inflation under imperfect insurance,” *Journal of Economic Dynamics and Control*, 16(1), 79–91.
- JAPPELLI, T. (1990): “Who is credit constrained in the US economy?,” *The Quarterly Journal of Economics*, 105(1), 219–234.
- KRUEGER, D., AND F. PERRI (2006): “Does income inequality lead to consumption inequality? Evidence and theory,” *The Review of Economic Studies*, 73(1), 163–193.
- KRUSELL, P., T. MUKOYAMA, A. ŞAHİN, AND A. A. SMITH JR (2009): “Revisiting the welfare effects of eliminating business cycles,” *Review of Economic Dynamics*, 12(3), 393–404.
- KRUSELL, P., AND A. A. SMITH (1997): “Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns,” *Macroeconomic dynamics*, 1(2), 387–422.
- KRUSELL, P., AND A. A. SMITH, JR (1998): “Income and wealth heterogeneity in the macroeconomy,” *Journal of political Economy*, 106(5), 867–896.
- LUCAS JR, R. E. (1987): “Models of Business Cycles,” .
- (2003): “Macroeconomic priorities,” *American economic review*, 93(1), 1–14.
- OTROK, C. (2001): “On measuring the welfare cost of business cycles,” *Journal of Monetary Economics*, 47(1), 61–92.
- PATINKIN, D. (1965): “Money, interest, and prices; an integration of monetary and value theory,” Discussion paper.
- QUADRINI, V. (2000): “Entrepreneurship, saving, and social mobility,” *Review of economic dynamics*, 3(1), 1–40.
- RAGOT, X. (2014): “The case for a financial approach to money demand,” *Journal of Monetary Economics*, 62, 94–107.
- RODRIGUEZ, S. B., J. DÍAZ-GIMÉNEZ, V. QUADRINI, J.-V. RÍOS-RULL, ET AL. (2002): “Updated facts on the US distributions of earnings, income, and wealth,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 26(3).
- SIDRAUSKI, M. (1967): “Rational choice and patterns of growth in a monetary economy,” *The American*

- Economic Review*, 57(2), 534–544.
- STERK, V., AND S. TENREYRO (2018): “The Transmission of Monetary Policy Operations through Redistributions and Durable Purchases,” *Journal of Monetary Economics*, 99, 124–137.
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2001): “The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk,” *European Economic Review*, 45(7), 1311–1339.
- WALSH, C. E. (2017): *Monetary theory and policy*. MIT press.
- WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.
- YOUNG, E. R. (2010): “Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations,” *Journal of Economic Dynamics and Control*, 34(1), 36–41.