Villanova School of Business Economics Working Paper #40

Preference Heterogeneity, Inflation, and Welfare*

Michael Curran Scott J. Dressler†

Villanova University Villanova University

February 2019

Abstract

This paper assesses the welfare implications of long-run inflation in an environment with essential money, a competing illiquid asset, and potential ex-ante heterogeneity of households with respect to their behavioral measures of risk aversion and elasticity of intertemporal substitution. The results show that the relative liquidity position of households' portfolio as well as potential inter-cohort transfers of resources can deliver fewer welfare costs to inflation than has been previously reported, and in some instances net welfare benefits to low levels of positive inflation. These results hold in versions of the model calibrated to both US and euro area data.

Keywords: Inflation; Welfare; Recursive Preferences

JEL: E21; E41; E50

*This paper has benefitted from comments by YiLi Chien, Chao He, and seminar participants at the Spring 2018 Midwest Macroeconomics Meetings and the 2018 Madison Monetary Conference. This paper uses data from the Eurosystem Household Finance and Consumption Survey.

†Corresponding author. Villanova University; 800 Lancaster Avenue; Villanova, PA 19085-1699. Phone: (610) 519-5934. Fax: (610) 519-6054. Email: scott.dressler@villanova.edu

1. Introduction

This paper examines how ex-ante heterogeneity in household preferences can impact the welfare implications of long-run monetary policy. Lucas (2003) concludes that households' attitude towards risk is crucial when assessing the welfare implications of fiscal and monetary policies, while most models studied in the monetary literature have quantified households' aversion to risk using either long-established values or via model calibration. If these models further assume utility functions of the constant relative risk aversion (CRRA) variety, then a household's degree of risk aversion (RA) and elasticity of intertemporal substitution (EIS) are in fact inverses of each other. There exists no empirical evidence nor theoretical reasoning behind this strict relationship between RA and EIS. What exists is empirical evidence that estimated degrees of RA and EIS can differ from commonly used *average* values, as well as exhibit large dispersion in the cross-section. As welfare calculations directly use household preferences for compensating variation, these behavioral measures can have a first-order impact on the welfare implications of long-run policy.

RA is defined as the degree to which a household prefers to reduce risk. A significantly risk-averse household would prefer an asset with a low average return over an asset with a higher average return if it provides some reduction in risk (e.g. lower return volatility or less liquidity risk). The central estimate of RA implied by labor supply studies is between 1 and 2, while estimates using asset or insurance markets have appeared significantly higher.[1] Barsky et al. (2007) and Kimball et al. (2008, 2009) provide estimates of RA distributions in the US using the Health and Retirement Study (HRS) and the Panel Study of Income Dynamics (PSID). These studies directly measure respondents' RA by utilizing a set of questions regarding hypothetical lotteries on lifetime income, and find large cross-sectional heterogeneity in RA as well as an average degree of RA significantly larger than 2.[2] In

[1] See Chetty (2006) for a detailed explanation on the relatively low estimates for risk aversion implied by labor supply studies. The larger estimates from asset markets stems from the classic equity premium puzzle of Mehra and Prescott (1986).

[2] Chiappori and Paiella (2011) performed a similar analysis using the Survey of Household Income and Wealth (SHIW) in Italy, but were able to indirectly measure risk aversion by examing the changes in respon-

particular, Kimball et al. (2009) report a mean of 4.26 and a median of 2.92.

EIS measures household responsiveness to the intertemporal price of consumption. A significantly low EIS implies that households possess a strong consumption-smoothing motive and do not alter consumption much in response to changes in the return on savings. The empirical literature reports a wide range of plausible values. Havránek et al. (2015) conduct a meta-analysis of 2,753 EIS estimates from 169 published studies covering over 104 countries. A broad weighted-average delivers an EIS estimate of roughly 0.5, but with a large dispersion across countries. For example, they find within-country average estimates of 0.9 for Japan, 0.6 for the US, 0.5 for the UK, 0.4 for Canada, 0.2 for Israel, and 0.1 for Sweden.[3] Gomes and Paz (2013) and Best et al. (2017) conclude that EIS is closer to Hall's (1988) original estimate of 0.1 for the US and UK.

This paper examines a model with rich household heterogeneity to determine the how these wide ranges of behavioral measures impact the welfare implications of long-run inflation. The model has several key features. First, households have generalized recursive preferences so they can possess measures of RA and EIS that are not strictly related to each other as under expected utility. Second, there may exist ex-ante heterogeneity among household cohorts that possess innately different measures of RA and EIS. Third, money is essential and not imposed as a medium of exchange.[4] Finally, there exists a productive and return-dominating asset (i.e. capital) that competes with money as a store of value, but this asset is relatively illiquid and incurs a fixed market entry cost. These features allow agents to self-insure against risk by choosing asset portfolios ranging from completely liquid to completely illiquid, depending upon their RA and EIS. Versions of the model are considered with either one household cohort possessing empirically plausible *average* values

dents' portfolio composition over time. Large and dispersed values of RA have also been found by Harrison et al. (2007) in a lab experiment involving Danish individuals, as well as von Gaudecker et al. (2011) in an on-line questionaire involving Dutch households.

[3] These within-country averages were calculated using over 50 EIS estimates for each country.

[4] Models featuring *essential money* generally satisfy two conditions: (i) the model allows for a nonmonetary equilibrium, and (ii) the set of allocations supported with money is larger (and possibly better) than without.

of RA and EIS, or multiple cohorts exhibiting ex-ante *dispersion* in their values of RA and EIS while possessing plausible average values across cohorts.

Traditional heterogeneous-agent models like the one considered here deliver a steady-state distribution of households, and welfare analyses involve a separation of the consumption distribution into households considering a particular long-run policy as a benefit and those who consider it a detriment. Changing household degrees of RA or EIS can deliver two impacts on the model (inter alia) which together deliver an ambiguous impact on welfare. The first impact is the change in the sensitivity of welfare calculations due changes in curvature of the utility function. Figure 1 compares a household's CRRA utility function with the preference parameter either set to the average RA estimate from Kimball et al. (2009) of 4.26 or a traditional value of 2. It is apparent that a household's marginal utility for any consumption level is positively related to the degree of RA, and any change in consumption would therefore deliver significantly different welfare results. The second impact is that increasing these behavioral measures will also increase households' desire to save in the form of money, capital, or both. An increase in the capital stock results in an increase in output, an increase in average income, and could therefore impact the moments of the consumption distribution. If the consumption distribution depends on households' behavioral measures, then so does the separation of households into policy winners and losers. It is therefore not obvious if a larger aversion to risk or desire to smooth consumption would necessarily deliver larger welfare costs of positive inflation on average.

The results of the model suggest that households who are either more risk averse or have higher motives to smooth consumption can potentially experience reduced welfare costs of long-run inflation and in some cases net welfare *benefits*. While a large portion of the explanation behind this result lies in a shifting of the consumption distribution due to an increase in savings mentioned above, another key driver is the liquidity position of the savings portfolio. Consider the results from a benchmark version of the model where there is only one household cohort exhibiting traditional values of RA $= 2$ and EIS $= 0.5$ (as under expected

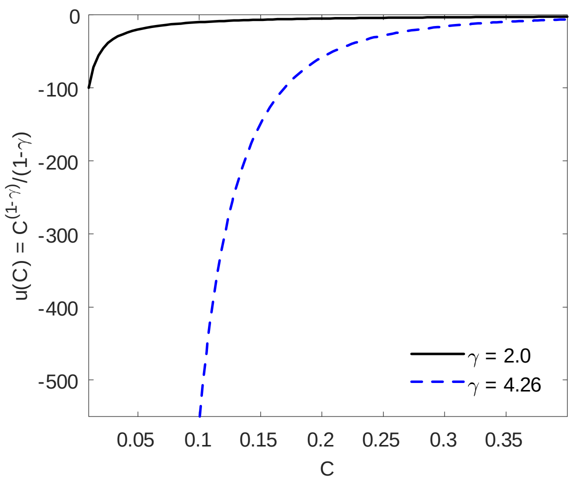

Figure 1: Comparison of CRRA Utility Functions

utility). When calibrated to traditional monetary measures such as velocity and moments of the monetary distribution, the welfare predictions of the model are in-line with previous analyses. Namely, welfare costs are monotonically decreasing in inflation, and 10 percent inflation costs around 2 percent of lifetime consumption. An increase in households' RA or EIS (ceteris paribus) increases their demand to save more liquid and illiquid assets, with a larger demand for liquid assets delivering a more liquid portfolio. At zero percent inflation, a household's liquidity position is highest because the liquid asset is paying the highest (zero percent) real return. When the environment is subjected to small amounts of inflation, households reduce the liquidity position of their portfolios by substituting the liquid asset for the illiquid asset, resulting in an increase in capital, output, and wages (which can also be used to smooth consumption). Ultimately, households find themselves with relatively higher average incomes, which shift the moments of the steady state consumption distribution. Depending on the degree of RA and EIS, the separation of the consumption distribution into policy winners and losers might find a larger proportion on the winning end.

The results described above for an environment with a single cohort of households are actually amplified in environments featuring multiple cohorts with innately different degrees of RA or EIS due to inter-cohort transfers. A multiple cohort version of the model can be thought of as an environment where each cohort has a distinct consumption distribution that together contribute to the same markets clearing. High risk-averse agents are now able to increase their ability to self-insure (i.e. save more) by transferring resources to low risk-averse agents who place less value on self-insurance. These transfers add additional welfare gains of inflation which cannot be displayed in environments with ex-ante homogeneous agents. A final application considers members of the euro area. Using country-average EIS estimates from Havránek et al. (2015) and assuming a common value for RA = 2 across countries, the results again suggest that ex-ante heterogeneity in preferences across countries can lead to stark differences in the welfare implications of long-run inflation when compared to a model with ex-ante homogeneous agents with common (average) behavioral measures.

The paper proceeds with a literature review in Section 2. Section 3 presents the model, and Section 4 reports the results from the computational analysis. Section 5 concludes.

2. Related Literature

While this is the first analysis examining the impact of RA and EIS on the welfare implications of long-run monetary policy, there is nonetheless a related monetary literature worthy of mention.[5]

Camera and Chien (2014) conduct a systematic analysis of welfare costs in a heterogeneous-agent environment with endogenous labor and a rich decision set among assets including money, illiquid bonds, and physical capital. They assume a role for money by means of a cash-in-advance constraint and detail how long-run inflation impacts welfare as well as inequality in wealth, income, and consumption. They find shock persistence, labor elasticity, and financial structure to have important welfare implications. A key result of their analysis is that low values of inflation deliver net welfare benefits when labor supply is sufficiently inelastic. The model analyzed below assumes an inelastic labor supply for two important

[5] This section focuses on computational analyses involving heterogeneous agents in Walrasian markets, and regretably omits a large portion of the monetary literature approaching this topic from a search-theoretic angle following Lagos and Wright (2005).

reasons. First, the model contains essential money which would require significant complications to the timing of markets or realizations of uncertainty in order for labor supply to be elastic.[6] Second, Swanson (2012) shows that elastic labor complicates a clear interpretation of risk aversion in many classes of preferences. While these complications are surmountable, the model considered below is an important first step.

Cao et al. (2018) consider cross-sectional heterogeneity on the welfare implications of inflation. They extend Erosa and Ventura (2002) into an overlapping generations environment and capture several observations of Canadian household survey data such as money holdings increasing (decreasing) with age (consumption level). Agents in their model have homogeneous preferences but are heterogeneous with respect to generational features such as age and access to credit markets, and they report a lower welfare cost of inflation than previously found in the literature. A key driver of their result is the calibration of discount rates across age cohorts as well as age-specific costs to entering a credit market. An interesting question beyond the scope of the current analysis is the extent to which household preferences change with age, possibly alleviating the requirement that households have age-specific costs to accessing credit.

Ex-ante heterogeneity has played an important role in papers outside of the monetary literature. Cozzi (2014) considers the distribution of risk aversion detailed by Kimball et al. (2009) in an Aiyagari (1994) environment without money. Assuming expected utility and a large number of cohorts differing in risk aversion, he shows that the model improves upon predicting wealth inequality relative to a single-cohort environment. Chien et al. (2016) analyze an environment with potential cross-sectional heterogeneity in household preferences, discount rates, and beliefs. They consider recursive preferences and focus on the rich cross-sectional heterogeneity of asset portfolios seen in the data, but do not address the welfare implications of long-run monetary policy.

[6]Examples of these environments include the search-theoretic analyses of Molico (2006) and Chiu and Molico (2011), as well as centralized-market analyses of Dressler (2016) and Wen (2015). Wen (2015) shows how one can have an elastic labor supply with essential money by assuming households make their labor supply decision before the realization of individual uncertainty.

3. The Model

3.1. Environment

Time is discrete with an infinite horizon. The environment is populated by a unit measure of infinitely lived households, sorted into J cohorts. Each household within cohort $j \in \{1, \ldots, J\}$ share the same innate measures of RA and EIS, and these behavioral measures are the only differences across cohorts. The proportion of each cohort is denoted φ_j, with $\Sigma_j \varphi_j = 1$. All other features of the environment are identical across cohorts such as endowments, technologies, and access to markets. The environment is therefore described in terms of household i from cohort j.

Period t begins with a household realizing the outcome of two sources of individual uncertainty. First, she receives an endowment of effective labor units $h_{ijt} \in \{h_L, h_H\}$ that she inelastically supplies to the firm. The endowment process is persistent, evolves according to a finite-state Markov process, and the long-run distribution of labor is invariant. Second, she receives an idiosyncratic preference shock $\theta_{ijt} \in \{\theta_L, \theta_H\}$ which is iid across households and time with $\Pr(\theta = \theta_q) = 0.5,\ q = \{L, H\}$. These shock processes are the only source of uncertainty (i.e. there is no aggregate uncertainty), they are faced by all households regardless of cohort, and they introduce ex-post heterogeneity *within* each cohort j.[7]

There exists a stock of fiat money that is perfectly divisible, costlessly storable, and unable to be produced or consumed by any private individual. Let a superscript N denote nominal values, so M_t^N denotes the nominal stock of money available at the beginning of period t. The law of motion for the money supply is given by $M_{t+1}^N = \mu_t M_t^N$ where μ_t denotes the period t growth rate. A household can hold any non-negative amount of money balances $\left(m_{ijt}^N \in R_+\right)$, and new money is injected into the economy via identical lump-sum transfers τ_t^N to all households at the beginning of the period. Money is a purely liquid asset, and

[7] In addition to providing useful degrees of freedom for model calibration, it will be made clear below that two dimensions of uncertainty are needed to shape the decisions of the household regarding the amount of income to consume versus save, as well as the savings portfolio proportions of liquid and illiquid assets.

there are no ad-hoc requirements on it being a medium of exchange.

There exists a stock of physical capital K_t which is owned by the households. A representative profit-maximizing firm accepts labor and rents physical capital from the households, produces a perishable consumption good, and pays a competitive market wage (w_t) and rental rate (r_t). Capital depreciates at the exogenous rate δ. The production technology is CRS and homogeneous of degree one

$$Y_t = f(K_t, H_t) = K_t^{\alpha} H_t^{1-\alpha}, \tag{1}$$

where K_t and H_t denote the aggregate supply of capital and labor. The profit-maximizing rates paid for capital and labor are standard.

$$r_t = \alpha \left(\frac{H_t}{K_t} \right)^{1-\alpha} \tag{2}$$

$$w_t = (1-\alpha) \left(\frac{K_t}{H_t} \right)^{\alpha} \tag{3}$$

A household enters period t with balances of real money (m_{ijt}) and capital (k_{ijt}), and receives an endowment (h_{ijt}), a preference shock (θ_{ijt}), and the real monetary transfer (τ_t). The market for production opens first, where a household supplies labor inelastically in exchange for a wage rate w_t and lends capital to firms at interest rate r_t. After the household receives all income, the goods and asset markets open so the household can purchase consumption goods (c_{ijt}) as well as choose their asset portfolio $(m_{ijt+1},\ k_{ijt+1})$ to bring into the next period.

Note the timing assumptions imply that households could use their contemporaneous income and capital assets to fully self-insure against risk, potentially leaving no meaningful role for money. A demand for money is induced by emphasizing its relative liquidity value. The environment makes two assumptions on the capital market, taken from analyses examining firm capital investment such as Abel and Eberly (1994, 1996) and Cooper and Haltiwanger

(2006). First, the capital market has a fixed entry cost that a household will incur in order to adjust her capital holdings. A household with capital holdings k_{ijt} must pay a proportional fixed fee $F \times (1 + k_{ijt})$ to enter the market and adjust her capital holdings.[8] If a household chooses to not enter, then she retains her current capital holdings net of depreciation. Second, there is a gap between the buying and selling prices of capital. This gap could capture a partial irreversibility of capital investment or a lemons problem. Taken together, these capital market frictions can deliver periods of investment inactivity and result in a wide array of asset portfolio options. Some households can choose zero money balances believing that their flow of income from labor and capital holdings (as well as their monetary transfer) will provide adequate insurance, while others choose the opposite extreme of entirely self-insuring with money and holding no capital. Households choosing portfolios between these two extremes will be holding a portfolio of both assets, and accumulate money until it is optimal to enter the capital market.

3.2. The Household's Problem

Given a market price P_t, all nominal variables are transformed into real terms (e.g. $m_{ijt} = m_{ijt}^N / P_t$ and $M_t = M_t^N / P_t$). Clearing of the money market in a stationary equilibrium implies $\pi_t = \frac{P_{t+1}}{P_t} = \mu_t$.

The current state of household i from cohort j consists of her current real asset portfolio $(m_{ij}, k_{ij} \in R_+)$ and shock realizations (h_{ij}, θ_{ij}). Let V_j be the value function of a household from cohort j after all shocks and transfers have been realized, and the household has decided whether or not to enter the capital market. Given state vector $\omega_{ij} = (m_{ij}, k_{ij}, h_{ij}, \theta_{ij})$, the household's problem is to choose $\left(c_{ij}, m'_{ij}, k'_{ij}\right)$ to maximize V_j,

$$V_j(\omega_{ij}) = \max\left\{V_j^I(\omega_{ij}), V_j^B(\omega_{ij}), V_j^S(\omega_{ij})\right\} \tag{4}$$

[8] Proportional entry costs are used to eliminate size effects. The environment uses $F \times (1 + k_t)$ as the entry cost and not a simple proportional cost of Fk_t so agents entering a period with zero capital will still face a positive cost of entry F.

where $V_j^I(\cdot)$, $V_j^B(\cdot)$, and $V_j^S(\cdot)$ are the values of a cohort$-j$ household choosing to be inactive, an active buyer, or an active seller in the capital market. The value of choosing action $A \in \{I, B, S\}$ is the solution to

$$V_j^A(\omega_{ij}) = \max_{\left\{c_{ij}, m'_{ij}, k'_{ij}\right\}} (1+\theta_{ij}) \frac{c_{ij}^{1-\gamma_j}}{1-\gamma_j} - \beta \left[E\left(-V_j\left(m'_{ij}, k'_{ij}, h'_{ij}, \theta'_{ij}\right)\right)^{1-\epsilon_j}\right]^{\frac{1}{1-\epsilon_j}}. \quad (5)$$

The constraint set for a household choosing inaction in the capital market $(A = I)$ is given by

$$\begin{aligned} c_{ij} + \pi m'_{ij} &\leq wh_{ij} + m_{ij} + rk_{ij} + \tau \\ 0 &\leq m'_{ij} \\ k'_{ij} &= (1-\delta) k_{ij} \end{aligned}$$

Note the decision to be inactive in the capital market imposes the capital stock next period to be $k'_{ij} = (1-\delta) k_{ij}$.

The constraint set for a household choosing to be an active buyer in the capital market $(A = B)$ is given by

$$c_{ij} + \pi m'_{ij} + p_B \left[k'_{ij} - (1-\delta) k_{ij}\right] + F(1+k_{ij}) \leq wh_{ij} + rk_{ij} + m_{ij} + \tau$$

$$0 \leq m'_{ij}$$

$$(1-\delta) k_{ij} \leq k'_{ij}$$

where p_B denotes the real purchase price of capital. Note the decision to be an active buyer in the capital market results in the cohort$-j$ household paying the fixed entry cost $F(1+k_{ij})$ and choosing $k'_{ij} \geq (1-\delta) k_{ij}$.

The constraint set for a household choosing to be an active seller in the capital market

$(A = S)$ is given by

$$c_{ij} + \pi m'_{ij} + p_S \left[k'_{ij} - (1-\delta)\, k_{ij}\right] + F\,(1 + k_{ij}) \leq wh_{ij} + rk_{ij} + m_{ij} + \tau$$

$$0 \leq m'_{ij}$$

$$k'_{ij} \leq (1-\delta)\, k_{ij}$$

where $p_S\,(\leq p_B)$ denotes the real selling price of capital. Note the decision to be an active seller in the capital market results in the cohort$-j$ household paying the fixed entry cost $F\,(1 + k_{ij})$ and choosing $k'_{ij} \leq (1-\delta)\, k_{ij}$.

The household's problem considers the generalized recursive specification of Epstein and Zin (1989) and Weil (1989), previously used by Rudebusch and Swanson (2012) and others in the finance literature. The Bellman equations in (5) have the same general form as expected utility preferences, but the expectation operator is *twisted* and *untwisted* by the coefficient $1 - \epsilon_j$. Swanson (2012) shows that the EIS of the cohort-j household is given by $\frac{1}{\gamma_j}$ as under expected utility, but RA is given by $\gamma_j + \epsilon_j\left(1 - \gamma_j\right)$.[9] When $\epsilon_j = 0$, the above specification reduces to standard expected utility maximization. When $\epsilon_j \neq 0$, the household's RA is amplified (or attenuated) by the additional curvature parameter ϵ_j. Assuming $\gamma_j > 1$, $\epsilon_j < 0$ $(\epsilon_j > 0)$ results in RA being higher (lower) than under expected utility.[10]

3.3. The Stationary Equilibrium

With the exception of market clearing conditions requiring aggregation across all J cohorts, the equilibrium can be defined independently for each cohort. The cohort-j notation is therefore suppressed. Let the state of a household be denoted $\omega := (m, k, h, \theta) \in \Omega := M \times K \times H \times \Theta$, with $M = [0, \infty)$, $K = [0, \infty)$, $H = \{h_L, h_H\}$, $\Theta = \{\theta_L, \theta_H\}$. Let $\mathcal{P}(H)$

[9] Swanson (2012) shows how this specification is equivalent to traditional Epstein-Zin recurrsive preferences, while delivering computational advantages such as standard dynamic programming regularity conditions.

[10] The specific form of preferences implicitly assume $\gamma_j > 1$, which is the case for all numerical experiments considered below.

denote the power set of H, $\mathcal{P}(\Theta)$ denote the power set of Θ, and $\mathcal{B}(M)$ and $\mathcal{B}(K)$ denote the respective Borel $\sigma-$algebra of M and K. Let $\mathcal{B}(\Omega) := \mathcal{B}(M) \times \mathcal{B}(K) \times \mathcal{P}(H) \times \mathcal{P}(\Theta)$ and define the subset of possible states $\mathbb{B}(\Omega) := (\mathbb{M}, \mathbb{K}, \mathbb{H}, \mathbb{T}) \subseteq \mathcal{B}(\Omega)$. Let $\{\Omega, \mathcal{B}(\Omega), \Phi\}$ define the probability space, where Φ is a probability measure. Finally, let ϕ denote the joint probability density associated with the probability space $\mathcal{B}(\Omega)$. Two discrete random variables (h and θ) and two continuous random variables (m and k) imply that ϕ is a mixed density.

Given the current realizations of the labor shock $(h \in H)$ and preference shock $(\theta \in \Theta)$, $p(h'|h)\, p(\theta')$ is the conditional probability of receiving (h', θ') next period. The evolution of the distribution of the state ω can be characterized using a transition function $\mathcal{Q} : \Omega \times \mathcal{B}(\Omega) \to [0, 1]$ defined by

$$\mathcal{Q}(\omega, \mathbb{B}(\Omega)) = \begin{cases} \sum\limits_{h' \in \mathbb{H}} \sum\limits_{\theta' \in \mathbb{T}} \text{ if } (m'(\omega), k'(\omega)) \in \mathbb{M} \times \mathbb{K} \\ \qquad 0 \text{ otherwise} \end{cases}$$

for all $\omega \in \Omega$ and all $\mathbb{B}(\Omega) \subseteq \mathcal{B}(\Omega)$.

A stationary monetary equilibrium is a time-invariant distribution of consumption, real money balances, and real capital holdings across the population of households (of each cohort-j), such that on each date the optimal plan of a household in state $\omega = (m, k, h, \theta) \in \Omega$ involves $c(\omega)$ consumption, $m'(\omega)$ monetary savings, and $k'(\omega)$ capital savings. These optimal plans solve the household problem given that the firm maximizes profit, all markets clear, and the distributions of states are stationary.

Definition 1 *(Recursive Equilibrium). A stationary monetary recursive competitive equilibrium is a constant inflation rate π, a wage rate w, an interest rate r, a set of policy functions $m' : \Omega \to \mathbb{R}_+$, $k' : \Omega \to \mathbb{R}_+$, $c : \Omega \to \mathbb{R}_+$, and an invariant probability measure Φ such that*

1. Given π, w, and r, the policy functions $m'(\omega), k'(\omega)$, and $c(\omega)$ for $\omega \in \Omega$ solve the household problem (4).

2. Given w and r, the firm demands labor and capital according to (2) and (3).

3. Markets for money and capital clear:

$$\sum_{j\in J}\left(\sum_{h\in H}\sum_{\theta\in\Theta}\int_{m\in M}\int_{k\in K}k'(\omega)\,\phi(\omega)\,dmdk\right)\varphi_j=K$$
$$\sum_{j\in J}\left(\sum_{h\in H}\sum_{\theta\in\Theta}\int_{m\in M}\int_{k\in K}m'(\omega)\,\phi(\omega)\,dmdk\right)\varphi_j=M$$

4. For each $j\in J$ and all subsets $\mathbb{B}(\Omega)\subseteq\mathcal{B}(\Omega)$, the cdf $\Phi'(\mathbb{B}(\Omega))$ satisfies

$$\Phi'(\mathbb{B}(\Omega))=\sum_{h\in H}\sum_{\theta\in\Theta}\int_{m\in M}\int_{k\in K}\mathcal{Q}(\omega,\mathbb{B}(\Omega))\,\phi(\omega)\,dmdk$$

and $\Phi'(\mathbb{B}(\Omega))=\Phi(\mathbb{B}(\Omega))$.

4. Quantitative Analysis

4.1. Preference Values

Model versions are assessed considering one or more cohorts of agents $(J\geq 1)$, and assume values of RA and EIS as follows.

Model versions assuming $J=1$ consider four cases using two parameters each for a household's RA and EIS. Risk aversion is assumed to be either 2, as in most CRRA analyses, or 4.26 which is the Kimball et al. (2009) estimated mean of the risk aversion distribution using the PSID. Given these two RA values, two cases naturally arise under expected utility. In particular, Case 1 considers RA = 4.26 and EIS = 0.24 ($=1/RA$), while Case 4 considers RA = 2 and EIS = 0.5. In order to assess the impact of breaking the link between RA and EIS made possible by recursive preferences, two additional cases are considered where the RA and EIS values are switched. The resulting preference parameters for these four cases are detailed in Table 1. It will be convenient to keep track of these cases by referring to their RA and EIS values as relatively *high* or *low*. In the results that follow, Case 1 will be referred to as the high RA - low EIS case or (R_H,E_L). Cases 2 through 4 will be respectively

Case	RA	EIS	$\gamma = \frac{1}{EIS}$	$\epsilon = \frac{RA-\gamma}{1-\gamma}$
$1\,(R_H, E_L)$	4.26	0.24	4.26	0
$2\,(R_H, E_H)$	4.26	0.5	2	-2.26
$3\,(R_L, E_L)$	2	0.24	4.26	0.69
$4\,(R_L, E_H)$	2	0.5	2	0

Table 1: Preference Parameters, Single Cohort Cases

referred to as (R_H, E_H), (R_L, E_L), and (R_L, E_H). It should be stressed that given the range of plausible values for behavior measures reported in the empirical literature, the parameter values considered are quite conservative.

Model versions considering multiple cohorts assume $J = 3$ where RA values are again taken from Kimball et al. (2009). In particular, all cases consider a distinct RA value for each cohort (1.62, 5.24, 7.93) which correspond to the 25th, 75th, and 87.5th percentiles of the estimated distribution. The proportion of each cohort in the economy is given by $\varphi_1 = 0.5$, and $\varphi_2 = \varphi_3 = 0.25$, resulting in an average RA to be roughly 4.26 as found by Kimball et al. (2009).[11]

Taking these RA values as fixed, three cases are considered. Case J1 assumes expected utility with $EIS_j = 1/RA_j$ for each cohort j. Case J2 assumes an identical $EIS_j = 0.5, \forall j$ and allows the impact of ex-ante heterogeneity in RA to be assessed in isolation. Finally, Case J3 adjusts the γ_j values in Case J1 proportionally so the cohorts have distinctly different values of EIS, but possess a weighted average $\overline{EIS} = 0.5$ as in Case J2. This case is considered because most empirical EIS estimates in the literature are in fact average values. This may be innocuous in a representative-agent environment, but a comparison of Case J2 with Case J3 will determine how much cross-sectional heterogeneity matters. The resulting preference parameters for these three cases are detailed in Table 2.

[11] The larger proportion placed on the first cohort is in line with the empirical distribution estimated by Kimball et al. (2019), while the exact RA value was conservatively chosen after some experimentation. These experiments showed that households with any RA value at or below the median of the distribution behaved similarly insofar as they were choosing to hold few assets for self insurance.

	EIS_j			$\gamma_j = \frac{1}{EIS_j}$			$\epsilon_j = \frac{RA_j - \gamma_j}{1-\gamma_j}$		
Case	$j = 1$	$j = 2$	$j = 3$	$j = 1$	$j = 2$	$j = 3$	$j = 1$	$j = 2$	$j = 3$
$J1$	0.62	0.19	0.13	1.62	5.24	7.93	0	0	0
$J2$	0.5	0.5	0.5	2	2	2	0.39	-3.25	-5.93
$J3$	0.80	0.25	0.16	1.26	4.07	6.15	-1.42	-0.38	-0.35

Table 2: Preference Parameterization, Multiple Cohort Cases

4.2. Other Parameter Values, Calibration Strategy, and Computation

The model frequency is annual and the parameters shared across model versions are as follows. The stochastic process for labor earnings is taken from Floden and Linde (2001), who use the PSID to estimate an AR1 model with a correlation $\rho(\ln h) = 0.92$ and standard deviation $std(\ln h) = 0.21$. The Rouwenhorst (1995) method was used to approximate a two-state Markov Process with resulting values for the endowment process of $(h_L, h_H) = (0.1, 0.3)$ and a transition matrix of $\pi_{LL} = \pi_{HH} = 0.96$. All models assume traditional values of $\beta = 0.96$, $\alpha = 0.36$, and $\delta = 0.10$, as well as $p_B = 1$ and $p_S = 0.75$ taken from Cooper and Hawltiwanger (2006).

The remaining parameters are the size of the preference shock θ and the capital entry cost F. These parameters are used to calibrate the models so they predict two monetary moments at two percent long-run inflation: a monetary velocity of 5, and a median-to-mean ratio in the equilibrium distribution of money holdings of roughly 0.3.[12] This ratio is the unweighted average of the median-to-mean holding of checking accounts (in 2016 dollars) taken from the Survey of Consumer Finances (SCF) for the years 2010, 2013, and 2016.[13]

The resulting parameter values for all cases are reported in Table 3. Note higher RA or

[12]The choice of money demand for calibration is a common practice in welfare analyses (see Lucas, 2000). This analysis takes a more direct approach of calibrating to one point on the money demand curve (as opposed to fitting the entire curve itself) due to computational feasibility.

[13]The use of the median-to-mean ratio in the SCF follows Dressler (2016) who show that capturing the number of agents with below-average money balances (i.e. those who stand to gain from money transfers associated with inflation) results in a majority of households voting for inflation above the level chosen by a social planner (i.e. the Friedman Rule).

Case	$1\,(R_H, E_L)$	$2\,(R_H, E_H)$	$3\,(R_L, E_L)$	$4\,(R_L, E_H)$	$J1$	$J2$	$J3$
$\pm\theta$	0.19	0.33	0.35	0.35	0.15	0.35	0.26
F	$1.35e^{-3}$	$3.57e^{-3}$	$1.99e^{-3}$	$5.23e^{-3}$	$1.60e^{-3}$	$3.70e^{-3}$	$2.00e^{-3}$

Table 3: Parameters via Calibration, Money and Capital Model

lower (or more dispersed) EIS results in lower values of entry cost to the capital market F.

The model was solved via standard value function iteration with a Howard improvement added to increase computational speed. Given the kinks in the value function due to the discrete choice of entering the capital market, as well as the non-monotonic behavior of the decision rule on money holdings, this solution method proved to be the most robust.

4.3. Single Cohort Results ($J = 1$)

This section reports the impact of changing preferences in the model when populated by one cohort of ex-ante homogeneous households. These models still provide stationary distributions of agents in a steady-state equilibrium, only these distributions are due to ex-post heterogeneity in shock histories and not due to innate differences in preferences across cohorts within the same environment. Before presenting the predictions of the model from changing both preferences and calibrated parameters, a comparative statics analysis is presented where the four cases use the same parameter values of θ and F recovered from the calibration exercise of the benchmark case (Case 4) and differ only in preferences.

4.3.1. Comparative Statics

The results of the comparative statics analysis are presented in Figure 2, and can be summarized as follows.

Result 1: At *low* levels of inflation, welfare costs are positively related to EIS and negatively related to RA only when EIS is low.

The top-left panel of Figure 2 illustrates the welfare result. Welfare is calculated throughout the analysis using a compensating variation method as is common in the literature and detailed in an appendix. For a positive amount of long-run inflation, the welfare measure is the percentage of consumption a household would need to give up or be compensated in order to be indifferent between that inflation rate and zero inflation. A positive number indicates that households prefer that level of inflation over zero inflation (i.e. net welfare benefits), while a negative number indicates net welfare costs of positive inflation.

Recall that the benchmark case (Case 4) considers $RA = 2$ and $EIS = 0.5$ as under expected utility. As common in the related literature, this case reports that inflation is always costly, welfare costs are monotonically increasing in inflation, and the welfare cost of ten percent inflation is 1.9% of consumption. The other cases indicate how changing the households' level of RA or EIS (ceteris paribus) can potentially impact the welfare implications of long-run inflation. Increasing the household's level of RA (Case 2) delivers almost no change in welfare, especially when inflation is relatively low. Increasing RA and decreasing EIS together (Case 1) delivers a stark change in the welfare results at low levels of inflation, and even delivers net welfare benefits (0.04 percent) of inflation up to 3 percent. The remaining case (Case 3) illustrates the impact of decreasing EIS while leaving RA at its benchmark level of 2. The welfare measures are smaller than in Case 1, but this case still predicts net welfare benefits (0.11 percent) of 1 percent inflation. Taken together, these four cases show that EIS has a robust impact on the welfare costs of inflation while RA only impacts welfare when EIS is low.

The remaining panels of Figure 2 provide the intuition behind the result. The top-right panel illustrates the average ratio of money to capital in households' portfolio. It is clear that increasing a household's aversion to risk or desire to smooth consumption (i.e. decreasing EIS) increases the liquidity of her portfolio, especially at zero percent inflation where the return on money is the highest. When subjected to positive amounts of inflation, households reduce their liquidity position across all cases. The reduction in liquidity is more dramatic depending on how much household preferences differ from the benchmark case. As inflation gets increasingly large, all households choose a similarly low liquidity position regardless of preferences. This suggests that increasing inflation can eventually make the liquid asset equally unattractive to households regardless of preferences.

While the change in the liquidity position of household assets are qualitatively similar across the four cases, the differences in magnitudes is what is driving the results. The second row of panels in Figure 2 illustrate steady state aggregate levels of money and capital. Note that at zero percent inflation, more risk aversion or a larger consumption smoothing motive increases the savings of both money and capital. The left panel illustrates that positive inflation induces all households to decrease their money holdings, while the right panel illustrates qualitative differences in changes in their capital holdings. Households in Cases four and two slightly decrease their average capital balances in response to positive inflation while households in Cases one and three have periods where they increase their capital holdings. This shows that the more dramatic reductions in the money-capital ratio can be achieved through substitution of the liquid asset for the illiquid asset. The third row of panels in Figure 2 illustrate the impact of the changes in capital holdings on wages and output, illustrated as percentage changes relative to their zero inflation levels. Due to inelastic labor supply, both wages and output track the changes in capital holdings. When households receive higher (lower) wages, they have more (less) income which can be used for self-insurance. Finally, the bottom row of panels in the figure illustrate the response to consumption. All cases show a reduction in average consumption (the left panel), which is due to lower levels of output in Cases two and four or due to higher rates of savings in Cases one and three. However, the right panel gives an indication of what is happening in the consumption distribution across the four cases. In particular, Case one illustrates an increase in median consumption at low levels of inflation which further supports that welfare gains are going to a larger proportion of the households at these positive inflation levels.[14]

Result 2: At *high* levels of inflation, changes in EIS and RA alone have little to no effect

[14]The median is displayed because it offers the most stark comparison across the four cases, but other measures of the consumption distribution (i.e. the standard deviation, Gini index, etc.) all support this result.

on welfare costs.

While the first result focuses on the impact of low levels of inflation (from one to five percent), the welfare impact of ten percent inflation appears roughly constant across the four cases. This is due to the eventual convergence of the households' liquidity position indicated in the top-right panel. Once households reach a minimal holding of liquid assets, they all face the same low level of self-insurance provided by these small balances. The welfare benefits (if any) of substituting from the liquid to illiquid asset gets eroded, and households must be compensated by roughly the same amount of their zero inflation consumption level to be at 10 percent long-run inflation.

4.3.2. Calibrated Results

While the comparative statics analysis isolates the impact of changing preferences, it should be noted that Cases one through three did not display any empirical features used for calibration. Figure 3 illustrates the same model comparisons, only now each case uses their calibrated parameter values of F and θ reported in Table 3.

Result 3: At *all* levels of inflation, welfare costs of inflation are positively related to EIS, and negatively related to RA only when EIS is low.

Compared to the comparative statics exercise above, the results presented here are quantitatively different due to the size of the entry cost F across the four cases. Note in Table 3 that as households are given increasingly higher amounts of risk aversion or incentives to smooth-consumption, the model requires a lower cost to enter the capital market. This directly follows from the fact that higher RA or lower EIS increases the household's desired liquidity position. In order for each case to match monetary velocity, this increased desire for liquidity needs to be suppressed by making capital easier to obtain.

At low levels of inflation, the welfare results are quite close to the comparative static results. Namely, increasing only risk aversion (Case 2) delivers almost no change in welfare costs relative to the benchmark, while decreasing EIS alone or with an increase in risk

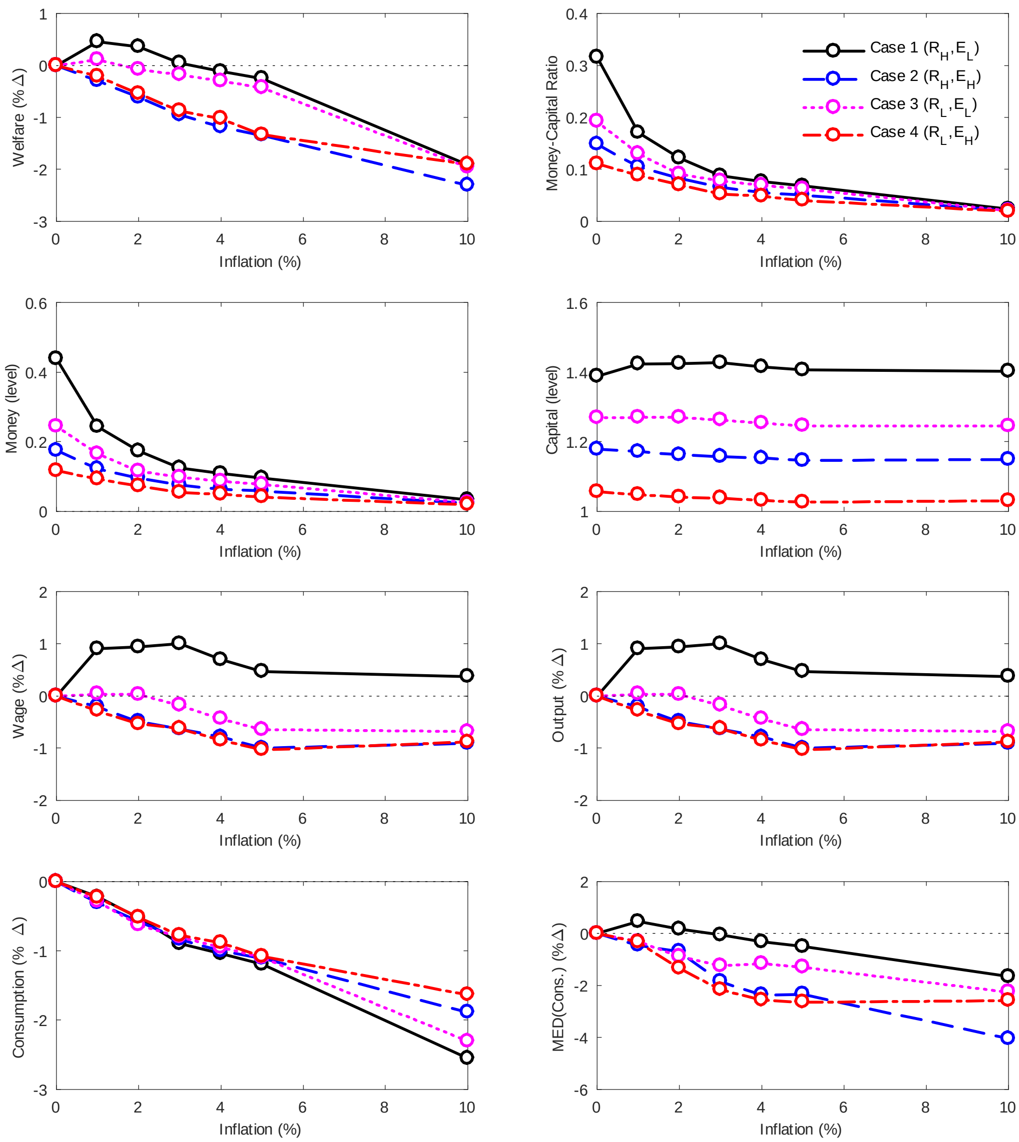

Figure 2: Single-Cohort Model: Comparative Static Results

aversion (Cases 1 and 3) deliver net welfare benefits. The intuition is also similar, but now the movements in the households' average liquidity positions are more drastic in the cases where the entry cost to the capital market is the lowest. The lower the entry cost, the easier households can shift from the relatively high liquidity position at zero inflation to a lower liquidity position when inflation becomes positive.

At high levels of inflation, although the liquidity positions of households across all cases converge to a minimal level as in the comparative statics analysis, the welfare benefits of increased inflation for some cases do not entirely erode away. This is because cases where households have the lowest capital entry costs bring down their money-capital ratio by drastically decreasing their money holdings and increasing their capital holdings. For example, the aggregate level of capital for Case 1 is roughly 4 percent higher at ten percent inflation than at zero percent, resulting in over a 1 percent increase in wages and output. This sustained increase in output prevents a drop in average consumption as large as in the benchmark case (Case 4) and delivers a reduced welfare cost of 10 percent inflation. When comparing the welfare costs across the four cases, it can be seen that increasing risk aversion alone (Case 4 vs Case 2) does not impact the welfare cost experienced at 10 percent inflation, while decreasing EIS only (Case 3) and together with an increase in RA (Case 1) reduces the welfare cost of ten percent inflation by 53 percent and 77 percent.

4.4. Multiple Cohort Results ($J = 3$)

This section presents results from three cases of the model populated by multiple cohorts of households. Recall that the three cases considered share the same distribution of RA, yet differ in the distribution of EIS.

Result 4: At all levels of inflation, cross-sectional heterogeneity in RA alone does not impact the welfare costs of inflation.

Result 5: At all levels of inflation, cross-sectional heterogeneity in EIS reduces the welfare costs of inflation.

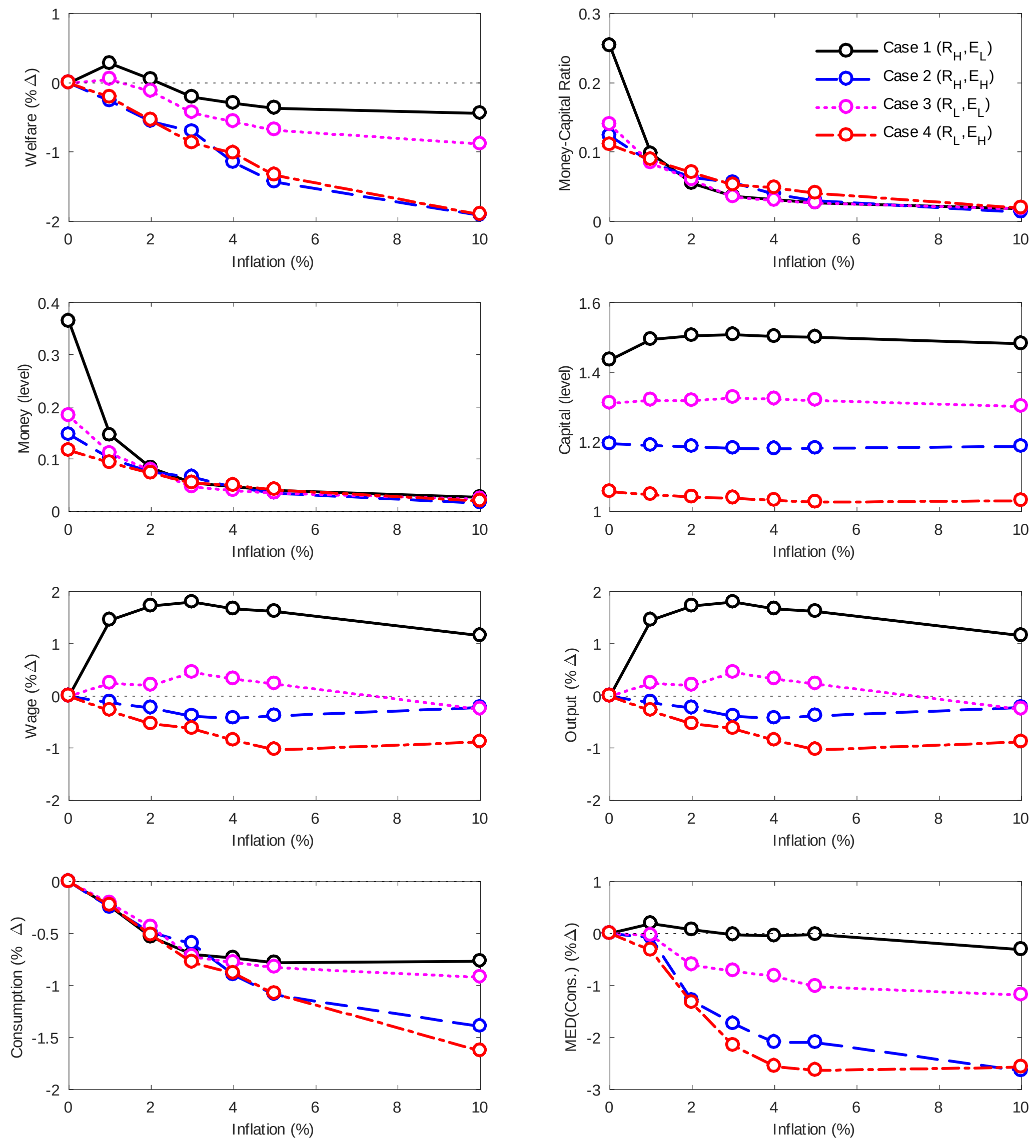

Figure 3: Single-Cohort Model: Calibrated Results

These welfare results are illustrated by comparing the multiple-cohort cases with the relevant single-cohort cases illustrated in Figure 4. Aggregate welfare is calculated as in Cao et al. (2018), where the independent welfare calculations of each cohort are aggregated according to each cohort's proportion of the population.[15] First, the single-cohort Case 2 is compared to the multiple-cohort Case 2 in the bottom panel of the figure. The single-cohort case considers ex-ante homogeneous households with relatively high RA (= 4.26) and EIS (= 0.5), while the multiple-cohort case considers ex-ante heterogeneous households with respect to their levels of RA, but all cohorts possess an EIS of 0.5. Note that the weighted average RA in the multiple-cohort model is similar to the RA of the single-cohort model. The welfare results are essentially identical: welfare costs are monotonically increasing, and the welfare costs of 10 percent inflation is roughly 1.9 percent of consumption. Second, the single-cohort Case 1 is compared to the multiple-cohort Case 1 in the top panel of the figure. The single-cohort case considers ex-ante homogeneous households with relatively high RA and a relatively low EIS (= 0.24) , while the multiple-cohort case considers ex-ante heterogeneous households with respect to both RA and EIS. The multiple-cohort case reports much larger welfare benefits of low levels of positive inflation up to 5 percent as well as a 72 percent reduction in the welfare cost of 10 percent inflation. Finally, compare the single-cohort Case 2 with the multiple-cohort Case 3 where households are ex-ante heterogeneous with respect to both their levels of RA and EIS, but the weighted average degree of EIS is equal to 0.5 (bottom panel). As in the multiple-cohort Case 1, this case reports welfare benefits to low levels of inflation as well as a 73 percent reduction in the welfare cost of 10 percent inflation. This final comparison is the most compelling, for it illustrates that a model with ex-ante heterogeneous households possessing average degrees of RA and EIS equal to those of a model with ex-ante homogeneous households can deliver decidedly different predictions on the welfare costs of inflation.

Figure 5 compares additional results for the multiple-cohort cases, and tells a similar

[15] Calculating a single welfare cost across cohorts ultimately results in some cohorts being over (under) compensated, which unneccesarily clouds the analysis.

story to the single-cohort cases presented above. At zero percent inflation, households with a larger (on average) or more dispersed motive to smooth consumption save more money and capital in levels, and hold a larger liquidity position on their assets. When faced with positive inflation, these households reduce their liquidity position by both decreasing their stock of money and increasing their stock of capital. This again results in increased wages and output. The increase in income for the cases with smaller or more dispersed EIS is still paired with lower consumption on average as in the benchmark case (Case 2), but these cases do experience an increase in median consumption.

Two additional figures help gain further insight to the multiple-cohort model. Figure 6 illustrates the individual cohort results for Case 3 together with their weighted average. As one would expect, households possessing larger risk aversion or self-insurance motives are the larger savers of both money and capital, but households possessing the lowest risk aversion or self-insurance motive are holding the most liquid portfolio at zero inflation by saving almost entirely in money.[16] All cohorts eventually converge to a similar money-capital ratio when faced with positive inflation, but only the lowest risk averse households are reducing this ratio by decreasing money while increasing capital. Notice in the middle-right panel illustrates that the weighted average level of capital remains almost unchanged as inflation increases. This is the result of inter-cohort transfers are not possible in the single-cohort model. When a decrease in demand for capital from a cohort is met (or exceeded) with an increase from other cohorts, then there are no general equilibrium factors placing downward pressure on wages or output. It still works out that the more risk averse households are holding a relatively more liquid portfolio, and this coincides with relative changes in consumption and welfare.

Figure 7 further illustrates how the average results displayed in Figure 6 can feed into welfare results by comparing the consumption distributions of the cohorts in Case 3 under zero and one percent inflation. One percent inflation was chosen for comparison because this is the inflation rate where the cohorts in Case 3 received the largest welfare gains of

[16] Due to this cohort holding almost no capital at zero inflation on average, the extremely high money-to-capital ratio is not illustrated in the figure.

positive inflation. The figure illustrates that increasing inflation results in an increase in consumption in the inter-quartile range of cohort 1 (the least risk averse), but decreases for cohorts 2 and 3. Although these changes appear slight, two important details should be noted. First, households in the lower portion of the consumption support in the least risk averse cohort are on the steepest portion of their utility function and experience the largest gains to consumption. Second, the least risk averse cohort makes up 50 percent of the population (as in the data). Taken together, these results show that there are levels of positive inflation where the gains to the winners are larger than the costs to the losers.[17]

4.5. Summary of Results

The results have shown that changes in preferences, within models populated by either ex-ante or ex-post heterogeneous agents, can deliver qualitatively different welfare implications of long-run inflation. The main driver of these results are due to the liquidity position of the households' asset holdings at zero percent, which is positively correlated with a household's degree of risk aversion or motive to smooth consumption. The higher the liquidity position at zero percent inflation, the larger the shift into illiquid assets at positive levels of inflation. This shift in the liquidity position can deliver increases in output and income, and can deliver reduced welfare costs and sometimes net welfare benefits relative to a benchmark model with traditional preferences. This section briefly discusses several additional takeaways from the analysis.

The model results challenge several widely held beliefs on what fuels the welfare costs of inflation. First, one common belief is that households who hold a more liquid portfolio are the ones who find inflation more costly. This model suggests that households with the highest liquidity position can be made better off by substituting their liquid assets for illiquid assets made available by households with a lower liquidity position. Since these households in the model were also the least risk averse and populated the lower end of the support of

[17]Note that the comparison of utility functions illustrated in Figure 1 provides intuition behind this outcome, the utility functions are not directly comparable because they do not represent life-time utility.

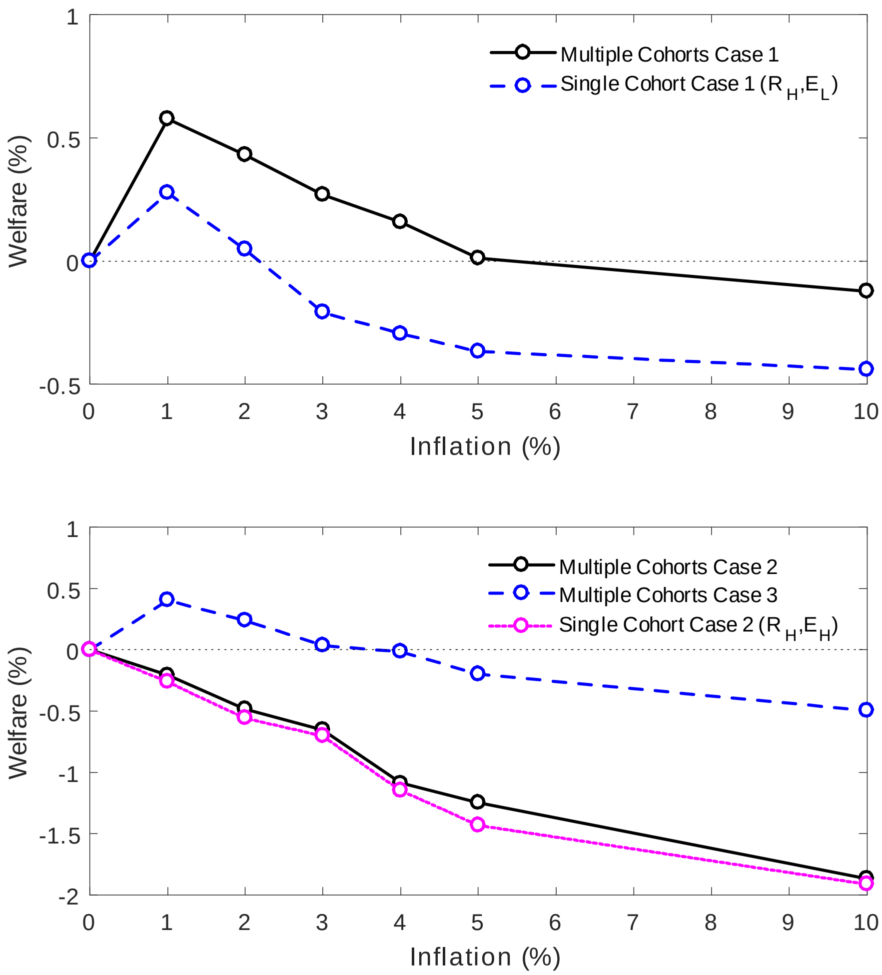

Figure 4: Welfare Results: Single vs Multiple Cohort Models

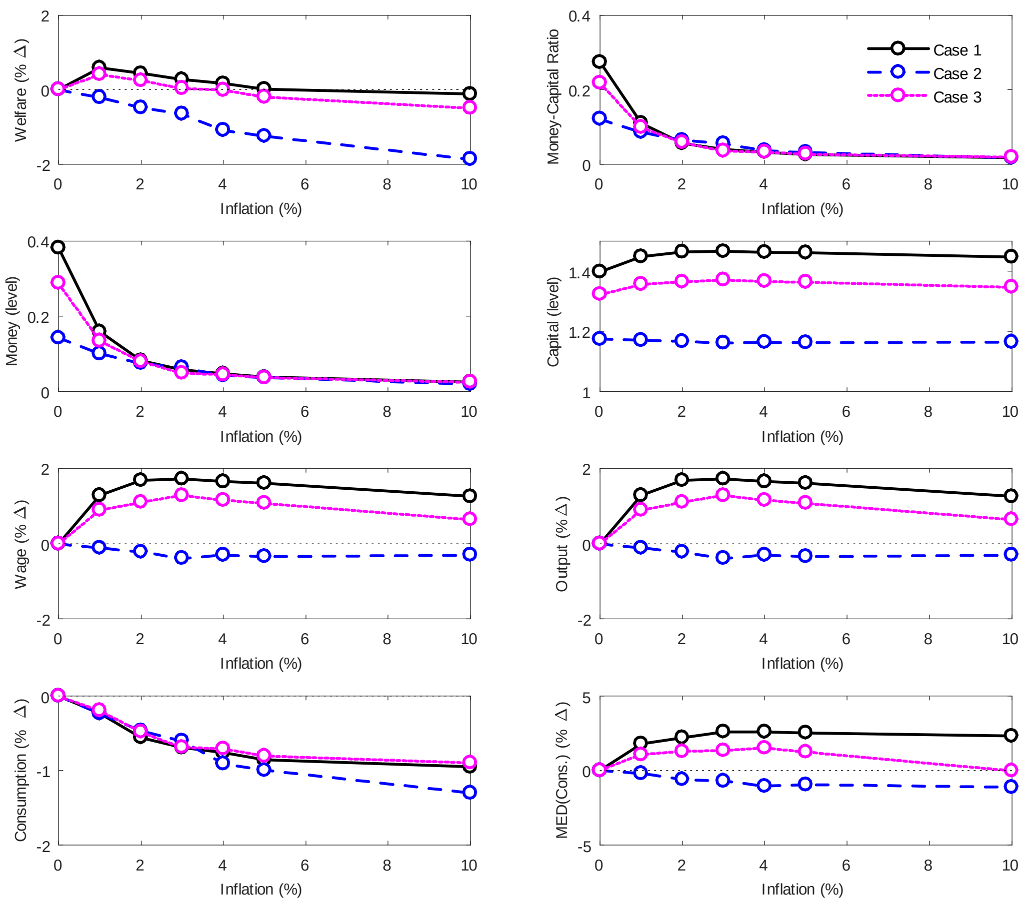

Figure 5: Multiple-Cohort Model: Calibrated Results

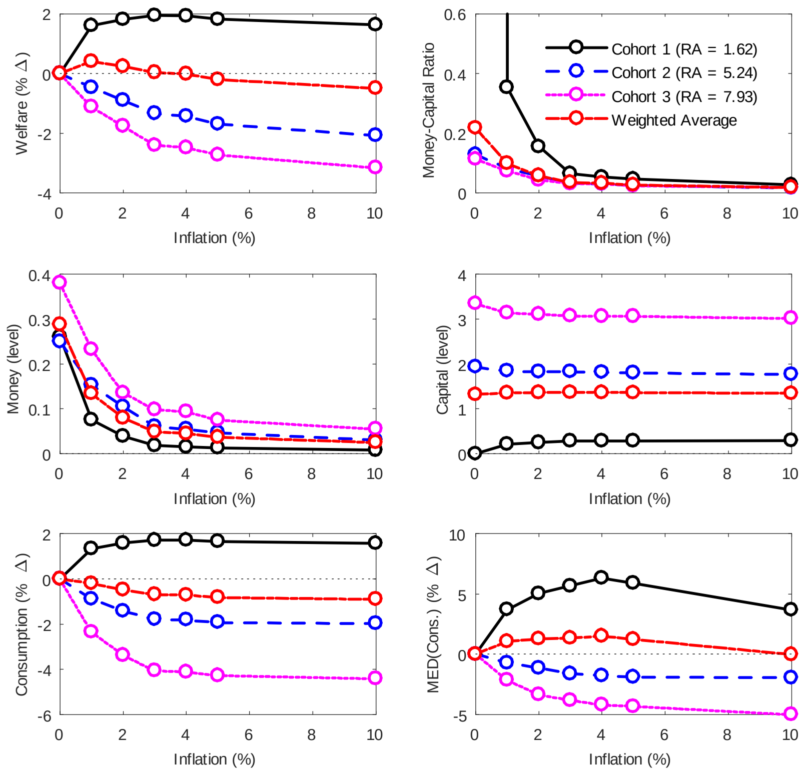

Figure 6: Multiple-Cohort Model: Individual Cohorts for Case 3

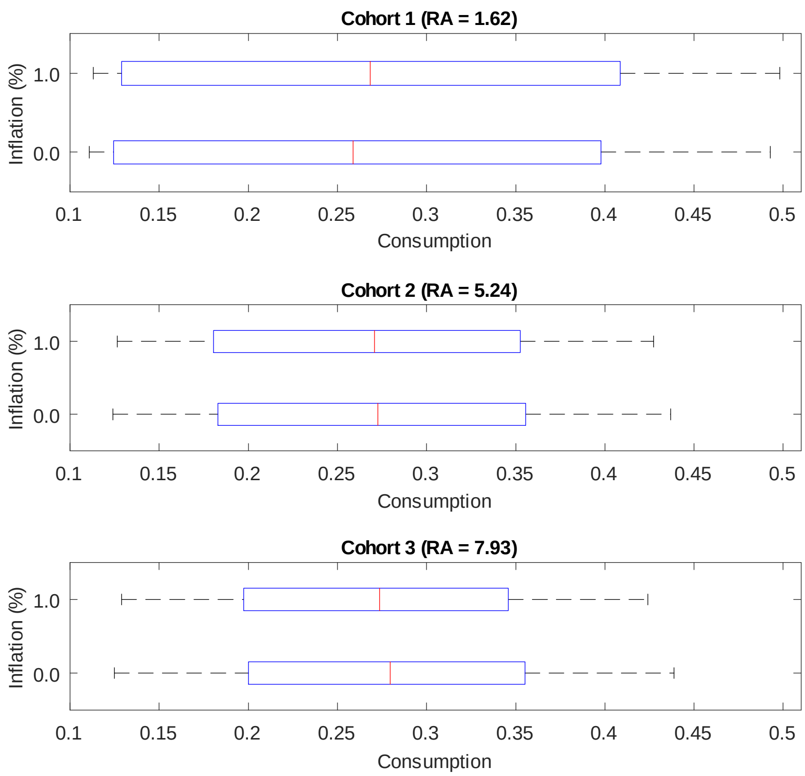

Figure 7: Multiple-Cohort Model: Individual Cohorts for Case 3

the consumption distribution, the net welfare gain from these households increasing their consumption can at times exceed the welfare loss of others. Second, results on the measures of inequality from the model were not displayed because there was no general link between the welfare implications and the impact on dimensions of inequality. Namely, the models above all predict a large increase in liquid wealth inequality and a much smaller decrease in illiquid wealth inequality in response to inflation. There is however little to no change in total wealth inequality, income inequality, and consumption inequality. While one reason for this result is that inelastic labor supply creates a large portion of household income (which is identical for all households), another reason is that the impact of inflation on inequality is netted out by some household cohorts experiencing an increase in inequality while others experience a decrease. When the cohorts are aggregated into one population, the overall changes in inequality are not directly relatable to changes in inflation.

This section concludes with a brief discussion of the robustness of the results. The comparative statics exercise performed in the single-cohort model suggests that the general mechanism behind the model is robust to changes in the key friction to the illiquid asset market as well as to the size of the preference shocks, at least for low levels of inflation. More formal sensitivity analyses were performed regarding the size of the resale value of capital (p_S) as well as the choices for the exact RA values to use from the distribution estimated using the PSID. Changes in p_S were found to not have much of an impact on the results, because while one can imagine that a lower cost of liquidating capital would place downward pressure on the liquidity position of households at zero inflation, it would also make households more willing to substitute into the illiquid asset when faced with positive inflation. When increasing the RA for the least risk averse cohorts, it was found that the liquidity position of this cohort was reduced and therefore reduced the benefits from inter-cohort transfers. However, this could be countered with adding more cohorts to the environment with even lower risk aversion that have a proportional size as in the data. While the results presented here used parameters arrived at under a great deal of conservatism and discipline, the main story goes through for a variety of parameters.

This analysis can also add to the debate on investment costs being fixed or convex. Versions of the model considering convex adjustment costs were unable to reproduce the welfare conclusions reported here for one key reason. Convex adjustment costs deliver frequent and small investment decisions, while the main mechanism of the model uses fixed adjustment costs to deliver infrequent and lumpy investment decisions. Convex adjustment costs constrain the inflation-induced shift in the liquidity position of asset portfolios, and therefore constrain the drastic difference in welfare implications due to changes in preferences. While the results presented here are not robust to the form of capital adjustment costs, Kahn and Thomas (2008) report that the majority of micro-evidence on firm-level investment behavior favors fixed costs.

4.6. Application: The Euro Area

The models thus far utilized empirical estimates for the US risk aversion distribution, and considered several possibilities for the EIS distribution due to lack of data. This section presents an application that uses various EIS estimates for members of the euro area reported by Havránek et al. (2015) to show that the general take away of the model holds when EIS levels have an empirical foundation. In other words, the cohorts described above having ex-ante heterogenous preferences but share a common monetary policy can be considered to be ex-ante heterogeneous countries participating in a common currency union.

The application considers three countries from the euro area (Germany, Italy, and Spain), as well as a single country (cohort) model using weighted average features of the countries for comparison.[18] Havránek et al. (2015) reports the average EIS of the households from each of these three countries to be 0.080, 0.290, and 0.504, and it is assumed that all households have the same RA of 2. The proportion sizes for each country were $\varphi_G = 0.4378$, $\varphi_I = 0.3154$,

[18] These three countries were chosen because they are among the most populous countries in the euro area. France was not considered due to the weighted average EIS estimate reported by Havránek et al. (2015) being negative.

and $\varphi_S = 0.2468$ using 2017 population data. The EIS for the single-country model (0.251) is a weighted average of the EIS from each country. All parameters with the exception of F and θ remain at the values reported above.

Calibrating the model for this application uses measures of monetary velocity and the median-mean ratio of the monetary distribution as in the previous sections. Households in the euro area reportedly hold larger money balances than in the US, and monetary velocity range between 0.5 and 1.5 at two percent inflation within these three countries.[19] These three countries, however, have an average median-mean ratio of the monetary distribution of 0.28 that is roughly similar to the US.[20] The multiple-country model is unfortunately unable to achieve these targets for monetary velocity and median-mean ratio by varying values for F and θ. The parameter θ was used in the above analysis to fine-tune the median-mean ratio while F was used to target velocity. A higher monetary velocity as in the euro area brings with it a higher average stock of money as well as a larger median-mean ratio. The best the multiple-country model can do is set $F = 0.0078$ and $\theta = 0.005$ in order to achieve a velocity of 2 and median-mean ratio of 0.33. In order to be consistent, the parameters of the single-country model were set to $F = 0.0162$ and $\theta = 0.10$ in order to match the same targets achieved by the multiple-country model.

The results are illustrated in Figure 8 for aggregate measures and in Figure 9 for the individual countries.[21] Note that even though these countries share the same level of RA, the same stories go through with respect to dispersion in EIS. In fact, the aggregate welfare results of the multiple country model suggest that all positive inflation rates considered are

[19] Dreger & Wolters (2009) find a velocity of around 1 for the overall euro area around 1999. Inflation in the euro area has been about 2 on average from 1991 to 2019 and was around 2 for the first decade of the euro area. Velocity calculations for Germany, Italy, and Spain in 1999 were 0.5, 1, and 1.5 using nominal GDP and M1. Nominal GDP was available through Bloomberg, and M1 was constructed from individual country statistical institutes and central banks.

[20] Data for this ratio is available through the Eurosystem Household Finance and Consumption Survey. It is the simple average of median-mean ratios for sight accounts (i.e. checking accounts) across Germany (0.25), Italy (0.17), and Spain (0.42).

[21] Note that the model assumes perfect mobility of labor and capital such that all countries share the same factor prices. While the European Single Market guaranteeing free movement for goods, services, people, and capital has been in effect since 1993, the model assumption of perfectly identical prices is counterfactual

preferred to no inflation. This is again due to the shift in the consumption distributions of all countries delivering more winners from inflation than losers. The single country version shows some welfare gains at low levels of inflation, but these gains do not persist due to the lack of inter-cohort transfers. The country-specific results presented in Figure 9 illustrate these transfers. First, Germany having the highest motive to smooth consumption displays the largest holdings of money and capital, but the lowest liquidity position at zero percent inflation. When faced with positive inflation, Germany reduces its holdings of both money and capital while the other two countries reduce their liquidity position by decreasing money holdings and increasing capital holdings. The gains in capital win out and result in an increase in overall output and wages. It follows from the previous discussions that Italy and Spain experience welfare benefits from inflation due to larger average and median consumption. Germany still experiences welfare benefits even though it experiences smaller consumption in both the average and the median. Even for the country with the highest motive to smooth consumption, there are more households benefitting from long-run inflation than not.

5. Conclusion

This paper examined the welfare implications of long-run inflation in an environment with essential money, a competing illiquid asset, and potential ex-ante heterogeneity of households with respect to their behavioral measures of risk aversion and elasticity of intertemporal substitution. The results show that the relative liquidity position of households' portfolios as well as potential inter-cohort transfers of resources can deliver fewer welfare costs to inflation than has been previously reported, and in some instances net welfare benefits to low levels of positive inflation. While not being definitive on the long-standing topic of the welfare costs of inflation, the contribution of this paper supports that households' attitudes towards consumption smoothing and risk are crucial when assessing the welfare

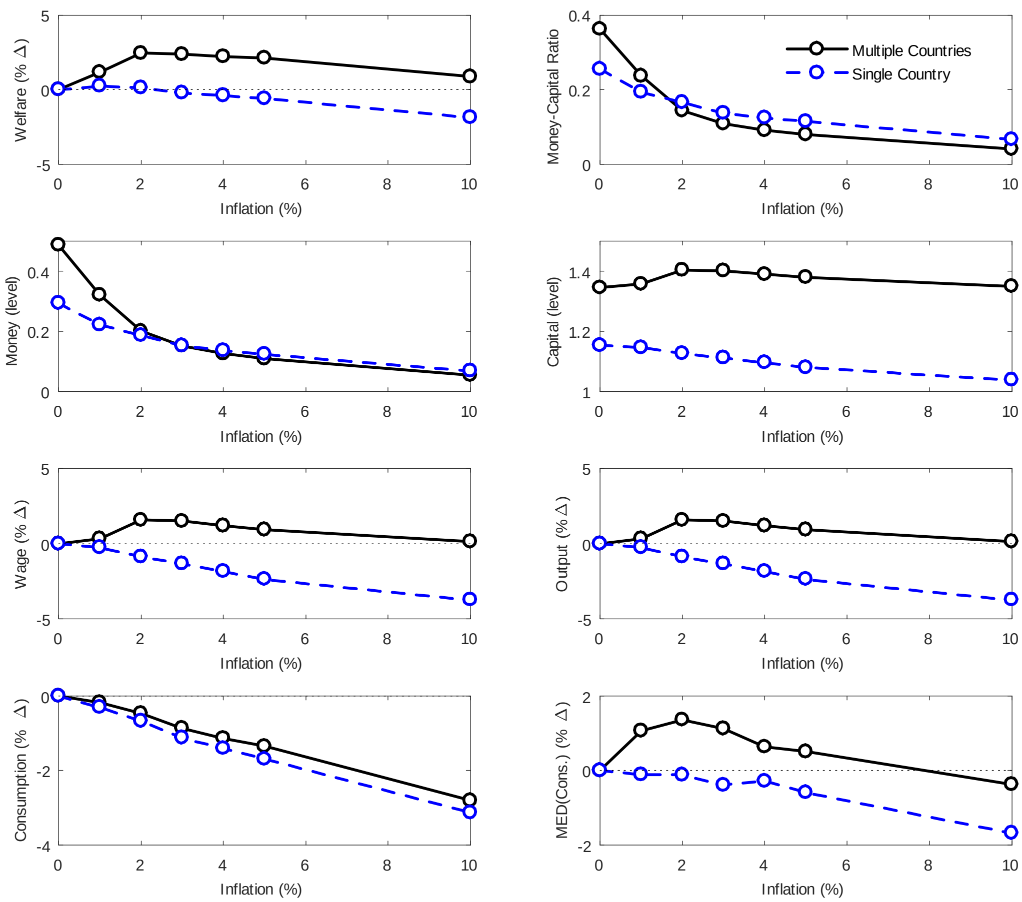

Figure 8: Multiple-Country Model: Aggregate Results

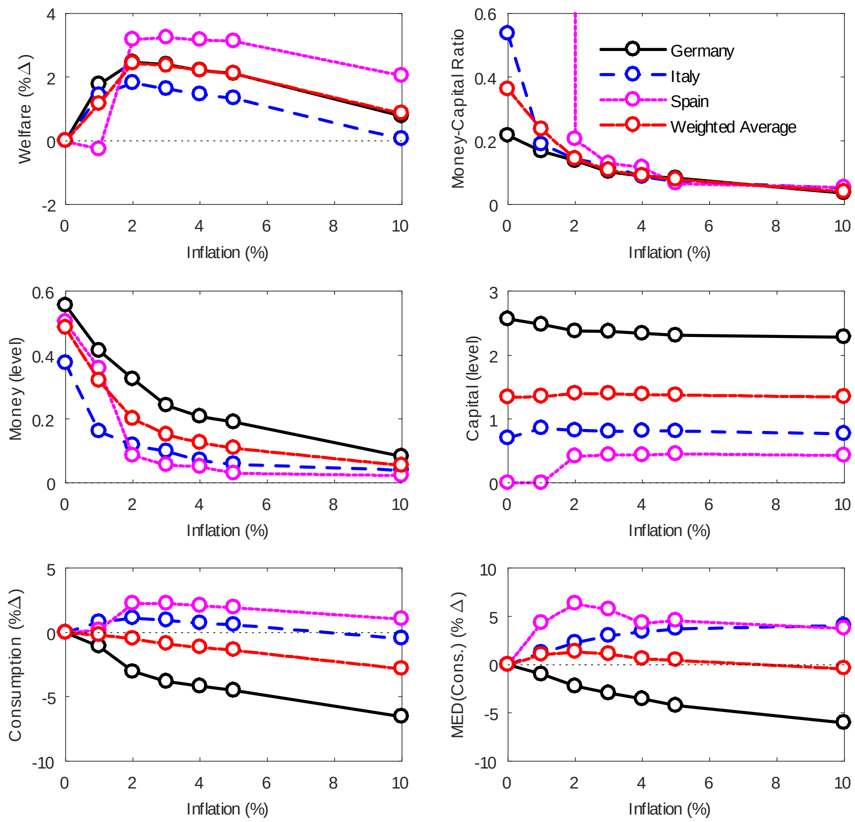

Figure 9: Multiple-Country Model: Individual Countries

implications of long-run monetary policies.

The model made two assumptions which may not be innocuous. First, the model assumed that households inelastically supplied their labor endowment. While this assumption allowed for the model to have essential money yet be as parsimonious as possible, it is not clear if allowing for elastically supplied labor will change the results. On one hand, households would be able to increase labor supply when faced with positive inflation, so an asset portfolio with a high liquidity position might not be as important. However, increasing labor in response to inflation would fuel the increase in output as well as make the illiquid asset more attractive. One way of addressing this issue would be to consider multiple cohorts in the model developed by Wen (2015), that combines essential money and elastically supplied labor through a specific timing realization of shocks and labor supply decisions. The second assumption made in the environment here is that all households receive an identical monetary transfer under positive inflation. While a standard assumption in most monetary models, it is again not clear if this is an important driver of the welfare results. There are many ways one can introduce monetary injections into the proposed environment, such as distributing new money to those active (or inactive) with the capital market, or proportional to their levels of capital, consumption, or even money balances. The last option is of interest because most models have shown that money injections proportional to a household's money balances have no real effects. This only holds when all households are holding *positive* money balances. If some households hold an entirely illiquid asset portfolio, then proportional monetary injections could have real effects. These issues are left for future research.

References

[1] Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. The Quarterly Journal of Economics, 109 (3), 659–684.

[2] Barsky, R. B., House, C. L. and Kimball, M. S. (2007). Sticky-Price Models and Durable Goods. American Economic Review, 97 (3), 984–998.

[3] Best, M. C., Cloyne, J., Ilzetzki, E. and Kleven, H. (2017). Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches. Tech. rep., Manuscript.

[4] Camera, G. and Chien, Y. (2014). Understanding the Distributional Impact of Long-Run Inflation. Journal of Money, Credit and Banking, 46 (6), 1137–1170.

[5] Cao, S., Meh, C. A., Ríos-Rull, J.-V., Terajima, Y. et al. (2018). The Welfare Cost of Inflation Revisited: The Role of Financial Innovation and Household Heterogeneity. Tech. rep., Bank of Canada.

[6] Chetty, R. (2006). A New Method of Estimating Risk Aversion. American Economic Review, 96 (5), 1821–1834.

[7] Chiappori, P.-A. and Paiella, M. (2011). Relative Risk Aversion is Constant: Evidence from Panel Data. Journal of the European Economic Association, 9 (6), 1021–1052.

[8] Chien, Y., Cole, H. and Lustig, H. (2016). Implications of Heterogeneity in Preferences, Beliefs and Asset Trading Technologies in an Endowment Economy. Review of Economic Dynamics, 20, 215–239.

[9] Chiu, J., & Molico, M. (2011). Uncertainty, Inflation, and Welfare. Journal of Money, Credit and Banking, 43, 487-512.

[10] Cooper, R. W. and Haltiwanger, J. C. (2006). On the Nature of Capital Adjustment Costs. The Review of Economic Studies, 73 (3), 611–633.

[11] Cozzi, M. (2014). Risk Aversion Heterogeneity, Risky Jobs and the Distribution of Wealth. Tech. rep., Working Paper.

[12] Díaz-Giménez, J., Glover, A. and Ríos-Rull, J.-V. (2011). Facts on the Distributions of Earnings, Income, and Wealth in the United States: 2007 update. Federal Reserve Bank of Minneapolis Quarterly Review, 34 (1), 2–31.

[13] Dreger, C. and Wolters, J., (2009). Money velocity and asset prices in the euro area. Empirica, 36(1), 51-63.

[14] Dressler, S. (2016). A Long-Run, Short-Run, and Politico-Economic Analysis of the Welfare Costs of Inflation. Journal of Macroeconomics, 47, 255–269.

[15] Epstein, L. and Zin, S. E. (1989). Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: a Theoretical Framework. Econometrica, 57, 937–969.

[16] Erosa, A. and Ventura, G. (2002). On Inflation as a Regressive Consumption Tax. Journal of Monetary Economics, 49 (4), 761–795.

[17] Floden, M. and Lindé, J. (2001). Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance? Review of Economic Dynamics, 4 (2), 406–437.

[18] Gomes, F. A. R. and Paz, L. S. (2013). Estimating the Elasticity of Intertemporal Substitution: Is the Aggregate Financial Return Free from the Weak Instrument Problem? Journal of Macroeconomics, 36, 63–75.

[19] Hall, R. E. (1988). Intertemporal Substitution in Consumption. Journal of Political Economy, 96 (2), 339–357.

[20] Harrison, G. W., Lau, M. I. and Rutström, E. E. (2007). Estimating Risk Attitudes in Denmark: A Field Experiment. Scandinavian Journal of Economics, 109 (2), 341–368.

[21] Havránek, T., Horvath, R., Irsova, Z. and Rusnak, M. (2015). Cross-country heterogeneity in intertemporal substitution, Journal of International Economics, 96 (1), 100-118.

[22] Khan A., Thomas J.K. (2008). Adjustment Costs. In: Durlauf S.N., Blume L.E. (eds) The New Palgrave Dictionary of Economics. Palgrave Macmillan, London

[23] Kimball, M. S., Sahm, C. R. and Shapiro, M. D. (2008). Imputing Risk Tolerance from Survey Responses. Journal of the American Statistical Association, 103 (483), 1028–1038.

[24] —, — and — (2009). Risk Preferences in the PSID: Individual Imputations and Family Covariation. American Economic Review, 99 (2), 363–68.

[25] Lucas, R. E. (2000). Some Macroeconomics for the 21st Century. Journal of Economic Perspectives, 14 (1), 159–168.

[26] Lucas Jr, R. E. (2003). Macroeconomic Priorities. American Economic Review, 93 (1), 1–14.

[27] Mehra, R. and Prescott, E. C. (1985). The Equity Premium: A Puzzle. Journal of Monetary Economics, 15 (2), 145–161.

[28] Molico, M. (2006). The Distribution of Money and Prices in Search Equilibrium. International Economic Review, 47 (3), 701–722.

[29] Rouwenhorst, G. K. (1995). Asset Pricing Implications of Equilibrium Business Cycle Moments. In T. Cooley (ed.), Frontiers of Business Cycle Research, 10, Princeton University Press.

[30] Swanson, E. (2018). Risk Aversion, Risk Premia, and the Labor Margin with Generalized Recursive Preferences. Review of Economic Dynamics, 28, 290–321.

[31] Swanson, E. T. (2012). Risk Aversion and the Labor Margin in Dynamic Equilibrium Models. American Economic Review, 102 (4), 1663–91.

[32] Von Gaudecker, H.-M., Van Soest, A. and Wengstrom, E. (2011). Heterogeneity in Risky Choice Behavior in a Broad Population. American Economic Review, 101 (2), 664–94.

[33] Weil, P. (1989). The Equity Premium Puzzle and the Risk-free Rate Puzzle. Journal of Monetary Economics, 24 (3), 401–421.

[34] Wen, Y. (2015). Money, Liquidity and Welfare. European Economic Review, 76, 1–24.

A Calculation of Welfare

Quantifying the welfare implications of long-run monetary policy follows Camera and Chien (2014), who define a *compensated* value function and then calculate its value via iteration. This calculation is done for each cohort independently, so welfare costs are calculated as a percentage of that cohort's own consumption. Aggregate welfare costs are calculated as a fraction of total consumption similar to the age-cohort welfare aggregation performed by Cao et al. (2018).

Let $\omega := (m, k, h, \theta)$ denote the state vector. Fix $\pi = 0$ and define $V_{j0}(\omega)$ as the value function for a cohort-j household with state vector under zero inflation. Using the distributions of asset holdings for each cohort j under zero inflation $\left(\phi_{j0}(\omega)\right)$, average welfare under zero inflation for cohort j (W_{j0}) and given by

$$W_{j0} := \sum_{h \in H} \sum_{\theta \in \Theta} \int_{m \in M} \int_{k \in K} V_{j0}(\omega) \phi_{j0}(\omega) dmdk. \tag{6}$$

The welfare cost of inflation $\pi > 0$ is a standard compensating variation measure. It is the percentage adjustment in consumption that leaves the household indifferent, ex ante, between inflation π and zero inflation. Define $c_{j\pi}(\omega)$, $m'_{j\pi}(\omega)$, and $k'_{j\pi}(\omega)$ to be the optimal decisions of the cohort-j household with state ω under π inflation. Let $\omega'_{j\pi} := \left(m'_{j\pi}(\omega), k'_{j\pi}(\omega), h', \theta'\right)$ denote the state resulting from current decisions $m'_{j\pi}(\omega)$ and $k'_{j\pi}(\omega)$ for a cohort-j household. Given that consumption is adjusted by the proportion $\Delta_{j\pi}$ each period, define the compensated value function $\hat{V}_{j\pi}(\omega)$ by

$$\hat{V}_{j\pi}(\omega) := (1+\theta) \frac{\left(\Delta_{j\pi} c_{j\pi}(\omega)\right)^{1-\gamma_j}}{1-\gamma_j} - \beta \left[E\left(-\hat{V}_{j\pi}\left(\omega'_{j\pi}\right)\right)^{1-\epsilon_j}\right]^{\frac{1}{1-\epsilon_j}}. \tag{7}$$

For a proposed compensation value $\Delta_{j\pi}$, the compensated value function can be determined using the cohort-j household's decision rules via iteration.[22] Once the compensated value functions are calculated over all possible states and all household types, average welfare under inflation π is defined as

$$W_{j\pi} := \sum_{h \in H} \sum_{\theta \in \Theta} \int_{m \in M} \int_{k \in K} \hat{V}_{j\pi}(\omega) \phi_{j\pi}(\omega) dmdk. \tag{8}$$

[22] Note that since the decision rules are taken as given, this is not a maximization step of a value function iteration. This can be thought of along the lines as a Howard updating iteration that is continued up to convergence to a specified tolerance.

Since $W_{j\pi}$ is a function of $\Delta_{j\pi}$, the last step involves finding the correct amount of compensation such that cohort-j households are indifferent between zero inflation and $\pi > 0$ inflation, on average $(W_{j0} = W_{j\pi})$. The resulting welfare costs are given in percentage terms by $\Delta^*_{j\pi} = (\Delta_{j\pi} - 1) \times 100\%$. When $\Delta^*_{j\pi} > 0$, the every cohort-j household is indifferent between zero inflation and π inflation after being compensated with $\Delta^*_{j\pi}$ percent more consumption on average.

When the economy is only populated by one cohort, the welfare calculation above is total welfare. When the economy features multiple cohorts, total welfare must be calculated as a fraction of total consumption. Let c^T_j denote total consumption of cohort j, and c^T denote total consumption aggregated across cohorts. Welfare as a percentage of total consumption is given by

$$\Delta^*_{\pi} = \frac{\sum_j \Delta^*_{j\pi} c^T_j}{c^T} \tag{9}$$