NBER WORKING PAPER SERIES

MONETARY POLICY AND THE REDISTRIBUTION CHANNEL

Adrien Auclert

Working Paper 23451
http://www.nber.org/papers/w23451

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2017

This paper is a revised version of Chapter 1 of my PhD dissertation at MIT. I cannot find enough words to thank my advisors Iván Werning, Robert Townsend and Jonathan Parker for their continuous guidance and support. I also thank many seminar participants for their insights. I have particularly benefited from the detailed comments of the editor (John Leahy), four anonymous referees, as well as Eduardo Dávila, Gauti Eggertsson, Xavier Gabaix, Adam Guren, Gregor Jarosch, Greg Kaplan, Guido Lorenzoni, Ben Moll, Makoto Nakajima, Matthew Rognlie, Yoko Shibuya, Alp Simsek, Christian Stoltenberg and Daan Struyven. Filippo Pallotti provided outstanding research assistance, and the Macro-Financial Modeling Group provided generous financial support. All remaining errors are my own. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Monetary Policy and the Redistribution Channel
Adrien Auclert
NBER Working Paper No. 23451
May 2017
JEL No. D31,D52,E21,E52

ABSTRACT

This paper evaluates the role of redistribution in the transmission mechanism of monetary policy to consumption. Three channels affect aggregate spending when winners and losers have different marginal propensities to consume: an earnings heterogeneity channel from unequal income gains, a Fisher channel from unexpected inflation, and an interest rate exposure channel from real interest rate changes. Sufficient statistics from Italian and U.S. data suggest that all three channels are likely to amplify the effects of monetary policy. A standard incomplete markets model can deliver the empirical magnitudes if assets have plausibly high durations but a counterfactual degree of inflation indexation.

Adrien Auclert
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305-6072
and NBER
aauclert@stanford.edu

An online appendix is available at http://www.nber.org/data-appendix/w23451

1 Introduction

There is a conventional view that redistribution is a side effect of monetary policy changes, separate from the issue of aggregate stabilization which these changes aim to achieve. Most models of the monetary policy transmission mechanism implicitly adopt this view by featuring a representative agent. By contrast, in this paper I argue that redistribution is a *channel* through which monetary policy affects macroeconomic aggregates, because those who gain from accommodative monetary policy have higher marginal propensities to consume (MPCs) than those who lose. The simple argument goes back to Tobin (1982):

> Aggregation would not matter if we could be sure that the marginal propensities to spend from wealth were the same for creditors and for debtors. But [...] the population is not distributed between debtors and creditors randomly. Debtors have borrowed for good reasons, most of which indicate a high marginal propensity to spend from wealth or from current income.

In this paper, I use consumer theory to refine Tobin's intuitions about aggregation. My analysis clarifies who gains and who loses from monetary policy changes, as well as the effect on aggregate consumption. Monetary expansions tend to increase real incomes, to raise inflation and to lower real interest rates. Not everyone is equally affected by these changes. This generates three distinct sources of redistribution.

First, monetary expansions induce gains in aggregate earnings from labor and profits. The distribution of these gains is unlikely to be equal: some agents benefit disproportionately, and conversely, some lose in relative terms. This is the *earnings heterogeneity channel* of monetary policy.

Second, unexpected inflation revalues nominal balance sheets, with nominal creditors losing and nominal debtors gaining: this is the *Fisher channel*, which has a long history in the literature since Fisher (1933). This channel has been explored by Doepke and Schneider (2006), who measure the balance sheet exposures of various sectors and groups of households in the United States to different inflation scenarios. Net nominal positions (NNPs) quantify the exposures to unexpected increases in the price level.

Real interest rate falls create a third, more subtle form of redistribution. These falls increase financial asset prices. But it is incorrect to claim that asset holders generally benefit: instead, we have to consider whether their assets have longer durations than their liabilities. Importantly, liabilities include consumption plans, and assets include human capital. Unhedged interest rate exposures (UREs)—the difference between all maturing assets and liabilities at a point in time—are the correct measure of households' balance-sheet exposures to real interest rate changes, just like net nominal positions are for price level changes. For example, agents whose financial wealth is primarily invested in short-term certificates of deposit tend to have positive UREs, while those with large long-term bond investments or adjustable-rate mortgage liabilities tend to have negative UREs. Real interest rate falls redistribute away from the first group towards the second group: this is what I call the *interest rate exposure channel*.

In this paper, I show how these three redistribution channels affect the transmission mechanism of monetary policy to consumption. My main theoretical result decomposes the consumption effect of a transitory change in monetary policy into a contribution from each of these channels, together with an *aggregate income* and a *substitution channel*. Representative-agent models only feature the latter two. My theorem shows that redistribution amplifies these effects, provided that winners from monetary expansions have higher MPCs than losers. The rest of the paper argues that this is likely to be the case in practice. Hence, the redistributive effects of monetary policy are important to understand its aggregate effect.[1]

In the first part of the paper, I establish my main decomposition by studying a general aggregation problem. In partial equilibrium, I consider an optimizing agent with a given initial balance sheet, who values nondurable consumption and leisure, and is subject to a transitory change in income, inflation and the real interest rate. I decompose his consumption response into a substitution effect and a wealth effect, and show that the latter is the product of his MPC out of income and a balance-sheet revaluation term in which NNPs and UREs appear. This result is robust to the presence of durable goods, incomplete markets, idiosyncratic risk, and (certain kinds of) borrowing constraints. In other words, the MPC out of a windfall income transfer is a key determinant of the response of optimizing consumers to inflation– or real interest rate–induced changes in their balance sheets. This result generalizes previous findings by Kimball (1990) on the importance of MPCs in incomplete-markets consumption models.

I then sum across the individual-level predictions and exploit the fact that financial assets and liabilities net out in general equilibrium to obtain the first-order response of aggregate consumption to simultaneous transitory shocks to output, inflation, and the real interest rate. This response is the sum of five terms, reflecting the contributions from the two aggregate and the three redistributive channels mentioned above. Moreover, the magnitudes of the redistributive channels are given by sufficient statistics: the cross-sectional covariances between MPCs and exposures to each aggregate shock. Since the pioneering work of Harberger (1964), sufficient statistics have been used in public finance to evaluate

[1] My theorem applies to a broad class of general equilibrium models with heterogeneous agents, so it can be used to understand consumption in other contexts than that of monetary policy. At the same time, I am leaving a number of redistributive channels out of my analysis. First, I abstract away from aggregate risk, so cannot handle changes in risk premia, as in Brunnermeier and Sannikov (2016). Second, I do not model limited participation, so monetary policy cannot differentially affect participants and nonparticipants, as in the studies of Grossman and Weiss (1983), Rotemberg (1984) and others. Finally, since I assume that all assets are remunerated at the risk-free rate, my analysis does not address the unequal incidence of inflation due to larger cash holdings by the poor (Erosa and Ventura 2002; Albanesi 2007). These are all interesting dimensions along which the theory could be extended.

the welfare effect of hypothetical policy changes in a way that is robust to the specifics of the underlying structural model (see Chetty 2009 for a survey). Mine are useful to evaluate the impact of hypothetical changes in macroeconomic aggregates on *aggregate consumption* in a similarly robust way. All that is required is information on household balance sheets, income and consumption levels, and their MPCs.

By further assuming that the elasticity of intertemporal substitution σ and the elasticity of relative income to aggregate income γ are constant in the population, I obtain a set of five estimable moments that summarize all we need to know about agents' heterogeneity to recover the aggregate elasticities of consumption to the real interest rate, the price level, and aggregate income. Contrary to σ (and perhaps γ), these sufficient statistics are not structural parameters: they are likely to vary over time and across countries. I set out to measure them in three separate surveys, covering different time periods, countries, and methods from the literature. I use a 2010 Italian survey containing a self-reported measure of MPC (Jappelli and Pistaferri 2014); the 1999-2013 waves of the U.S. Panel Survey of Income Dynamics, together with semi structural approach to identify the MPC out of transitory income shocks (Blundell, Pistaferri and Preston 2008); and the 2001–2002 waves of the U.S. Consumer Expenditure Survey, together with a method that exploits the randomized timing of tax rebates as a source of identification for MPC (Johnson, Parker and Souleles 2006).

Consider first the elasticity of consumption to the real interest rate. In a representative-agent world, this elasticity is due to intertemporal substitution. It is negative, and its magnitude depends on σ. I define a method for measuring UREs, and show that, in each of my three datasets, their covariance with MPCs is also negative. Through the lens of my theorem, this implies that the interest rate exposure channel acts in the same direction as the substitution channel, and with comparable magnitude provided that σ is around 0.1-0.2. Hence representative-agent analyses that abstract from redistribution fail to capture an important reason why real interest rates affect consumption.

Similarly, across datasets, the covariance between MPCs and NNPs is negative. This implies that consumption tends to rise with inflation as a result of the Fisher channel. However, when cast in terms of elasticities, the magnitude is small: an unexpected permanent 1% increase in the price level raises consumption today by no more than 0.1%. This suggests that nominal redistribution could be important in explaining why aggregate consumption increases in monetary expansions, though its contribution is likely to be modest.

Finally, in line with previous literature, I estimate the covariance between MPCs and incomes to be negative in the data. If, in addition, low-income agents disproportionately benefit from increases in aggregate income—as suggested by Coibion et al. 2016—the earnings heterogeneity channel also amplifies the effects of monetary policy. Future work can build on my results by providing more precise estimates of these sufficient statistics, as well as keeping track of their evolution over time.

A nascent literature analyzes the effects of monetary policy in dynamic stochastic general equilibrium models with rich heterogeneity, matching various aspects of the cross-section such as the wealth distribution. Prominent examples include Gornemann, Kuester and Nakajima (2012), McKay, Nakamura and Steinsson (2016), and Kaplan, Moll and Violante (2016). These structural models overcome a number of important limitations of my sufficient statistics approach. They can study the role of investment, analyze the precise interaction between monetary and fiscal policy, and explore the effect of shocks that are persistent and/or announced in advance. Yet, as highlighted by Kaplan, Moll and Violante (2016), a version of my main decomposition survives in these more complex models, shedding light on the underlying mechanisms. Moreover, sufficient statistics can discipline the construction of these models. By making sure that the model's sufficient statistics match the data, researchers can ensure that, even if the model is misspecified, its predictions for the response of consumption to shocks are consistent with the empirical evidence.

I illustrate this procedure by considering the sufficient statistics generated by a standard partial-equilibrium incomplete markets model, similar to the one used as a building block by the literature. Mine is a Bewley-Huggett-Aiyagari model with nominal, long-term, circulating private IOUs (as in Huggett 1993). Such a model features rich heterogeneity in MPCs, UREs, NNPs and incomes. I calibrate it to the U.S. economy and quantitatively evaluate, in its steady state, the size of my sufficient statistics. This exercise delivers three main insights.

First, in the model, the interest rate exposure channel has the same sign and comparable magnitude as it does in the data. However, this result relies crucially on long asset durations. If instead all assets are short term (a typical assumption in the literature), changes in real interest rates have very large redistributive effects. The intuition is as follows: under a shorter maturity structure, debtors—the high-MPC agents in the economy—roll over a larger fraction of their liabilities each period, and their consumption plans are therefore very sensitive to changes in those rates. The role of asset durations I uncover holds under any degree of shock persistence. It is consistent with the results of Calza, Monacelli and Stracca (2013), who find that consumption reacts much more strongly to identified monetary policy shocks in countries where mortgages predominantly have adjustable rates.[2]

Second, I find that the benchmark calibration of the model in which all assets are nominal displays a Fisher channel with the same sign as in the data, but with a much larger magnitude. This is because inflation redistributes along the asset dimension, which in this class of models is highly correlated with MPC. As a result, Bewley models with nominal

[2]See also Rubio (2011) and Garriga, Kydland and Sustek (2016).

assets tend to overstate the correlation between MPCs and NNPs that exists in the data. A model with real assets, or in which assets have a high degree of inflation indexation, is more consistent with the empirical evidence.[3]

Finally, in the model with short-term debt, changes in real interest rates have asymmetric effects. The sufficient statistic approach correctly predicts the effect of any increase in the real rate, but it overpredicts in the other direction. This asymmetry comes from the differential response of borrowers at their credit limit to rises and falls in income: while these borrowers save an important fraction of the gains they get from low interest rates, they are forced to cut spending steeply when interest rates rise. This could help explain the empirical finding that interest rate hikes tend to lower output by more than falls increase it (Cover 1992; de Long and Summers 1988; Tenreyro and Thwaites 2016).

This paper is motivated by an an extensive empirical literature documenting that MPCs are large and heterogenous in the population (see Jappelli and Pistaferri 2010 for a survey), and that they depend on household balance sheet positions.[4] Recently, di Maggio et al. (2017) have measured the consumption response of households to changes in the interest rates they pay on their mortgages. My theory shows that this paper quantifies an important leg of the redistribution channel of monetary policy.

Several papers have focused on the redistributive channels of monetary policy I highlight in isolation. Coibion et al. (2016) propose an empirical evaluation of the earnings heterogeneity channel by measuring how identified monetary policy shocks affect income inequality in the Consumer Expenditure Survey. The Fisher channel has received a great deal of attention in the literature following the work of Doepke and Schneider (2006). For example, on the normative side, Sheedy (2014) asks when the central bank should exploit its influence on the price level to ameliorate market incompleteness over the business cycle. On the positive side, Sterk and Tenreyro (2015) show that the Fisher channel can be a source of effects of monetary policy under flexible prices in a non-Ricardian model. The interest rate exposure channel has, by contrast, not received much attention in the context of monetary policy.[5]

The importance of MPC differences in the determination of aggregate demand is well understood by the theoretical literature on fiscal transfers.[6] MPC differences between borrowers and savers, in particular, have been explored as a source of aggregate effects from shocks to asset prices or to borrowing constraints.[7] In Farhi and Werning (2016b),

[3]My model also replicates the empirical covariance between MPCs and incomes but, since I assume that incomes are exogenous, I stop short of a full assessment of the earnings heterogeneity channel. A previous version of this paper examined this channel in general equilibrium.

[4]See for example Mian, Rao and Sufi 2013; Mian and Sufi 2014; Baker 2017 and Jappelli and Pistaferri 2014.

[5]Redistribution through real interest rates does play a prominent role, for example, in Bassetto (2014)'s study of optimal fiscal policy or in Costinot, Lorenzoni and Werning (2014)'s study of dynamic terms of trade manipulation.

[6]See Galí, López-Salido and Vallés 2007; Oh and Reis 2012; Farhi and Werning 2016a; McKay and Reis 2016.

[7]See King 1994; Eggertsson and Krugman 2012; Guerrieri and Lorenzoni 2015; Korinek and Simsek 2016.

MPCs enter as sufficient statistics for optimal macro-prudential interventions under nominal rigidities. None of these studies, however, focus on the role of MPC differences in generating aggregate effects of monetary policy.

The remainder of the paper is structured as follows. Section 2 presents a partial equilibrium decomposition of consumption responses to shocks into substitution and wealth effects. Section 3 provides my aggregation result and discusses the monetary policy transmission mechanism with and without heterogeneity. Section 4 assesses the quantitative magnitudes of each of my redistribution channels by measuring sufficient statistics in three surveys. Finally, section 5 compares the sufficient statistics from the data to those of a Huggett model, shedding light on their structural determinants. Section 6 concludes.

2 Household balance sheets and wealth effects

In this section, I show how households' balance sheets shape their consumption and labor supply adjustments to a transitory macroeconomic shock. I first highlight the forces at play in a life-cycle labor supply model (Modigliani and Brumberg 1954; Heckman 1974) featuring perfect foresight and balance sheets with an arbitrary maturity structure. Balance sheet revaluations and marginal propensities to consume and work play a crucial role in determining both the welfare and the wealth effects of the shock (theorem 1). Under certain conditions, the positive results from theorem 1 survive the addition of idiosyncratic income uncertainty (theorem 2) and therefore apply to a large class of microfounded models of consumption behavior.

2.1 Perfect-foresight model

Consider a household with separable preferences over nondurable consumption $\{c_t\}$ and hours of work $\{n_t\}$.[8] I assume no uncertainty for simplicity: the same insights obtain when markets are complete, except with respect to the unanticipated initial shock. The household is endowed with a stream of real unearned income $\{y_t\}$. He has perfect foresight over the general level of prices $\{P_t\}$ and the path of his nominal wages $\{W_t\}$, and holds long-term nominal and real contracts. Time is discrete, but the horizon may be finite or infinite, so I do not specify it in the summations. The agent solves the following utility

[8] I present results for separable preferences because expressions for substitution elasticities take simple and familiar forms in this case, but many of my results extend to arbitrary non satiable preferences (see Appendix A.3). I assume that both u and v are increasing and twice continuously differentiable, with u concave and v convex.

maximization problem:

$$\begin{aligned}\max \quad & \sum_t \beta^t \left\{u\left(c_t\right) - v\left(n_t\right)\right\} \\ \text{s.t.} \quad & P_t c_t = P_t y_t + W_t n_t + \left({}_{t-1}B_t\right) + \sum_{s\geq 1}\left({}_tQ_{t+s}\right)\left({}_{t-1}B_{t+s} - {}_tB_{t+s}\right) \\ & \qquad + P_t\left({}_{t-1}b_t\right) + \sum_{s\geq 1}\left({}_tq_{t+s}\right)P_{t+s}\left({}_{t-1}b_{t+s} - {}_tb_{t+s}\right) \end{aligned} \tag{1}$$

The flow budget constraint (1) views the consumer, in every period t, as having a portfolio of zero coupon bonds inherited from period $t-1$, and determining consumption c_t, labor supply n_t, as well as the portfolio of bonds he chooses to carry into the next period.[9] Specifically, ${}_tQ_{t+s}$ is the time-t price of a nominal zero-coupon bond paying at $t+s$, ${}_tq_{t+s}$ the price of a real zero-coupon bond, and ${}_tB_{t+s}$ (respectively ${}_tb_{t+s}$) denote the quantities purchased. This asset structure is the most general one that can be written for this dynamic environment with no uncertainty. To keep the problem well-defined, I assume that the prices of nominal and real bonds prevent arbitrage profits. This implies a Fisher equation for the nominal term structure:

$${}_tQ_{t+s} = \left({}_tq_{t+s}\right)\frac{P_t}{P_{t+s}} \quad \forall t,s$$

I focus on the period $t=0$. The environment allows for a very rich description of the household's initial holdings of financial assets, denoted by the consolidated claims, nominal $\{{}_{-1}B_t\}_{t\geq 0}$ and real $\{{}_{-1}b_t\}_{t\geq 0}$, due in each period. The former could represent deposits, long-term bonds and most typical mortgages. The latter could represent stocks (which here pay a riskless real dividend stream and therefore are priced according to the risk-free discounted value of this stream), inflation-indexed government bonds, and price-level adjusted mortgages. I write the real wage at t as $w_t \equiv \frac{W_t}{P_t}$, the initial real term structure as $q_t \equiv \left({}_0q_t\right)$, and impose the present-value normalization $q_0 = 1$.

Using either a terminal condition if the economy has finite horizon, or a transversality condition if the economy has infinite horizon, the flow budget constraints consolidate into an intertemporal budget constraint:

$$\sum_{t\geq 0} q_t c_t = \underbrace{\sum_{t\geq 0} q_t\left(y_t + w_t n_t\right)}_{\omega^H} + \underbrace{\sum_{t\geq 0} q_t\left(\left({}_{-1}b_t\right) + \left(\frac{{}_{-1}B_t}{P_t}\right)\right)}_{\omega^F} \equiv \omega \tag{2}$$

Equation (2) states that the present value of consumption must be equal to wealth ω: the sum of human wealth ω^H (the present value of all future income) and financial wealth ω^F. Since $\{{}_{-1}B_t\}$ and $\{{}_{-1}b_t\}$ only enter (2) through ω^F, it follows that financial assets with the same initial present value deliver the same solution to the consumer problem. For in-

[9] He may, of course, just decide to roll over his position from the previous period. This corresponds to the costless trade that sets ${}_{t-1}b_{t+s} = {}_tb_{t+s}$ and ${}_tB_{t+s} = {}_{t-1}B_{t+s}$ for all s.

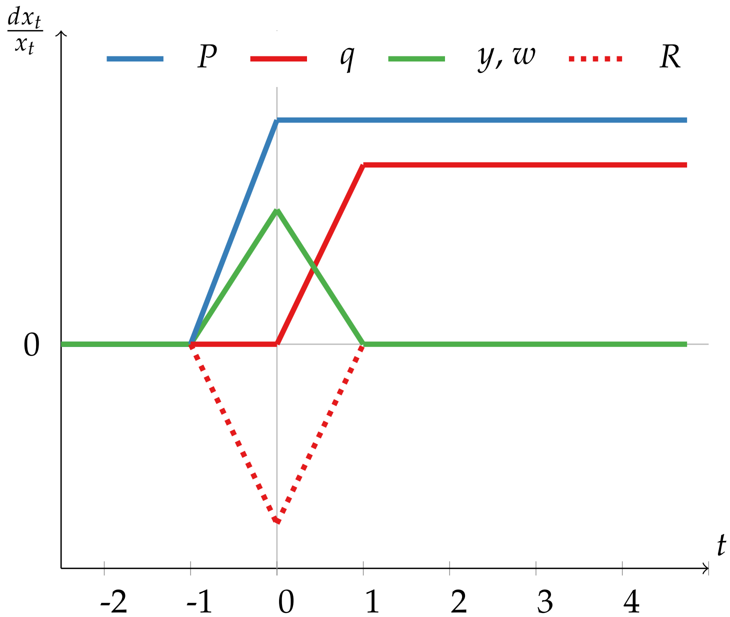

Figure 1: The experiment

stance, this framework predicts that a household with an adjustable-rate mortgage (ARM), with ${}_{-1}B_0 = -L$, chooses the same plan for consumption and labor supply as an otherwise identical household with a fixed-rate mortgage (FRM), ${}_{-1}B_t = -M$ for $t = 0 \ldots T$, provided the two mortgages have the same outstanding principal, i.e. $L = \sum_{t=0}^{T} Q_t M$. In this sense, the composition of balance sheets is irrelevant. But this composition matters following a shock, as the next section shows.

2.2 Adjustment after a transitory shock

I now consider an exercise where, keeping balance sheets fixed at $\{{}_{-1}B_t\}_{t\geq 0}$ and $\{{}_{-1}b_t\}_{t\geq 0}$, the paths of variables relevant to the consumer choice problem change in the following way:

a) all nominal prices rise in proportion, $\frac{dP_t}{P_t} = \frac{dP}{P}$, for $t \geq 0$

b) all present-value real discount rates rise in proportion, $\frac{dq_t}{q_t} = -\frac{dR}{R}$, for $t \geq 1$

c) the Fisher equation holds at the new sequence of prices: $\frac{dQ_t}{Q_t} = -\frac{dR}{R}$ for $t \geq 1$

d) the agent's unearned income at $t = 0$ rises by dy, and his real wage by dw.

This particular variation, depicted in figure 1, captures in a stylized way the major changes in a consumer's environment that usually follow a temporary change in monetary policy: over a period labelled $t = 0$, incomes and wages increase, the price level rises due to inflation between $t = -1$ and $t = 0$, and the real interest rate $R_0 = \frac{q_0}{q_1}$ falls.[10] As I show formally in Appendix A.1, these are the changes that occur in the standard representative-agent New Keynesian model following a one-period change in monetary policy. Hence this variation is a natural starting point for an analysis of the effects of monetary policy on individual households.

I am interested in the first-order change in initial consumption $dc \equiv dc_0$, labor supply $dn \equiv dn_0$, and welfare dU that results from this change in the environment.

Let σ and ψ be the local Frisch elasticities of substitution in consumption and hours.[11] Define the marginal propensity to consume as $MPC = \frac{\partial c_0}{\partial y_0}$ along the initial path. When a consumer exogenously receives an extra dollar of income, he increases consumption by MPC dollars, but, to the extent that labor supply is elastic ($\psi > 0$), he also reduces hours by $MPN = \frac{\partial n_0}{\partial y_0} < 0$, leaving only $MPS = 1 - MPC + w_0 MPN$ dollars for saving. [12]

These behavioral responses to income changes turn out to also matter for the response to the real interest rate, wage, and price level changes, as the following theorem shows.

Theorem 1. *To first order, dropping $t = 0$ subscripts whenever unambiguous,*

$$dc = MPC\left(d\Omega + \psi n dw\right) - \sigma c MPS \frac{dR}{R} \tag{3}$$

$$dn = MPN\left(d\Omega + \psi n dw\right) + \psi n MPS \frac{dR}{R} + \psi n \frac{dw}{w} \tag{4}$$

$$dU = u'\left(c\right) d\Omega \tag{5}$$

where $d\Omega$, the net-of-consumption wealth change, is given by

$$d\Omega = dy + ndw + \underbrace{\left(y + wn + \left(\frac{{}_{-1}B_0}{P_0}\right) + \left({}_{-1}b_0\right) - c\right)}_{\text{Unhedged interest Rate Exposure } (URE)} \frac{dR}{R} - \underbrace{\sum_{t\geq 0} Q_t \left(\frac{{}_{-1}B_t}{P_0}\right)}_{\text{Net Nominal Position } (NNP)} \frac{dP}{P} \tag{6}$$

The theorem, proved in appendix A.2, follows from an application of Slutsky's equations—separating the wealth and the substitution effects that result from the shock. The relative price changes dR and dw generate substitution effects on consumption and labor supply with familiar signs, and magnitudes given by a combination of Frisch elasticities and marginal propensities. All wealth effects get aggregated into a net term, $d\Omega$, which affects consumption and labor supply after multiplication by the marginal propensity to consume and work, respectively.

[10] The assumption that balance sheets are fixed implies that coupon payments are not contingent on the macroeconomic changes dw, dy, dP or dR. This is an incomplete markets assumption. If assets payoffs are state contingent, my results go through provided the change in payoff is counted towards dy.

[11] Formally, $\sigma \equiv -\frac{u'(c_0)}{c_0 u''(c_0)} > 0$ and $\psi \equiv \frac{v'(n_0)}{n_0 v''(n_0)} \geq 0$.

[12] Indeed, separable utility guarantees that $MPC \in (0,1)$, $MPS \in (0,1)$ and $MPN \leq 0$: in other words, consumption, saving and leisure are 'normal'. Below I provide an alternative definition of the marginal propensity to consume that corresponds to the more familiar split between consumption and savings alone.

Note that theorem 1 makes no assumption on horizon or the form of u and v. In appendix A.3, I show that it extends to general utility functions and to persistent shocks.

Net wealth revaluation: determinants and implications. The net wealth change $d\Omega$ in (6) is the key expression determining the sign and the magnitude of the welfare and the wealth effects in theorem 1. This term is a sum of products of *balance-sheet exposures* by *changes in aggregates*. The exposure to a one-off, immediate increase in the price level is the negative of the present value of the household's net nominal assets, also known as their net nominal position (NNP). This term can be computed directly from a survey of the household's finances. Doepke and Schneider (2006) conduct this exercise for various groups of U.S. households and show that NNPs are large and heterogenous in the population: they are very positive for rich, old households and negative for the young middle class with mortgage debt. Theorem 1 shows that these numbers are not only relevant for welfare, but also for the consumption response to this inflation scenario. Clearly, the composition of balance sheets matters. A household with a positive NNP loses when the price level increases. This exposure can be avoided by investing all wealth in inflation-indexed instruments, that is, by letting ${}_{-1}B_t = 0$ for all t.[13]

Just as an change in the price level 'acts' upon the consumer's net nominal position, equation (6) shows that a change in the real interest rate acts upon what I call his *unhedged interest rate exposure*, or URE. URE is the difference between *all* maturing assets (including income) and liabilities (including planned consumption) at time 0. It represents the net saving requirement of the household at time 0, from the point of view of date -1. Because it includes the stocks of financial assets that mature at date 0 rather than interest flows, it can significantly diverge from traditional measures of savings, in particular if investment plans have short durations.

Why does URE determine the wealth effect following a real interest rate change dR at time 0? To fix ideas, suppose $dR > 0$. This increase in the discount rate reduces the present value of assets, but also the present value of liabilities, with consumption being one such liability. The result is a net wealth loss if future assets exceed future liabilities. But this can only happen if currently-maturing liabilities exceed currently-maturing assets, i.e. if $URE < 0$. Indeed, equation (2) implies that

$$\sum_{t\geq 1} q_t \left(y_t + w_t n_t\right) + \sum_{t\geq 1} q_t \left(\left({}_{-1}b_t\right) + \left(\frac{{}_{-1}B_t}{P_t}\right)\right) - \sum_{t\geq 1} q_t c_t = -URE$$

[13]If prices adjust more sluggishly, the Fisher exposure measure changes. For example, if prices adjust only after T (so that $\frac{dP_t}{P_t} = \frac{dP}{P}$ for $t \geq T$), the formulas hold if NNP is replaced by $\sum_{t\geq T} Q_t \left(\frac{{}_{-1}B_t}{P_0}\right)$, the present value of assets maturing after T. In this case, short-maturity nominal assets maintain constant value, while long-maturity assets decline in value due to the increase in nominal discount rates that follows the expected rise in inflation. The general expression for any given path of price adjustment is given by formula (A.37) in appendix A.3.

The intuition here is that a fall in the price of future consumption relative to current consumption is the same as an increase in the price of current consumption relative to future consumption. But a rise in the price of current goods benefits those consumers that are supplying more goods than they demand at that date, and conversely, it hurts the net buyers of current goods. URE is the measure of the net exposure to this price change. Note that URE is *also* measurable from a survey of household finances that has information on income and consumption.[14]

This observation has the important implication that the *duration of asset plans* matters to determine what happens after a change in real interest rates. Fixed rate mortgage holders and annuitized retirees usually have income and outlays roughly balanced, and hence a URE of about zero. By contrast, ARM holders tend to have negative URE, and savers with large amounts of wealth invested at short durations tend to have positive URE. Hence the theory predicts that the former tend to lose and the latter tend to gain from a temporary increase in real interest rates. In response, consumption falls whenever the substitution effect dominates the wealth effect. Equation 3 allows us to quantify these two effects, and shows that this happens whenever $\sigma c MPS \geq MPC \cdot URE$.

Monetary policy and household welfare. Theorem 1 shows that asset value changes give incomplete information to understand the effects of monetary policy on household welfare. In the model just presented, monetary policy can be thought of as influencing asset values through three channels: a risk-free real discount rate effect (dR), an inflation effect (dP), and an effect on dividends (dy). But these asset value changes do not enter $d\Omega$ directly, so they are not relevant *on their own* to understand who gains and who loses from monetary policy, contrary to what popular discussions sometimes imply. For example, it is sometimes argued that accommodative monetary policy benefits bondholders by increasing bond prices. Yet theorem 1 shows that, while increases in dividends do raise welfare, lower real risk-free rates have ambiguous effects on savers. They have no effect on bondholders whose dividend streams initially match the difference between their target consumption and other sources of income. They benefit households who hold long-term bonds to finance short-term consumption, through the capital gains they generate. And they hurt households who finance a long consumption stream with short-term bonds, by lowering the rates at which they reinvest their wealth. Unhedged interest rate exposures, not asset price changes, constitute the welfare-relevant metric for the impact of real interest rate changes on households. This is why it is important to measure them, which I do in section 4.

[14] By contrast, measuring the exposure to real interest rate changes at any future date requires the knowledge of future income and consumption plans.

The response of consumption to overall income changes. Theorem 1 draws a distinction between exogenous changes in income and changes in wages, since the latter have substitution effects on consumption. However, since preferences are separable, it is possible to rewrite the consumption response as a function of the *total* income change, inclusive of the labor supply response, as shown in appendix A.4.

Corollary 1. *Given an overall change in income $dY = dy + ndw + wdn$, the household's consumption response is given by*

$$dc = \hat{MPC}\left(dY + URE\frac{dR}{R} - NNP\frac{dP}{P}\right) - \sigma c\left(1 - \hat{MPC}\right)\frac{dR}{R} \tag{7}$$

where $\hat{MPC} = \frac{MPC}{MPC+MPS} = \frac{MPC}{1+wMPN} \geq MPC$.

Hence, once we have factored in the endogenous response of income to transfers, the relevant marginal propensity to consume becomes $\hat{MPC}$, the number between 0 and 1 that determines how the remaining amount of income is split between consumption and savings. This corresponds more closely to the textbook measure of the marginal propensity to consume. It is also what empirical measures tend to pick up, since these are usually regressions of observed consumption on observed income.[15]

Durable goods. So far I have restricted my analysis to nondurable consumption. However, durable expenditures tend to account for a substantial share of the overall consumption response to monetary policy shocks, so it is important to understand their behavior. Understanding how durable goods fit into the theory is also important to deliver an accurate map to consumption data. As I show formally in Appendix A.5, adding durable goods to the model does not alter the substantive conclusions from Theorem 1, but there are some subtleties.

The most straightforward case is the one in which the relative price of durable goods and nondurable goods is constant. In this case, formulas (3) or (7) continue to hold, provided that c is interpreted as overall expenditures, MPC is the marginal propensity to spend on all goods, URE subtracts durable expenditures, and σ is adjusted upwards to reflect the fact that durable goods allow more opportunities for intertemporal substitution.

In multi-sector New Keynesian models with durable goods, a constant relative price of durable goods obtains when the prices of durables and nondurables are equally sticky (Barsky, House and Kimball 2007). However, there is some evidence that durables have more flexible prices (Klenow and Malin 2010), in which case the models imply a negative comovement between the relative price of durables p and the nondurable real interest rate

[15] Note that, if hours affected the marginal utility of consumption, it would not be generally possible to obtain an expression such as (7). Instead, dw would enter separately, with a sign reflecting the degree of complementarity between consumption and labor supply.

R. Let $\epsilon = -\frac{\partial p}{p}\frac{R}{\partial R}$ be the corresponding elasticity—in this case, nondurables and durables matter separately, so there no longer exists a straightforward notion of aggregate demand. Instead, in Appendix A.5 I derive separate expressions for the change in nondurable and durable consumption as a function of ϵ. These resemble equations (3) or (7), except for the fact that the expression for URE only subtracts a share $1-\epsilon$ of durable expenditures.[16] Consider a given size increase in the nondurable real interest rate dR. As ϵ rises, durable prices fall by more, and durable demand tends to expand. This is counterfactual, as argued by Barsky, House and Kimball (2007). Hence, in practice, elasticities closer to 0 may be more reasonable. In the empirical section, I will assume $\epsilon = 0$ as a benchmark from computing UREs, but I will also consider robustness of my results to the value of ϵ.

Even though all the results presented in this section assume no uncertainty and perfect foresight, they apply directly to environments with uncertainty provided that markets are complete, *except* for the shock that is unexpected (all summations are then over states as well as dates). An important feature of all these environments is that the marginal propensity to consume, MPC, is the same out of all forms of wealth ($\frac{\partial c_0}{\partial y_0} = \frac{\partial c_0}{\partial \omega}$). The next section relaxes this assumption.

2.3 The consumption response to shocks under incomplete markets

I now consider a dynamic, incomplete-market partial equilibrium consumer choice model. The consumer faces an idiosyncratic process for real wages $\{w_t\}$ and unearned income $\{y_t\}$. He chooses consumption c_t and labor supply n_t to maximize the separable expected utility function

$$\mathbb{E}\left[\sum_t \beta^t \left\{u\left(c_t\right) - v\left(n_t\right)\right\}\right] \tag{8}$$

The horizon is still not specified in the summation. As in the previous section, it will only influence behavior through its impact on the MPC. To model market incompleteness in a general form, I assume that the consumer can trade in N stocks as well as in a nominal long-term bond. In period t, stocks pay real dividends $\mathbf{d}_t = (d_{1t} \ldots d_{Nt})$ and can be purchased at real prices $\mathbf{S}_t = (S_{1t} \ldots S_{Nt})$; the consumer's portfolio of shares is denoted by θ_t. Following the standard formulation in the literature, I assume that the long-term bond can be bought at time t at price Q_t and is a promise to pay a geometrically declining nominal coupon with pattern $\left(1, \delta, \delta^2, \ldots\right)$ starting at date $t+1$. The current nominal coupon, which I denote Λ_t, then summarizes the entire bond portfolio, so it is not necessary to separately

[16] When $\epsilon = 1$, durable purchases are not counted at all in URE, for the same reason that purchases of bonds or shares aren't: in this case, durables completely hedge real interest rate movements.

keep track of future coupons. The household's budget constraint at date t is now

$$P_t c_t + Q_t \left(\Lambda_{t+1} - \delta \Lambda_t\right) + \theta_{t+1} \cdot P_t \mathbf{S}_t = P_t y_t + P_t w_t n_t + \Lambda_t + \theta_t \cdot \left(P_t \mathbf{S}_t + P_t \mathbf{d}_t\right) \tag{9}$$

A borrowing constraint limits trading. This constraint specifies that real end-of-period wealth cannot be too negative: specifically,

$$\frac{Q_t \Lambda_{t+1} + \theta_{t+1} \cdot P_t \mathbf{S}_t}{P_t} \geq -\frac{\overline{D}}{R_t} \tag{10}$$

for some $\overline{D} \geq 0$, where R_t is the real interest rate at time t. The constraint in (10) is a standard specification for borrowing limits[17] and we will see that it generates reactions of constrained agents to balance sheet revaluations that are closely related to those of unconstrained agents. Given that the extent to which borrowing constraints react to macroeconomic changes is an open question, (10) provides an important benchmark.

Provided that the portfolio choice problem just described has a unique solution at date $t-1$, the household's net nominal position and his unhedged interest rate exposure are both uniquely pinned down in each state at time t. This contrasts with the environment in section 2.2, where the consumer was indifferent between all portfolio choices. Here, these quantities are defined as

$$\begin{aligned} NNP_t &\equiv (1 + Q_t \delta) \frac{\Lambda_t}{P_t} \\ URE_t &\equiv y_t + w_t n_t + \frac{\Lambda_t}{P_t} + \theta_t \cdot \mathbf{d}_t - c_t \end{aligned}$$

As before, NNP_t is the real market value of nominal wealth: the sum of the current coupon, Λ_t, and the value of the bond portfolio if it were sold immediately, $Q_t \delta \Lambda_t$. Similarly, URE_t is maturing assets (including income, real coupon payments and dividends) net of maturing liabilities (including consumption).

Consider the predicted effects on consumption resulting from a simultaneous unexpected change in his current unearned income dy, his current real wage dw, the general price level dP and the real interest rate dR, for one period only. Assume that this variation leads asset prices to adjust to reflect the change in discounting alone : $\frac{dQ}{Q} = \frac{dS_j}{S_j} = -\frac{dR}{R}$ for $j = 1 \ldots N$.[18] If $MPC = \frac{\partial c}{\partial y}$, and both MPN and MPS are similarly defined as the responses to current *income* transfers, then the positive results from theorem 1 carry through.

[17] For example, with short-term debt and no stocks ($N = \delta = 0$), $Q_t = \frac{1}{R_t} \frac{P_t}{P_{t+1}}$ and (10) reads $\frac{\Lambda_{t+1}}{P_{t+1}} \geq -\overline{D}$, as in Eggertsson and Krugman (2012).

[18] This is a natural assumption that obtains if asset prices are determined in a general equilibrium with incomplete markets. Absence of arbitrage in such a model implies the existence of a probability measure $\mathbb{Q}$ such that the price of each stock j at date 0 is $S_{0j} = \frac{1}{R_0} \mathbb{E}^{\mathbb{Q}} \left[\sum_{t \geq 1} \frac{1}{R_1 \ldots R_{t-1}} d_{jt}\right]$, where R_t is the sequence of risk-free rates. My variation affects R_0 but does not affect future interest rates, dividends, or risk-neutral probabilities, so results in $\frac{dS_{0j}}{S_{0j}} = -\frac{dR}{R}$. The argument for $\frac{dQ_0}{Q_0} = -\frac{dR}{R}$ is identical.

Theorem 2. *Assume that the consumer is at an interior optimum, at a binding borrowing constraint, or unable to access financial markets (in the latter two cases, let MPS=0). Then his first order change in consumption dc and labor supply dn continue to be given by equations (3) and (4). In particular, writing* $\hat{MPC} \equiv \frac{MPC}{MPC+MPS}$*, the relationship between dc and the total change in income* $dY = dy + ndw + wdn$ *is still given by equation (7).*

The proof is given in appendix A.6. The intuition for why MPC, MPN and MPS are relevant to understand the response of all agents to changes in the real interest rate and the price level is simple: when the consumer is locally optimizing, these quantities summarizes the way in which he reacts to all balance-sheet revaluations, income being only one such revaluation. When the borrowing limit *is* binding, consumption and labor supply adjustments depend on the way the borrowing limit changes when the shock hits. Under the specification (10), the changes in dR and dP free up borrowing capacity[19] exactly in the amount $URE\frac{dR}{R} - NNP\frac{dP}{P}$. Finally, when the consumer is unable to access financial markets, he lives hand-to-mouth so $NNP = URE = 0$. In these latter two cases, $\hat{MPC} = 1$ so we can interpret the consumption response as a pure wealth effect.

By showing that the marginal propensity to consume out of transitory income shocks, which has been the focus of a large empirical literature, remains a key sufficient statistic for predicting behavior with respect to other changes in consumer balance sheets, theorem 2 provides important theoretical restrictions. The rest of the paper takes these restrictions as given and uses them to predict aggregate consumption responses to changes in R or P. But these restrictions are also directly testable empirically: given independent variation in dP, dy and dR as well as individual balance sheet information, one could check that individual consumption responds in accordance with equations (3) or (7). This provides an interesting avenue for future empirical work on consumption behavior.

3 Aggregation and the redistribution channel

This section shows how the microeconomic demand responses derived in section 2 aggregate in general equilibrium to explain the economy-wide response to shocks in a large class of heterogenous-agent models (theorem 3).

3.1 Environment

Consider a closed economy populated by I heterogenous types of agents with separable preferences (8). Each agent type i has its own discount factor β_i, period utility functions u_i

[19] The form of the borrowing constraint is clearly important for this result. For example, if the constraint on the level of wealth (10) is replaced by a constraint on the flow of income received from financial markets, $\frac{Q_t \Lambda_{t+1} + \theta_{t+1} \cdot P_t \mathbf{S}_t}{P_t} - \frac{\delta Q_t \Lambda_t + \theta_t \cdot P_t \mathbf{S}_t}{P_t} \geq -\overline{D}$, then the result collapses to $dc = dY$.

and v_i, and time horizon. To accommodate idiosyncratic uncertainty, assume that within each type i there is a mass 1 of individuals, each in an idiosyncratic state $s_{it} \in S_i$. I write $\mathbb{E}_I [z_{it}]$ for the cross-sectional average of any variable z_{it}, taken over individual types I and idiosyncratic states S_i. I write all aggregate variables in per capita units, so for example aggregate (per capita) consumption C_t is equal to average individual consumption $\mathbb{E}_I [c_{it}]$.

Agents and asset structure. Each agent type i in state s_{it} has a stochastic endowment of $e_i (s_{it})$ efficient units of work, and receives a wage of $w_{it} = e_i (s_{it}) w_t$ per hour, where w_t is the real wage per efficient hour. By choosing n_{it} hours of work, he therefore receives $w_t e_{it} n_{it}$ in earned income. The agent also receives unearned income $y_{it} = d_{it} - t_{it}$, the total dividends on the trees he owns d_{it} net of taxes from the government t_{it}. Let the agent's overall gross-of-tax income be

$$Y_{it} \equiv w_t e_{it} n_{it} + d_{it}. \tag{11}$$

The economy has a fixed supply of aggregate capital K. A set of N trees constitute claims to firm profits and the capital stock. Each tree delivers dividends which, in the aggregate, add up to the sum of aggregate capital income and profits: $\mathbb{E}_I [d_{it}] = \rho_t K + \pi_t$. Agents can also trade nominal government bonds in net supply B_t, as well as a set of $J - 1$ additional assets in zero net supply that can be nominal or real. Each agent of type i can trade a subset N_i of the trees and a subset J_i of the other assets. If both N_i and J_i are empty, agents of type i live hand-to-mouth. In other cases, I assume that trading is subject to a type-specific borrowing constraint $\overline{D_i}$, which takes the form in (10) and may be infinite.

Firms. There exists a competitive firm producing the unique final good in this economy, in quantity Y_t and nominal price P_t, by aggregating intermediate goods with a constant-returns technology. These intermediate goods are produced by a unit mass of firms j under constant returns to scale, using the production functions $X_{jt} = A_{jt} F \left(K_{jt}, L_{jt}\right)$. Markets for inputs are perfectly competitive, so firms take the real wage w_t and the real rental rate of capital ρ_t as given. These firms sell their products under monopolistic competition and their prices can be sticky. Firm j therefore sets its price P_{jt} at a markup over marginal cost and make real profits π_{jt}.[20] Summing across firms $j \in J$, aggregate production is equal to aggregate income:

$$Y_t = \mathbb{E}_J \left[\frac{P_{jt}}{P_t} X_{jt} \right] = w_t \mathbb{E}_J \left[L_{jt} \right] + \rho_t \mathbb{E}_J \left[K_{jt} \right] + \mathbb{E}_J \left[\pi_{jt} \right] \tag{12}$$

[20] Specifically, if μ_{jt} is firm j's markup at time t, then $\pi_{jt} = \left(\mu_{jt} - 1\right) \left(w_t L_{jt} + \rho_t K_{jt}\right)$.

Government. A government has nominal short-term debt B_t, spends G_t, and runs the tax-and-transfer system. Its nominal budget constraint is therefore:

$$Q_t B_{t+1} = P_t G_t + B_t - P_t \mathbb{E}_I [t_{it}] \tag{13}$$

where $Q_t = \frac{1}{R_t} \frac{P_t}{P_{t+1}}$ is the one-period nominal discount rate. The consequences of price-induced redistributive effects between households and the government depend crucially on the fiscal rule. I assume a simple rule in which the government targets a constant real level of debt $\frac{B_t}{P_t} = \overline{b} > 0$ and spending $G_t = \overline{G} > 0$. I also assume that the government balances its budget at the margin by adjusting all transfers in a lump-sum manner. Hence, unexpected increases in P_t (which create ex-post deviations of $\frac{B_t}{P_t}$ from $\overline{b}$) and reductions in the real interest rate R_t result in immediate lump-sum rebates.

Market clearing. In equilibrium, the markets for capital, labor and goods all clear. This implies that at all times t

$$\mathbb{E}_J [K_{jt}] \equiv K \tag{14}$$

$$\mathbb{E}_I [e_{it} n_{it}] = \mathbb{E}_J [L_{jt}] \tag{15}$$

$$\mathbb{E}_I [Y_{it}] = Y_t = C_t + G_t \tag{16}$$

Equilibrium also implies market clearing in all $J + N$ asset markets. This environment nests a large class of one-good, closed economy general equilibrium models. It can accommodate many assumptions about population structure, asset market structure and participation, heterogeneity in preferences, endowments and skills, as well as the nature of price stickiness. With some minor modifications, it would accommodate wage stickiness as well.

3.2 Aggregation result

I am interested in the aggregate consumption response to a perturbation of this environment in which individual gross incomes dY_i, nominal prices dP and the real interest rate dR change at $t = 0$ only. This exercise is useful to understand the effect of an *unexpected shock* that has *no persistence.* Let $dY \equiv \mathbb{E}_I [dY_i]$ be the aggregate change in gross income. Assuming labor market clearing after the shock, this is also the aggregate output change.

Aggregation is simplified by several restrictions from market clearing at $t = 0$. Market clearing for nominal assets implies that all nominal positions net out except for that of the government,

$$\mathbb{E}_I [NNP_{it}] = \overline{b} = -NNP_{gt} \quad \forall t \tag{17}$$

and market clearing for all assets, combined with (11)—(16) implies[21] that

$$\mathbb{E}_I\left[URE_{it}\right] = Y_t - \mathbb{E}_I\left[t_{it}\right] + \frac{B_t}{P_t} - C_t = G_t + \frac{B_t}{P_t} - \mathbb{E}_I\left[t_{it}\right] = -URE_{gt} \tag{18}$$

where NNP_{gt} and URE_{gt} are naturally defined as the net nominal position and the unhedged interest rate exposure of the government sector. Equations (17) and (18) are crucial restrictions from general equilibrium: since one agent's asset is another's liability, net nominal positions and interest rate exposures must net out in a closed economy. Aggregation of consumer responses as described by theorem 2 shows that the per capita aggregate consumption change can be decomposed as the sum of five channels:

Theorem 3. *To first order, in response to dY_i, dY, dP and dR, aggregate consumption changes by*

$$dC = \underbrace{\mathbb{E}_I\left[\frac{Y_i}{Y}M\hat{P}C_i\right]dY}_{\text{Aggregate income channel}} + \underbrace{\mathrm{Cov}_I\left(M\hat{P}C_i, dY_i - Y_i\frac{dY}{Y}\right)}_{\text{Earnings heterogeneity channel}} - \underbrace{\mathrm{Cov}_I\left(M\hat{P}C_i, NNP_i\right)\frac{dP}{P}}_{\text{Fisher channel}}$$

$$+\left(\underbrace{\mathrm{Cov}_I\left(M\hat{P}C_i, URE_i\right)}_{\text{Interest rate exposure channel}} - \underbrace{\mathbb{E}_I\left[\sigma_i\left(1 - M\hat{P}C_i\right)c_i\right]}_{\text{Substitution channel}}\right)\frac{dR}{R} \tag{19}$$

The proof is given in appendix A.7. The key step is to aggregate predictions from theorem 2, decomposing i's individual income change as $dY_i = \frac{Y_i}{Y}dY + dY_i - \frac{Y_i}{Y}dY$ (the sum of an aggregate component and a redistributive component), and using market clearing conditions, the fiscal rule, and the fact that $\mathbb{E}_I\left[dY_i - \frac{Y_i}{Y}dY\right] = 0$ to transform expectations of products into covariances.

Theorem 3 shows that, in the class of environments I consider, a small set of sufficient statistics is enough to understand and predict the first-order response of aggregate consumption to a macroeconomic shock. Equation (19) holds irrespective of the underlying model generating MPCs and exposures at the micro level, as well as the relationship between dY, dP and dR at the macro level. Most of the bracketed terms are cross-sectional moments that are measurable in household level micro-data and are informative about the economy's macroeconomic response to a shock, no matter the source of this shock. The two exceptions are the EISs σ_i, which need to be obtained from other sources, and $dY_i - Y_i\frac{dY}{Y}$, which in general depends on the driving force behind the change in output.

I now use this theorem to discuss the channels of monetary policy transmission under heterogeneity. Alternative applications, for example to short-term redistributive fiscal policy or open-economy models, are also possible.

[21] To see this, note that if b_{it} denotes the asset coupons that mature at time t for household i, we have $URE_{it} = Y_{it} - t_{it} + b_{it} - c_{it}$. Using market clearing in the $J-1$ zero net supply assets, all these coupons net out except for the government coupon, which here is $\mathbb{E}_I\left[b_{it}\right] = \frac{B_t}{P_t}$. The result then follows from goods market clearing and the government budget constraint.

3.3 Monetary policy shocks with and without a representative agent

Consider a transitory, accommodative monetary policy shock that, as in figure 1, lowers the real interest rate and raises aggregate income for one period ($dR < 0$, $dY > 0$), and permanently raises the price level ($\frac{dP}{P} > 0$). Since these are the changes implied by the textbook New Keynesian model with sticky prices and flexible wages after a transitory monetary policy shock, we can apply theorem 3 to understand the consumption response in that model.

The textbook model features a representative agent ($I = 1$) with separable preferences and EIS σ. Hence all covariance terms in (19) are zero, and we are left with

$$dC = \hat{MPC} dY - \sigma \left(1 - \hat{MPC}\right) C \frac{dR}{R} \tag{20}$$

The first term in (20) is a general-equilibrium income effect, and the second term is a substitution effect.[22] Solving out for $dC = dY$ gives the textbook response, $\frac{dC}{C} = -\sigma \frac{dR}{R}$. Intuitively, a Keynesian multiplier $\frac{1}{1-\hat{MPC}}$ amplifies the initial 'first-round' effect from intertemporal substitution. Here this multiplier is entirely microfounded, and in particular takes into account the substitution and wealth effects on labor supply that play out in the background.

Heterogeneity implies a role for redistributive channels in the monetary transmission mechanism, except under special conditions. For example, if aggregate income is distributed proportionally to individual income, so that $dY_i = \frac{Y_i}{Y} dY$; if no equilibrium asset trade is possible, so that agents consume all their incomes $Y_i = c_i$ and $NNP_i = URE_i = 0$; and if all agents have the same elasticity of intertemporal substitution $\sigma_i = \sigma$, then the representative-agent response $\frac{dC}{C} = -\sigma \frac{dR}{R}$ obtains even under heterogeneity. This important neutrality result is studied in Werning (2015).

Away from this benchmark, the redistributive channels of monetary policy can be signed and quantified by measuring the covariance terms in equation (19), either directly in micro data or within a given model. I follow each of these routes in the next two sections to obtain a sense of the plausible magnitudes. As I will show, both the data and my model suggest that all three of the following covariances are negative:

$$\text{Cov}_I \left(\hat{MPC}_i, URE_i\right) < 0 \tag{21}$$

$$\text{Cov}_I \left(\hat{MPC}_i, NNP_i\right) < 0 \tag{22}$$

$$\text{Cov}_I \left(\hat{MPC}_i, Y_i\right) < 0 \tag{23}$$

suggesting that redistribution *amplifies* the transmission mechanism of monetary policy.

Inequality (21) says that agents with unhedged borrowing requirements have higher

[22] Since the typical calibration of the representative-agent model implies a low $\hat{MPC}$, the substitution component is typically dominant in this decomposition, as noticed by Kaplan, Moll and Violante (2016).

marginal propensities to consume than agents with unhedged savings needs. In addition to being supported by the data, in section 5 I will show that it is naturally generated by models with uninsured idiosyncratic risk, with a magnitude that depends on asset durations. Because of this interest rate exposure channel, aggregate consumption is more responsive to real interest rates than measures of intertemporal substitution alone would suggest. In other words, the first-round effect of monetary policy is larger that what the representative-agent model predicts.

Inequality (22) says that net nominal borrowers have higher marginal propensities to consume than net nominal asset holders. This inequality is also both supported by the data and generated endogenously by my model in section 5. It implies that, through its general equilibrium effect on inflation, monetary policy can increase aggregate consumption via a Fisher channel.[23]

Inequality (23) says low-income agents have high MPCs, echoing a finding in much of the empirical literature. On its own, this fact is not enough to sign the earnings heterogeneity channel: we need to know how increases in aggregate income affect agents at different levels of income. More specifically, let

$$\gamma_i \equiv \frac{\partial\left(\frac{Y_i}{Y}-1\right)}{\left(\frac{Y_i}{Y}-1\right)}\frac{Y}{\partial Y} \tag{24}$$

be the elasticity of agent i's relative income to aggregate income. Assume that this is well approximated by a constant γ. Then the earnings heterogeneity channel term in equation (19) simplifies to $\gamma \text{Cov}_I\left(M\hat{P}C_i, Y_i\right)\frac{dY}{Y}$. There is empirical evidence that income risk is countercyclical (for example Storesletten, Telmer and Yaron 2004 or Guvenen, Ozkan and Song 2014) and that monetary policy accommodations reduce income inequality (Coibion et al. 2016). These studies all suggest that γ is negative. Combining this fact with (23), it is likely that monetary expansions increase aggregate consumption because of their endogenous effect on the income distribution.[24]

Independently of the sign of the covariance terms in (19), theorem 3 provides an organizing framework for future research on the role of heterogeneity in the monetary policy transmission mechanism.[25]

[23] Note that this effect from redistribution is conceptually distinct from the effect of future inflation lowering real interest rates, which has nothing to do with nominal redenomination and is present in representative-agent models with persistent shocks to inflation.

[24] Away from separable preferences, an additional *complementarity channel* of monetary policy can arise, even with a representative agent, when preferences are such that increases in hours worked increase the marginal utility of consumption.

[25] An early generation of papers in the heterogeneous agent New Keynesian literature analyzed the transmission of monetary policy under limited heterogeneity. In 'saver-spender' models, such as Bilbiie (2008), 'spender' agents live hand-to-mouth and consume their incomes, so they have $M\hat{P}C = 1$; while 'saver' agents have access to financial markets, with a low $M\hat{P}C$. This has the effect of increasing the aggregate MPC in the economy, raising the importance of income effects relative to substitution effects in equation (19). In 'borrower-saver' models, as in

3.4 Discussion

I now provide a discussion of my result, highlighting its limitations and possible generalizations.

Interactions between the household and other sectors. The market clearing equations (17) and (18) respectively state that the net nominal positions and the unhedged interest rate exposure of the combined household and government sectors are zero. This is a theoretical restriction that must hold in a closed economy, provided firms are correctly consolidated as part of the household sector. In practice there are two challenges: actual economies are open, and it is difficult to accurately take into account the indirect exposures through firms when measuring NNPs and UREs.

In an open economy, (17) and (18) are no longer true, so price-level and real interest rate changes redistribute between the domestic economy and the rest of the world. For example, Doepke and Schneider (2006) find that the net nominal position of the United States is negative, implying that unexpected inflation redistributes towards the U.S. Given a positive average MPC, consumption should rise by more than what equation (19) predicts. Similarly, Gourinchas and Rey (2007) find that the United States borrows short and lends long on its international portfolio, suggesting that it has a negative unhedged interest rate exposure. Hence, U.S. households benefit *on average* from lower real interest rates, further contributing to the expansionary effects of monetary accommodations on consumption.[26]

The assumption that households and firms are consolidated is also important. For example, the household sector tends to be maturity mismatched, holding relatively short-term assets (deposits) and relatively long-term liabilities (fixed-rate mortgages), but this is to a large extent a counterpart to the reverse situation in the banking sector. In principle, household UREs and NNPs should take into account the indirect exposure to interest rates that each household has through all the firms it has a stake in. In practice this is quite challenging, just as it is challenging to estimate indirect exposures of households to the government balance sheet. Undercounting household exposures to negative-URE sectors will imply a positive $\mathbb{E}_I\left[URE_i\right]$, as in equation (18). However, the logic of theorem 3 shows that this imperfect measurement does not matter to the extent that all marginal rebates from other sectors are immediate and lump-sum. In this sense, the covariance terms provide an important benchmark. In practice, rebates might be delayed, and they might target higher or lower MPC agents, so that the precise numbers may depart from

Iacoviello (2005), the high-MPC agents are also borrowers. The literature usually assumes short-term debt, implying (21) and sometimes also nominal debt, implying (22). However, whether (23) holds crucially depends on the assumptions these paper make about the distribution of wages and profits across savers vs spenders.

[26] To the extent that these gains are evenly distributed across the population, these effects can be quantified, respectively, by evaluating $\mathbb{E}_I\left[\hat{MPC}_i\right]\cdot NNP_{US}$ and $\mathbb{E}_I\left[\hat{MPC}_i\right]\cdot URE_{US}$.

the covariance expression in either direction.

One way to assess the importance of all these effects is to directly measure in the data expressions such as $\mathbb{E}_I\left[M\hat{P}C_i URE_i\right]$ and to compare them to the covariance numbers. These 'no-rebate' numbers replace the covariance terms in (19) under the extreme assumption that none of the outside sectors rebate gains to the household sector. In this context, it is interesting to note the theoretical possibility that the interest rate exposure term $\mathbb{E}_I\left[M\hat{P}C_i URE_i\right]$ may not only be positive, but *larger* than the substitution term in (19). Hence, in a world in which outside rebates are highly delayed or benefit low-MPC agents, real interest rate cuts could *lower* aggregate consumption demand, significantly altering the conventional understanding of how monetary policy operates.[27]

General equilibrium and persistent shocks. Theorem 3 provides the response of consumption to a *transitory* shock to R, P and Y. While this exercise provides an insightful decomposition that has the significant merit of involving measurable sufficient statistics, it has two major limitations.

First, the exercise is partial equilibrium in nature: in general, theorem 3 does not permit us to solve for the general equilibrium consumption effect of a given exogenous shock. This is because even transitory exogenous shocks tend to have long-lasting effects on agent behavior and the wealth distribution, which in general equilibrium tends to generate adjustments in future interest rates and/or income. Equation (19) does characterize the full equilibrium in my leading case of the benchmark New Keynesian model, but in more general heterogeneous-agent models it will typically only hold as an approximation of the consumption response to a transitory monetary policy shock.[28]

Second, empirically, monetary policy changes tend to be persistent. Persistent shocks make the derivation of sufficient statistics much more difficult: for example, to characterize the effect of future changes in R, one needs to know the distribution of future consumption and income plans.

In the context of a given structural model, a decomposition such as (19) can be performed for any degree of exogenous and endogenous persistence (see section 5.4 and Kaplan, Moll and Violante 2016).[29] As models grow in complexity and realism, the importance of the channels identified in Theorem 3 can be assessed and refined using such

[27] This theoretical possibility is sometimes mentioned in economic discussions of monetary policy. See Raghuram Rajan ("Interestingly [...] low rates could even hurt overall spending"), "Money Magic", Project Syndicate, November 11, 2013

[28] For instance, the theorem cannot accommodate capital investment, where a current fall in the real interest rate $dR < 0$ comes together with a future fall in capital income, $d\rho_1 < 0$. A previous working paper version of this paper showed the quality of the approximation $dC \simeq dY$ in the context of a model without investment.

[29] See also Auclert, Rognlie and Straub (2017) for an alternative approach that generalizes theorem 3 to shocks with arbitrary persistence and general equilibrium.

Table 1: Seven cross-sectional moments that determine consumption in (25)

	Definition	Name	Channel
$\mathcal{E}_R$	$\text{Cov}_I\left(MPC_i, \frac{URE_i}{\mathbb{E}_I[c_i]}\right)$	Redistribution elasticity for R	Interest-rate exposure
$\mathcal{E}_R^{NR}$	$\mathbb{E}_I\left[MPC_i \frac{URE_i}{\mathbb{E}_I[c_i]}\right]$	—, No Rebate	—
$\widehat{S}$	$\mathbb{E}_I\left[(1-MPC_i)\frac{c_i}{\mathbb{E}_I[c_i]}\right]$	Hicksian scaling factor	Substitution
$\mathcal{E}_P$	$\text{Cov}_I\left(MPC_i, \frac{NNP_i}{\mathbb{E}_I[c_i]}\right)$	Redistribution elasticity for P	Fisher
$\mathcal{E}_P^{NR}$	$\mathbb{E}_I\left[MPC_i \frac{NNP_i}{\mathbb{E}_I[c_i]}\right]$	—, No Rebate	—
$\mathcal{M}$	$\mathbb{E}_I\left[MPC_i \frac{Y_i}{\mathbb{E}_I[c_i]}\right]$	Income-weighted MPC	Aggregate income
$\mathcal{E}_Y$	$\text{Cov}_I\left(MPC_i, \frac{Y_i}{\mathbb{E}_I[c_i]}\right)$	Redistribution elasticity for Y	Earnings heterogeneity

a procedure.[30] I believe that my key finding that redistribution amplifies the effects of monetary policy is likely to remain robust, but it will certainly need to be qualified.

Estimable moments. As discussed above, some of the terms in equation (19) require knowledge of additional information before they can be taken to the data. I make two further assumptions on these structural parameters so as to turn the equation into a full set of estimable moments. For convenience, I also rewrite the decomposition in terms of elasticities.

Corollary 2. *Assume that individuals have common elasticity of intertemporal substitution, $\sigma_i = \sigma$, and common elasticity of relative income to aggregate income, $\gamma_i = \gamma$ for all i. Then,*

$$\frac{dC}{C} = (\mathcal{M} + \gamma\mathcal{E}_Y)\frac{dY}{Y} - \mathcal{E}_P\frac{dP}{P} + (\mathcal{E}_R - \sigma S)\frac{dR}{R} \tag{25}$$

where $\mathcal{M}$, $\mathcal{E}_Y$, $\mathcal{E}_P$, $\mathcal{E}_R$ and S are all measurable cross-sectional moments summarized in table 1.

The proof is in appendix A.8. The assumption of a constant γ parametrizes the incidence of increases in aggregate output dY using a convenient functional form.[31] As is clear from equation (24), when $\gamma > 0$, agents with income above the mean benefit disproportionately from such an increase. The opposite happens when $\gamma < 0$. As discussed above, the evidence on the cyclicality of income risk tends to suggest that the latter case is plausible, though a constant γ is obviously a very strong assumption.

[30] One open question in this literature, in particular, is whether the redistribution channel can help account for the empirical impulse response to a monetary policy shock documented in Christiano, Eichenbaum and Evans (2005) and many others.

[31] Such a specification appears, for example, if labor supply is inelastic ($\psi = 0$) and all income is labor income ($d = 0$). In this case, agent i's gross earnings are $e_i Y$, the product of his skills e_i and aggregate output Y. Suppose that the government taxes these earnings at a rate $\tau(Y)$ and rebates them lump-sum. Then post-redistribution earnings are $Y_i = ((1-\tau(Y))e_i + \tau(Y)\mathbb{E}[e_i])Y$. A constant γ_i follows if the net-of-tax rate has constant elasticity with respect to output, i.e. $\frac{\tau'(Y)}{1-\tau(Y)} = -\gamma$.

Table 1 summarizes the definitions of the moments entering equation (25). I call $\mathcal{E}_P$, $\mathcal{E}_R$ and $\mathcal{E}_Y$ the *redistribution elasticities* of consumption with respect to the price level, the real interest rate and income, since these terms enter explicitly as elasticities in equation (25).[32] The next section measures these numbers in the data.

4 Measuring the redistribution elasticities of consumption

This section turns to data from three surveys to get a sense of the empirical magnitudes of each of the terms in table 1. This exercise is not intended as definitive and will need to be refined in future work. Yet it already paints a fairly consistent picture. With these moment estimates in hand, only two parameters in equation (25) remain unknown. σ can be obtained from the vast literature studying the elasticity of intertemporal substitution, and γ can be obtained from studies on the cyclicality of the distribution of income.

4.1 Three surveys, three identification strategies

In order to compute my key cross-sectional moments, I need household-level information on income, consumption, and balance sheets. Several recent household surveys have collected all this information, both in the United States and abroad, with varying degrees of precision. I also need information on $\hat{MPC}$, the marginal propensity out of transitory income shocks.[33] The literature has used various techniques to estimate these MPCs (see Jappelli and Pistaferri 2010 for a survey). Three of the most influential approaches are implementable using public survey data. I compute my moments using all three approaches, each in a different survey. Since I build on standard references in the literature, I restrict myself to a brief description of the methods, and refer the reader to Appendix B and to the original sources for further detail.

My first source of data is the Italian Survey of Household Income and Wealth (SHIW). In 2010, the survey asked households to self-report the part of any hypothetical windfall that they would immediately spend (Jappelli and Pistaferri 2014). The benefit of this approach is that the windfall can be taken as exogenous for all agents, so in principle this empirical measure of MPC is the number that matters for the theory. This approach also provides MPCs at the household level, making it easy to compute covariances with individual balance-sheet information. These are significant advantages, but the numbers only correspond to one specific setting, that of Italy in 2010. Moreover, a concern with

[32] Calling $\mathcal{E}_Y$ an elasticity is a slight abuse of terminology, since the actual elasticity is $\gamma\mathcal{E}_Y$.

[33] Recall that the theory makes a distinction between $\hat{MPC}$, which takes into account the endogenous response of labor supply, and MPC which does not. The methods used to compute MPC either regress observed consumption on observed income, or ask a question to respondents without mentioning a potential labor supply adjustment, so from now on I assume that they measure $\hat{MPC}$, and sometimes write it MPC for convenience.

self-reported answers to hypothetical situations is that they may not be informative about how households would actually behave in these situations. For these reasons I also turn to other datasets, and to settings where MPCs are estimated from actual behavior.

My second source of data is the U.S. Panel Study of Income Dynamics (PSID) and uses a 'semi-structural' approach to compute MPCs out of transitory income shocks. The procedure is due to Blundell, Pistaferri and Preston (2008) and has since been popularized by Kaplan, Violante and Weidner (2014) and others in the context of macroeconomics. The idea here is to postulate an income process and a consumption function, and to use restrictions from the theory to back out the MPC out of transitory shocks from the joint cross-sectional distribution of consumption changes and income changes. Since this procedure can only recover an estimate at the group level, I compute my redistribution elasticities by first grouping households into different bins, then estimating MPCs within bins and covariances across bins. One drawback of such a procedure is therefore that it generates significantly larger error bands.

My third and final source of data is the U.S. Consumer Expenditure Survey (CE), in which MPC is identified using exogenous income variation following Johnson, Parker and Souleles (2006). These authors estimate the MPC out of the 2001 tax rebate by exploiting random variation in the timing of the receipt of this rebate across households. As the policy was announced ahead of time, they identify the MPC out of an increase in income that is expected in advance. This is, in general, different from the theoretically-consistent MPC out of an unexpected increase. However, to the extent that borrowing constraints are important, or if households are surprised by the receipt despite its announcement, the estimation gets closer to the MPC that is important for the theory. This procedure also yields an MPC at a group level, so I again estimate covariances across groups.

As discussed in appendix B, each of the three techniques has its own limitations, and no survey contains perfect information on all components of household balance sheets. Notably, the consumption data in the SHIW and the PSID is limited, as are the income and the asset data in the CE. In addition, none of these surveys samples very rich households whose consumption behavior may be an important determinant of aggregate expenditures. Hence, the exercise in this section is tentative and intended to give a sense of magnitudes based on the current state of knowledge in the field. As administrative-quality household surveys become available and more sophisticated identification methods for *MPC*s arise, a priority for future work is to refine the estimates I provide here.[34]

[34]Using administrative Norwegian data and the MPCs of lottery winners, Fagereng, Holm and Natvik (2016) provide estimates of redistribution covariances that are broadly consistent with mine.

4.2 Conceptual measurement issues

Even though my analysis is term of elasticities, which are unit-less numbers, the choice of time units is important: MPC needs to be measured over a period of time consistent with the time unit for income, consumption, and maturing elements of the balance sheet. I follow the structure of the datasets, and measure each at an annual rate in the SHIW and the PSID, and at a quarterly rate in the CE.

MPC. As discussed above and in Appendix A.5, my ideal measure of MPC would be one that encompasses both nondurable and durable goods, since this would correspond to the concept that matters for predicting changes in aggregate consumption spending. The question in the SHIW refers to 'spending' without distinguishing between types of purchases, so it is safe to assume that it refers to both durables and nondurables. For my U.S. exercises, I prefer to follow the baseline estimates from Blundell, Pistaferri and Preston (2008) and Johnson, Parker and Souleles (2006), neither of which include durable goods in MPC estimation. My PSID estimate includes only nondurables, while my CE estimate only includes food. Appendix 4.3 considers robustness to using a broader set of goods into the MPC estimation.

URE. As defined in section 2.2, URE_i measures the total resource flow that a household i needs to invest over the first period of his consumption plan. In each survey, I construct URE_i as

$$URE_i = Y_i - T_i - C_i + A_i - L_i \tag{26}$$

where Y_i is gross income, T_i is taxes net of transfers, C_i is consumption, and A_i and L_i represent, respectively, assets and liabilities that mature over the period over and above the amounts already included in Y_i or C_i. Specifically, C_i includes expenditure on durable goods, rents and interest payments, while Y_i includes income from all sources: labor, dividend, and interest income. Therefore, Y_i comprises the maturing portion of equities and bonds in household portfolios. Moreover, $Y_i - C_i$ includes the 'maturing' portion of housing, which I treat as a special asset that pays a dividend equal to its consumption by owner-occupiers.[35] In L_i, I count principal payments on all loans, notably all mortgages, since these are not included in consumption.

The remaining maturing asset and liabilities that I include in A_i and L_i consist of short-term and adjustable-rate assets and liabilities. For these assets, I only observe the stocks,

[35] This differential treatment of housing relative to other durable goods is consistent with the assumption made in national income and product accounts. In the language of section 2.2, the implicit assumption is that the relative price of housing has an elasticity $\epsilon^h = 1$ with respect to the real interest rate, while the relative price of all other durable goods is $\epsilon^d = 0$. I consider robustness to alternative values of ϵ^d in appendix B.5.2.

and detailed maturity information is typically absent. I therefore define a benchmark scenario in each survey, and perform an extensive sensitivity analysis in Appendix B.5.3. I assume in this benchmark that, for every agent i, a) time and savings deposits have a duration of two quarters, b) adjustable-rate mortgages have a duration of three quarters, and c) debt outstanding on credit cards has duration of two quarters. Appendix 4.3 contains detailed information on the asset and liability classes reported in each survey, as well as a comparison across surveys.

NNP and Income. I compute net nominal positions as the difference between directly held nominal assets (mainly deposits and bonds) and directly held nominal liabilities (mainly mortgages and consumer credit). When assets are clearly indicated as shares of a financial intermediary that mostly owns nominal assets (for example, money market mutual funds), I also include the value of these shares in the households' nominal position. However, relative to Doepke and Schneider (2006), I do not calculate the indirect nominal positions arising from holdings of equity or other financial intermediaries, since my data is not sufficiently detailed for this purpose. For my income measure, in keeping with the theory, I use pre-tax income in the PSID and the CE where it is available; in the SHIW I use post-tax income.

Measurement error. Measurement error is a very important issue in this exercise. These errors can stem from many sources: poor data quality, imperfect coverage, underreporting of consumption, or timing differences in the reporting of consumption and income. As discussed in the appendix, each survey has its own strengths and weaknesses. The CE has excellent information on consumption and liabilities, but very poor information on assets. Both the PSID and the SHIW appear to considerably undermeasure consumption. My covariance estimates are unbiased provided that the measurement errors in in MPC and its cross-term (URE, NNP or Y) are additive and uncorrelated. Economically, this assumption corresponds to the presence of a 'mismeasurement' sector that rebates gains and losses lump-sum, just as the government does in the setting of theorem 3.[36] Keeping this important caveat in mind, I proceed to my measurement exercise.

4.3 Estimating the redistribution elasticities

Table 2 reports the main summary statistics from each survey, with the appendix containing further detail. Each line is normalized by average consumption in the survey, which

[36] For example, by abstracting away from indirect exposures to the banking sector, I tend to overstate the aggregate URE. If gains to the banking sector disproportionately favor low-MPC households, my estimate of the MPC/URE correlation would be biased downwards.

Table 2: Main summary statistics from the three surveys

Survey	SHIW		PSID		CE	
Variable	mean	s.d.	mean	s.d.	mean	s.d.
Income after tax ($Y_i - T_i$)	1.31	0.92	2.13	2.63	1.16	1.03
Consumption (C_i)	1.00	0.61	1.00	0.63	1.00	0.83
Maturing assets (A_i)	0.98	2.64	1.46	6.38	0.48	1.70
Maturing liabilities (L_i)	0.34	1.55	0.81	2.11	0.53	1.55
Unhedged interest rate exposure (URE_i)	0.95	3.13	1.78	7.60	0.16	2.36
Nominal assets	0.82	2.61	1.41	5.00	1.90	7.50
Nominal liabilities	0.55	1.65	2.72	3.95	4.97	7.73
Net nominal position (NNP_i)	0.27	2.92	-1.31	6.10	-2.79	10.06
Income before tax (Y_i)	1.31	0.92	2.67	4.11	1.25	1.11
Marginal propensity to consume (MPC_i)	0.47	0.35				
Number of households	7,951		9,620		4833	

In each survey, 'mean' and 's.d.' represent the sample mean and standard deviation.
All statistics are computed using sample weights.
All variables except for MPC are normalized by average consumption in the sample.

facilitates comparability and corresponds to the normalization behind my elasticities in table 1. Note that the average URE is positive all three surveys. One reason, in addition to those highlighted in section 3.4, is that consumption is below income at the mean, especially in the PSID and the SHIW—likely because of underreporting and coverage issues. This is another reason why my preferred estimate of the redistribution elasticity is $\mathcal{E}_R$ rather than $\mathcal{E}_R^{NR}$, which is mechanically pushed up by the high average URE. The average net nominal position is quite negative in CE and PSID—possibly reflecting poor asset measurement—and moderately positive in the Italian survey, where far fewer households own mortgages.

Figure 2 reports the distribution of MPC by URE, NNP and income across the three surveys. Columns correspond to datasets, and rows to redistribution channels. The first column displays data from the SHIW, where individual MPC information is available. The graphs report the average value of MPC in each percentile of the x-axis variable. On the other hand, in the PSID (second column) and the CE (third column), I estimate the MPC by stratifying the population in terciles of the x-axis variable, and then report the point estimate together with confidence intervals within each bin.

Starting with the interest exposure channel, looking across the first row, all three surveys show a negative correlation between MPC and URE. This is particularly apparent in the SHIW data, but the pattern is there in the U.S. surveys as well. A direct implication is that $\mathcal{E}_R < 0$ in each of these datasets: falls in interest rates increase consumption demand via the redistribution channel. Turning to the Fisher channel, we also observe an overall

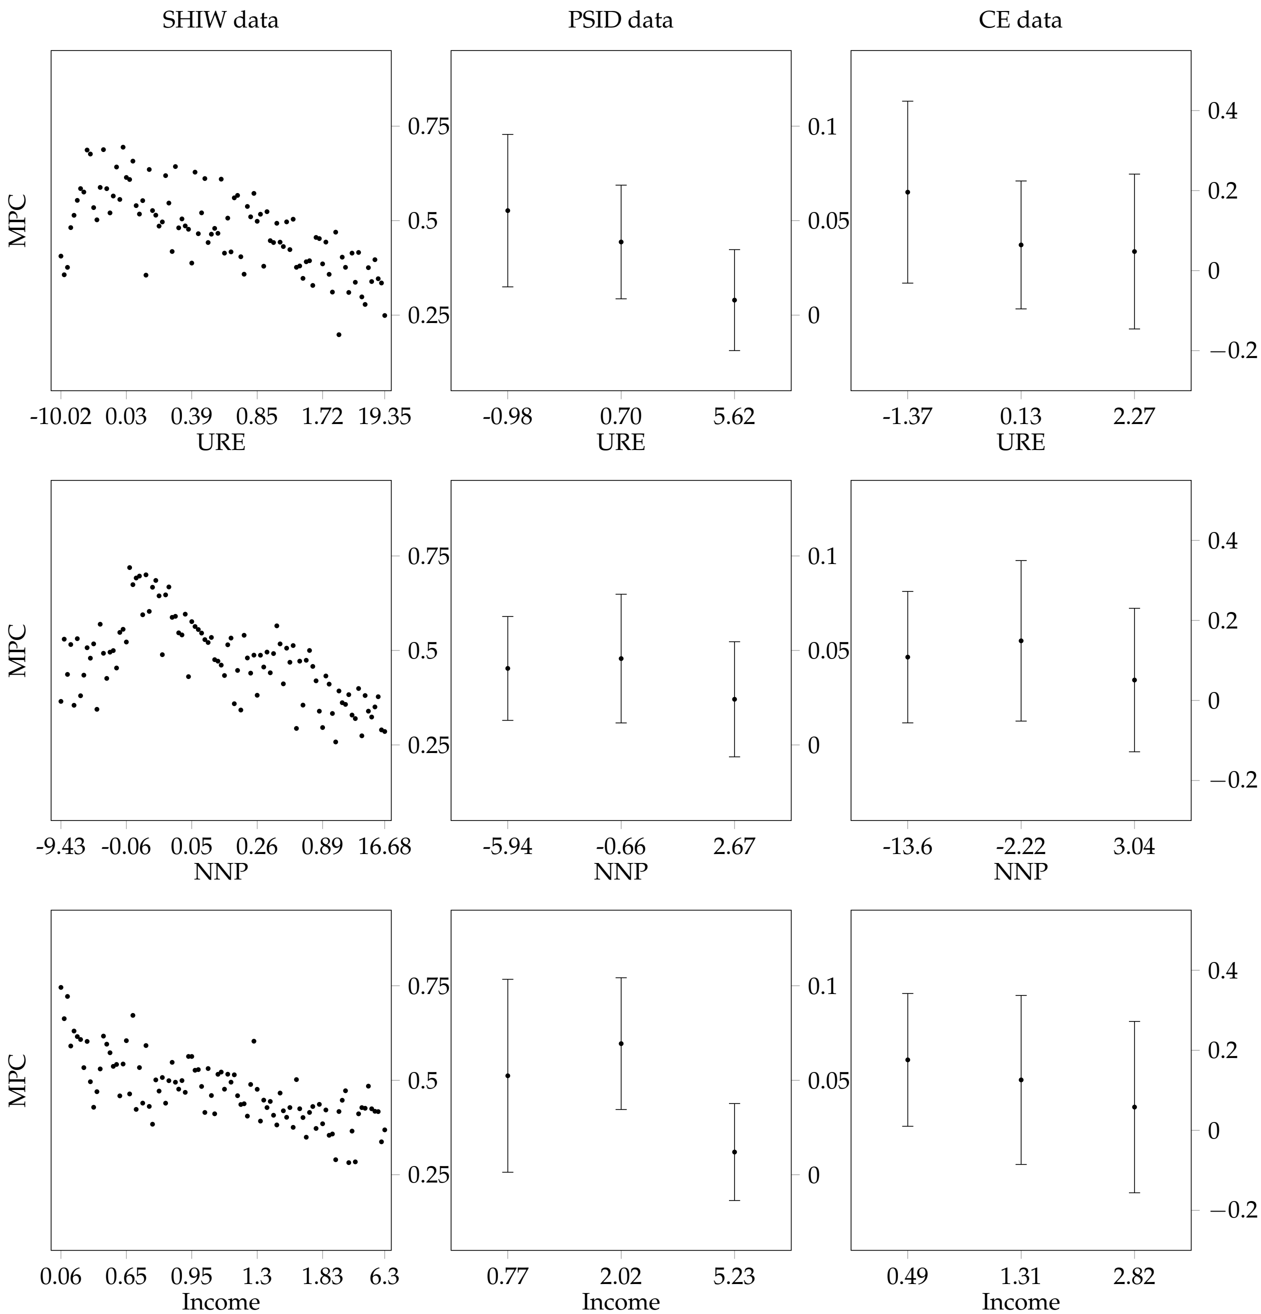

Figure 2: Marginal propensities to consume and the redistribution channels.

negative correlation, though it is somewhat less pronounced. In particular, MPCs tend to be slightly higher in the center of the NNP distribution than at the extremes, potentially consistent with a 'wealthy hand-to-mouth' explanation as in Kaplan and Violante (2014). This suggests that $\mathcal{E}_P < 0$, consistent with Fisher's hypothesis—unexpected increases in prices tend to increase consumption overall, though this effect is quantitatively small. Finally, across all three surveys, the covariance between MPCs and gross incomes is also negative. Again, this pattern is not entirely clear across the income distribution, consistent with the presence of some wealthy hand-to-mouth individuals. Combined with $\gamma < 0$, a negative $\mathcal{E}_Y$ implies an amplification role for the earnings heterogeneity channel in the transmission of monetary policy.

Moving on to magnitudes, table 3 computes my seven key cross-sectional moments, together with 95% confidence intervals. For the PSID and the CE, the estimation is done across bins by using three bins, just as in figure 2.[37]

Confirming the visual impression from figure 2, the point estimates for the redistribution elasticities $\widehat{\mathcal{E}_R}$, $\widehat{\mathcal{E}_P}$ and $\widehat{\mathcal{E}_Y}$ are negative in all three surveys. However, the magnitudes are relatively small—in particular, the confidence bands in the CE always include zero.[38]

To put these numbers in the context of standard representative-agent analyses, consider that many macroeconomists believe 0.1 to 0.5 as plausible values for the elasticity of intertemporal substitution σ. Equation (25) shows that σ should be compared to $-\mathcal{E}_R/S$ to gauge the relative strength of the redistribution effect. According to the point estimates from table 3, this number is between 0.05 and 0.20. Hence, the data suggests that the redistribution effect might be as important as the substitution effect in explaining why aggregate consumption responds to changes in real interest rates. Similarly, the magnitudes of $\widehat{\mathcal{E}_P}$ and $\widehat{\mathcal{E}_Y}$ are fairly small, so that (unless γ is very negative) neither channel can account on its own for very large movements in consumption. But their combined effect may nevertheless be substantial, and further research is needed to refine the precision of these estimates.

As more sources of joint consumption, income and asset data become available, a better empirical understanding of UREs and NNPs will become possible, helping to shape our understanding of the winners and losers from changes in real interest rates and inflation. Real-time estimates of the redistribution covariances will also provide useful information about the dynamic evolution of the monetary policy transmission mechanism.

[37] Appendix B.5.4 reports a sensitivity analysis using four to eight bins. The results are little changed.

[38] Moreover, the estimated value of $\mathcal{E}_R^{NR}$ is usually positive, implying that the negative covariance is not strong enough to overwhelm the positive value of URE at the mean. As argued above, taking $\mathcal{E}_R^{NR}$ at face value requires an extreme view of outside-sector rebates.

Table 3: Estimates of table 1's cross-sectional moments using SHIW, CE and PSID

Survey	SHIW		PSID		CE	
	Estimate	95% C.I.	Estimate	95% C.I.	Estimate	95% C.I.
$\widehat{\mathcal{E}_R}$	-0.11	[-0.16,-0.06]	-0.05	[-0.10,-0.00]	-0.09	[-0.26,0.09]
$\widehat{\mathcal{E}_R^{NR}}$	0.34	[0.29,0.39]	0.01	[-0.05,0.06]	-0.05	[-0.23,0.13]
$\widehat{S}$	0.55	[0.53,0.58]	0.97	[0.95,0.98]	0.90	[0.77,1.03]
$\widehat{\mathcal{E}_P}$	-0.07	[-0.12,-0.03]	-0.02	[-0.08,0.04]	-0.11	[-0.83,0.60]
$\widehat{\mathcal{E}_P^{NR}}$	0.05	[0.01,0.10]	-0.07	[-0.13,-0.01]	-0.55	[-1.33,0.23]
$\widehat{\mathcal{M}}$	0.57	[0.55, 0.59]	0.08	[0.03,0.13]	0.14	[-0.12,0.39]
$\widehat{\mathcal{E}_Y}$	-0.05	[-0.07, -0.03]	-0.04	[-0.08,-0.00]	-0.05	[-0.15,0.06]

All statistics are computed using survey weights. In the CE and the PSID, confidence intervals are bootstrapped by resampling households 100 times with replacement.

4.4 Empirical drivers of the redistribution covariances

While the sufficient statistic approach suggests that only the population-level redistribution elasticities matter to determine an overall effect, in practice it is interesting to understand the empirical drivers of these covariances. For example, is the covariance between MPC and URE negative because older households tend to have lower MPCs and higher UREs? In order to shed light on this and related questions, I perform a covariance decomposition, projecting each covariance onto observable components such as age or education. This procedure is inspired by the law of total covariance: focusing on URE for ease of notation, for any covariate Z_i we know that

$$\text{Cov}\left(MPC_i, URE_i\right) = \underbrace{\text{Cov}\left(\mathbb{E}\left[MPC_i|Z_i\right], \mathbb{E}\left[URE_i|Z_i\right]\right)}_{\text{Explained fraction of covariance}} + \underbrace{\mathbb{E}\left[\text{Cov}\left(MPC_i, URE_i|Z_i\right)\right]}_{\text{Unexplained fraction of covariance}} \quad (27)$$

We can then implement this decomposition using an OLS regression, which performs a linear approximation to the conditional expectation function.[39] For any observable covariate Z_i, I run two OLS regressions

$$\begin{aligned} MPC_i &= \alpha_M + \beta_M Z_i + \epsilon_{Mi} \\ URE_i &= \alpha_R + \beta_R Z_i + \epsilon_{Ri} \end{aligned}$$

and compute the covariance between the fitted values $\widehat{MPC_i}$ and $\widehat{URE_i}$ to get an empirical counterpart of the explained component in (27). This gives me the part of the covariance

[39] This is similar to implementing the law of total variance using R^2.

Table 4: Covariance decomposition for URE, NNP and income in the SHIW

			$\mathcal{E}_R$		$\mathcal{E}_P$		$\mathcal{E}_Y$	
Z_i	Var (Z_i)	$\widehat{\beta_M}$	$\widehat{\beta_R}$	% expl.	$\widehat{\beta_P}$	% expl.	$\widehat{\beta_Y}$	% expl.
Age bins	0.77	-0.027	0.459	9%	0.521	15%	0.062	3%
Male	0.24	-0.055	0.396	5%	0.285	5%	0.282	7%
Married	0.18	-0.016	0.116	0%	-0.070	-0%	0.417	2%
Years of ed.	18.8	-0.005	0.064	6%	0.031	4%	0.088	17%
Family size	1.71	0.023	-0.107	4%	-0.215	12%	0.122	-10%
Res. South	0.22	0.198	-0.481	19%	-0.255	15%	-0.561	48%
City size	1.21	0.037	0.029	-1%	0.053	-3%	0.068	-6%
Unemployed	0.04	0.189	-0.728	5%	-0.308	3%	-0.624	10%

that can be explained by Z_i, since

$$\begin{aligned}\text{Cov}\left(MPC_i, URE_i\right) &= \text{Cov}\left(\widehat{MPC_i} + \widehat{\epsilon_{Mi}}, \widehat{URE_i} + \widehat{\epsilon_{Ri}}\right) \\ &= \text{Cov}\left(\widehat{\beta_M} Z_i + \widehat{\epsilon_{Mi}}, \widehat{\beta_R} Z_i + \widehat{\epsilon_{Ri}}\right) \\ &= \text{Var}\left(Z_i\right) \widehat{\beta_M}\widehat{\beta_R} + \text{Cov}\left(\widehat{\epsilon_{Mi}}, \widehat{\epsilon_{Ri}}\right) \qquad (28)\end{aligned}$$

where the last line follows because, by construction, $\text{Cov}\left(\widehat{\epsilon_{Mi}}, Z_i\right) = \text{Cov}\left(\widehat{\epsilon_{Ri}}, Z_i\right) = 0$. For example, in table 4, when Z_i is age, $\widehat{\beta_M}$ is negative and $\widehat{\beta_R}$ is positive, so older agents do tend to have lower *MPC* and higher *URE*. However, on its own, age can only explain 9% of the total covariance.

This procedure is straightforward to implement in the SHIW, where *MPC* is available at the individual level. Table 4 reports these results using as control variables all those that Jappelli and Pistaferri (2014) use to explain *MPC*, one covariate at a time. For each of my three redistributive channels, I report each of the terms in the decomposition (28), as well as the fraction of the variance explained. In Appendix B.6, I generalize this approach to multiple covariates, and also report graphs of *URE* and *NNP* by age and income bins in each survey. All of these tend to give a consistent message: age, education and income are all negatively correlated with MPC and positively correlated with URE and NNP, so they help explain the negative covariance overall.

5 Sufficient statistics in a Huggett model

This section puts some additional structure to the model of section 3 to connect the empirical magnitudes estimated in the previous section back to theory. It answers the following four main questions: 1) Can a simple model rationalize the empirical signs and magnitudes obtained in the previous section? 2) What are some key theoretical determinants of these redistribution elasticities? 3) How robust are the sufficient statistics predictions to large shocks? 4) What about persistent shocks such as those likely to prevail in practice?

Since I explicitly specify the heterogeneity and the driving processes, the sufficient statistics now become endogenous. However, I do not explicitly model the endogenous determination of incomes, and instead assume an endowment economy with exogenous labor supply. Endogenizing the earnings heterogeneity channel remains a significant challenge for the burgeoning literature on New Keynesian Heterogeneous-Agent models, so I do not attempt to do this here, and provide a discussion at the end of the section.

5.1 Environment

The economy is now populated by a continuum of infinitely-lived, ex-ante identical but ex-post heterogenous households indexed by $i \in [0, 1]$. Agents do not work, but face face idiosyncratic uncertainty with respect to their endowment of goods $\{y_{it}\}$ and their discount factor $\{\beta_{it}\}$. The process for the exogenous idiosyncratic state $\mathbf{s}_{it} = (y_{it}, \beta_{it})$ is uncorrelated across agents and follows a Markov chain $\Gamma(\mathbf{s}'|\mathbf{s})$ over time. This Markov chain is assumed to have a stationary distribution $\varphi(\mathbf{s})$, which I take to be the cross-sectional distribution of idiosyncratic states at $t = 0$. There is no aggregate uncertainty: the path for all macroeconomic variables is perfectly anticipated.

Labor supply is exogenous so all households value consumption streams only. They do so with separable preferences, as in (8), with inelastic labor supply ($\psi = 0$) and common elasticity of intertemporal σ.

I assume that there are two assets available for trade, both risk-free, long-term bonds with identical rates of decay δ as in section 2.3. One of these assets is a nominal asset and one is a real asset. Prices are expected to remain constant forever, and therefore households are completely indifferent between both types of bonds. To break indifference, I assume that each household allocates a fraction κ of his portfolio to the nominal asset. A borrowing constraint limits the size of bond issuances so that the market value of real end-of-period liabilities is bounded by a limit $\overline{D}$, which takes the form in (10).

Readers will recognize the standard incomplete markets model, taken here in partial equilibrium as in Deaton (1991), Carroll (1997), and which forms the basis of general equilibrium variants such as Bewley (1980), Huggett (1993) and Aiyagari (1994). The only difference is that assets may be nominal and have long maturity, two features that are crucial characteristics of household balance sheets. I have abstracted away from many additional important features in household finance, such as portfolio choice, to focus on the key determinants of sufficient statistics.

In this environment, Theorem 1 applies to every individual agent, with $\hat{MPC} = MPC$ since labor supply does not enter preferences. I consider a calibration of the steady-state of this model in which aggregate income and consumption are equal ($C = \mathbb{E}_I[y_i]$). This can

Table 5: Calibration parameters, targets, and main sufficient statistics

Parameters		Value	Targets	
Elasticity of intertemporal substitution	σ	0.5		
Impatient discount factor	β^I	0.93	Average MPC	0.25
Patient discount factor	β^P	0.99	Real interest rate (annual)	3%
Borrowing limit (% of per capita annual C)	$\overline{D}$	195%	Household debt (% of C)	113%
Outcomes				
Redistribution elasticity for R ($\delta = 0.95$)	$\mathcal{E}_R$	-0.09	See Figure 3a	
Hicksian scaling factor	S	0.84		
Redistribution elasticity for P ($\kappa = 0$)	$\mathcal{E}_P$	-1.8	See Figure 3b	
Income-weighted MPC	$\mathcal{M}$	0.17		
Redistribution elasticity for Y	$\mathcal{E}_Y$	-0.08		

be interpreted as the general equilibrium a closed economy with no government spending or taxes. Starting from such a steady state, Theorem 3 applies, allowing me to ask my four main questions of this section.

5.2 Steady-state calibration and solution method

I perform my calibration at quarterly frequency. I assume a steady state real interest rate of 3% and a household debt to consumption ratio 113%—the U.S. level for 2013. Since there is no net savings, total assets are also 113% of consumption, which is consistent with data on interest-paying assets held by the household sector.[40] As already noted, I assume that there is no inflation at steady-state ($\Pi = 1$), and consider a range of calibrations for bond durations, from 1 quarter to 10 years ($\delta \in [0, 0.95]$), and for inflation indexation ($\kappa \in [0, 1]$).[41] As a benchmark, Doepke and Schneider (2006) report that the average duration of U.S. household assets and liabilities is 4.5 years (see their figure 3), which falls in the middle of my range. On the other hand, the household assets that my calibration includes are entirely nominal, making $\kappa = 0$ a useful reference point.

I follow the vast majority of the literature in postulating an income process that follows an $AR(1)$ process in logs at quarterly frequency, and follow Guerrieri and Lorenzoni (2015) to calibrate this process.[42] I normalize $\overline{y}$ such that $\mathbb{E}[y] = 1$.

[40] According to the U.S. Financial Accounts, in 2013 households held interest-paying liabilities worth \$13trn and interest-paying assets worth \$12.8trn. I define the former as the sum of mortgages and consumer credit, the latter as time and savings deposits and credit market instruments.

[41] The duration of the nominal bond is $\frac{R\Pi}{R\Pi-\delta} = \mathcal{D}$, so $\frac{\delta}{\Pi} = R\left(1 - \frac{1}{\mathcal{D}}\right)$. $\delta = 0$ and $\mathcal{D} = 1$ for short-term debt.

[42] Specifically, the process is $\log y_t^i - \log \overline{y} = \rho_y \left(\log y_{t-1}^i - \log \overline{y}\right) + \sigma_y \sqrt{1-\rho^2}\epsilon_t^i$ with $\epsilon_t^i \sim \mathcal{N}(0,1)$, with a coefficient of mean reversion of $\rho_y = 0.96$ at quarterly frequency, and a cross-sectional standard deviation of log income of $\sigma_y = 0.52$. I discretize the process using a ten state Markov chain.

Since the moments of the redistribution channel all feature a prominent role for MPCs, I make sure that my model generates *average* marginal propensities to consume that are in line with the empirical evidence. The empirical literature replicated above, and reviewed in more detail in Jappelli and Pistaferri (2010), consistently finds numbers between 0.1 to 0.4 at an annual rate. I settle for a number in the middle of this range, and target an average quarterly marginal propensity to consume of 0.25. It is well-known that the benchmark incomplete markets model cannot generate this level of average MPCs (see for example Kaplan and Violante 2014). I follow the simple Krusell and Smith (1998)–Carroll et al. 2017 solution of assuming slow-moving time preference heterogeneity, with agents alternating between patience (discount factor β^P) and impatience (discount factor β^I).[43] I then jointly calibrate β^P, β^I and the borrowing limit $\overline{D}$ to achieve my targets for the average MPC, household debt, and a closed current account at $R = 3\%$.[44] The top row of table 5 summarizes my benchmark parameters.

The model is solved using a version of Carroll (2006)'s endogenous gridpoints methods. Details are provided in appendix C.

5.3 Sufficient statistics in the model and the data

The bottom row of table 5 shows the five key endogenous sufficient statistics in the model. The Hicksian scaling factor S and the income-weighted MPC $\mathcal{M}$ take on values consistent with the empirical evidence. $\mathcal{M}$ is below the calibrated average MPC because of the negative covariance between MPCs and income, which is given by $\mathcal{E}_Y = -0.08$. The fact that this covariance is both negative and small is consistent with the empirical evidence from table 3, and is one success of the model. It results from the fact that, in the model, MPCs are strongly negatively correlated with liquid wealth ('cash-on-hand'), but income and cash on hand are not that highly correlated precisely because liquid wealth is used to smoothe income fluctuations.

Next, when I choose $\delta = 0.95$ to match an average asset duration of 4.5 years, I obtain a redistribution elasticity for real interest rates, $\mathcal{E}_R = -0.08$, that is negative but small, as it is in the data. This is because long durations imply endogenously small unhedged interest rate exposures, since households roll over only a fraction of their wealth every quarter. As a consequence, falls in real interest rates imply a limited amount of redistribution—though this redistribution does favor high-MPC households on average. However, the left panel of figure 3 shows that this result is very sensitive to the assumed duration of assets. Typical calibrations of Bewley models assume that debt is short-term ($\delta = 0$). As

[43] My Markov process is such that the stationary population distribution contains patient and impatient agents in equal numbers, and that consumers stay in their patience state for 50 years on average.

[44] Intuitively, β^P controls net asset accumulation, β^I the average MPC and $\overline{D}$ the debt-to-GDP ratio, so a global solver has no difficulty finding a solution to this system of three equations in three unknowns.

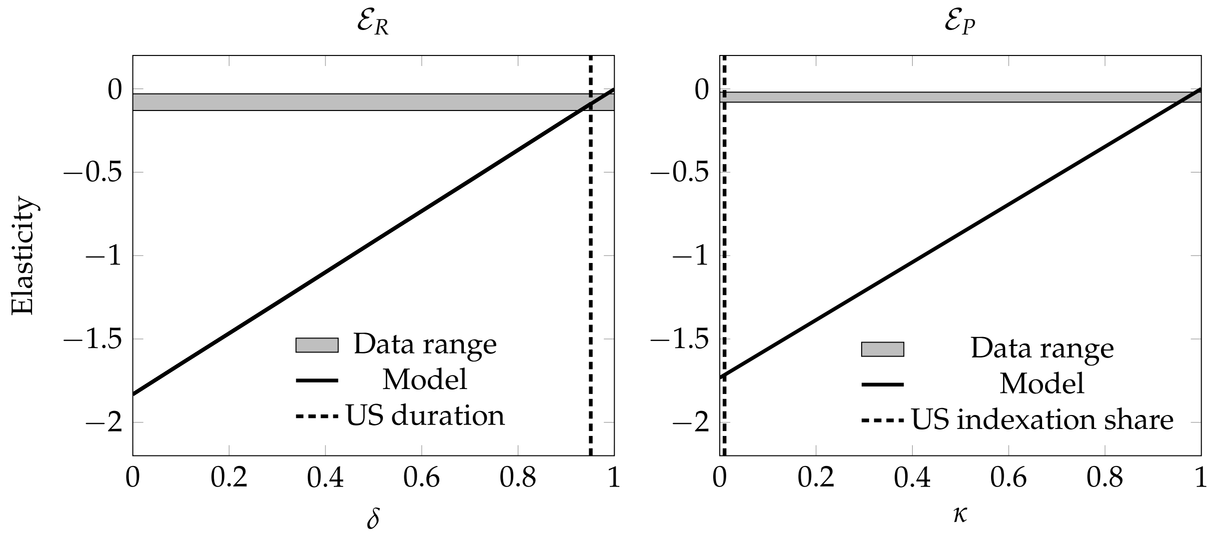

Figure 3: Sufficient statistics: model vs data

the figure shows, assuming short-term debt implies a very large and negative $\mathcal{E}_R$. In such a calibration, the redistributive effects of real interest rate changes completely swamp the substitution effects: $-\mathcal{E}_R/S$ is more than four times larger than the elasticity of intertemporal substitution σ, implying that more than 80% of the consumption effects of real interest rate changes come from redistribution. This finding implies a very important role for the maturity structure in determining the aggregate effects of monetary policy changes.

There are two equivalent ways of interpreting the more muted response of the economy to monetary policy shocks under longer asset durations. The first is that long durations reduce the endogenous amount of unhedged interest rate exposures—making everyone's consumption less sensitive to changes in real interest rates. A second and more subtle interpretation is that under longer asset maturities, expansionary monetary policy creates more capital gains for asset holders and additional upward revaluation of liabilities for borrowers. These capital gains and losses redistribute against the economy's MPC gradient, and therefore make monetary policy less potent in affecting consumption.

Such a role for the maturity structure in monetary policy transmission is consistent with the cross-country structural VAR evidence presented in Calza, Monacelli and Stracca (2013). It suggests that wealth redistribution is the primary reason why monetary policy affects consumption in a country like the United Kingdom, where mortgages have adjustable rates.

Turning to the redistributive role of inflation, the model with purely nominal assets ($\kappa = 0$) implies a counterfactually large elasticity of aggregate consumption with respect to increases in the price level, $\mathcal{E}_P = -1.77$. In contrast to the data, this version of the

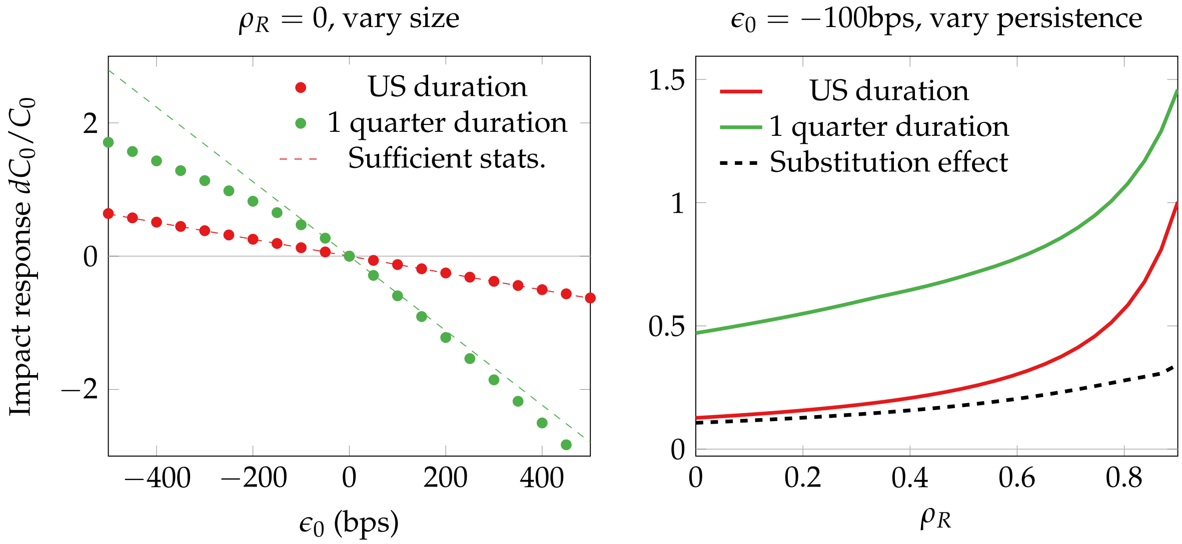

Figure 4: Effect of shock persistence and shock size

model implies instead very powerful redistribution through the Fisher channel. There is a strong intuition for this result. Inflation redistributes along the asset dimension, which in this class of models is highly correlated with MPC (a consequence of the concavity of the consumption function). If all assets are nominal, inflation directly affects agents' real asset positions. Of course, as the right panel of figure 3 shows, more inflation indexation brings the model closer to the empirical results. However, matching the data requires assuming at least 90% inflation indexation, which is clearly counterfactual. Understanding this discrepancy between model and data, and evaluating more carefully the role of the Fisher channel in monetary policy transmission, is a priority for future research.

Whether as cross-checks or as direct targets for calibration, sufficient statistics play a promising role in disciplining heterogeneous-agent general equilibrium models going forward.

5.4 Effects of shock size and persistence

I now explore the role of persistence, focusing on real interest rate changes. Specifically, in this exercise, I maintain a constant price level $P_t = \overline{P}$ and change the real interest rate by

$$R_t - R^* = \rho_R \left(R_{t-1} - R^*\right) - \epsilon_t, \quad t \geq 1 \tag{29}$$

This allows me to answer my two final questions: how robust are the predictions of Theorem 1 for large transitory shocks? And do the main intuitions regarding the role of UREs survive with persistent shocks?

Asymmetric effects from increases and cuts. My first exercise maintains transitory shocks ($\rho_R = 0$) but varies the size of ϵ_0. The left panel of figure 4 shows the result of this exercise as a function of ϵ_0, for increases and cuts of up to 500bps. Two conclusions emerge. First, in my benchmark calibration with 4.5 year durations, the sufficient statistic approximation is excellent in both directions, including for large increases and cuts. Second, in the economy with one quarter duration, where the elasticity is much larger, an asymmetry emerges between the effects of large increases and large cuts: large enough *cuts* in interest rates do not stimulate consumption as much as the sufficient statistics predicts. This asymmetric effect can be traced back to the asymmetric behavior of borrowing-constrained agents to increases and falls in *income*. While these agents have to cut consumption one for one in response to income falls, their MPC out of moderate increases is below 0.3. Because their debt is short term in the ARM calibration, falls in interest rates effectively act as reductions in payments on their credit limit, and therefore as increases in income. In the aggregate, this generates an effective reduction in MPC differences that is strong enough to affect the quantitative magnitude of the redistribution channel. Increases in interest rates do not have the same feature, since the MPC of borrowers out of increases in interest payments is exactly one, as captured by the sufficient statistics.

This type of asymmetric effect of monetary policy has received empirical support (see for example Cover 1992; de Long and Summers 1988 and recently Tenreyro and Thwaites 2016). My explanation, which has to do with asymmetric MPC differences in response to policy rate changes, provides an alternative to the traditional Keynesian interpretation of this fact, which relies on downward nominal wage rigidities.[45]

Robust predictions from persistent shocks. The right panel of figure 4 displays the impact response of the economy as a function of ρ_R, for a shock to ϵ_0 of 100 basis points, under a short duration calibration and a long duration calibration (the degree of indexation of contracts is of course irrelevant). The graph decomposes the response as the sum of an income effect and a substitution effect. Consider first the output effect from a transitory shock ($\rho_R = 0$, to the left of the graph). As we already know, the benchmark calibration has a limited role for the redistribution effect relative to the substitution effect, whereas the redistribution effect is much more important in the economy with short durations. The key message of this graph is that this pattern continues to hold no matter what the persistence of the shock ρ_R is. As shock persistence grows, the substitution effect grows but the redistribution effect grows as well. However, the graph shows that it is not quite right to hold the relative sizes of these two effects fixed: in fact, in the model, redistribution

[45]While my U.S. benchmark calibration does not feature asymmetric effects of interest rates, in practice, the refinancing option embedded in fixed rate mortgages in the United States is likely to create an asymmetric effect in the opposite direction from the one I stress here. See Wong (2015) for theory and empirical evidence along these lines.

becomes more important as the shock becomes more persistent.

This result shows both the benefits and the costs of using the sufficient statistic approach. Measured sufficient statistics for transitory shocks can be informative about what would happen under more persistent shocks, at least in terms of direction. Yet they are not structural objects, so they cannot be used as elasticities that stay constant as persistence changes.

5.5 General equilibrium effects on income

The model presented above takes incomes as exogenous. This distinguishes it from the recent wave of Heterogeneous-Agent New Keynesian models which endogenize aggregate income and its distribution. My analysis shows that the endogenous distribution of income can matter a great deal for monetary policy transmission because of the earnings heterogeneity channel. As illustrated in Appendix A.1, models with sticky prices tend to generate procyclical wages, procyclical capital income and countercyclical profits. On the contrary, models with sticky wages tend to generate countercyclical wages and procyclical profits. Hence both the nature of nominal rigidities and the way in which labor, capital and profits are distributed across the population matters for the results. A successful model needs to match the empirical evidence on the cyclicality of the distribution of income by income type. Models such as those of Gornemann, Kuester and Nakajima (2012) and Kaplan, Moll and Violante (2016) make progress along these lines.

6 Conclusion

This paper contributes to our understanding of the role of heterogeneity in the transmission mechanism of monetary policy. I identified three important dimensions along which monetary policy redistributes income and wealth, and argued that each of these dimensions was likely to be a source of aggregate effects on consumption. My classification holds in many environments and provides a simple, reduced-form approach to computing aggregate magnitudes. Hence it can guide future work on the topic, both theoretical and empirical.

An important finding of my paper is that capital gains and losses, both nominal and real, matter for understanding monetary policy transmission. This finding has broad implications for monetary policy. A change in the inflation target can create large redistribution in favor of high MPC agents and be expansionary over and beyond its effect on real interest rates. With long asset maturities, lower real interest rates can benefit asset holders with lower MPCs and make interest rate cuts less effective at increasing aggregate demand than they would otherwise be. Monetary policy becomes intertwined with fiscal policy, but also with government debt maturity management and mortgage design policies.

These are just some of the macroeconomic consequences of the presence of large and heterogeneous marginal propensities to consume, which are a robust feature of household micro data. My investigation was very much a first pass, and opens up many avenues for future research on monetary policy with heterogeneous agents.

References

Aiyagari, S. Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, August 1994, *109* (3), 659–684.

Albanesi, Stefania, "Inflation and Inequality," *Journal of Monetary Economics*, May 2007, *54* (4), 1088–1114.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub, "The Intertemporal Keynesian Cross," *Manuscript*, March 2017.

Baker, Scott, "Debt and the Consumption Response to Household Income Shocks," *Journal of Political Economy*, March 2017, (forthcoming).

Barsky, Robert B., Christopher L. House, and Miles S. Kimball, "Sticky-Price Models and Durable Goods," *American Economic Review*, June 2007, *97* (3), 984–998.

Bassetto, Marco, "Optimal Fiscal Policy with Heterogeneous Agents," *Quantitative Economics*, November 2014, *5* (3).

Bewley, Truman, "The Optimum Quantity of Money," in John H. Kareken and Neil Wallace, eds., *Models of Monetary Economies*, Minneapolis: Federal Reserve Bank of Minneapolis, 1980, pp. 169–210.

Bilbiie, Florin O., "Limited Asset Markets Participation, Monetary Policy and (inverted) Aggregate Demand Logic," *Journal of Economic Theory*, May 2008, *140* (1), 162–196.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, December 2008, *98* (5), 1887–1921.

Brunnermeier, Markus K. and Yuliy Sannikov, "The I Theory of Money," *Manuscript*, August 2016.

Calza, Alessandro, Tommaso Monacelli, and Livio Stracca, "Housing Finance and Monetary Policy," *Journal of the European Economic Association*, 2013, *11*, 101–122.

Carroll, Christopher D., "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, February 1997, *112* (1), 1–55.

—, "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 2006, *91* (3), 312–320.

— **, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White**, "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, January 2017, *forthcoming*.

Chetty, Raj, "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, January 2009, *1* (1), 451–488.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, February 2005, *113* (1), 1–45.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia, "Innocent Bystanders? Monetary Policy and Inequality in the U.S.," *Manuscript*, February 2016.

Costinot, Arnaud, Guido Lorenzoni, and Iván Werning, "A Theory of Capital Controls as Dynamic Terms-of-Trade Manipulation," *Journal of Political Economy*, February 2014, *122* (1), 77–128.

Cover, James Peery, "Asymmetric Effects of Positive and Negative Money-Supply Shocks," *Quarterly Journal of Economics*, November 1992, *107* (4), 1261–1282.

de Long, J. Bradford and Lawrence H. Summers, "How Does Macroeconomic Policy Affect Output?," *Brookings Papers on Economic Activity*, January 1988, *1988* (2), 433–494.

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, 1991, *59* (5), 1221–1248.

di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao, "Monetary Policy Pass-Through: Mortgage Rates, Household Consumption and Voluntary Deleveraging," *Manuscript*, 2017.

Doepke, Matthias and Martin Schneider, "Inflation and the Redistribution of Nominal Wealth," *Journal of Political Economy*, December 2006, *114* (6), 1069–1097.

Eggertsson, Gauti B. and Paul Krugman, "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach," *Quarterly Journal of Economics*, August 2012, *127* (3), 1469–1513.

Erosa, Andrés and Gustavo Ventura, "On Inflation as a Regressive Consumption Tax," *Journal of Monetary Economics*, May 2002, *49* (4), 761–795.

Fagereng, Andreas, Martib B. Holm, and Gisle J. Natvik, "MPC Heterogeneity and Household Balance Sheets," Working Paper 2861053, SSRN, October 2016.

Farhi, Emmanuel and Iván Werning, "Chapter 31 - Fiscal Multipliers: Liquidity Traps and Currency Unions," in Harald Uhlig and John B. Taylor, eds., *Handbook of Macroeconomics*, Vol. 2, Elsevier, 2016, pp. 2417–2492.

_ **and** _ , "A Theory of Macroprudential Policies in the Presence of Nominal Rigidities," *Econometrica*, September 2016, *84* (5), 1645–1704.

Fisher, Irving, "The Debt-Deflation Theory of Great Depressions," *Econometrica*, October 1933, *1* (4), 337–357.

Galí, Jordi, J. David López-Salido, and Javier Vallés, "Understanding the Effects of Government Spending on Consumption," *Journal of the European Economic Association*, 2007, *5* (1), 227–270.

Garriga, Carlos, Finn E. Kydland, and Roman Sustek, "Mortgages and Monetary Policy," *Review of Financial Studies*, May 2016, *forthcoming.*

Gornemann, Nils, Keith Kuester, and Makoto Nakajima, "Monetary Policy with Heterogeneous Agents," Working Paper 2147841, SSRN, September 2012.

Gourinchas, Pierre-Olivier and Helene Rey, "From World Banker to World Venture Capitalist: U.S. External Adjustment and the Exorbitant Privilege," in "G7 Current Account Imbalances: Sustainability and Adjustment," University of Chicago Press, 2007, pp. 11–66.

Grossman, Sanford and Laurence Weiss, "A Transactions-Based Model of the Monetary Transmission Mechanism," *American Economic Review*, December 1983, *73* (5), 871–880.

Guerrieri, Veronica and Guido Lorenzoni, "Credit Crises, Precautionary Savings, and the Liquidity Trap," *Manuscript*, January 2015.

Guvenen, Fatih, Serdar Ozkan, and Jae Song, "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, June 2014, *122* (3), 621–660.

Harberger, Arnold C., "The Measurement of Waste," *American Economic Review Papers and Proceedings*, May 1964, *54* (3), 58–76.

Heckman, James, "Life Cycle Consumption and Labor Supply: An Explanation of the Relationship between Income and Consumption Over the Life Cycle," *American Economic Review*, March 1974, *64* (1), 188–194.

Huggett, Mark, "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, September 1993, *17* (5–6), 953–969.

Iacoviello, Matteo, "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, June 2005, *95* (3), 739–764.

Jappelli, Tullio and Luigi Pistaferri, "The Consumption Response to Income Changes," *Annual Review of Economics*, 2010, *2* (1), 479–506.

_ **and** _ , "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics*, October 2014, *6* (4), 107–136.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, December 2006, *96* (5), 1589–1610.

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, July 2014, *82* (4), 1199–1239.

_ **, Benjamin Moll, and Gianluca Violante**, "Monetary Policy According to HANK," *Manuscript*, January 2016.

_ **, Giovanni L. Violante, and Justin Weidner**, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, *2014* (1), 77–138.

Kimball, Miles S., "Precautionary Saving and the Marginal Propensity to Consume," Working Paper 3403, National Bureau of Economic Research, July 1990.

King, Mervyn A., "Debt Deflation: Theory and Evidence," *European Economic Review*, April 1994, *38* (3–4), 419–445.

Klenow, Peter J. and Benjamin A. Malin, "Chapter 6 - Microeconomic Evidence on Price-Setting," in Benjamin M. Friedman and Michael Woodford, ed., *Handbook of Monetary Economics*, Vol. 3, Elsevier, 2010, pp. 231–284.

Korinek, Anton and Alp Simsek, "Liquidity Trap and Excessive Leverage," *American Economic Review*, 2016, *106* (3), 699–738.

Krusell, Per and Anthony A. Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, October 1998, *106* (5), 867–896.

McKay, Alisdair and Ricardo Reis, "The Role of Automatic Stabilizers in the U.S. Business Cycle," *Econometrica*, January 2016, *84* (1), 141–194.

__ , **Emi Nakamura, and Jón Steinsson**, "The Power of Forward Guidance Revisited," *American Economic Review*, October 2016, *106* (10), 3133–3158.

Mian, Atif and Amir Sufi, "What Explains the 2007–2009 Drop in Employment?," *Econometrica*, November 2014, *82* (6), 2197–2223.

__ , **Kamalesh Rao, and Amir Sufi**, "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, November 2013, *128* (4), 1687–1726.

Modigliani, Franco and Richard Brumberg, "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in Kenneth K. Kurihara, ed., *Post Keynesian Economics*, 1954.

Oh, Hyunseung and Ricardo Reis, "Targeted Transfers and the Fiscal Response to the Great Recession," *Journal of Monetary Economics*, December 2012, *59, Supplement*, S50–S64.

Rotemberg, Julio J., "A Monetary Equilibrium Model with Transactions Costs," *Journal of Political Economy*, February 1984, *92* (1), 40–58.

Rubio, Margarita, "Fixed- and Variable-Rate Mortgages, Business Cycles, and Monetary Policy," *Journal of Money, Credit and Banking*, 2011, *43* (4), 657–688.

Sheedy, Kevin D., "Debt and Incomplete Financial Markets: A Case for Nominal GDP Targeting," *Brookings Papers on Economic Activity*, 2014, pp. 301–373.

Sterk, Vincent and Silvana Tenreyro, "The Transmission of Monetary Policy Operations through Redistributions and Durable Purchases," *Manuscript*, May 2015.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, June 2004, *112* (3), 695–717.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, October 2016, *8* (4), 43–74.

Tobin, James, *Asset Accumulation and Economic Activity: Reflections on Contemporary Macroeconomic Theory*, University Of Chicago Press, August 1982.

Werning, Iván, "Incomplete Markets and Aggregate Demand," Working Paper 21448, National Bureau of Economic Research, August 2015.

Wong, Arlene, "Population Aging and the Aggregate Effects of Monetary Policy," *Manuscript*, December 2015.