HANKSSON*

Florin O. Bilbiie   Sigurd Mølster Galaasen   Refet S. Gürkaynak
Mathis Mæhlum   Krisztina Molnar

November 2024

Abstract

HANK Sufficient Statistics Out of Norway (HANKSSON) answers a core question of the heterogeneity in macroeconomics literature theoretically and empirically: does heterogeneity *amplify* the aggregate effects of demand shocks and policies. We provide two sufficient statistics (SS) and test these using individual-level matched data for personal characteristics, income, wealth and consumption for the Norwegian population. The first SS gauges whether heterogeneity drives a wedge between the (representative-agent) average MPC and a model-consistent (heterogeneous-agent) aggregate MPC. The second SS elicits whether the consumption of constrained, "hand-to-mouth" agents is more exposed to aggregate fluctuations. Our robust key finding is that to analyze aggregate behavior, one does not need to keep track of heterogeneity: the average and the aggregate are about the same. Along the way we show that the amplification result currently prevalent in the literature is due to using labor earnings and is overturned when using model-consistent disposable income. This is due to the strong insurance effect of taxes and transfers; even the much less progressive US tax and transfer system produces no amplification due to heterogeneity. The same "close to irrelevance" conclusion arises based on the second statistic using consumption data directly. Not even during the Great Recession do we see heterogeneity contribute meaningfully to demand shock amplification.

JEL Codes:

Keywords: heterogeneity; inequality; HANK; MPCs; hand-to-mouth; multipliers; consumption; income; wealth; aggregate demand.

*Bilbiie: University of Cambridge and CEPR; Galaasen: Norges Bank; Gürkaynak: Bilkent University and CEPR; Molnar: Norwegian School of Economics NHH; Mæhlum: Norges Bank. We are grateful for comments to Adrien Auclert, Anmol Bhandari, Corina Boar, Tobias Broer, Marco d'Amico, Davide Debortoli, Axelle Ferrière, Joel Flynn, Jordi Gali, Diego Känzig, Simon Mongey, Christina Patterson, Fabrizio Perri, Paolo Surico, Gianluca Violante. Bilbiie gratefully acknowledges without implicating the support of the Janeway Institute and Keynes Fund at Cambridge, and the hospitality of CREi Barcelona and the Paris School of Economics. The views do not represent Norges Bank.

1 Introduction

Does household heterogeneity contribute to an amplification of the aggregate effects of demand shocks and policies? This is a core question of a recent flourishing literature that studies the interaction of distributions and aggregates in shaping macro fluctuations and the transmission and design of monetary and fiscal policies. A wide range of such "HANK" (heterogeneous-agent New Keynesian) models makes predictions regarding the mechanisms driving amplification, which are relevant for monetary and fiscal policy. At the core of all models in the HANK class, including their simpler two-agent (TANK) versions, lies a two-way feedback between aggregate macro variables and distributions of individual variables, or "inequality", which shape each other and are determined jointly in general equilibrium through a variety of economic mechanisms. Several theoretical and/or quantitative contributions shed light on such mechanisms and derived implications for certain "sufficient statistics" that drive some of the key co-movements.

In this project, we tackle the *measurement* of some of these key statistics, informed by some simple organizing theory that captures several key mechanisms of the larger heterogeneous-agent model class: are these mechanisms salient in the data, and if so which heterogeneity dimensions are most important. The purpose of this measurement is to document and test the empirical relevance of these channels, in isolation and jointly, using a unique dataset from Norway that contains individual-level information on consumption, income, and wealth.[1] The key predictions and transmission channels in the majority of this literature concern the distribution and evolution of *consumption*. Yet notorious data issues make the measurement of this based on the traditional data sets challenging. The literature often thus uses earnings or, at best, income data, even though the relevant variable for both transmission, welfare, and optimal policy prescriptions is, in fact, consumption. It is thus of first-order importance to test whether these mechanisms are salient using actual consumption data; making progress on this front is one of the contributions of this paper.

The main transmission channels we focus our measurement on are as follows. A key amplification channel is that individual variables (consumption, income, earnings) vary over the cycle in a manner that is correlated with that individual's (or household) marginal propensity to consume (MPC). Specifically, if the income or consumption of high-MPC individuals falls more in a recession, this amplifies the recession itself through a Keynesian-cross mechanism common to much of the heterogeneous agent (HA) literature. This is described theoretically in Bilbiie (2008), Bilbiie (2020), Bilbiie (2018), Auclert (2019), and at work in quantitative HANK such as i.a. Oh

[1] Despite the explosion of research using administrative data, examining consumption and income links proved more difficult and relied mainly on survey data. We solve this data challenge by linking transaction-based consumption data to administrative data in Norway. The advantage of our data is twofold. First, our consumption data is representative and covers realized debit card and transfer expenditures. Second, we are able to link this data to high-quality administrative data that allows us to follow disposable income, saving choices, and wealth of households over time.

and Reis (2012), Kaplan et al. (2018), Bayer et al. (2019), Gornemann et al. (2016). Patterson (2023) provides a novel measurement of this amplification of business cycle shocks but crucially based on earnings, not disposable income or consumption. A key contribution of this paper is to bring both disposable total income (including capital income and taxes and transfers) and consumption data to inform that measurement.[2]

We therefore provide two sufficient statistics and measure them using the rich consumption, income, and wealth data through the model lens. The first statistic consists of computing the "Aggregate MPC" (of Norway), which combines individual MPCs and–to use a terminology that has been borrowed from finance–individual "betas" (elasticities of the individual variables to the aggregate variable). We estimate the *MPC distribution* of Norway and find its most salient determinants, which turn out to be *liquid* wealth and education. We then compute their **betas** for labor earnings and total disposable (post tax and transfer) income. Combining the two delivers the first sufficient statistic, the **Aggregate MPC**, an object we define formally below that takes into account not only the average MPC level but also the general-equilibrium cyclical individual interactions. This in turn allows to compute the "multiplier" and ascertain whether, through the lens of this "aggregate MPC" based on disposable income, there is *Aggregate Demand amplification* in the Norwegian economy: have recessions been amplified through these inequality/heterogeneity mechanisms?

The second sufficient statistic uses consumption betas directly, together with agents' hand-to-mouth status based on their wealth data, as a way to ascertain whether agents are on their Euler equations or not. This allows us to answer the same question more directly, as we review theoretically below.. These two statistics determine whether heterogeneity amplifies or dampens aggregate fluctuations. They are equivalent when net savings are zero, but differ when assets are in positive supply. In the latter case, the sufficient statistic based on consumption betas is the correct measure rather than the one based on income betas. With our comprehensive data we study the dynamics of the Norwegian economy through the lens of these sufficient statistics.

In a nutshell, we find that measurement based on labor earnings (betas) would lead us to conclude through the model's lens that there is aggregate-demand amplification at work in Norway. However, measurement based on total disposable income instead completely undoes that amplification and points to near-irrelevance of these heterogeneity mechanisms for aggregate fluctuations. Part of this undoing is due to the role of capital income (which is cyclical and accrues to low-MPC individuals), and part to the tax-and-transfer system whose in-built progressivity insures particularly high-MPC individuals. Since the latter plays a large role in our quasi-irrelevance result, it is legitimate to ask whether "is this just Norway": would using a counterfactual US tax system instead undo our result, i.e. preserve the amplification found when using labor earnings. We conduct such a counterfactual and show that while using the tax system qualifies our result quantitatively, to a first order and qualitatively the quasi-neutrality still obtains.

Even with perfectly insured disposable incomes, individual consumptions may still fluctuate differently over the cycle as households make different investment and precautionary saving and

[2] Patterson (2023) imputes consumption using the Blundell et al. (2008) method.

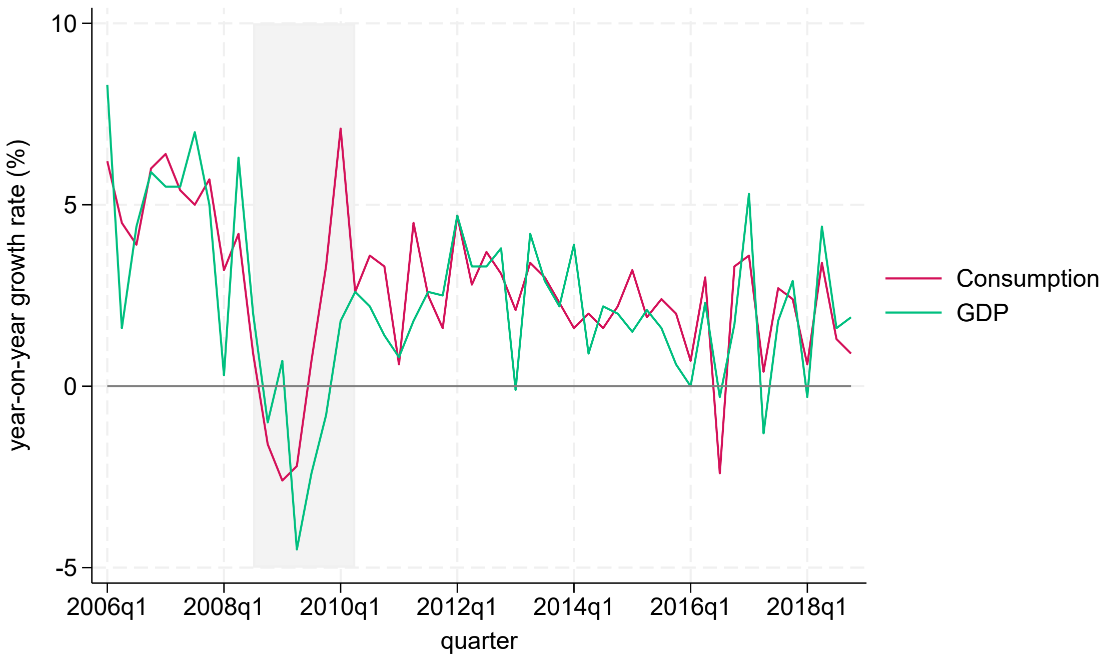

liquidity choices. Thus finally, we leverage our high-quality consumption data to compute the second, complementary but different, direct sufficient statistic. Namely, we elicit whether the consumption of agents who are more likely to be "hand-to-mouth" (their consumption does not obey an Euler equation) is more "cyclical", or more responsive to aggregate consumption—be it over the whole sample or focusing on the Great Recession episode; we illustrate the magnitude of the Great Recession (for GDP and consumption) in Norway in Figure 1. Using a wide range of definitions of "hand-to-mouth" based on wealth, liquidity, and education, we find that this is hardly the case: consumption of both the HtM and non-HtM groups, regardless of the splitting criterion, behave in very similar ways. The most unequal incidence in consumptions across groups can be seen when zooming in on the Great Recession episode: but even then, through the lens of our simple model, an upper bound on the degree of amplification yielded by heterogeneity is of the order of 10%.

This paper is most related to—but significantly further builds on—the seminal paper by Patterson (2023). We share the focus on an attempt at measuring the correlation between estimated MPCs and individual "cyclicalities" as a way to ascertain whether the theoretical amplification mechanisms are salient in the data. There are several differences. First, the coverage and quality of our transactions-based consumption data allows a more precise estimation of the MPCs (which are nevertheless reassuringly in line with those estimated by Patterson). Then, the ability to match this with wealth data allows us to ascertain whether among MPC determinants wealth plays a major role: we find that it does, but it in in fact "liquid wealth" which is a predictor of MPCS. Indeed, looking at net worth to predict an individual's MPC can even be misleading; this informs the "wealthy hand to mouth" literature pioneered by Kaplan et al. (2014) and of the essence in HANK models. Third, the availability of administrative disposable-income data, i.e. both capital and taxes and transfers, allows us to study the cyclicalities of individual incomes: we find that even though—in line with Patterson—the cyclicality of earnings is higher for high-MPC individuals, this is overturned when looking at net income. Partly due to capital income (which is very cyclical and accrues mostly and naturally to low-MPC savers) and partly due to transfers (and less so to taxes) which act as an automatic-stabilizer insurance device, the cyclicality of disposable incomes is in fact slightly lower for high-MPC individuals. To a first order, nevertheless, the "aggregate MPC" (which takes into account general-equilibrium feedback) is very close to a pure average MPC. Finally, our data allows us to ascertain more directly whether aggregate-demand amplification of the type predicted by HANK and TANK-type models is of the essence empirically; the (observable, given our data and estimation) consumption of agents who participate in the market for liquid assets obeys an Euler equation, which prices those assets. Whereas the consumption of non-participants (comprising the "wealthy" HtM) does not. The inequality between the two groups' consumption is a sufficient statistic for measuring whether the amplification mechanism is operating in the data. We perform this measurement, by estimating at the disaggregated level "consumption betas" across the MPC distribution: while there is somewhat more incidence for the HtM agents, the difference in the cross-section is not very large; the most difference in incidence can be found when focusing on the Great Recession, but it is still nowhere near what would be concluded based on the incidence of labor earnings.

Therefore, to a first order and through a variety of channels, consumption insurance seems to be at work in Norwegian data in a way that precludes heterogeneity-based, distributional channels to act as amplifiers of business cycles.[3] The reverse side of this is that this also implies that AD management, monetary and fiscal policies are less effective; nevertheless, other more direct tools such as targeted transfers to high-MPC individuals still have high AD effects even in such economies.

A final remark pertains to the applicability of our analysis beyond Norway. First, it can be reasonably argued that Norwegian consumption is a lower bound of consumption inequalities in other countries, such as the US or developing nations. While Norway consistently ranks among the countries with the lowest income inequality, we show that consumption is nevertheless unequal (thus, one can expect consumption to be even more unequal in higher-income-inequality countries). Furthermore, Norway's welfare system significantly reduces inequality through unemployment benefits and lower income taxes, especially for the poorest; and it offers extensive social services like healthcare and education to those unable to afford them. Lack of such social services in countries like the US would naturally lead to higher consumption inequality.

Related literature A first literature we contribute to is the burgeoning literature on TANK and HANK models to analyze fluctuations and policies. These include earlier TANK models (Galí et al. (2007), Bilbiie (2008)), rich-heterogeneity HANK models (i.a. Oh and Reis (2012), McKay and

[3] To draw on a famous quote by Enrico Fermi (*"There are two possible outcomes: if the result confirms the hypothesis, then you've made a **measurement**. If the result is contrary to the hypothesis, then you've made a **discovery**."*), one could argue that instead of a measurement we ended up with a "discovery".

Reis (2016), Kaplan et al. (2018), Gornemann et al. (2016), Auclert (2019), Ravn and Sterk (2017), Den Haan et al. (2017), Luetticke (2021), Bayer et al. (2019), Auclert et al. (2018), Hagedorn et al. (2019)), and tractable-HANK models leveraging insights from the former class to capture a subset of the channels of the latter (i.a. Bilbiie (2020), Bilbiie (2018), Broer et al. (2020), Debortoli and Galí (2024), Acharya and Dogra (2020), Ravn and Sterk (2020), Challe et al. (2017), Bilbiie et al. (2022b), Debortoli and Galí (ming)). Our contribution to this literature is to provide measurement using high-quality Norwegian data.

In the realm of measurement purely, we also contribute to the vast empirical literature on MPC estimation, by using representative actual consumption and wealth data, where the latter allows us to elicit the determinants of MPC heterogeneity. Earlier literature used survey responses, imputed consumption, or, more recently, data from a particular bank, or experimental data. An incomplete list includes the seminal studies of Souleles et al. (2006), Shapiro and Slemrod (1995, 2009); Sahm et al. (2010), Gruber (1997); Jacobson et al. (1993), Blundell et al. (2008), as well as more recent attempts such as Parker et al. (2013), Broda and Parker (2014), Christelis et al. (2019); Jappelli and Pistaferri (2014, 2020), Coronado et al. (2005), Fuster et al. (2021), Misra and Surico (2014), Kueng (2018), Commault (2022), Lewis et al. (2024), Fagereng et al. (2021, 2024a), Orchard et al. (2025), and Boehm et al. (2025). A related literature studies the link between MPC heterogeneity and "hand-to-mouth" behavior, including attempts to elicit the determinants of the latter based on net worth, liquid wealth, housing ownership status, or liquid assets; key references include Campbell and Mankiw (1989), Kaplan et al. (2014); Kaplan and Violante (2014), Cloyne et al. (2020), Aguiar et al. (2023), and Ganong et al. (2020).

Finally, we contribute to the literature providing empirical evidence pertaining to inequality over the cycle, and to the related notion of income risk. Seminal empirical contributions include Heathcote et al. (2010), Heathcote et al. (2023), Storesletten et al. (2004), Attanasio and Pistaferri (2016), Guvenen et al. (2014), Coibion et al. (2017), Guvenen et al. (2021). Estimated HANK models including such cross-sectional time series to elicit the role of heterogeneity channels include Bayer et al. (2020), Auclert et al. (2020), and Bilbiie et al. (2022a). Different measurement exercises using micro data include the Patterson paper reviewed above and Pekanov (2024); the latter uses EU data from several sources and obtains similar results to the former, pointing to amplification through the estimated earnings betas. Related exercises focusing on the unequal incidence of monetary policy shocks using micro data include Holm et al. (2021) for Norway and Coglianese et al. (2024) for Sweden. Berger et al. (2023) instead measure "heterogeneity wedges" to elicit the amount of risk sharing in US CEX data. In the realm of consumption risk sharing, our paper is also related to the literature pioneered by Cochrane (1991); Mace (1991); Townsend (1994) . To these literatures, what our paper brings is the use of (representative) transactions consumption data and matched administrative gross and net income data, as well as wealth data.

2 Some Organizing Theory: Defining Sufficient Statistics

In this section, we define the sufficient statistics that inform a large class of heterogeneous-agent NK models. We do so by drawing on the class of tractable, two-agent models reviewed in the introduction, and in particular on the analytical model(s) in the two-agent class in Bilbiie (2008, 2020), generalized by Auclert (2019) to HANK models with richer heterogeneity. The insights nevertheless transcend this simple framework when it comes to measurement of the key (for aggregate amplification) objects, as we discuss below.

The key question in the literature in terms of the aggregate effects of heterogeneity is whether any given impulse is amplified or dampened due to the presence of heterogeneity. In a simple Keynesian framework, which extends naturally into NK settings, the effect on impact of any shock is the direct effect times a multiplier, which depends on the MPC. In a RANK model the relevant MPC is the average one, in a model with non-trivial heterogeneity (not having full insurance), the model consistent aggregated MPC will depend on the distribution of individual MPCs, as we will see below.

This is a good place to emphasize the usage of "aggregated," a term we will employ often. In theoretical settings we will use aggregated to denote model consistent aggregates rather than averages, similar to the usage of the correct price index in New Keynesian models given the intratemporal consumption aggregator. In empirical settings we will use it to denote aggregates correctly built up from the micro data that we will use, rather than corresponding quantities in national accounts.

2.1 Sufficient Statistic 1: The Aggregate MPC and Multipliers

To set the stage, it is useful to start from the basics and classics: namely, the "Keynesian Cross" representation of aggregate demand (AD) pioneered by Samuelson (1948). This starts by postulating a generic aggregate consumption function, or "Planned Expenditure":

$$C = C\left(Y; R, ...\right)$$

where C is consumption and Y a measure of total disposable income—thus, the slope of this curve is the marginal propensity to consume $MPC \equiv \frac{\partial C}{\partial Y}$. R is the gross real interest rate; there are several other possible arguments to this function, such as other aggregate-demand management tools (government spending, transfers) or indeed exogenous shocks (impatience, deleveraging, financial disruptions). We take a change in interest rates as a prototypical AD shifter, but our conclusions translate to any AD shocks and policies.

The effect of an interest rate change dR (or indeed any AD shift) is then found by differentiating the consumption function:

$$\frac{dC}{dR} = MPC\frac{dY}{dR} + \frac{\partial C}{\partial R},$$

where $\frac{\partial C}{\partial R}$ is the "autonomous expenditure", PE curve shifter. In other words, it is the partial-equilibrium, "direct" effect of the shock/policy on aggregate demand.

To find the general-equilibrium, total aggregate effect, we need to make extra assumptions and impose market clearing. In the simplest case where there is no saving in equilibrium, we thus only need to impose that consumption equals total disposable income $C = Y$. Adding this, we can see **General-Equilibrium** AD amplification (which again will apply to any demand shock or policy):

$$GE \text{ Multiplier } \frac{dC}{dR} = \frac{1}{1 - MPC} \frac{\partial C}{\partial R} \tag{1}$$

The shift in the PE curve $\frac{\partial C}{\partial R}$only gives us the impulse and initial, direct impact. But since this initial expansion creates income, of which the agent consumes MPC, this creates a further expansion of $MPC * \frac{\partial C}{\partial R}$ and so on for an infinity of rounds with decaying powers of MPC: the classic Keynesian-cross *multiplier*.

The "*New*" Keynesian cross (Bilbiie (2020)) operating in heterogeneous-agent, HANK models adds the following layer of complexity: to find the aggregate, economy-wide MPC of an economy populated by many agents we need to aggregate them carefully. Start from an individual j "MPC" out of *aggregate* income:

$$\frac{dC^j}{dY} = \frac{dC^j}{dY^j} \frac{dY^j}{dY} = MPC^j \frac{Y^j}{Y} \frac{d \ln Y^j}{d \ln Y} = MPC^j \frac{Y^j}{Y} \beta_y^j, \tag{2}$$

where the first equality uses only the chain rule, the second uses the definition of $\frac{dC^j}{dY^j}$ as the MPC out of *own* income of agent j and the third uses $\frac{Y^j}{Y}$ as the income share and rewrites the change in individual income to aggregate income in elasticity form. In particular, it introduces a crucial object, denoted by β^j, that we refer to as "beta", the elasticities of individual variables to aggregate ones, in this case elasticity of individual incomes to aggregate income:

$$\beta_y^j \equiv \frac{\partial y_t^j}{\partial y_t}, \tag{3}$$

where y_t is log aggregate income and y_t^j is log individual income.

The **aggregate MPC**, i.e. the general-equilibrium object relevant for macro amplification, which takes into account how the income distribution changes when aggregate income changes, is then found by summing these individual MPCs out of aggregate income, that is:

$$\mathbf{A}MPC \equiv \frac{dC}{dY} = \sum_j \frac{dC^j}{dY} = \sum_j MPC^j \frac{Y^j}{Y} \beta^j. \tag{4}$$

This is the relevant object for macroeconomic amplification of aggregate demand shocks and policy impulses, shown to capture the aggregate demand amplification in general equilibrium—in a manner that is conceptually similar to Samuelson's staple "Keynesian Cross" representation studied in TANK and HANK models by Bilbiie (2008, 2020) and Auclert (2019), respectively. Patterson (2023) was the first to attempt to empirically measure this important object.

Hence, the aggregate MPC measured as above is a sufficient statistic for the general equilibrium multiplier. We use these simple expressions to discipline our estimation and to compute the AMPC and the GE multiplier as sufficient statistics below, after we estimate the individual components. The sufficient statistic to have an amplification effect due to heterogeneity here is:

$$Cov(MPC^j, \beta^j) > 0 \Longleftrightarrow \mathbf{A}MPC > aMPC|_{\beta_y^j=1}. \tag{5}$$

Amplification occurs when the aggregate MPC just defined is larger than the simple "average" MPC, obtained by merely aggregating individual MPCs using the population shares, i.e. assuming that all individual incomes are proportional to the aggregate one, $\beta_y^j = 1$ for all j, $aMPC \equiv \sum_j MPC^j \frac{Y^j}{Y}$. Naturally, when the income of higher-MPC agents is more cyclical, an expansion is amplified as an aggregate income change is distributed disproportionately to higher-MPC agents, who shift their demand further engendering a further expansion round and thus a Keynesian-cross-like multiplier.

Therefore, the general-equilibrium amplification for *any* demand shock (monetary, fiscal, financial, etc.) is given by a multiplier-like object that is remarkably similar to the old-fashioned Keynesian cross representation, namely:

$$GE \text{ Multiplier} = \frac{1}{1 - \mathbf{A}MPC}. \tag{6}$$

To ascertain the degree of amplification relative to a case with uniform betas, where the aggregate MPC is equal to the simple weighted average (denoted $aMPC$ above), we need to simply compare the above multiplier with its counterpart in that case, given by $1/\left(1 - aMPC\right)$.

To fix ideas by means of a simple example, in an analytical TA(NK) framework with measure λ hand to mouth (HtM, denoted H) consumers and $(1-\lambda)$ savers (S), who follow Euler equations and may in principle save although we do not yet allow non-zero net assets, the aggregate MPC is:

$$\mathbf{A}MPC = \lambda \times MPC^H \times \frac{Y^H}{Y} \times \beta_y^H + (1-\lambda) \times MPC^S \times \frac{Y^S}{Y} \times \beta_y^S$$

Since in this model $MPC^H = 1$ and MPC^S is a very low number (S are permanent-income consumers with an MPC of $1 - e^{-\rho}$, where ρ is the discount rate), assuming without loss of generality that in a long-run steady state $Y^H = Y^S = Y$, we infer that the general-equilibrium, multiplier-like amplification will be governed my an aggregate MPC of the form $\lambda \beta_y^H$. Indeed, it can be easily shown (see Bilbiie (2008, 2020)) that in this model the general-equilibrium multiplier of an interest rate cut is given by (where small letters denote log deviations from steady state):

$$\frac{dc_t}{d\left(-r_t\right)} = \sigma \frac{1-\lambda}{1-\lambda\beta_y^H}. \tag{7}$$

Indeed, in this expression σ is the elasticity of intertemporal substitution, which is also the general-equilibrium multiplier of a one-time interest-rate cut in the RANK model. The "direct" (in the Kaplan et al. (2018) sense) effect in TANK is scaled down, since only $1-\lambda$ agents are directly exposed to interest rates (have a demand for assets). Therefore, the equivalent of the "autonomous expenditure" term defined above as $\frac{\partial C}{\partial R}$ is $\sigma(1-\lambda)$. However, there are also indirect effects amplifying this initial direct effect through the (new) Keynesian-cross mechanism described above, manifesting in the expression in the denominator, with $\lambda\beta_y^H$ as the aggregate MPC. There is amplification when these indirect effects governed by $1/\left(1-\lambda\beta_y^H\right)$ dominate the direct effects $1-\lambda$, i.e. whenever $\beta_y^H > 1$, which trivially makes the aggregate MPC $\lambda\beta_y^H$ larger than the average MPC λ, because the high-MPC agent's income reacts disproportionately to aggregate income. The covariance condition 5 is a many-agent generalization of this simple logic.

2.2 Sufficient Statistic 2: "Consumption betas"

When high-quality consumption data is available at the micro level as a panel, there is a direct way to ascertain the amplification properties, relying on a sufficient statistic that captures consumption inequality and its dynamics in connection with the aggregate, i.e. its cyclicality. To illustrate this, consider again the simplest analytical TANK framework with hand-to-mouth (H) and potentially saving agents (S) as above. The loglinearized Euler equation governing consumption of the latter group (who hold and price all the assets in this simple economy) is:

$$c_t^S = E_t c_{t+1}^S - \sigma r_t, \tag{8}$$

where $c_t^S$is the log consumption of savers, c_t^Hthat of HtM agents, σ the elasticity of intertemporal substitution and $r_t$the real interest rate. Define consumption inequality as $\Gamma_t \equiv \frac{C_t^S}{C_t^H}$,and in log deviations:

$$\gamma_t = c_t^S - c_t^H. \tag{9}$$

Rewriting the aggregation equation (loglinearized) $c_t = \lambda \frac{C^H}{C} c_t^H + (1-\lambda) \frac{C^S}{C} c_t^S$ using this definition, we have:

$$c_t = c_t^S - \tilde{\lambda}\gamma_t, \tag{10}$$

where we defined the re-normalized (inequality-weighted) population share of H by$\tilde{\lambda} \equiv \frac{\lambda}{1+(1-\lambda)(\Gamma-1)}$

We are now ready to show that consumption inequality is indeed a sufficient statistic for demand amplification due to heterogeneity in this class of models. To do so, replace the last equation in the savers' Euler equation to obtain the aggregated Euler equation:

$$c_t = E_t c_{t+1} - \tilde{\lambda}\left(\gamma_t - E_t\gamma_{t+1}\right) - \sigma r_t. \tag{11}$$

This has been used as an organizing device to summarize how an analytical TANK model can approximate different departures from RANK of richer, quantitative HANK models, e.g. in Bilbiie (2020) and Debortoli and Galí (2024).

Proposition: There is AD *amplification* iff consumption inequality is *countercyclical, i.e.:*

$$\gamma_c \equiv \frac{d\gamma_t}{dc_t} < 0 \iff \beta_c^H > \beta_c^S, \tag{12}$$

where consumption betas β_c^j are defined, similarly to earnings betas and income betas defined in 3, as the elasticity of individual j's consumption to aggregate consumption, i.e. $\beta_c^j \equiv \frac{\partial c_t^j}{\partial c_t}$.

Now, by virtue of our data set containing both consumption and *wealth* data we can split the population along various dimensions, and assess the amplification properties through the lens of the model. Notice that by this strategy no MPC estimation is, in fact, needed (other than if we wish to split the population according to their MPC levels). In fact, no shares are needed either (Γ only matters for the magnitude, not for the sign of the response). The only object we need to calculate are "consumption betas".

Taking again a one-time cut in interest rates as the prototypical AD impulse (and recalling that a similar logic applies to any AD shock or policy), the GE multiplier is then given by:

$$\frac{dc_t}{d(-r_t)} = \sigma \frac{1}{1 + \tilde{\lambda}\gamma_c} \tag{13}$$

We can therefore determine whether there is amplification through heterogeneity channels, i.e. relative to a representative-agent benchmark (or a complete-market, perfect-insurance benchmark whereby betas are uniformly unity, $\beta_c^H = \beta_c^S = 1$) where the multiplier is merely σ, by directly computing $\left(1 + \tilde{\lambda}\gamma_c\right)^{-1}$ and assessing whether it is larger than 1.

It is important to emphasize that the key statistic "cyclicality of consumption inequality γ_c" is model-dependent and shaped by whatever drives both the cyclicality of the income distribution (inequality) and the mapping from income to consumption at the individual level. Yet it is useful to notice that in the simplest case when there is no saving—consumption is equal to income, $C_t^j = Y_t^j$ for each agent and on aggregate—this sufficient-statistic is exactly equivalent to the sufficient-statistic derived above based on incomes and MPCs. In particular, assuming without loss of generality a symmetric steady state $\Gamma \equiv \frac{C^S}{C^H} = 1$, we have $\lambda = \lambda$ and $\gamma_c = \frac{1-\beta_y^H}{1-\lambda}$. Replacing this in the multiplier expression (13) shows directly that it is identical to the multiplier derived based on sufficient-statistic 1, found in expression (7). In other words, the statistic derived based on disposable income is still "sufficient" to learn about the properties of consumption, under the zero-savings assumption.

2.3 Income vs Consumption: The role of savings

The brief derivation above assumes no net assets and has consumption equal to income in equilibrium, i.e. there are no net savings. In that model, the statistic based on disposable income is indeed "sufficient" to draw inference for the properties of consumption. This is no longer true (that is, the statistic will be "insufficient") when consumption is not equal to income. In a richer model, and certainly in the data, savings are non-zero and the cyclicalities of various forms of saving can play a key role by driving a wedge between the cyclicalities of income and consumption at the individual level. To illustrate this, consider the following simple formalization (Appendix A provides a more complete analysis, including the case with idiosyncratic risk).

Denoting net saving/investment by X_t, the savers' log-linearized budget constraint (with $Z_Y \equiv Z/Y$ the share of variable Z in total income, $Z \in \{C, X, Y^S\}$):

$$C_Y c_t^S + \frac{X_Y}{1-\lambda} x_t = Y_Y^S y_t^S, \tag{14}$$

while for the HtM we still have:

$$c_t^H = y_t^H. \tag{15}$$

Through some simple algebra outlined in the Appendix A, and using individual income elasticities β_y^H, we obtain that consumption inequality now depends not only on aggregate income, but also on net saving:

$$\gamma_t = c_t^S - c_t^H = \frac{1-\beta_y^H C_Y}{(1-\lambda)C_Y} y_t - \frac{X_Y}{(1-\lambda)C_Y} x_t. \tag{16}$$

This illustrates transparently that there can be **amplification** (γ_t countercyclical) despite $\beta_y^H < 1$ if and only if net savings (be they in the form of productive investment or liquidity, etc.) are procyclical *enough*. We will now use these theoretical insights to study Norwegian data and empirically asses to what extent heterogeneity may be creating demand amplification.[4]

3 Data Description

Our study draws on several administrative records, all merged with a panel of individual-level spending based on electronic transactions. In this section we describe the various data sources and sample selection criteria.

3.1 Administrative Data

The administrative data is collected from Statistics Norway and covers the universe of Norwegian residents aged 16 and above over the period 1993-2018. The uniqueness of the Norwegian data lies in its combination of comprehensive coverage, granularity, and high quality. The administrative

[4]This simple derivation captures the insights of a part of the HANK literature that studied the role in shaping amplification of either cyclical liquidity (Werning 2015, Bilbiie (2018)) or investment in physical capital (Auclert et al. (2020), Bilbiie et al. (2022b))

records provide near-complete population coverage, along with detailed demographic information and data on various sub-components of income and wealth. Since the balance sheet data are derived from tax records, the potential for measurement issues is minimal compared to survey data.[5]

We rely on data from multiple records, all linked using unique anonymized personal identifiers. From various population registers, we obtain demographic information such as age, gender, and education. These records also contain data on family relationships, allowing us to link individuals living together in a household.

The demographic information is then combined with detailed income data from third-party reported annual tax records. This source provides a breakdown of individual annual after-tax income into labor earnings, capital income, transfers, and taxes. We consider several income measures based on this data. Our baseline measure, similar to Heathcote et al. (2010), is individual gross labor earnings (salaries), which includes zeroes (no labor earnings) as well. As a result, our baseline inequality estimates are influenced by movements in and out of employment. We also consider four, more comprehensive, income measures: (i) the sum of gross labor earnings and income (net of costs) from self-employment,[6] (ii) the sum of gross labor earnings, income (net of costs) from self-employment and capital income, (iii) total income pre-tax (equal to (ii) + government transfers), and finally (iv) total income after tax.

Due to the presence of a wealth tax, the tax records also include detailed balance sheet information on wealth and its components, such as deposits, financial assets, real assets including real estate, cars, and private business wealth.

Appendix B contains detailed information on data construction, definitions and sample restrictions. In particular, following the literature (Heathcote et al. (2010) (XXX Others? XXX)), we focus on the prime working age population and restrict our sample to individuals between 25-55 years.

3.2 Consumption data

Our consumption spending measure is constructed from an electronic transactions database, detailed in Ahn et al. (2024). The data, sourced from the Norwegian retail clearing institution Nets Branch Norway (hereafter referred to as Nets), covers all Norwegian residents during the time period 2006-2018. Having access to this granular, observed rather than imputed consumption data is one of the ways that this paper differs from the rest of the literature asking similar questions.

The dataset includes two primary payment types for consumption for each individual: debit card transactions processed through the BankAxept system and online bank wire transfers cleared via the Norwegian Interbank Clearing System (NICS). Although credit card transactions are not directly recorded, the granularity of the data enables us to infer credit card spending from online

[5] Fagereng et al. (2020); Holm et al. (2021); Ring (2024) are some recent applications of these data on household savings and consumption behavior.

[6] Self-employment is defined as reporting income from typically sole proprietorship firms, and includes net income from agriculture and forestry, fishing and hunting, income from other business activities, and sick pay in business activities during the calendar year. Deducted from this is the annual loss in the business.

wire transfers where the recipient is a bank.[7] While our card and transfer data cover the majority of electronic payments, it nevertheless excludes debit card payments processed by VISA and Mastercard (i.e. online and abroad debit card payments) as well as wire transfers not cleared by NICS. Nonetheless, as shown in Ahn et al. (2024) our payment sources account for roughly 80 percent of all electronic payments made by the household sector between 2006-2018 and there is no reason to think that the excluded 20 percent has any systematic difference.

When measuring spending based on electronic payments, a potential concern is that unobserved cash payments may vary systematically over time or across individuals. However, this concern is mitigated in the context of Norway, which has long been a near-cashless society. Survey data from 2017 indicates that cash transactions accounted for only 10 percent of all point-of-sale transactions in Norway Norges Bank (2023). In contrast, across the Euro area the average was 80 percent Esselink and Hernández (2017).[8] Additionally, our data include cashback transactions (withdrawals made at points of sale), which account for 18 percent of total cash withdrawals Norges Bank (2023). We include this in our consumption spending measure.

Overall, our aggregate consumption measure closely aligns with the household consumption growth reported in the Norwegian national accounts. In Figure 3.2 we compare the quarterly growth of nominal household consumption in the national accounts with the growth of total electronic spending. The national accounts series is constructed as domestic household consumption, excluding imputed owner-occupied housing. Details on the cleaning of the electronic transactions data are presented in Appendix B.[9] The correlation of the two time series is 0.83 over the whole sample period.

We merge the electronic transactions data with the administrative records presented above. This gives us a data set that has individual-level information on characteristics such as gender, education, employment, cohabitation status; income; and consumption for essentially the whole country. We perform a series of sample selections

3.3 Institutional details of Norwegian welfare state

Before estimating MPCs, we outline key features of the Norwegian setting. Norway's welfare state relies on a mandatory, comprehensive system of taxes and transfers.

Norwegian income is taxed progressively and with a broad tax base. Pensions, social benefits, rental income, and even lottery winnings, gifts, and inheritances are taxed. Wealth is taxed above a threshold, and capital gains are taxed upon realization. Figure 3.3 shows the tax system's progressivity in our data. The gap between pre- and post-tax income widens across deciles, reflecting

[7] As explained in Appendix B, in the raw data, transactions are observed at weekly frequency across all Norwegian postal codes (address of recipient) and separated into 26 different categories. One such category is payments made to banks.

[8] In terms of value of transactions, the corresponding averages was 3 and 53 percent in Norway and the Euro area respectively.

[9] As explained in B, we exclude imputed mortgage payments, very large single transactions (above 12,500 USD in 2018 dollars), and person-to-person online transfers.

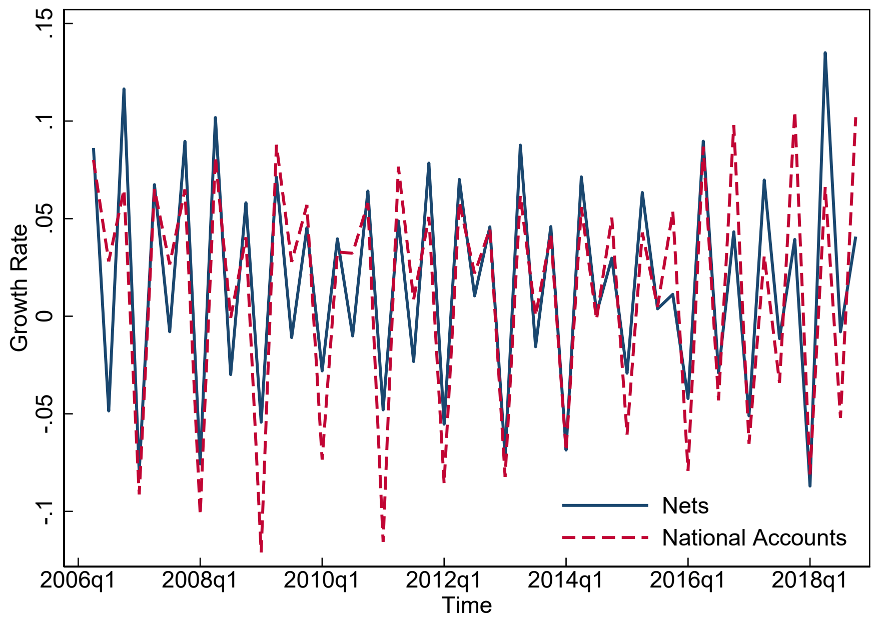

rising average tax burdens at higher incomes.\footnote{A flat 22% tax applies to labor and capital income, with additional bracketed marginal taxes reaching 17.7%.} By international standards, taxation is high, consistently above the OECD average. Since the 1970s, tax revenue has remained $40 - 45\%$ of GDP.

A core feature of Norway's welfare state is its extensive transfer system, including unemployment insurance (UI), pensions, disability pensions, child allowances, and parental leave. Figure 3.3shows how transfers vary across income groups. Government support plays a larger role for low-income households, reflecting the system's redistributive nature. Between 2006 and 2016, households in the bottom income decile received, on average, half of their income from transfers. For most, labor income remains the primary source of earnings, while at higher income levels, transfers largely consist of pensions. As part of the welfare state, education and healthcare are free.

Despite Norway's extensive welfare system, unemployment leads to substantial income losses. The unemployed lose one-third to one-half of their income, with losses that persist—recovery typically begins only after two years (see Fagereng et al. (2024b) estimations). The reason for this is that UI is less generous than many Western European schemes. The 62.4% replacement rate is the lowest in the Nordics, comparable to Germany, but Germany offers a higher benefit ceiling and Norwegian benefits are taxed unlike in Germany.[10]

Eligibility is stricter than in Western Europe, requiring one year of prior employment and an income threshold. However, Norway offers a longer benefit duration than most European countries, and UI payments are easy to calculate with an online tool.

[10]The maximum monthly UI benefit as of 2024 is approximately EUR 3,200 in Norway versus EUR 4,530 in Germany for individuals without children. The U.S. system is even less generous: 53% replacement rate, capped at 41% of the average wage, with benefits typically lasting only six months.

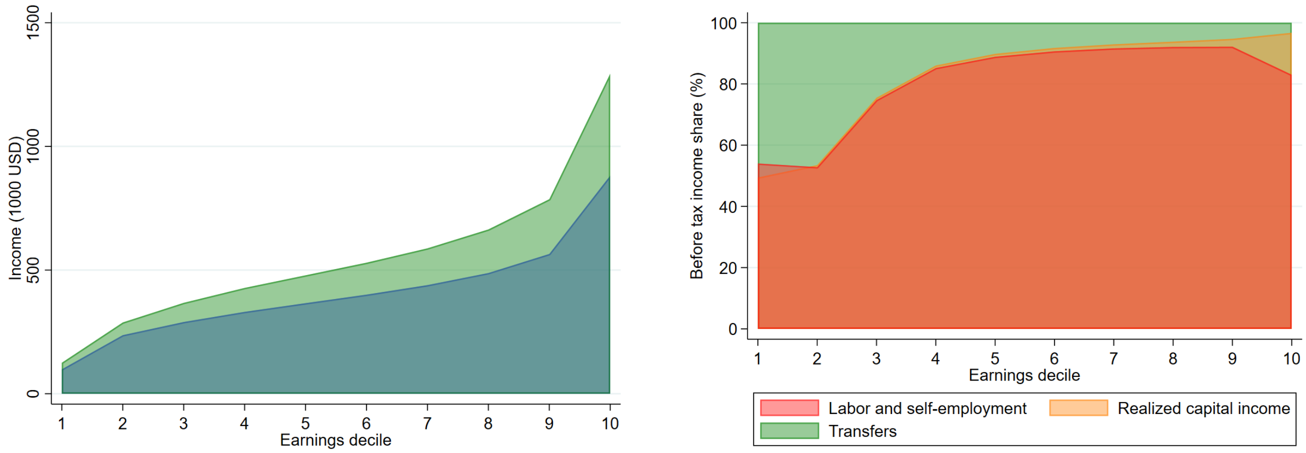

Before and After Tax Income

Components of Before Tax Income

Figure 1: Norwegian Income, Taxation and Transfers: By Income Deciles

High earners face the largest income losses due to a UI cap at 120% of lost wage. Households with significant wealth or debt see sharp declines, as wealth taxes and debt payments persist during unemployment. Child benefits, however, help smooth income for those with children. In our empirical analysis, we exploit these income fluctuations while we carefully control for wealth, debt levels

4 Inequality over Time in Norway: Descriptive Statistics

Before we turn to our main questions, we begin with a first look at the distributional dynamics in Norway, focusing on two descriptive statistics. First, we report the distributional dynamics for net income and consumption along the lines of Heathcote et al. (2010, 2020, 2023). Second, we estimate "worker betas", the cyclicalities of individual earnings in the cross-section, along the lines of Guvenen et al. (2017, 2021). We later estimate the distribution of MPCs and then connect it to the individual-level variables in order to compute the general equilibrium, "aggregate MPC" object. The main takeaway is that these distributional dynamics in Norway are surprisingly similar to the ones documented for the U.S. by the aforementioned studies.

4.1 Inequality over time

We begin by plotting the distribution of income and consumption over time, in log change across the *income* distribution for the deciles of our sample. Figure 2 shows how the deciles of income and consumption vary over time, across the income distribution. It is apparent that, similarly to the US dynamics documented in the seminal paper of Heathcote et al. (2010), inequality has been increasing over time in both income and consumption (the distributions fan out); more importantly for our purpose, the recession (shaded are in 2007-2009) does seem to affect the bottom

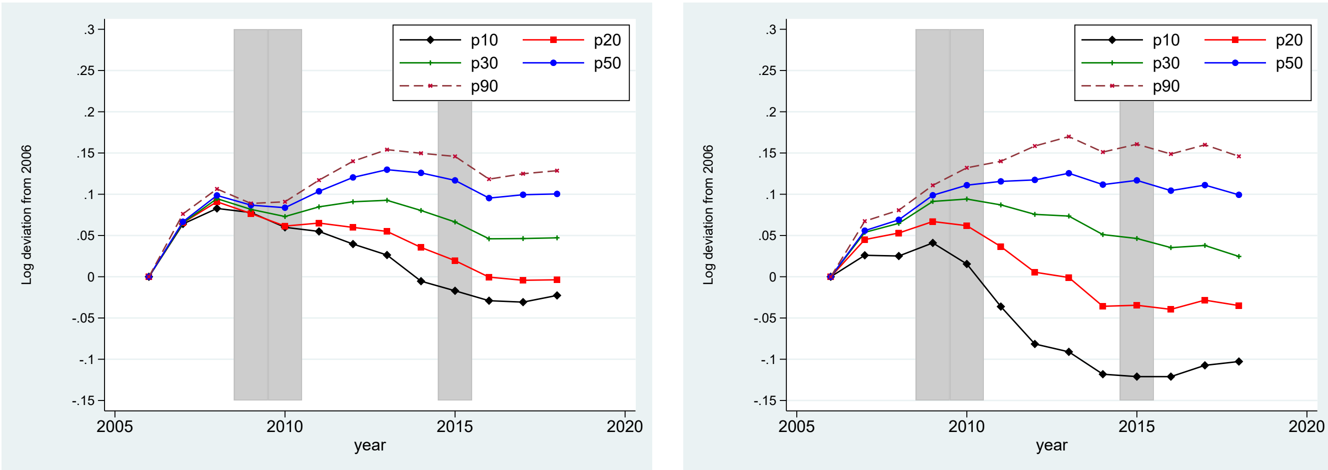

Total Income After Tax   Consumption

Figure 2: Household Income and Consumption: By Income Level

deciles relatively more.

4.2 Norwegian "worker betas"

We now compute so-called worker betas as in Guvenen et al. (2017) for different population groups. These betas measure the elasticity of individual's labor earnings with respect to aggregate income, and as such reflect cross-sectional variation in income risk. In this section we replicate the study in Guvenen et al. (2017) on Norwegian data. That is, we estimate the labor earnings betas for the same partitioning of the population as in their study.

We divide the population into groups based on a combination of 12 bins of average earnings (calculated over the previous six years), gender, and three age categories. For each group g we then estimate:

$$\Delta y_{i,t} = \alpha_g + \beta_g \Delta Y_t + \varepsilon_{i,t} \tag{17}$$

where $y_{it}$and Y_t denote log of individual real labor earnings and real GDP, respectively. The GDP-elasticities $\widehat{\beta}_g$, plotted in Figure 3, are in line with the US evidence from Guvenen et al. (2017): The elasticity is higher at the bottom and the top of the permanent earnings distribution.

How the inequality documented in this section behaves when one looks at disposable income and consumption, how it manifests itself as MPC heterogeneity, how the betas and MPCs relate to each other and empirically aggregate into the aggregated MPC will be the topics of the rest of the paper. We will do this analysis guided by the theory discussed in section 2.

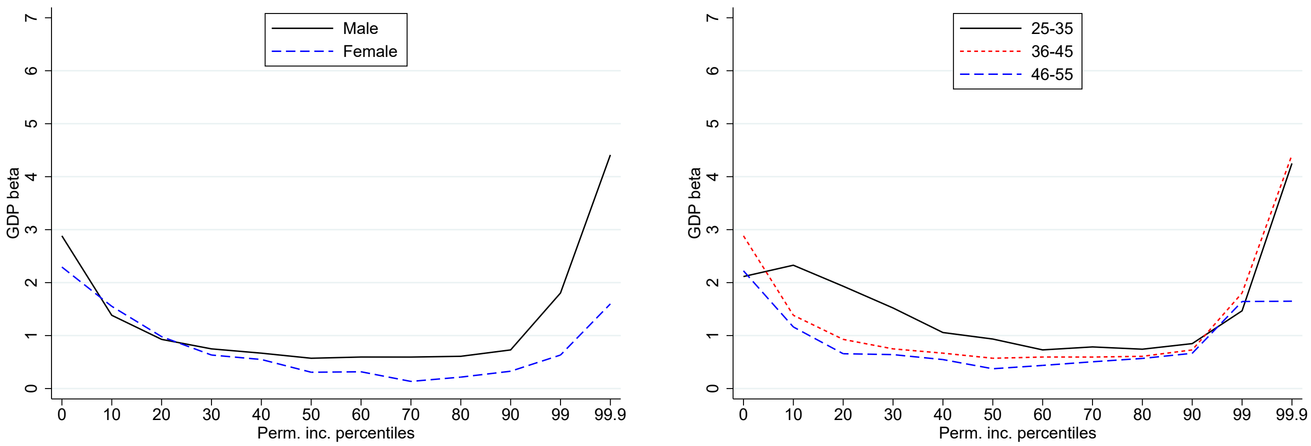

Gender Age

Figure 3: Worker Betas (Labor Earnings)

5 The MPC Distribution of Norway and its Determinants

Other than being a key object for our measurement exercise of eliciting sufficient statistics for HANK, the distribution of MPCs in Norway is of interest in and of itself. Existing attempts, reviewed in the literature section, have employed a variety of estimation strategies and data types from a variety of countries.Our approach is different from existing studies in that we have actual, transactions-level consumption data; yet as we will see our estimates are well in line with those obtained using different estimation techniques on other Norwegian data or in other countries using similar (or different) estimation techniques. One important contribution of this paper is therefore to justify the use of (much easier to obtain) imputed consumption data in answering related questions. Having employed actual, granular consumption data, we are able to show that what was assumed in earlier studies (that partial and/or imputed consumption is a valid proxy for consumption) is indeed the case.

5.1 Estimating the MPC distribution of Norway

To estimate the MPCs, we employ the strategy used in the influential study in Patterson (2023), itself building on a seminal Gruber (1997) paper. Unlike essentially all previous work on this question, we have the luxury of using actual consumption data covering about all of the population. We estimate:

$$\Delta C_{i,t} = \sum_{x}(\beta_x \Delta Y_{i,t} \times x_{i,t-1} + \alpha_x x_{i,t-1}) + \delta_{t,s} + \varepsilon_{i,t}, \tag{18}$$

where $C_{i,t}$ is real total household consumption of individual i at time t, $Y_{i,t}$ the real after tax total income of individual i, and $x_{i,t}$ are individual characteristics.[11] The empirical specification in (18)

[11] Specifically, in x we include dummies for gender and education (<high school, high school, university (lower), university (higher)), partnership status (single or cohabitant), five quintiles of average after-tax income

allows these individual characteristics to affect consumption changes directly, but crucially for us, they also shape consumption sensitivity to income changes. In particular, the estimated MPC for a given individual i can be calculated as $\widehat{MPC}_{i,t} = \sum \hat{\beta}_x x_{i,t-1}$.

Since both consumption and income are endogenous and jointly influenced by many factors, a naive OLS of (18) would likely be biased. We therefore proceed as in Patterson (2023) and employ an instrumental variable approach, using unemployment as the source of exogenous income change. We focus on workers who were employed in $t-1$ and construct an unemployment dummy $u_{i,t}$ equal to 1 if the worker reports being unemployed in t (measured as having received UI benefits on the tax return). Hence, we perform an IV regression in which:[12]

$$\sum_x \beta_x \Delta Y_{i,t} \times x_{i,t-1}$$

is instrumented with:

$$\sum_x \beta_x u_{i,t} \times x_{i,t-1}.$$

The advantage of this approach is that this is a typical business-cycle shock experienced by workers, is large, and therefore generates identifying variation in income. The disadvantage is that it may not always be an exogenous shock and is often a persistent (rather than purely transitory) income shock. The latter qualification implies that it is important to ensure that the theoretical object is consistent with what we actually estimate (and thus adjust for shock persistence). Furthermore, we will also study robustness by applying the method of Blundell et al. (2008), and in particular the variant put forth by Commault (2022) to deal with the potential bias induced by the shock persistence .

Figure 4 presents the estimated MPC distribution. The results are remarkably in line with Patterson (2023) and those independently estimated with other data, other estimation strategies, and for other countries, as in Parker et al. (2013),Fagereng et al. (2021), Kueng (2018), Lewis et al. (2024), Ganong et al. (2020),Fagereng et al. (2024a). The average MPC is high (0.38), varying substantially in the cross section, with a large fraction of households having a relatively elevated MPC. As shown below in Figures 5 and 6 our estimates, consistent with findings elsewhere, exhibits a negative relationship between liquid assets and MPC.

An important caveat to our MPC estimation is that the income shock is potentially persistent and has different persistence across agents depending on the duration of unemployment, which

in the past three years, and dummies for low liquid and low net wealth. In contrast to the other covariates, the wealth dummies are constructed at the level of individual i's household. Low liquid wealth is defined as the household having gross liquid wealth below two weeks net income (annualized), while low net wealth is defined as below two months net income (annualized). These wealth dummies are related to often used proxies for hand-to-mouth status used in the literature.

[12]In addition to restricting the estimation sample to individuals being employed in t-1 (i.e. positive labor earnings), we also exclude individuals who's partnership status changes between t-1 to t. We also follow Patterson (2023) and handle outliers by restricting the estimation sample to individuals whose one-year income $\Delta Y_{i,t}$ and consumption $\Delta C_{i,t}$ changes are less than twofold.

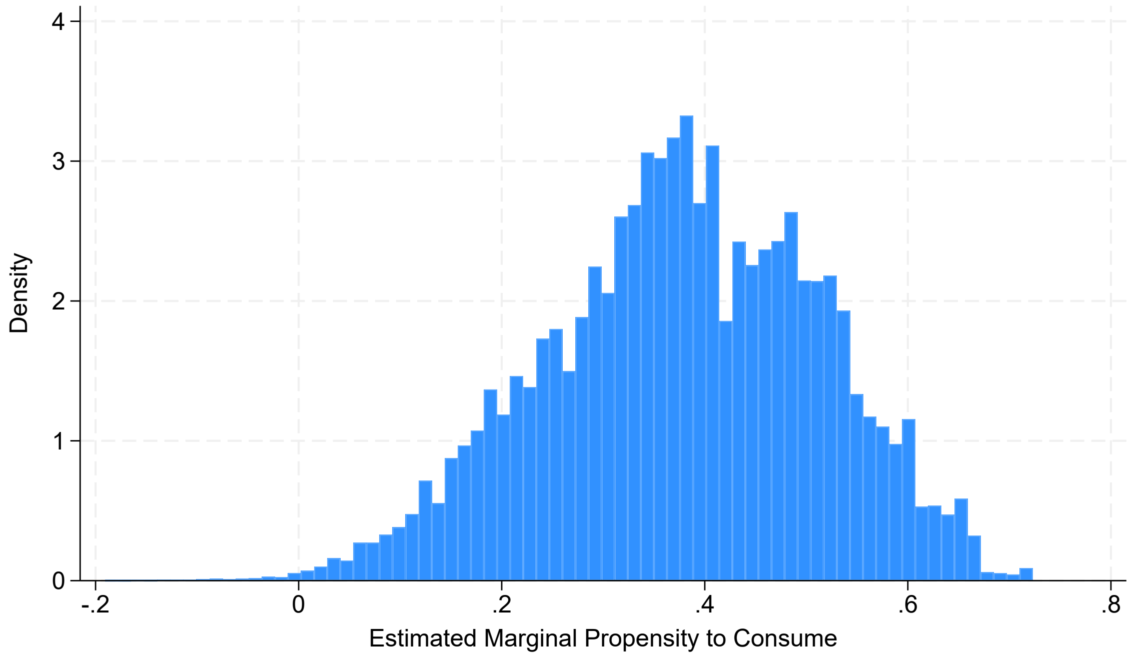

Figure 4: The MPC Distribution of Norway

may bias the MPC estimates. This issue is common to the literature where unemployment is used to identify exogenous variation in income. To address this, we employ an alternative MPC method to solve the shock persistence problem, namely the Commault (2022) variant of the Blundell et al. (2008) methodology.

This alternative method produces MPCs that are quantitatively in line with our benchmark method. To anticipate, we will also see below that one of our important results, the correlation of MPCs with income and consumption betas, is also invariant to using BPP methodology rather than the benchmark estimation.

5.2 Determinants of MPC heterogeneity

To make progress towards estimating an empirical counterpart to the theoretical models, we next turn to studying the dimensions along which MPC heterogeneity manifests itself. To that end, in Figure 5 we report the results of bi-variate regressions corresponding to the set of characteristics captured by the covariates in equation (18): gender, partnership status, and age in the left panel; lagged earnings, hand-to-mouth HTM status (defined below) and education in the right one. The figure shows coefficient estimates obtained from a bi variate IV version of the regression in section 5.1 Each dot represents the estimated MPC for the bi variate regression using the characteristics on the the x-axis. In Table 4 we report the full results of the multivariate regression containing all determinants jointly.

In doing analysis of this nature, while most categories are quite obvious (gender, cohabitation

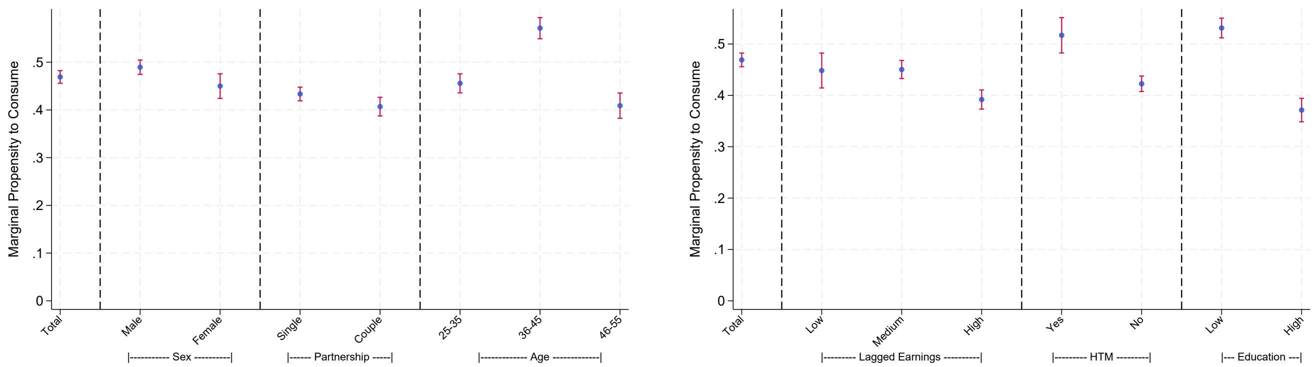

a. b.

Figure 5: MPC Determinants

status etc. are directly observable in the registry data), hand-to-mouth status is not. In many HANK-type models being HtM and having a high MPC are by assumption the same. To test whether this is indeed the case, we need a HtM measure that does not depend on the estimated MPC.

This is one of the instances where having data on wealth holdings enables us to properly study and important question empirically. As we have data on wealth that is separated into liquid and illiquid components, we are able to ask whether wealth holdings as a whole, as in Aguiar et al. (2023), who classify agents are HtM if they have total wealth that is less than two months' after-tax income; or liquid wealth, as in Kaplan et al. (2014), who classify agents are HtM if they have liquid assets that are less than half of monthly after-tax income, have any bearing on estimated MPCs.[13] This latter measure allows wealthy hand-to-mouth agents, those having high total wealth and low liquid wealth.

We have three important findings to report. One is that while income by itself is not an indicator of MPC level, education is[14] Second is that HtM status is also an important predictor of MPC. Third is that HtM status *only* works as a predictor if it is conditional on liquid wealth. Total wealth is not informative about MPC status. Figure 5 shows these findings, with the liquidity-based HtM measure shown in the right panel.

In figure 6 we illustrate how certain characteristics of individuals varies along the MPC distribution. In particular we consider liquid wealth and life-cycle income calculated using a Mincerian regression, where we employ predicted income obtained from regressing the log of either after-tax income or labor earnings on gender, age, education and year dummies. We consider a flexible specification in which the age-income profile is allowed to vary with gender and education.

The main takeaway is that the upper tail of the MPC distribution (high-MPC agents) is char-

[13] All variables are measured at yearly level, so monthly of income is measured as one twelfth of yearly income in our data.

[14] Nevertheless, when it comes to the determinants of income and consumption cyclicalities above, education is not a significant factor. We have chosen not to report this for the sake of brevity, but the illustration of this is available upon request.

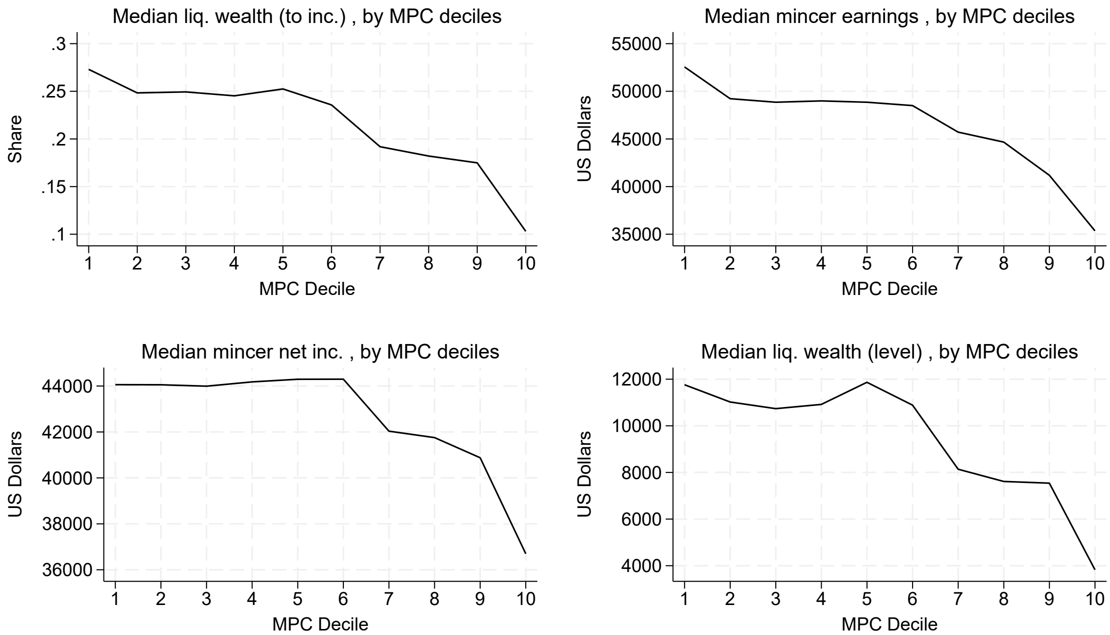

Figure 6: Liquid Wealth and Permanent Income in the MPC distribution

acterized by a markedly lower level of liquid assets and low Mincerian income. This provides support for the often-used assumption in HANK models that high MPC, hand-to-mouth agents are the liquid-asset poor. We leverage this below in our computation of sufficient statistics based directly on consumption.

6 Sufficient Statistic 1: The Aggregate MPC

We are now in a position to compute the first sufficient statistic from section 2, the aggregate MPC. Along the way, we first show how the income betas for each measure of income vary across the distribution of MPCs. We find that the betas become flatter in the MPC dimension when we move from labor income to adding capital income, transfers and taxes. In a counterfactual exercise, we find that a substantial degree of flattening would also arise even in a less redistributive tax and transfer system such as that of the United States. Finally, we find the aggregate MPC directly by combining the individual MPC estimates with estimated income betas for each individual (or finely-disaggregated groups of individuals) in our sample.

6.1 Income Betas Across the MPC Distribution

To move towards computing our object of interest, we estimate the betas as in section 4.2, but we make two changes. First, we use *disposable income* instead of earnings. And second, we group households by their estimated MPC instead of by their permanent income. Informed by the theory in section 2, what ultimately interests us is to what extent people with higher MPCs have more procyclical *disposable income*.

Figure 7 plots the estimates of income betas for both labor earnings and disposable income for each decile of the MPC distribution. We plot the estimates of betas both with respect to GDP growth (left-hand side), as in Guvenen et al. (2021), and with respect to the growth rate of the aggregated income variable itself (right-hand side).[15] Consistent with the results for the United States reported by Patterson (2023), we find that the betas for labor earnings are higher for individuals with higher MPCs. The beta for the top MPC decile is almost double that of the bottom decile. Hence, if used to draw conclusions regarding the model properties, labor earnings betas imply amplification, of a magnitude we will quantify momentarily. However, we find that this pattern disappears when we move to estimating betas for disposable incomes. If anything, incidence is lower for high-MPC groups, which suggests a *dampening* of aggregate demand shocks and policies rather than an amplification. This holds regardless of whether we estimate the betas with respect to aggregate GDP or aggregated income.

This is one of our important results. MPCs are functions of disposable income, yet in empirical work earnings are often used to proxy for the theoretically appropriate disposable income measure because earnings data are easier to come by. We see that earnings is not a good enough proxy for disposable income when the question at hand is MPC heterogeneity and its aggregate effects. In Norwegian data, earnings-based betas would have us think heterogeneity in MPCs creates an aggregate demand amplification effect, whereas the proper disposable income measure shows that there is in fact a dampening effect.

6.2 A Decomposition: The Role of Capital Income and the Tax and Transfer System

The income beta and MPC relationship changes sign and becomes negatively sloped (from a pronounced positive slope) as we move from labor earnings to disposable income for the income measure. We decompose disposable income into its constituent components to understand the reason. Figure 8 shows the results of estimating betas by MPC decile for labor earnings, then step-by-step adding capital income, adding government transfers, and subtracting taxes.[16]

Adding capital income to labor earnings mildly flattens the beta curve, in particular for individuals with very high and very low MPCs. Since capital income is highly cyclical and low MPC individuals–who on average have higher income and net worth than high MPC individuals–on

[15]In these and subsequent figures that plot betas across the MPC distribution, we use a slightly different specification of the regression than the one in section 4.2. Specifically, since we are interested in the group-level beta – not the average beta within the group – we define the left-hand side variable as the individual-level change in income relative to mean income within the decile group. See details in section C.3. As we also show in the appendix, this specification also ensures that the group-level betas will approximately sum to one when weighted by the group income shares.

[16]Figure 17 in the appendix displays beta estimates for each income measure with confidence intervals.

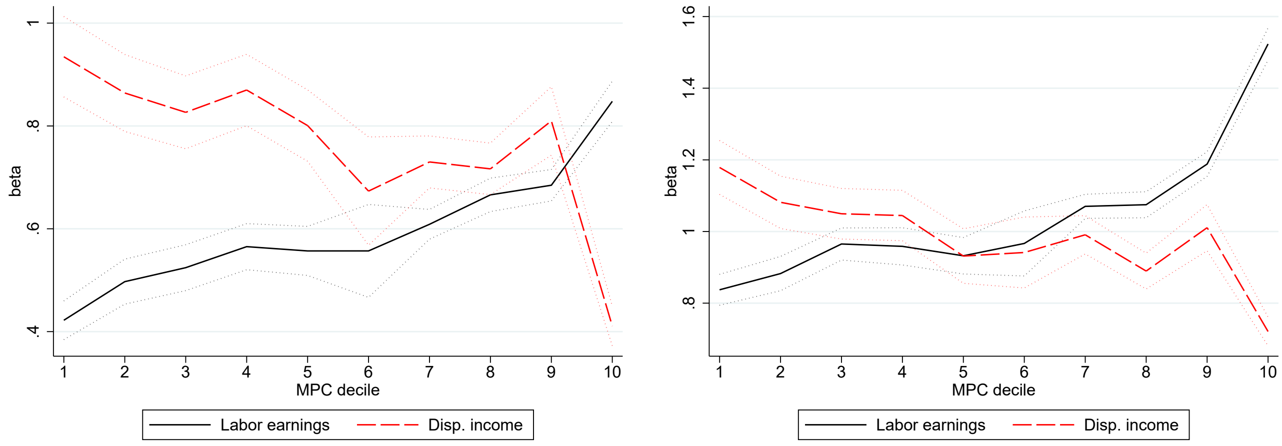

(a) GDP betas

(b) Aggregated income betas

Notes: The figures show estimates of income betas for labor earnings and disposable income for deciles of the MPC distribution, when the right-hand side variable is either GDP growth (left) or the growth rate of aggregated individual income in the sample (right). Dotted lines indicate 95% confidence bands based on White standard errors.

Figure 7: Estimated tax and transfer functions for US and Norway.

average get a larger share of their income from capital income, adding capital income to labor earnings naturally lowers the betas in the top deciles relative to the betas in the lowest deciles. But it is clear that this is not the main contributor to the difference between labor-earnings- and disposable income-based measures.

Adding government transfers changes the beta curve substantially, making it downward sloping. There are two reasons for this. First, transfers are less cyclical than market income, and high MPC individuals on average get a larger share of their income from transfers. Second, as detailed in section 6.2, the transfer system is highly progressive, in particular at the bottom of the income distribution. As a result, a large drop in labor income translates into a smaller drop in income inclusive of transfers for low-income individuals than high-income individuals. For instance, while a low-income person is more likely to lose her job in a recession, unemployment benefits make up for a larger share of pre-employment income than for a person with higher initial income. Since low-income individuals are more likely to have high MPCs, these effects contribute to a substantial dampening of betas at the top of the MPC distribution. Finally, subtracting taxes leads to a further negative slope of the beta curve.

The positive slope of the beta curve across the MPC deciles, known in the literature for other countries (in particular the US) when labor earnings is the income measure, is due to the behavior of the betas in the in the lowest and highest thirds of the MPC deciles, with the middle third being essentially flat. Capital income, taxes, and transfers differentially affect the disposable incomes of agents with the highest and lowest incomes and hence lowest and highest MPCs. That these effects are large enough to change the slope of the relationship is a new empirical finding that reverses the received wisdom about amplification due to heterogeneity.

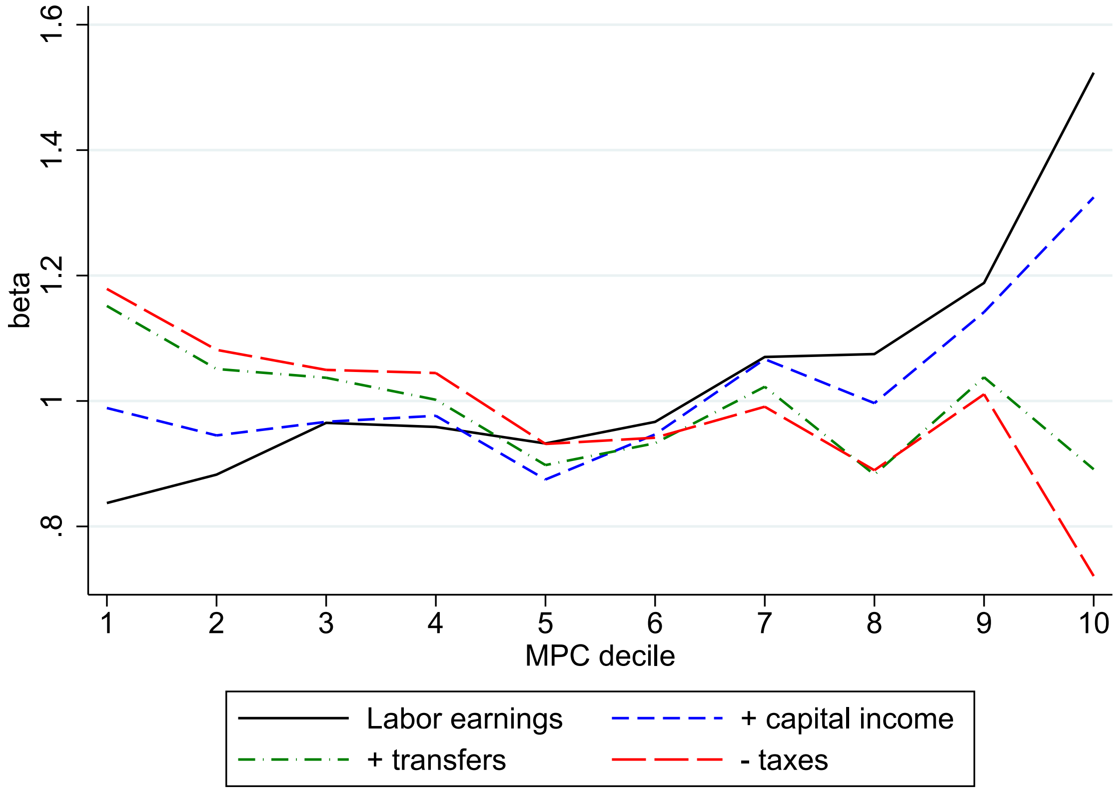

Notes: The figure shows the point estimates for income betas across 10 MPC deciles based on labor earnings, labor and capital income, total income before tax, and total income after tax (disposable income).

Figure 8: Betas Decomposition: From Earnings to Net Income

6.3 A Counterfactual:The US Tax and Transfer System

The results in the previous section point to an important role of automatic stabilizers in mitigating the impact of aggregate shocks on the incomes of high-MPC individuals. Given the important role of the tax and transfer system, it is important to ask whether the same flattening of the beta curve would occur in countries with less redistributive systems. To answer this question, we counterfactually apply an estimated tax and transfer system for the United States to our Norwegian micro data, then re-estimate the beta curve for disposable income imputed from the US system.[17] We employ the specification estimated by Ferriere et al. (2023) on data from the Current Population Survey, consisting of a flat tax on capital income, a labor tax rate that varies with labor income, and a transfer that varies with total labor and capital income. The functional forms are specified in appendix C.4.1.

Our counterfactual exercise proceeds in two steps. First, to insure that the results we get based on the US tax and transfer system is not due to the approximation of a complicated system by a

[17]We are grateful to Adrien Auclert and Axelle Ferriere for suggesting this exercise and to Axelle Ferriere for sharing her codes.

parsimonious set of functions, we re-estimate the same set of functions on Norwegian data and estimate betas across the MPC distribution for both actual and imputed disposable income in Norway. Second, we switch to the parameters estimated by Ferriere et al. (2023) and re-estimate the betas for disposable income imputed based on the US tax and transfer system.[18]

Figure 9 compares estimated taxes and transfers in the US and Norway. Labor taxes are more progressive in Norway, with an average tax rate that is slightly lower for low-income individuals and noticeably higher for high-income individuals. Of more consequence, low income individuals in Norway receive substantially more government transfers, equal to more than 50% of average population income for the lowest income individuals, compared to less than 10% in the US.[19] Figure 10 plots disposable income by quintiles of labor income both in the data and estimated based on the Norwegian and US tax and transfer systems, respectively. Our estimates indicate that households with labor income below around $50'000 have the most to gain from the Norwegian system, primarily due to the substantially higher levels of transfers received by this group. The Norwegian welfare system is particularly generous at the bottom of the income distribution.

Now we turn to the estimates of disposable income betas under the two tax and transfer systems. The betas along the MPC distribution are shown in figure 11. The betas based on imputed disposable income from the estimated Norwegian system are very close to the betas based on *actual* disposable income, indicating that imputing taxes and transfers in itself does not change our main result. Perhaps more surprisingly, there is also substantial flattening of the beta curve—relative to the betas based on labor earnings, or even labor and capital income—even under the counterfactual US tax and transfer system. Although there is less dampening of the betas for high-MPC individuals under the US system, the counterfactual brings us closer to the betas based on disposable income in Norway than to those based on pre-tax and transfers income. The exercise thus suggests that our main result of substantially less amplification than suggested by pre-tax labor earnings (in the Norwegian case a switch to dampening) would still hold in countries with less progressivity of taxes and transfers.

6.3 Measuring Amplification: The Aggregate MPC and Implied Multipliers

We see that the general-equilibrium amplification of demand shocks due to heterogeneity based on *labor earnings* is overturned when estimating betas based on *disposable income*. We can get at this result more directly by estimating the aggregate MPC in equation 4. To do so, we estimate individual-specific income betas $\hat{\beta}_j$ for every individual j in our sample who is observed in at least two consecutive years.[20] The aggregate MPC is then $\sum_j \hat{MPC}_j \times \frac{Y_j}{Y} \times \hat{\beta}_j$, where Y_j is individual j's

[18] Ferriere et al. (2023) estimate the tax and transfer functions on household-level data, while we estimate them on individual-level data for Norway. In appendix C.4.3 we show that re-estimating the betas along the MPC distribution on household-level data does not alter the conclusions in this section.

[19] The exact parameter values, as well as measures of fit, can be found in table 5.

[20] These betas are estimated based on running regression 17 at the individual level.

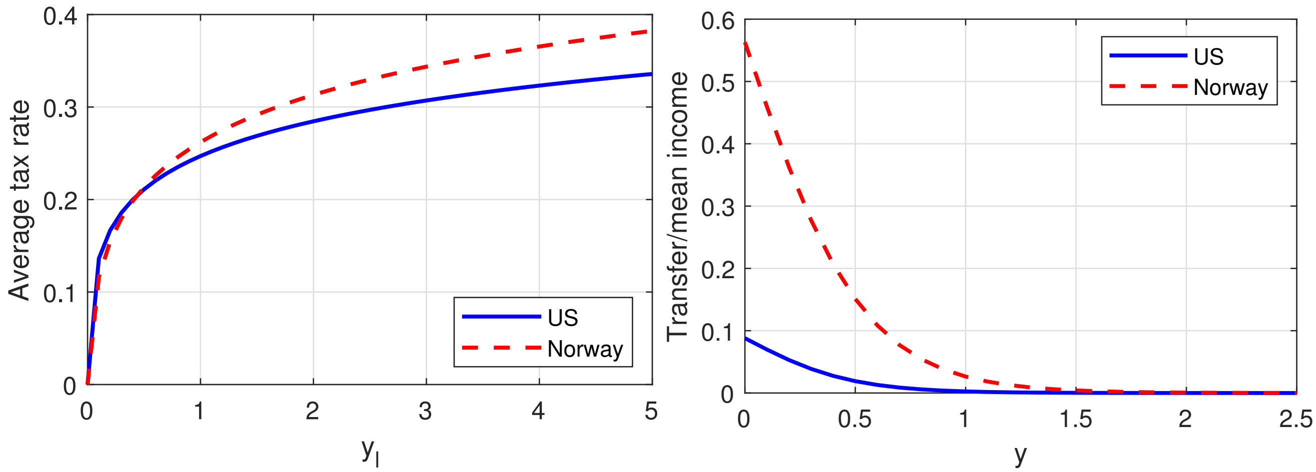

(a) Labor tax function

(b) Transfer function

Notes: The figures illustrate the estimated functions for average labor tax and transfer relative to mean income, by labor income relative to mean income (y_l) and disposable income relative to mean income (y). These figures are based on figure C1 in Ferriere et al. (2023). The functional forms are specified in equations 37 and 38, and the parameter values are found in table 5.

Figure 9: Estimated tax and transfer functions for US and Norway.

income and Y is total income in the sample.

Table 1 reports the aggregate MPC based on each of the four income measures, as well as the *average* MPC – assuming that $\hat{\beta}_j = 1$ for every individual j. The difference between the aggregate and average MPC is reflected in the covariance between individual MPCs and betas, as in Patterson (2023). These covariances are reported in the fourth column. Finally, we also report the GE multiplier based on the AMPC (labeled Multiplier A), and the multiplier based on the uniform-betas, average MPC (labeled "Multiplier a").

Consistent with the group-level results reported in section 6.1, we find a large aggregate MPC and multiplier based on labor income betas. This is due to in part to a positive covariance between earnings betas and MPCs, resulting in an aggregate MPC that is higher than the simple average MPC in the sample. Adding capital income decreases the implied amplification, but the aggregate MPC is still higher than the simple average. However, adding transfers and subtracting taxes from the measure of income essentially inverts the gap between the aggregate and average MPC, with a slightly negative covariance. When betas are estimated based on disposable income, the multiplier is in fact slightly lower than it would be in a representative agent model where the agent has–for whatever reason–an MPC of 0.366.[21]

The significance of this result is hard to over emphasize. Policy analysis employing HANK models will lead to erroneously large aggregate multipliers if income is measured only for labor earnings and this is used as a proxy for disposable income. Taxes and transfers *qualitatively* change the cyclical behavior of incomes and their relationship with MPCs.

[21]Notice that since the post-tax and transfer income measures redistribute income to the high-MPC agents on average, their average income shares increase and this also automatically increases the average MPC.

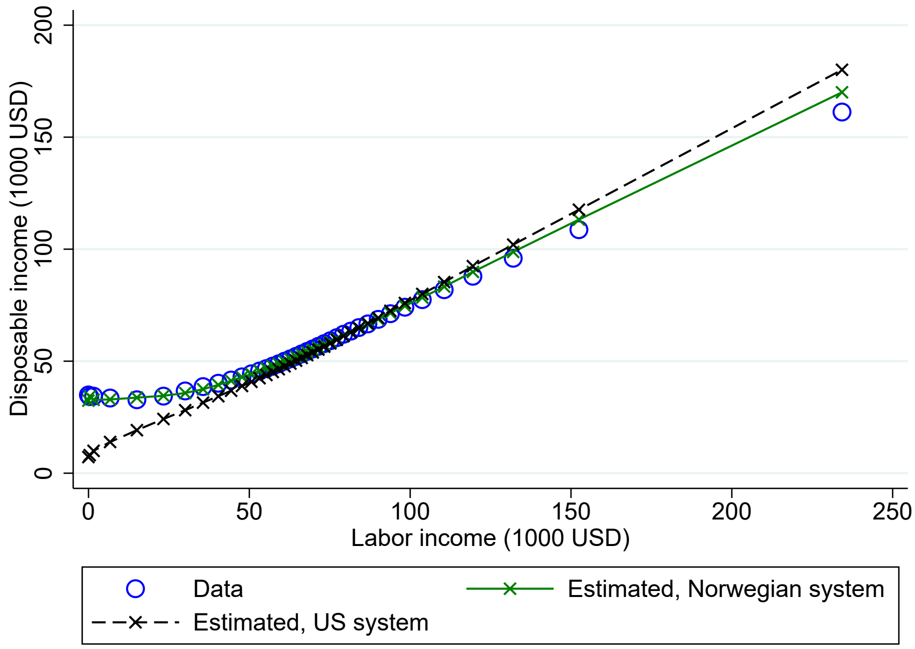

Notes: Disposable income by 2.5% quantiles of labor income, in data and based on estimated tax and transfer system for Norway and the United States. NOK values are converted to USD using the average exchange rate over the period 2006-2018.

Figure 10: Imputed disposable income by labor income, based on estimated tax and transfer system for Norway and the US.

6.4 Dealing with shock persistence

Since our main result concerns the ordering of betas along the MPC distribution, it is important to address the persistence issues in MPC estimation discussed earlier. To do so, we group households into income beta quintiles and estimate an alternative MPC for each using Blundell et al. (2008), extended by Commault (2022) (BPP-C), a semi-structural approach that accounts for shock persistence. Following BPP-C, we restrict the sample to stable married couples. Appendix C.6 details the sample selection and estimator.

Our main result is that the magnitude and pattern of estimates align with our benchmark findings. First, MPC estimates using BPP-C remain flat across beta quartiles (Figure 12).[22] Second, since correcting for persistent income shocks lowers average MPCs, the BPP-C estimates are lower compared to our baseline but not significantly so. In addition, correcting for income persistence lowers MPCs uniformly in the income betas dimension, consistent with the notion that the persistence of income shocks does not vary considerably in the relevant cross-sectional dimension we focus on. These findings align with Fagereng et al. (2024a), who find that while unemployment reduces income persistently, MPCs remain homogenous across income groups, though a direct comparison is less straightforward since they do not focus on income betas.[23]

[22]Due to the short time sample, BPP-C estimates for smaller beta quantiles are noisier but also flat.

[23]Fagereng et al. (2024a) find that after becoming unemployed, Norwegians' earnings decline by $20\%–

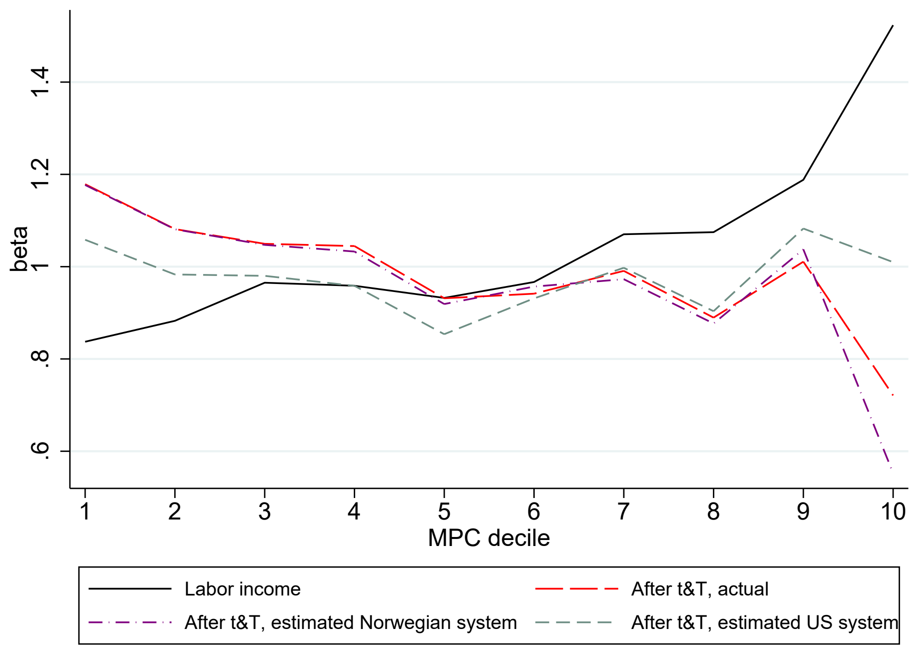

Notes: The figure shows the point estimates for income betas across 10 MPC deciles based on labor earnings, actual disposable income (after taxes and transfers), and disposable income imputed from the estimated taxes and transfer systems for Norway and the United States, respectively.

Figure 11: Income betas with US counterfactual tax and transfer system.

Our overall estimate for the full sample closely matches both our benchmark and U.S. estimates. We find an MPC of 0.3 in response to transitory income shocks over the following year, slightly below our benchmark estimate. In the U.S., Commault (2022) estimates an MPC of 0.32 over the same horizon, remarkably close to our Norwegian estimate.

7 HANK Sufficient Statistics 2: Consumption Betas

We are now in a position to compute the second sufficient statistic, based on the consumption betas. This is a direct way of determining the amplification or dampening effects due to heterogeneity by comparing the cyclical behavior of consumption of different groups in the population. We will be considering two groups, one following an Euler equation and the other hand to mouth.

7.1 Consumption betas: The extent of consumption risk sharing

In Figure 13 we report the betas computed directly using our rich consumption data. As before, we compute the betas both with respect to a national income accounts measure of aggregate con-

30\%$ before starting to recover after two years. They estimate MPCs from inputed consumption and find modest variation across male or household earnings.

Table 1: Aggregate MPC of Norway:

Income measure	AMPC	avg. MPC	$Cov(MPC, \beta)$	Multiplier A	Multiplier a
Labor income	0.374	0.357	0.017	1.6	1.557
+ capital	0.364	0.357	0.07	1.57	1.556
+ transfers	0.356	0.363	-0.007	1.55	1.571
- taxes	0.354	0.366	-0.0112	1.549	1.578
US counterfact.	0.358	0.359	-0.001	1.558	1.561

Notes: The table displays, for each of the four income measures, the aggregate MPC (AMPC) based on equation 4, the average MPC (when individuals are weighted by their income shares), the covariance between MPC and income betas, and the general equilibrium multiplier based on the AMPC and equation 6.

sumption and with respect to the aggregated consumption measure itself. The picture is the same regardless of the measure used. As the figure shows, the consumption betas are close to flat over the MPC distribution, indicating a high degree of consumption risk sharing, or insurance, in a manner similar to Townsend (1994), Mace (1991), and Cochrane (1991).

7.2 Measuring amplification using consumption betas

We now study amplification directly—–as described in the theory section—–by comparing the consumption betas of groups based on their hand-to-mouth status. Since hand-to-mouthness is subject to a variety of possible definitions, we report this for a wide range of splits, based on liquid wealth, net worth, stock market participation, and slicing the respective distribution at various thresholds. Table 2 reports the various measures for these different ways of dividing the distribution.

Table 2 reinforces the conclusion we drew based on the computation of the "aggregate MPC" using the MPC distribution and the betas of net total disposable income at the individual level. While HtM agents have a somewhat higher cyclicality of consumption, the magnitude is negligible. In other words, regardless of the classification criteria, the betas based on consumption of HtM and non-HtM groups are not that different. This is fully consistent with Figure 13, which showed that all MPC deciles have about the same consumption beta. Here, we learn that grouping the population in other ways does not change the consumption beta invariance result. Clearly, any aggregate-demand amplification—because an aggregate-demand shock that shifts the Euler equation would be amplified relative to a representative-agent or perfect-insurance (complete-market) benchmark when HtM agents have higher betas—will have to be modest. To give some concrete examples, the multipliers (computed using the expression in 13) would be merely **1.009** (for the

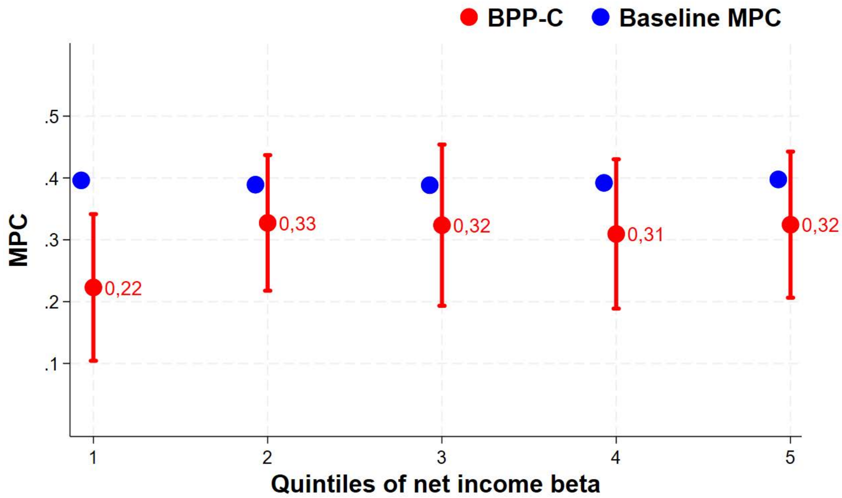

Notes: The figure shows (in red) the point estimates and standard errors for the MPCs estimated using the BPP-C method, alongside our baseline estimates (blue dots) by quintiles of disposable income betas

Figure 12: MPCs estimated by the BPP-C method as a function of net income betas.

first split), up to **1.033** (for the fourth split) or **1.045** (for the eighth and last split).[24] We provide a further summary calculation of the implied multiplier for the case study of the Great Recession below.

Note that consumption betas could have implied amplification, despite disposable income betas having indicated dampening. In an economy with saving, these two metrics do not overlap, with the consumption-beta-based measure being the correct sufficient statistic.

The fact that the direct, consumption-based sufficient statistic also does not point to multiplier amplification due to heterogeneity suggests the aggregate irrelevance of yet another layer of heterogeneity mechanisms. Recall that HtM (or high-MPC) agents' consumption can be relatively more cyclical (than non-HtM/low MPC agents'), while their disposable income is less cyclical, through two forces. First, investment is naturally, almost by definition, concentrated among the low MPC (high saving) group. Insofar as investment is procyclical, this automatically implies that the consumption of this latter group (which, together with investment, sums to their disposable income) will be less cyclical than their income. For instance, at similar income cyclicalities, the low-MPC would have more subdued procyclical consumption because they invest and investment is strongly procyclical. The second force has to do with precautionary savings in liquid assets, which may be correlated with investment but is distinct (much of investment is in illiquid assets such as housing). Our findings of flat consumption betas suggest that overall, the lower cyclicality of disposable incomes of high-MPC consumers (shown in the previous subsection) is offset by the consumption smoothing ability of low-MPC agents. This relates to the cyclicality of income risk and its concentration. The measurement and teasing out of these channels and their

[24]Grouping by education leads to an interesting mild dampening result, despite the education being a key determinant of MPC heterogeneity. If anything, more educated agents, who have lower MPCs, seem to have more cyclical consumption.

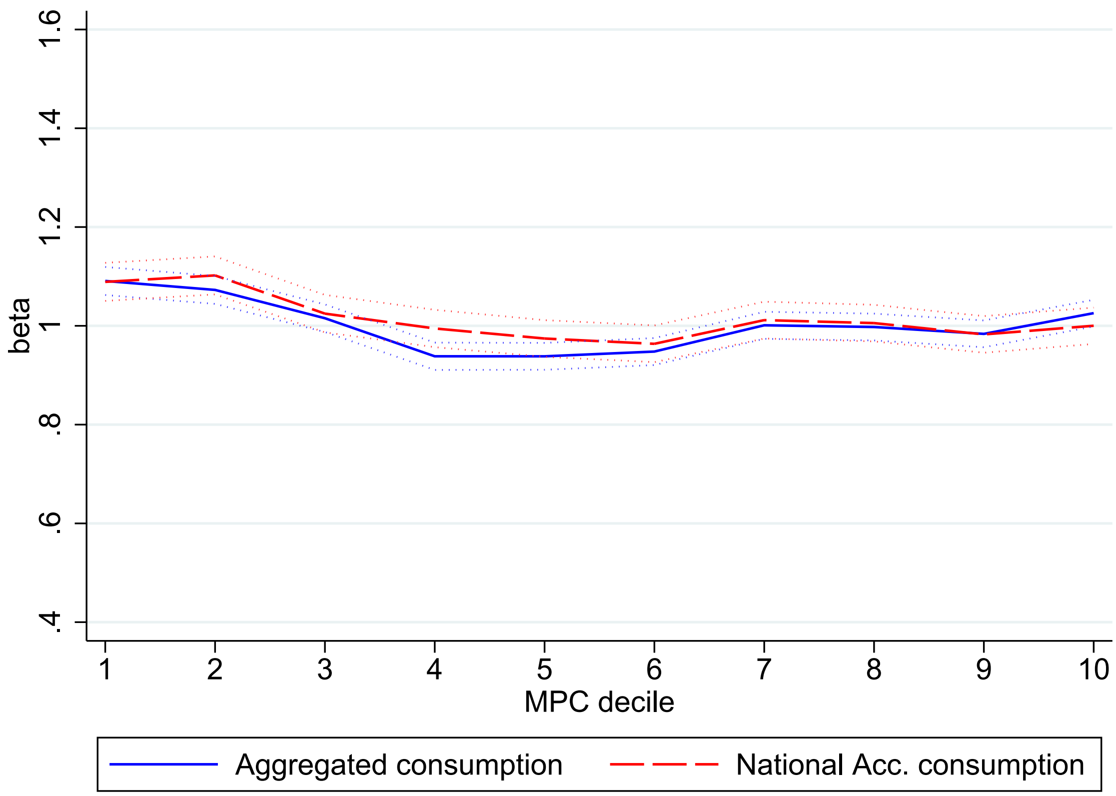

Figure 13: Consumption betas.

Table 2: Consumption Betas Norway:

HtM Definition	Share in		Consumption betas	
	Pop.	C	Agg'D c_is	NA C
Liq.<2 wks perm. y				
HtM	0.16	0.126	1.06	1.11
Non-HtM	0.84	0.874	0.99	1.01
Liq.<2 wks of y_{-1}				
HtM	0.17	0.14	1.078	1.11
Non-HtM	0.83	0.86	0.989	0.999
Liq. bottom 30% perm. y				
HtM	0.31	0.27	1.074	1.116
Non-HtM	0.69	0.73	0.971	0.982
Liq. bottom 30% y_{-1}				
HtM	0.30	0.26	1.09	1.12
Non-HtM	0.70	0.74	0.965	0.97
Liq. bottom 30% (level)				
HtM	0.27	0.24	1.049	1.081
Non-HtM	0.71	0.77	0.985	0.994
Liq. bottom 10% (level)				
HtM	0.08	0.06	0.984	1.085
Non-HtM	0.92	0.94	1.001	1.011
Not owning stocks				
HtM	0.58	0.5	1.023	1.008
Non-HtM	0.42	0.5	0.984	1.02
Net wealth < 2 months y_{-1}				
HtM	0.32	0.28	1.0\ref{fig:BPP}82	1.078
Non-HtM	0.68	0.72	0.98	0.992
Net wealth bottom 30%				
HtM	0.27	0.24	1.036	1.147
Non-HtM	0.73	0.76	0.994	0.966

separate roles is the subject of a separate, follow-up paper.

7.3 Case Study: The Great Recession in Norway

As an alternative to gauge amplification, we report a decomposition of the Great Recession in Norway: whose income and consumption were affected more, based on hand-to-mouth status? And does the measured unequal incidence imply through the lens of the model that the Great Recession itself was amplified by these mechanisms, relative to an equal incidence, perfect-insurance case? The estimates of income and consumption betas and their cyclical behavior suggest that amplification due to heterogeneity is unlikely to be materially important. Here, we will calculate betas for the Great Recession to learn whether this large shock created a more pronounced heterogeneity

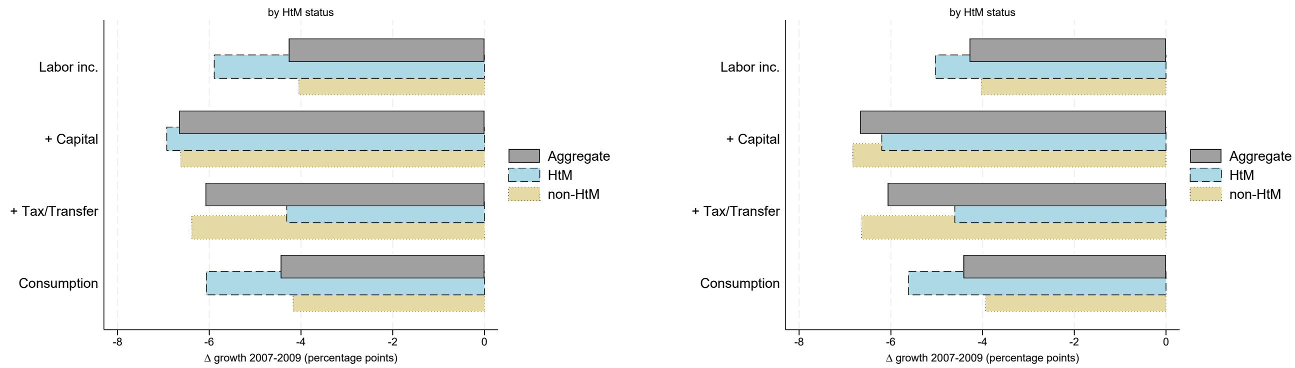

Liq.<2 weeks perm. y | Liq. (by y_{-1}) in bottom 30%

Figure 14: The Great Recession and Inequality in Norway: A Decomposition

in consumption betas. Figure 14 provides the decomposition based on the change in group consumption relative to the aggregate, for two alternative hand-to-mouth definitions. Remember that the perfect insurance case will have betas of unity for both groups and a corresponding multiplier that is also unity.

Based on these estimates, we conclude that at most, with Great Recession betas of the HtM and non-HtM groups (1.355 vs 0.948), we would have observed a mere **1.054** multiplier based on the first HtM definition in the left panel. To move towards obtaining an *upper bound* and give these mechanisms the best chance at delivering substantial amplification, we consider the second HtM categorization using the bottom 30% of the liquid assets distribution (normalized by previous-period individual income). Even in this case, at most, the Great Recession betas (1.30 vs 0.89) delivers a multiplier of **1.12**, similar to what was reported in Table 1. [25] Even in the Great Recession period we thus find scant evidence of amplification of shocks due to heterogeneity.

8 The Reflection Problem

Estimates such as ours may be subject to what Manski (1993) called the *reflection problem*: estimates in a linear regression of individual elasticities to the aggregate of that same variable are biased towards one.[26] To address this, we re-estimate all our beta regressions instrumenting with the past aggregate variable (earnings, net income, and consumption) instead of using the current one

[25]Comparing these multipliers to the (dampening-inducing) ones based on disposable income allows us to tease out the relative contribution of "cyclical savings" (be it in investment or liquidity). Thus, the multipliers relative to that benchmark double for each case, being respectively equal to 1.1 and 1.22 (implying that cyclical savings in and of themselves amplify fluctuations relative to a uniform-income hypothetical scenario by 10 and 22 percent, respectively).

[26]We thank Joel Flynn for raising this point

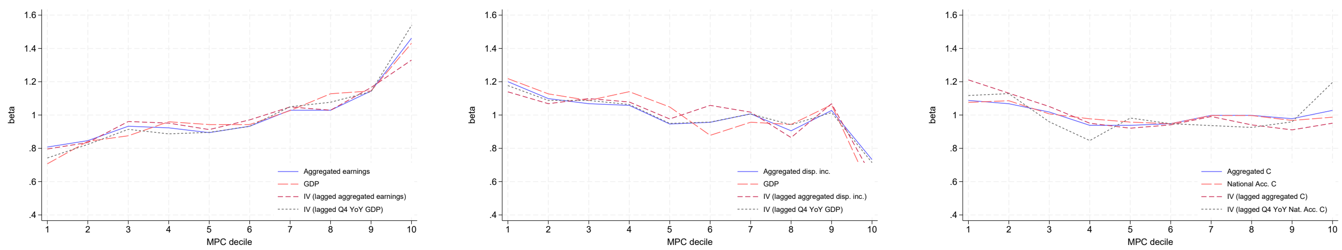

Earnings Betas — Income After Tax Betas — Consumption (transactions)

Figure 15: Betas with current vs (instrumented) lagged aggregate variables.

directly, as in Flynn et al (2024). We do this for both the national accounts variables and for the micro-aggregated variables. The results are summarized in Figure 15, with one panel for each of earnings, net income, and consumption. Each panel includes two lines reproducing the baseline betas, i.e. with respect to both the contemporaneous aggregated and national-accounts version of the respective variable, as well as two lines for the instrumental-variable estimates.

The IV is performed as a two-stage least-squares where in the first stage we regress the contemporaneous aggregate growth on the instrument and in the second stage we regress the individual growth rates on the predicted contemporaneous growth, where the predicted value is the fitted value from the first stage. We use two IV's mirroring our alternatives for aggregate variables: either (i) the $t-1$ aggregated growth (for the internally aggregated elasticities) or (ii) the $t-1$ fourth-quarter year-on-year growth rate (for the national account elasticities).

Figure 15 shows clearly that the general pattern of the betas in the cross-section of MPCs stays unchanged: the betas are still increasing for earnings, and are flat to decreasing for disposable income and consumption. We are therefore reassured that the flatness of the last two measures is not merely an artifact of the reflection problem.

9 Conclusion

Do distributional dynamics interact with aggregate fluctuations in a way that leads to amplification, or dampening of macroeconomic shocks? Based on very detailed data on Norway, we find that heterogeneity is close to irrelevant when the object of interest is aggregate moments. It seems that the Campbell and Mankiw (1989) assumption of different (by MPC) agents' incomes being proportional to aggregate income is about right, which then implies about no effects of heterogeneity on aggregate dynamics, as in the theoretical works of Bilbiie (2008, 2020) and Werning (2015) in various settings.

We first proposed two sufficient statistics that determine whether dynamics due to heterogeneity will contribute to aggregate fluctuations and showed that these are equivalent in a world where assets are in net zero supply, but will differ when net aggregate savings are possible. In that case, sufficient statistic 2, based on the distribution of consumption, will be the correct measure rather than (*in*)sufficient statistic 1, based on heterogeneous income dynamics. With our comprehensive data we study the dynamics of the Norwegian economy through the lens of these sufficient statistics.

We find that, based on *earnings*, income cyclicality (income betas) and MPCs are positively correlated and produce an aggregate MPC of 0.374, similar to the influential Patterson (2023) estimate for the U.S., also based on earnings. Based on this estimate, there is amplification of aggregate shocks through heterogeneity: a simple weighted average of individual MPCs that does not correct for the individual exposures is lower, i.e. 0.357. However, our data allows us to work with disposable income as well, and using that theory-consistent measure, we find that heterogeneity if anything mildly *dampens* aggregate fluctuations, with an aggregate MPC of 0.354 which is now *lower*, not higher than the simple weighted average, 0.366.

This is an important change in our understanding of aggregate implications of heterogeneity. We show that it is taxes and transfers, rather than capital income, that change the cyclical dynamics of income across MPC groups. Furthermore, we show by means of a counterfactual that even the much less progressive US tax and transfer system would have essentially offset any heterogeneity-induced amplification of shocks.

Turning to the better sufficient statistic, based directly on the behavior of consumption—–which we are able to empirically study for the first time thanks to our granular consumption data—–our substantive finding is that consumption cyclicalities of different groups are quite similar, regardless of how these groups are defined. Hence, heterogeneity does not add much to aggregate fluctuations based on this sufficient statistic either. Even in response to the substantial Global Financial Crisis shock, in a calibration that gives heterogeneity the best chance to play a role (i.e. with the largest justifiable share of hand-to-mouth agents), heterogeneity contributed only a modest 12% to amplification.

This paper focused on the effect of heterogeneity on amplifying shocks on impact and found about no evidence of such amplification. There are other ways heterogeneity may be consequential, above and beyond the importance of distributional issues per se, as well as those of targeted transfers. Not all response to shocks are on impact and intertemporal consumption responses may be affected by heterogeneity in ways contemporaneous ones are not. Measuring the distribution of relevant intertemporal MPCs, as theorized by Auclert et al. (2018) and Bilbiie (2018) is one item on our future research plan. Another is to properly measure the effect of heterogeneity-specific mechanisms on equilibrium outcomes due to heterogeneous income risk and associated self-insurance behavior. The challenging task of measuring (perhaps unrealized but perceived) individual income risk in micro data and studying how that affects aggregate behavior is waiting to be done. These are all questions central to macroeconomics of heterogeneous agents for which measurement is now catching up with theory.

References

Acharya, S. and Dogra, K. (2020). Understanding hank: Insights from a prank. *Econometrica*, 88(3):1113–1158.

Aguiar, M., Bils, M., and Boar, C. (2023). Who are the hand-to-mouth?

Ahn, S., Galaasen, S. M., and M, M. (2024). The cash-flow channel of monetary policy - evidence from billions of transactions. Technical report, Norges Bank mimeo.

Attanasio, O. P. and Pistaferri, L. (2016). Consumption inequality. *Journal of Economic Perspectives*, 30(2):3–28.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–2367.

Auclert, A., Rognlie, M., and Straub, L. (2018). The intertemporal keynesian cross.

Auclert, A., Rognlie, M., and Straub, L. (2020). Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model.

Bayer, C., Born, B., and Luetticke, R. (2020). Shocks, frictions, and inequality in us business cycles. CEPR Discussion Papers 14364.

Bayer, C., Lütticke, R., Pham-Dao, L., and Tjaden, V. (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica*, 87(1):255–290.

Berger, D., Bocola, L., and Dovis, A. (2023). Imperfect risk sharing and the business cycle. *The Quarterly Journal of Economics*, 138(3):1765–1815.

Bilbiie, F., Primiceri, G. E., and Tambalotti, A. (2022a). Inequality and business cycles. Working paper, University of Cambridge.

Bilbiie, F. O. (2008). Limited asset market participation, monetary policy and (inverted) aggregate demand logic. *Journal of Economic Theory*, 140(1):162–196.

Bilbiie, F. O. (2018). Monetary policy and heterogeneity: An analytical framework.

Bilbiie, F. O. (2020). The new keynesian cross. *Journal of Monetary Economics*, 114:90–108.

Bilbiie, F. O., Kaenzig, D. R., and Surico, P. (2022b). Capital and income inequality: An aggregate-demand complementarity. *Journal of Monetary Economics*, 126:154–169.

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921.

Boehm, J., Fize, E., and Jaravel, X. (2025). Five facts about mpcs: Evidence from a randomized experiment. *American Economic Review*, 115(1):1–42.

Broda, C. and Parker, J. A. (2014). The economic stimulus payments of 2008 and the aggregate demand for consumption. *Journal of Monetary Economics*, 68(S):S20–S36.

Broer, T., Harbo Hansen, N.-J., Krusell, P., and Öberg, E. (2020). The new keynesian transmission mechanism: A heterogeneous-agent perspective. *The Review of Economic Studies*, 87(1):77–101.

Campbell, J. Y. and Mankiw, N. G. (1989). Consumption, income, and interest rates: Reinterpreting the time series evidence. *NBER Macroeconomics Annual*, 4:185–216.

Challe, E., Matheron, J., Ragot, X., and Rubio-Ramirez, J. F. (2017). Precautionary saving and aggregate demand. *Quantitative Economics*, 8(2):435–478.

Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., and van Rooij, M. (2019). Asymmetric consumption effects of transitory income shocks. *The Economic Journal*, 129(622):2322–2341. Published: 17 May 2019.

Cloyne, J., Ferreira, C., and Surico, P. (2020). Monetary policy when households have debt: new evidence on the transmission mechanism. *The Review of Economic Studies*, 87(1):102–129.

Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy*, 99(5):957–976.

Coglianese, J., Olsson, M., and Patterson, C. (2024). Monetary policy and the labor market: A quasi-experiment in sweden. Mimeo.

Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2017). Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics*, 88:70–89.

Commault, J. (2022). Does consumption respond to transitory shocks? reconciling natural experiments and semistructural methods. *American Economic Journal: Macroeconomics*, 14(2):96–122.

Coronado, J. L., Lupton, J. P., and Sheiner, L. M. (2005). The household spending response to the 2003 tax cut: Evidence from survey data. Technical Report 2005-32, Federal Reserve Board, Washington, DC.

Debortoli, D. and Galí, J. (2024). Heterogeneity and aggregate fluctuations: Insights from tank models.

Debortoli, D. and Galí, J. (Forthcoming). Idiosyncratic income risk and aggregate fluctuations. *American Economic Journal: Macroeconomics*. Forthcoming in American Economic Journal: Macroeconomics.

Den Haan, W. J., Rendahl, P., and Riegler, M. (2017). Unemployment (fears) and deflationary spirals. *Journal of the European Economic Association*, 16(5):1281–1349.

Esselink, H. and Hernández, L. (2017). The use of cash by households in the euro area. *ECB Occasional Paper*, (201).

Fagereng, A., Guiso, L., Malacrino, D., and Pistaferri, L. (2020). Heterogeneity and persistence in returns to wealth. *Econometrica*, 88(1):115–170.

Fagereng, A., Holm, M. B., and Natvik, G. J. (2021). Mpc heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics*, 13(4):1–54.

Fagereng, A., Onshuus, H., and Torstensen, K. N. (2024a). The consumption expenditure response to unemployment: Evidence from norwegian households. *Journal of Monetary Economics*, page 103578.

Fagereng, A., Onshuus, H., and Torstensen, K. N. (2024b). The consumption expenditure response to unemployment: Evidence from norwegian households. *Journal of Monetary Economics*, 146:103578.

Ferriere, A., Grübener, P., Navarro, G., and Vardishvili, O. (2023). On the optimal design of transfers and income tax progressivity. *Journal of Political Economy Macroeconomics*, 1(2):276–333.

Fuster, A., Kaplan, G., and Zafar, B. (2021). What would you do with $500? spending responses to gains, losses, news, and loans. *The Review of Economic Studies*, 88(4):1760–1795.

Galí, J., López-Salido, J. D., and Vallés, J. (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association*, 5(1):227–270.

Ganong, P., Jones, D., Noel, P. J., Greig, F. E., Farrell, D., and Wheat, C. (2020). Wealth, race, and consumption smoothing of typical income shocks. (w27552).

Gornemann, N., Kuester, K., and Nakajima, M. (2016). Doves for the rich, hawks for the poor? distributional consequences of monetary policy.

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review*, 87(1):192–205.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What do data on millions of u.s. workers reveal about life-cycle earnings dynamics? *Econometrica*, 89(5):2303–2339.

Guvenen, F., Ozkan, S., and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, 122(3):621–660.

Guvenen, F., Schulhofer-Wohl, S., Song, J., and Yogo, M. (2017). Worker betas: Five facts about systematic earnings risk. *American Economic Review*, 107(5):398–403.

Hagedorn, M., Manovskii, I., and Mitman, K. (2019). The fiscal multiplier. Technical report, National Bureau of Economic Research.

Heathcote, J., Perri, F., and Violante, G. L. (2010). Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006. *Review of Economic Dynamics*, 13(1):15–51.

Heathcote, J., Perri, F., and Violante, G. L. (2020). The rise of us earnings inequality: Does the cycle drive the trend? *Review of economic dynamics*, 37:S181–S204.

Heathcote, J., Perri, F., Violante, G. L., and Zhang, L. (2023). More unequal we stand? inequality dynamics in the united states 1967–2021. *Review of Economic Dynamics*. forthcoming.

Holm, M. B., Paul, P., and Tischbirek, A. (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy*, 129(10):2861–2904.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *American Economic Review*, 83(4):685–709.

Jappelli, T. and Pistaferri, L. (2014). Fiscal policy and mpc heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–136.

Jappelli, T. and Pistaferri, L. (2020). Reported mpc and unobserved heterogeneity. *American Economic Journal: Economic Policy*, 12(4):275–297.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Kaplan, G., Violante, G. L., and Weidner, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*.

Kueng, L. (2018). Excess sensitivity of high-income consumers. *The Quarterly Journal of Economics*, 133(4):1693–1751.

Lewis, D., Melcangi, D., and Pilossoph, L. (2024). Latent heterogeneity in the marginal propensity to consume. Technical report, National Bureau of Economic Research.

Luetticke, R. (2021). Transmission of monetary policy with heterogeneity in household portfolios. *American Economic Journal: Macroeconomics*, 13(2):1–25.

Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy*, 99(5):928–956.

McKay, A. and Reis, R. (2016). The role of automatic stabilizers in the us business cycle. *Econometrica*, 84(1):141–194.

Misra, K. and Surico, P. (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics*, 6(4):84–106.

Norges Bank (2023). Retail payment services 2022. *Norges Bank Papers*, (1).

Oh, H. and Reis, R. (2012). Targeted transfers and the fiscal response to the great recession. *Journal of Monetary Economics*, 59:S50–S64.

Orchard, J., Ramey, V., and Wieland, J. (2025). Micro mpcs and macro counterfactuals: The case of the 2008 rebates. *Quarterly Journal of Economics*. Forthcoming.

Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer Spending and the Economic Stimulus Payments of 2008. *American Economic Review*, 103(6):2530–2553.

Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review*, 113(4):982–1012.

Pekanov, A. (2024). The earnings elasticities to aggregate fluctuations in the euro area. Mimeo.

Ravn, M. O. and Sterk, V. (2017). Job uncertainty and deep recessions. *Journal of Monetary Economics*, 90:125–141.

Ravn, M. O. and Sterk, V. (2020). Macroeconomic fluctuations with hank & sam: an analytical approach. *Journal of the European Economic Association*.

Ring, M. A. (2024). Wealth taxation and household saving: Evidence from assessment discontinuities in norway. *Review of Economic Studies*, page rdae100.

Sahm, C. R., Shapiro, M. D., and Slemrod, J. (2010). Household response to the 2008 tax rebate: Survey evidence and aggregate implications. *Tax Policy and the Economy*, 24.

Shapiro, M. D. and Slemrod, J. (1995). Consumer response to the timing of income: Evidence from a change in tax withholding. *American Economic Review*, 85(1):274–283.

Shapiro, M. D. and Slemrod, J. (2009). Did the 2008 tax rebates stimulate spending? *American Economic Review*, 99(2):374–379.

Souleles, N. S., Parker, J. A., and Johnson, D. S. (2006). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review*, 96(5):1589–1610.

Storesletten, K., Telmer, C., and Yaron, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717.

Townsend, R. M. (1994). Risk and insurance in village india. *Econometrica*, 62(3):539–591.

A Some Theoretical Derivations

We outline some details of the algebra for the case with net saving in equilibrium. The budget constraint of savers (in log-linear form) reads:

$$C_Y c_t^S + \frac{X_Y}{1-\lambda} x_t = Y_Y^S y_t^S, \tag{19}$$

where y_t^S is the (post-transfer) income of the savers and $X_Y \equiv X/Y$ denotes the steady-state share of variable X in GDP (income) Y; same for C and Y^S.[27]

Spenders just consume all their income in every period, i.e.:

$$c_t^H = y_t^H. \tag{20}$$

Goods market clearing requires that income equal saving (which can be investment, or "liquidity"):

$$y_t = C_Y c_t + X_Y x_t. \tag{21}$$

Aggregate consumption and income are given by:

$$c_t = \lambda c_t^H + (1-\lambda) c_t^S \tag{22}$$

$$y_t = \lambda Y_Y^H y_t^H + (1-\lambda) Y_Y^S y_t^S. \tag{23}$$

Replacing in the budget constraints (19) and (20):

$$c_t^H = \beta_y^H y_t \tag{24}$$

$$C_Y c_t^S + \frac{X_Y}{1-\lambda} x_t = \frac{1-\lambda \beta_y^H Y_Y^H}{1-\lambda} y_t.$$

We solve again for savers' consumption:

$$c_t^S = \frac{1-\lambda \beta_y^H \Omega}{1-\lambda} c_t, \tag{25}$$

where

$$\Omega \equiv \frac{1-X_Y}{1-\eta X_Y} > 1 \text{ when } \eta > 0$$

that is when saving (investment, or liquidity) is procyclical. This governs how much more volatile

[27]We focus on a case with equal consumption in steady state across households, i.e. $C^S = C^H = C$, achieved by a fixed steady-state transfer . This simplifies the analytics but is not needed. Furthermore, average consumption shares in the data of the two groups are pretty uniform (close to population shares).

total income is relative to consumption, solved from $y_t = C_Y c_t + X_Y x_t$

$$y_t = \frac{1 - X_Y}{1 - \eta X_Y} c_t$$

We obtain the aggregate Euler equation and our next Proposition:

$$c_t = E_t c_{t+1} - \frac{1 - \lambda}{1 - \lambda \beta_y^H \Omega} r_t. \tag{26}$$

The multiplier of an interest-rate cut when both channels are active is:

$$\frac{\partial c_t}{\partial \left(-r_t\right)} = \frac{1 - \lambda}{1 - \lambda \beta_y^H \Omega}. \tag{27}$$

There can be amplification even if $\beta_y^H < 1$ (disposable-income betas inversely correlated with MPCs).

A.1 Adding risk (for liquidity)

We use the simplest version of the THANK model in Bilbiie (2018) and refer the reader to that paper for details. Savers have a risk of becoming hand-to-mouth, which follows a Markov chain with transition probability $1 - s$. The savers' loglinearized Euler equation for (now, *liquid*) bonds takes into account the risk of transitioning to the constrained H state next period with probability $1 - s$:

$$c_t^S = s E_t c_{t+1}^S + \left(1 - s\right) E_t c_{t+1}^H - r_t. \tag{28}$$

Replacing individual consumptions (24) and (25) in (28) delivers the next Proposition.

The aggregate Euler equation with idiosyncratic risk and saving is:

$$c_t = \Theta_c E_t c_{t+1} - \Theta_r r_t,\ \Theta_c \equiv 1 + \left(1 - s\right) \frac{\beta_y^H \Omega - 1}{1 - \lambda \beta_y^H \Omega}; \tag{29}$$

There is aggregate-Euler **compounding** $\Theta > 1$ (for $s < 1$), that is an additional source of intertemporal amplification, if and only if saving/investment/liquidity is procyclical **enough**, specifically:

$$\beta_y^H \Omega > 1 \rightarrow \eta > 1 + \left(1 - \beta_y^H\right) \frac{1 - X_Y}{X_Y}, \tag{30}$$

This additional source of intertemporal amplification has stark implications, e.g. making the Taylor principle insufficient for determinacy and aggravating the forward guidance puzzle (Bilbiie et al. (2022b) for the case with investment). Procyclical *enough* investment in the sense of (30) generates Euler compounding even when income inequality is procyclical $\beta_y^H < 1$ and would *by itself* generate discounting $\Theta < 1$. The compounding intuition is similar to the one stemming from countercyclical inequality and risk, previously emphasized by Bilbiie (2018, 2020), Acharya and Dogra (2020), and Ravn and Sterk (2020).

A.2 Consumption inequality

Here we derive our main proposition in text. Subtracting budget constraints from one another and using the cyclicality of individual incomes β_y^H :

$$c_t^S - c_t^H = \frac{1-\beta_y^H C_Y}{(1-\lambda)C_Y} y_t - \frac{X_Y}{(1-\lambda)C_Y} x_t. \tag{31}$$

To show this, use $y_t^S = \left(1 - \lambda \beta_y^H Y_Y^H\right) y_t / \left((1-\lambda) Y_Y^S\right)$ to rewrite:

$$C_Y c_t^S - C_Y c_t^H + \frac{X_Y}{1-\lambda} x_t = Y_Y^S y_t^S - C_Y y_t^H \tag{32}$$

$$= \frac{1-\lambda\beta_y^H Y_Y^H}{1-\lambda} y_t - C_Y \beta_y^H y_t \tag{33}$$

$$C_Y c_t^S - C_Y c_t^H + \frac{X_Y}{1-\lambda} x_t = \frac{1-\beta_y^H \left(\lambda Y_Y^H + (1-\lambda) C_Y\right)}{1-\lambda} y_t \tag{34}$$

Use $C_Y + \frac{X_Y}{1-\lambda} = Y_Y^S$

$$\frac{1-\beta_y^H \left(\lambda Y_Y^H + (1-\lambda) C_Y\right)}{1-\lambda} = \frac{1-\beta_y^H \left(1 + (1-\lambda)\left(C_Y - Y_Y^S\right)\right)}{1-\lambda} = \frac{1-\beta_y^H \left(1 - X_Y\right)}{1-\lambda} = \frac{1-\beta_y^H C_Y}{1-\lambda}$$

Substitute

$$C_Y \left(c_t^S - c_t^H\right) = \frac{1-\beta_y^H C_Y}{1-\lambda} y_t - \frac{X_Y}{1-\lambda} x_t \tag{35}$$

Consumption inequality can be countercyclical iff, even with income-betas skewed towards H ($\beta_y^H < 1$), saving is procyclical enough.

B Data

B.1 Administrative Data(XXX TO BE WRITTEN XXX)

Things to include here

- What exactly is our income variables measuring?
  - Eg. labor earnings includes sickness and parental leave benefits

    - Capital income is mainly dividends?
    - How do we go from tax value to market value (not for business wealth), and issue with housing wealth.

- How do we define households

B.2 Electronic Transactions Data (XXX TO BE WRITTEN XXX)

Draw extensively from Ahn et al (2024).

Things to include here

- Boring details on the data
- Back up the 80 percent claim
- Imputation of credit card, and basic pre-cleaning
    - Mortgage payments, large transactions, person to person transfers

B.3 Sample Selection and Summary Statistics (XXX TO BE BE_WRITTEN XXX)

Things to include here

- Why and which individulas to we exclude.
- Summary statistics on whole population and final analysis sample.

We restrict the sample to individuals between the ages of 25 and 55. We further exclude individuals for which consumption expenditures are measured less well.

Summary statistics are presented in Table **??** and sample size by year in Figure **??**. In Table **??** we also report summary statistics for the individual income components. All nominal values are deflated to 2018 real terms using the CPI index and expressed in US dollars using the 2018 exchange rate.

Summary statistics are presented in Table **??** and sample size by year in Figure **??**. In Table **??** we also report summary statistics for the individual income components. All nominal values are deflated to 2018 real terms using the CPI index and expressed in US dollars using the 2018 exchange rate.

Need to note how we handle household level variables (divide by two, but could have done some equivalization too). But important, mpc is estimated without normalization, and elasticities are performed at the invidiual level (but here consumption is household consumption divided by 2).

To study consumption inequality, we need to exclude individuals with very low or even zero consumption, as this likely reflects missing data. For the transaction-based consumption measure, an additional issue is that the missing information seems to be time-varying, which could potentially influence our assessment of time-varying inequality.

In addition to the full sample, we consider an alternative sample. The first alternative sample is based on restrictions derived from the transaction-based consumption measure. Specifically, to be included in this sample we require household. We require the transaction-based consumption-to-income ratio to be within the 2.5-97.5 percentile. Households should have at least one debit card transaction a week for 50 percent of the year. Household members should be Norwegian residents the whole year (i.e., registered residents on both Jan 1st and Dec 31st) throughout the sample. We exclude self-employed.

The restriction on non-migration and non-self-employed is related to only partially observing consumption for the former group and contamination from business-related expenditures for the latter group. In total, imposing the residential and self-employed restriction means that we drop around 4.8 million out 16.4 million observations. The restriction on debit card activity and consumption-to-income ratio implies a further dropping 1.4 million observations. The final sample has 10.2 million household-year observations.

Summary statistics for this restricted sample is shown in Panel (C) in Table 3

C Further Estimation Results and Details

This Appendix groups several robustness exercises and describes details of some of the estimation results.

C.1 MPC estimation results

Table XXX reports the estimation results from a regression including all the determinants jointly:

Table 3: Income and Consumption Statistics Households 2006-2018

	N (mill.)	mean	sd	p10	p50	p90
Panel (A): Unrestricted sample 2006-2018:						
Salaries	16,4	54600	41312	0	54899	97477
Salaries, Self empl.	16,4	57964	62019	616	57123	101149
Salaries, Self empl., Capital	16,4	60789	87639	1316	57941	105560
Salaries, Self empl., Capital, Transfers	16,4	68232	85761	29092	61853	108730
Total after-tax income	16,4	50375	69387	24959	47101	76738
Cons (transact.)	16,4	36057	34057	5577	31995	64802
pst. w Cons>0 (transact.)	16,4	95	22	100	100	100
Deposits	16,4	20891	78692	330	7188	48137
Debt	16,4	122293	194399	0	94991	266434
Securitites	16,4	28555	1183108	0	0	16964
Housing	16,4	148655	177679	0	115883	355898
Age (head)	16,4	41	9	29	41	52
Panel (B): Unrestricted 2006-2015:						
Total after-tax income	12,4	50236	62252	25242	46898	76309
Cons (transact.)	12,4	35440	34899	3952	31567	64050
Cons (imputed)	12,4	47296	231604	13771	39996	86907
pst. w Cons>0 (transact.)	12,4	94	24	100	100	100
pst w Cons>0 (imputed)	12,4	94	24	100	100	100
Panel (C): Restricted sample 2006-2018:						
Salaries	10,2	59803	38710	7881	58889	99789
Salaries, Self empl.	10,2	59803	38710	7881	58889	99789
Salaries, Self empl., Capital	10,2	62171	51376	8562	59593	103474
Salaries, Self empl., Capital, Transfers	10,2	69751	47801	35135	63225	106532
Total after-tax income	10,2	51911	30774	29453	48157	75574
Cons (transact.)	10,2	35056	19504	13559	32408	58537
pct. w Cons>0 (transact.)	10,2	100	0	100	100	100
Deposits	10,2	18954	47982	456	7068	44196
Debt	10,2	120516	140466	879	99830	258226
Securitites	10,2	18080	426907	0	0	14269
Housing	10,2	157320	167728	0	133115	359744
Age (head)	10,2	40	9	28	40	52

Notes. This table shows summary statistics for our sample of households where both the head and partner is of age 25-55. All nominal values are deflated to 2018 real terms using the CPI index and expressed in US dollars using the 2018 exchange. For households couples values are divided by two. Panel (A) reports values for the unrestricted sample for the years 2006-2018, while panel (B) for the years 2006-2015. In panel (C) the sample is restricted to include households with transaction based consumption-to-income ratio within the 2.5-97.5 percentile and have at least one debit card transaction a week in 50 percent of the years. We further require household members to be Norwegian residents the whole year (i.e. registered residents both Jan 1st and Dec 31st) and to never be observed as self-employed.

Table 1: Coefficient Estimates for Individual Marginal Propensities to Consume

Dep. Var.	ΔC
ΔI	-1.068
	(0.188)
Age	0.088
	(0.010)
Age sq.	-0.001
	(0.000)
Male	0.000
	(.)
Female	0.034
	(0.017)
1st Quintile (lagged income)	0.000
	(.)
2nd Quintile (lagged income)	-0.059

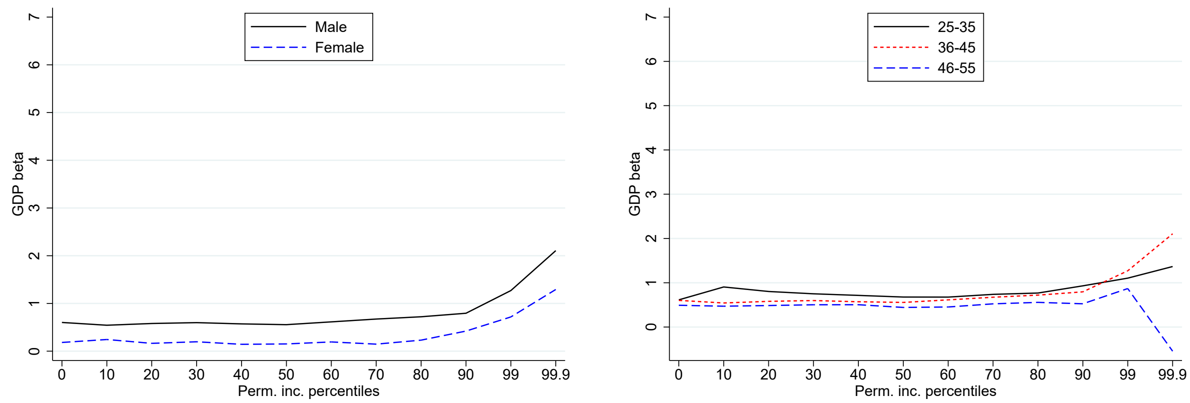

Gender Age

Figure 16: Disposable Income Betas.

In Table: 4 we estimate drivers of household MPCs, by regressing MPC estimates on household characteristics, income quartiles, and HTM status.

C.2 Appendix: Disposable-Income Betas

The results in Figure 16 show that, perhaps as might be expected especially in a country like Norway, much of the cyclicality present for labor earnings is eliminated when looking at net income–presumably mostly due to the insurance inherent in automatic stabilizers, through taxes and transfers (which is something we decompose and ascertain in the data in text, in the MPC dimension). This confirms and is another way to visualize one of our main findings in text, showing that the pattern of income betas flattens out in the MPC distribution when passing from labor earnings to total income to post-tax-and-transfer, net disposable income.

C.3 Appendix: Income Betas By MPC Decile

In this section we explain how we estimate income betas along the MPC distribution and provide estimates of each income measure with confidence intervals. In section XXX, we used specification XXX from Guvenen et al., where the earnings beta is the coefficient on GDP growth in a regression of the log change in earnings on the log change in GDP. The beta coefficient within a permanent income group can then be interpreted as the average elasticity of earnings growth to GDP growth within that group. In section XXX, it is more relevant to estimate the betas weighted by the income shares of individuals within the group, which brings us closer to our main object of interest, the aggregate MPC. Hence, we estimate the regression

$$\frac{Y_{i,t} - Y_{i,t-1}}{\bar{Y}_{g,t-1}} = \alpha_g + \beta_g \left(\frac{Y_t - Y_{t-1}}{Y_{t-1}} \right) + \varepsilon_{n,t}, \tag{36}$$

Table 4: Coefficient Estimates for Individual Marginal Propensities to Consume

Dep. Var.	ΔC
ΔI	-1.094
	(0.288)
Age	0.092
	(0.014)
Age sq.	-0.001
	(0.000)
Male	0.000
	(.)
Female	0.013
	(0.021)
1st Quintile (lagged income)	0.000
	(.)
2nd Quintile (lagged income)	-0.056
	(0.039)
3rd Quintile (lagged income)	-0.135
	(0.036)
4th Quintile (lagged income)	-0.123
	(0.034)
5th Quintile (lagged income)	-0.082
	(0.034)
< High School	0.000
	(.)
High School	-0.019
	(0.028)
Uni (lower)	-0.153
	(0.031)
Uni (higher)	-0.177
	(0.035)
Missing education	-0.192
	(0.044)
Single	0.000
	(.)
Couple	0.010
	(0.018)
Not HTM	0.000
	(.)
HTM (Net Wealth)	-0.153
	(0.025)
HTM (Liquid)	0.033
	(0.049)
HTM (Net Wealth & Liquid)	0.051
	(0.037)
No. Observations	7823252
Year X Municip. FEs	X

Notes:

where $Y_{i,t}$ is a measure of income for individual i, $\bar{Y}_{g,t}$is the average income for individuals within group g at time t, and $Y_t = \sum_i Y_{i,t}$is aggregated income for all individuals in the sample at time t. The coefficient β_g can now be interpreted as the elasticity of group g's income with respect to aggregated income.

We estimate equation 36 using pooled OLS, with White heteroskedasticity robust standard errors. The sample is balanced year by year. This ensures that the independent variable in the regression is the total income growth of all individuals in the regression sample.It also ensures that the betas approximately sum to one across all MPC deciles when weighted by each decile's income share. To see why this is the case, consider the general regression

$$y_{n,t} = \alpha_g + \beta_g x_t + \varepsilon_{n,t},$$

where n is an individual, g is a group (here: deciles of MPC) and t is time. Let I_g be the set of individuals in group g, which we assume to be fixed over time, and let T_g be the number of individuals in each group.When we estimate the regression with pooled OLS, we get the estimator

$$\hat{\beta}_g = \frac{\sum_t \sum_{n \in I_g} \hat{y}_{n,t} \hat{x}_t}{\sum_t \sum_{n \in I_g} \hat{x}_t \hat{x}_t},$$

where hats denote deviations from mean. Suppose that there are weights s_g such that

$$\sum_g s_g \left(\frac{\sum_{n \in I_g} y_{n,t}}{T_g} \right) = x_t$$

for every t. In other words, the weighted sum of y across all individuals and groups equals the variable x, with equal weights for every year. Then we have

$$\sum_g s_g \hat{\beta}_g = 1.$$

Figure 17 contains the results of this regression for the four measures of income.

C.4 Appendix: The US Tax and Transfer System

C.4.1 Parameterization of the US tax and transfer system

Following Ferriere et al. (2023), our parameterization of the US tax and transfer system has three elements. First, we assume that there is a flat tax rate on capital income, τ_c. Second, we assume that labor income y_l is taxed at rate

$$\tau_l(y_l) = \exp\left(\log(\lambda) \left(\frac{y_l}{\bar{y}} \right)^{-2\theta} \right), \tag{37}$$

where y_l is labor income and $\bar{y}$ is average labor and capital income. The parameter θ determines the progressivity of the tax system; marginal tax rates increase in labor income when θ is positive.

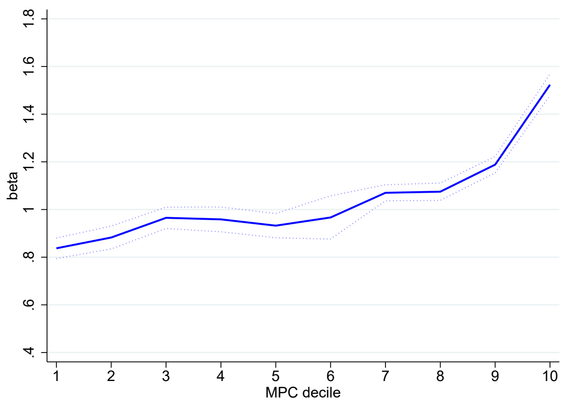

(a) Labor earnings

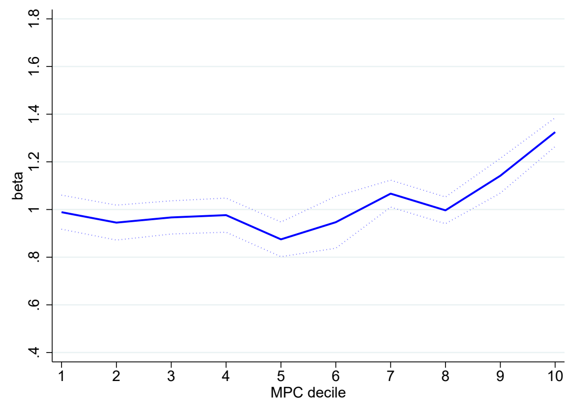

(b) Labor earnings and capital income

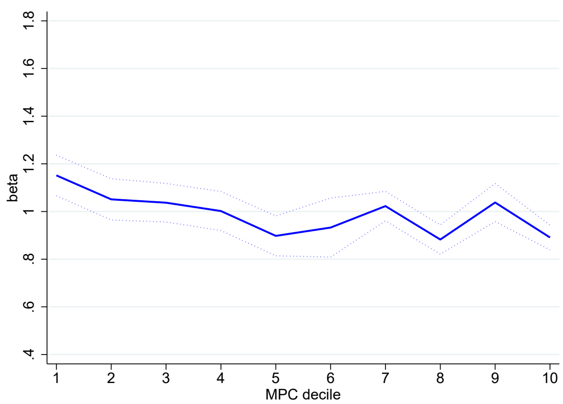

(c) Labor earnings, capital income and transfers

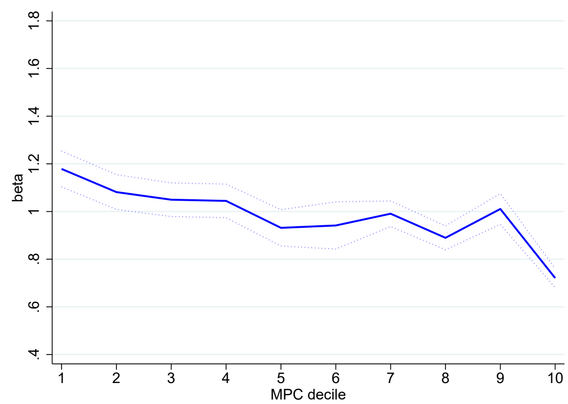

(d) Disposable income

Notes: Estimated income betas from regression XXX, by MPC decile. Dotted lines indicate 95% confidence bands based on heteroskedasticity robust standard errors.

Figure 17: Estimated income betas by MPC decile.

The parameter λ determines the level of tax rates; when $\theta = 0$, everyone pays a flat tax equal to λ.

Third, transfers depend on total labor and capital income y. Specifically, we assume that the level of transfers received is given by the function

$$T(y) = m\bar{y}\frac{2\exp\left(-\xi\left(\frac{y}{\bar{y}}\right)\right)}{1+\exp\left(-\xi\left(\frac{y}{\bar{y}}\right)\right)}. \tag{38}$$

This function implies that households with zero income receive transfers equal to m times average income in the population. The parameter ξ determines the rate at which transfers phase out with increasing income; when $\xi = 0$, each households receive a lump sum.

Our parameterization of the US tax and transfer system is based on the estimates of these functions on Current Population Survey (CPS) household-level data for the year 2013, from Ferriere et al. (2023). The estimates for the five parameters of interest are reproduced on the first line of table 5, which also shows the square root of the mean squared deviation between actual income after taxes and transfers and the imputed level based on applying the estimated functions to income before taxes and transfers. We report the deviation in both levels and logs. Ferriere et al. (2023) report other measures of fit, finding that this set of functions provide a good approximation to actual transfers and taxes across the income distribution.

When applying the US tax and transfer system to our Norwegian data, we follow Ferriere et al. (2023) as closely as possible. We let y_l be labor earnings inclusive of unemployment benefits, while y is y_l plus capital income.

C.4.2 Estimation of the Norwegian tax and transfer system

In this section we describe how we estimate the tax and transfer functions described in section C.4.1 on our Norwegian micro data. Ferriere et al. (2023) estimate the US tax and transfer system on household-level data. Since all Norwegian tax returns are filed jointly, and because we are concerned with individual-level outcomes, we estimate the functions in section C.4.1 on individual data.

In Norway, individuals pay two types of tax rates. First, general income (labor earnings, pensions, capital income, self-employment income and taxable government transfers) net of deductions are taxed at a flat rate.[28] Second, personal income (labor earnings, taxable transfers and pensions) are also taxed at an additional bracket tax rate, such that income in higher brackets are taxed at higher marginal rates. In addition, a wealth tax is paid on net wealth above a threshold. Since the wealth tax is not based on income, we subtract it from total taxes.

When estimating the functions in section C.4.1, we first let capital and self-employment income be taxed at the flat rate τ_c. We calibrate this rate at the average of the tax rate on general income over our sample period 2006-2018.[29] Second, since we only observe total taxes, we construct a measure of labor taxes paid at the individual level by subtracting imputed capital taxes from total taxes net of the wealth tax. We define y_l as personal income, which also gives us a value for the average labor tax rate paid by each individual. Then we estimate equation 37 by non-linear least squares. Third, we define y as labor earnings and capital income, and we estimate equation 38 using non-linear least squares.

The second line of table 5 contains the parameter estimates on Norwegian data. The flat rate on capital is substantially higher in Norway than in the United States. While a similar level for the parameter λ indicates that levels of labor taxes are similar in the two countries, a higher level of θ in Norway shows that the Norwegian tax system is more progressive. Finally, the level parameter for transfers, m, is substantially higher in Norway than in the US, indicating that low-income individuals in Norway receive more government transfers as a fraction of average income.

[28]Most people take a flat minimum deduction. Taxable government transfers include pensions and unemployment benefits.

[29]The rate was fixed at 28% from 2006 to 2013, decreasing to 27% in 2014 and then increasing to 29% in 2016, 30% in 2017 and 31% in 2018.

Table 5: Estimated coefficients for tax and transfer functions

	Parameters					Deviations	
	τ_c	λ	θ	m	ξ	In levels	In logs
US System (Ferriere et al., 2023)	0.133	0.247	0.077	0.088	4.22	7543	0.158
Norwegian System	0.283	0.262	0.103	0.563	3.73	9616	0.237

Notes: The table shows the estimated parameters for the tax and transfer functions in equations 37 and 38, respectively, as well as the flat capital tax rate τ_c. The parameters for the US system are from table B3 in Ferriere et al. (2023). The two rightmost columns summarize the fit of the functions as the square root of the mean squared difference between fitted values and actual values for disposable income. NOK values are converted to USD using the average exchange rate over the period 2006-2018.

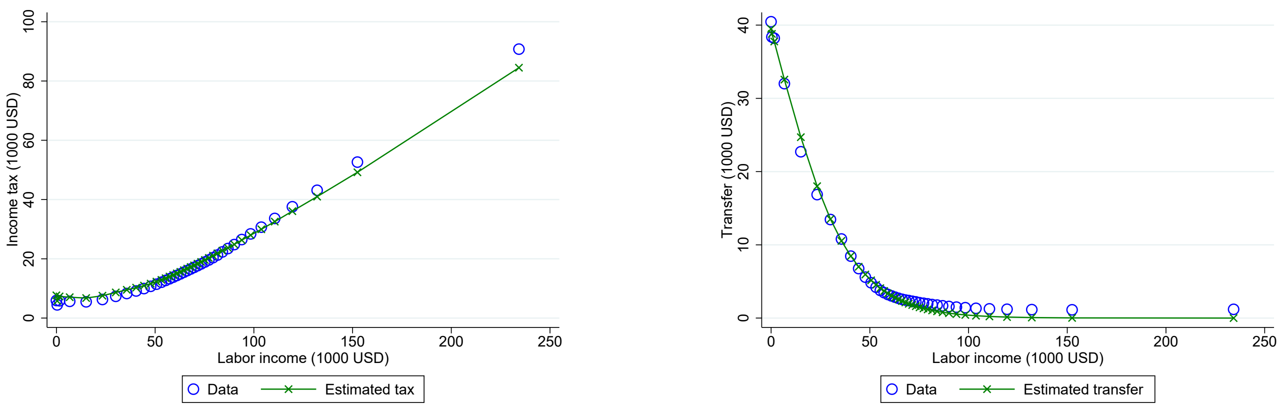

(a) Fit of labor tax function

(b) Fit of transfer function

Notes: The figures demonstrate the fit of the estimated tax and transfer functions in equations XXX and XXX, by 2.5% quantiles of labor income. NOK values are converted to USD using the average exchange rate over the period 2006-2018.

Figure 18: Fit of tax and transfer functions for Norway.

Table 5 also shows that the fit of the transfer and labor tax functions are on par with the same set of functions estimated on US data. Figure 18 shows the fit of the estimated average taxes and transfers along the distribution of labor income. [30]

C.4.3 Income betas at the household level

Figure 20shows the results of re-estimating the income betas along the MPC distribution with household-level data. The income betas show a similar pattern across MPC deciles both for la-

[30]Part of the reason for the slightly worse fit in Norway might be that while we only require individuals in our dataset to have non-negative income, Ferriere et al. (2023) only keep individuals with income and labor income above $5000 in their estimation dataset. We find that estimating on an equivalently censored dataset gives a better fit. Results are available upon request.

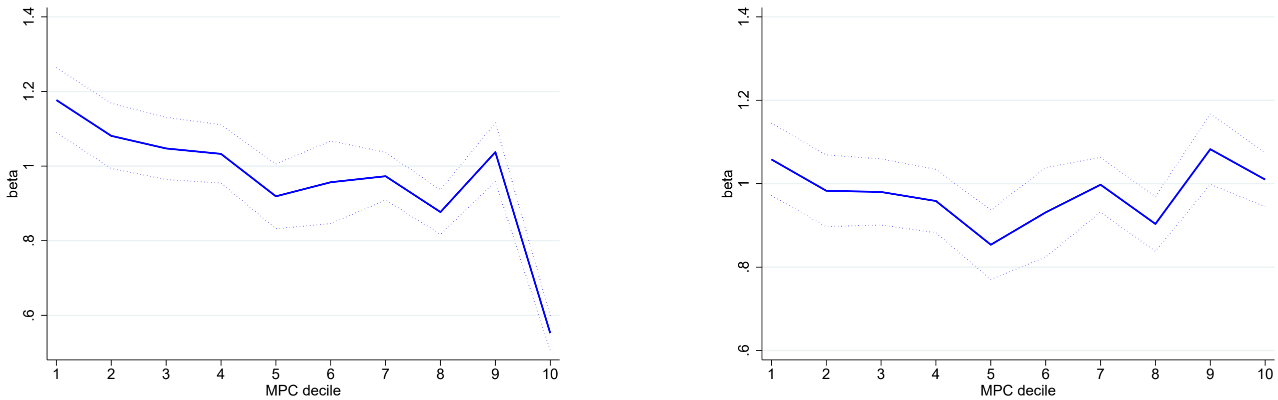

(a) Norwegian tax and transfer system

(b) US tax and transfer system

Notes: Estimated income betas from regression XXX, by MPC decile, for disposable income imputed from estimated tax and transfer system of Norway (left) and the United States (right). Dotted lines indicate 95% confidence bands based on heteroskedasticity robust standard errors.

Figure 19: Estimated disposable income betas based on estimated tax and transfer system.

bor income, actual disposable income and imputed disposable income based on the US tax and transfer system.

C.5 Appendix: External Validation/Sanity Check: Beta-Implied MPCs

We use the two betas to compute "beta-implied MPCs" for j:

$$\frac{\beta_C^j}{\beta_Y^j} = \frac{\partial c_t^j}{\partial c_t} / \frac{\partial y_t^j}{\partial y_t} \rightarrow \tilde{MPC}^j = \frac{\beta_C^j}{\beta_Y^j} \frac{\partial c_t}{\partial y_t} \frac{C_t^j}{Y_t^j}$$

C.6 Robust Blundell-Pistaferri-Preston MPC estimator

To address the role of persistence of income shocks we estimate the robust semi-structural estimator proposed by Commault (2022) (BPP-C). Building on Blundell et al. (2008) (BPP) uses panel data on income and consumption to leverage the variance of income and its covariance with consumption to estimate MPC with respect to transitory income shocks. The key difference is that while BPP assumes i.i.d. transitory income shocks, BPP-C allows for persistent transitory income shocks. The statistical model of log household rezidualized income $ln(y_{i,t})$ is assumed to be the sum of a random walk permanent $p_{i,t}$ and an $MA(k)$ transitory $\mu_{i,t}$ income component

$$ln(y_{i,t}) = p_{i,t} + \mu_{i,t}$$

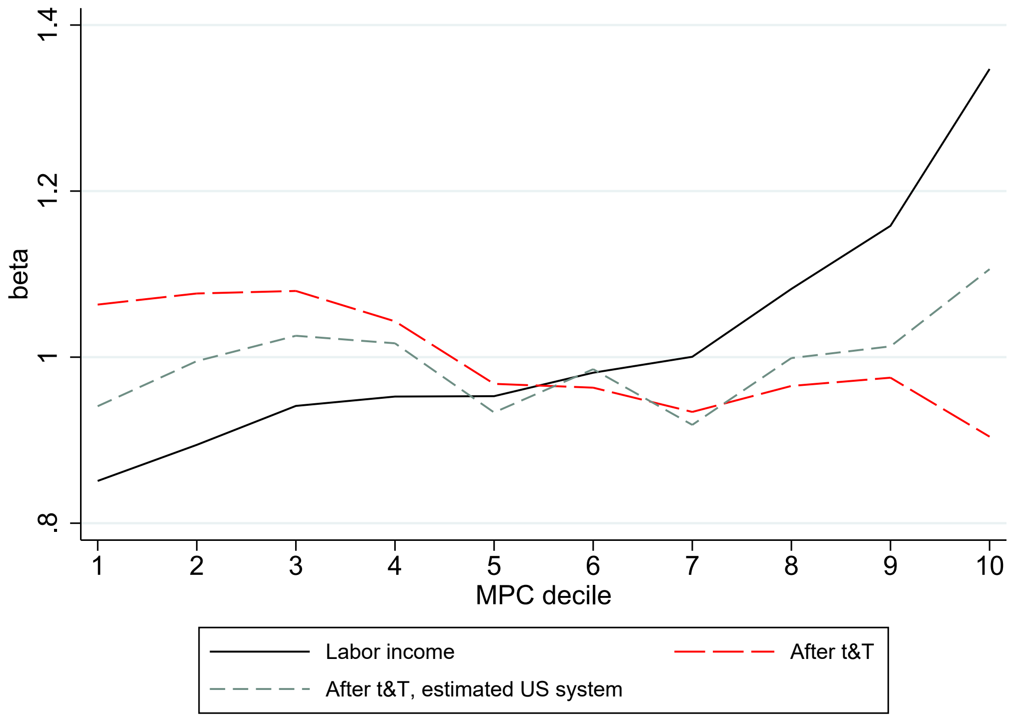

Notes: Estimated income betas at the household level for labor income, actual disposable income (after tax and transfers) and imputed disposable income based on the US tax and transfer system.

Figure 20: Estimated income betas at household level.

$$p_{i,t} = p_{i,t-1} + \nu_{i,t}$$

$$\mu_{i,t} = \varepsilon_{i,t} + \theta_1 \varepsilon_{i,t-1} + \theta_2 \varepsilon_{i,t-2} + ... + \theta_k \varepsilon_{i,t-k},$$

where ν and ε are drawn independently from one another and over time.

Allowing for persistence has implications for identification of the factor loading of temporary income shocks on consumption growth. Both BPP and BPP-C instrument temporary income shocks with future income changes. BPP uses next-period income, while BPP-C uses $\Delta ln(y_{i,t+k})$ as an instrument. The idea of BPP-C is to isolate the effect of the current transitory income shock by the value of future log-income growth, that correlates with the realization of the current transitory shock at t but is orthogonal to any of the other current of past income shocks. The resulting BPP-C estimator is the following

$$\phi^{\varepsilon} = \frac{cov(\Delta ln(c_{i,t),\varepsilon_{i,t}}}{var(\varepsilon_{i,t})} = \frac{cov(\Delta ln(c_{i,t}), -\Delta ln(y_{i,t+k+1})}{cov(\Delta ln(y_{i,t}), -\Delta ln(y_{i,t+k+1})}.$$

Using panel data on both net income and consumption, we can estimate BPP-C with the generalized method of moments.

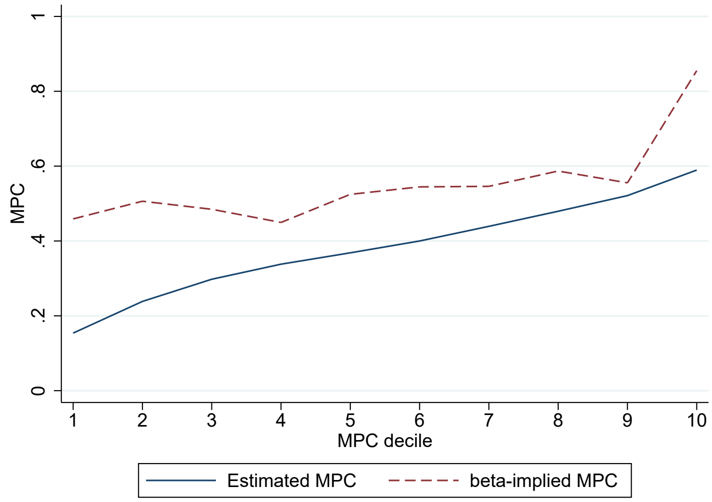

Figure 21: Betas-Implied MPCs.

Our sample selection is similar to BPP and BPP-C. We exclude households that were not continuously married between 2006 and 2016 or experienced a change in household head. We retain the life cycle age restriction from our main sample, keeping only households with heads aged 25–55. Additionally, we drop outliers in income, consumption, and individual income betas. The final sample includes 1,044,994 household-year observations from 107,660 households.

We begin by residualizing income and consumption using controls for year, family size, number of children under 18, household head's age and education, and place of residence. We then compute covariances of residualized log income growth to verify that the transitory component aligns with an $MA(1)$ process. Accordingly, we instrument transitory income at t using income changes at $t+2$. A summary of our estimation results for the full sample is provided below.

Table 6: BPP-C MPC estimate for the whole sample

MPC	0.300***
	(0.0273)
Observations	859321

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$