HANKSSON*

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

May 2025

Abstract

HANK Sufficient Statistics Out of Norway (HANKSSON) answers a core question of the heterogeneity in macroeconomics literature: does heterogeneity amplify the *aggregate* effects of demand policies and shocks? We provide two sufficient statistics (SS) and test them using individual-level matched data for personal characteristics, income, wealth and non-imputed, actual 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 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. 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; when applied to our data, 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: E21, E24, E32, E52, E64, H31

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, Jeanne Commault, Marco d'Amico, Davide Debortoli, Axelle Ferrière, Joel Flynn, Jordi Gali, Diego Känzig, Ben Moll, Simon Mongey, Christina Patterson, Fabrizio Perri, Paolo Surico, Gianluca Violante, to conference participants at the 2024 SED in Barcelona, "Frontiers in Macro" Cambridge, CEPR Paris Symposium, Paris Dauphine Heterogeneity, Pavia workshop, and seminar participants at the EUI, Imperial, and PSE. Bilbiie gratefully acknowledges the support of the Janeway Institute and Keynes Fund at Cambridge, and the hospitality of CREi and PSE. The views do not represent Norges Bank.

1 Introduction

Does household heterogeneity amplify the aggregate effects of demand policies and shocks? This is a core question of a flourishing literature that studies the interaction of distributions and aggregates in shaping macro fluctuations and the transmission of monetary and fiscal policies. In a wide range of recent models, such as those in the HANK (Heterogeneous Agent New Keynesian) class, heterogeneity drives a two-way feedback between macro variables and distributions of individual variables, or "inequality", which shape each other and are determined jointly in equilibrium.

In this paper, we tackle the measurement of some key amplification channels in these models. Informed by straightforward organizing theory and using a new and unique dataset of individual-level income, wealth and consumption from Norway, we ask whether these channels are salient in the data, and if so, which dimensions of heterogeneity are most important.[1] The key predictions and transmission channels in the majority of models concern the distribution and evolution of *consumption*, which also determines welfare and optimal policy prescriptions. Yet notorious data issues in traditional datasets make the measurement of consumption challenging. Hence, the literature has often resorted to using labor earnings or, at best, income data, and their relationships to marginal propensities to consume (MPC) based on imputed consumption measures. Depending on how well various proxies work, these measures may or may not be model consistent. It is thus of first-order importance to test whether these transmission mechanisms due to heterogeneity are salient using actual consumption data.

A key amplification channel in models of the HANK class is that individual variables (consumption, income, earnings) vary over the cycle in a manner that is correlated with that individual's (or household's) marginal propensity to consume. 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, 2020, 2024) and Auclert (2019), and is at work in quantitative HANK models such as Oh and Reis (2012), Kaplan et al. (2018), Bayer et al. (2019), Gornemann et al. (2016). Patterson (2023), in recent but justly influential work, provides a novel measurement of this amplification mechanism based on earnings, not disposable income or actual consumption (consumption expenditures for MPC estimation is imputed using the Blundell et al. (2008) method). A key contribution of our paper is to bring both disposable total income (including capital income, as well as taxes and transfers) and actual consumption data in combination with data on wealth to inform that measurement.

We provide two sufficient statistics and measure them using our rich consumption, income, and wealth data. The first sufficient statistic is the difference between the aggregate MPC and the

[1] Despite the explosion of research using administrative data, examining consumption and income links proved 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. It is important to identify households for our analysis due to their important role in insurance against income fluctuations, see e.g. Halvorsen et al. (2024).

average MPC. To measure the former, we first estimate an MPC for every individual in Norway, using a similar methodology to Patterson (2023). Consistent with much of the literature, we find that liquid asset holdings are an important determinant of an individual's MPC. We then estimate individual "income betas" (using a terminology borrowed from finance) — the elasticity of individual income with respect to aggregate income — for both labor earnings and total disposable (post tax and transfer) income. Combining MPCs and income betas, we calculate the aggregate MPC of Norway. This in turn allows us to compute the multiplier and ascertain whether, through the lens of the aggregate MPC, there is *aggregate demand amplification* through heterogeneity in the Norwegian economy, compared to a multiplier based on the average MPC that would be relevant in a representative agent model.

In most HANK models, agents who live hand-to-mouth (HtM)—and thus are not on their Euler equation—are those who have low levels of liquid assets and/or net wealth relative to income. The second sufficient statistic combines betas for consumption, rather than income, with agents' HtM status based on their wealth position. This allows us to measure amplification due to heterogeneity more directly, without the need to estimate MPCs. The two sufficient statistics determine whether heterogeneity amplifies or dampens aggregate fluctuations. But while they are equivalent when saving rates are time invariant (or, trivially, when savings are zero), they differ when savings have cyclical dynamics of their own. In the latter case, the sufficient statistic based on consumption betas is the correct measure of amplification.

We find that estimating betas based on labor earnings would lead us to conclude that there is aggregate demand amplification at work in Norway, qualitatively similar to the findings for the US (Patterson, 2023), albeit of a smaller magnitude. However, measurement based on total disposable income 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 agents), and part to the tax and transfer system whose progressivity insures high-MPC agents to a larger extent than low-MPC ones. Since the tax and transfer system plays an important role for our quasi-irrelevance result, it is legitimate to ask whether this would also be the case in an economy with a less progressive system.We therefore ask whether the US tax and transfer system counterfactually applied to Norway would undo our result. Applying the counterfactual system estimated by Ferriere et al. (2023) to the labor and capital income data for Norway, we find that while the results are quantitatively slightly different, the quasi-neutrality still remains. Therefore, in light of these results, it seems that the assumption made by Campbell and Mankiw (1989) of different (by MPC) agents' incomes being proportional to aggregate income over the business cycle is about right; this instead implies about near-irrelevance of heterogeneity for aggregate dynamics, as in the theoretical works of Bilbiie (2008, 2020) and Werning (2015) in two-agent and heterogeneous-agent models, respectively, whereby consumption equals income.

Yet even with perfectly insured disposable income, individual consumption may still fluctuate differently over the cycle, as households make different investment, precautionary saving, and liquidity choices. Thus finally, we leverage our high-quality consumption data to compute the second sufficient statistic, studying whether the consumption of agents who are more likely to be hand-to-mouth is more responsive to aggregate consumption — be it over the whole sample or focusing only on the Great Recession episode. Measuring hand-to-mouth status in a variety of ways, based on either total wealth, liquidity or education, we find that this is hardly the case: consumption of both the HtM and non-HtM groups, regardless of the classification criterion, behave in very similar ways. The most unequal cyclicality of consumption across groups can be seen when zooming in on the Great Recession episode, but even then, an upper bound on the degree of amplification yielded by heterogeneity is only of the order of ten percent.

Our work builds on the seminal paper by Patterson (2023). We share the focus on measuring the correlation between estimated MPCs and individual "cyclicalities" as a way to learn whether the theoretical amplification mechanisms are salient in the data. There are several important differences that form the basis of this paper's contributions. First, the coverage and quality of our transactions-based consumption data allow a more precise estimation of the MPCs. Second, the ability to match this with wealth data allows us to study whether wealth variables are important determinants of MPCs. Consistent with the literature pioneered by Kaplan et al. (2014), we find that *liquid* wealth predicts a person's consumption response to an income shock. Third, the availability of administrative disposable income data, including both capital and taxes and transfers, allows us to study the cyclicality of individual income. We find that even though, in line with Patterson (2023), the cyclicality of earnings is higher for high-MPC consumers, this is not the case for disposable income. Partly due to capital income (which is very cyclical and accrues mostly 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 income is in fact slightly lower for high-MPC consumers. Thus, a key result is that the aggregate MPC is very close to the average MPC.

Finally, our data allows us to look directly at whether aggregate demand amplification of the type predicted by HANK and TANK-type models is of economic significance empirically. In these models, the consumption of agents who participate in the market for liquid assets obeys an Euler equation that prices those assets, whereas the consumption of non-participants does not. The cyclicality of inequality between the two groups' consumptions is a sufficient statistic for measuring whether the amplification mechanism is operating in the data. We perform this test by estimating the consumption betas across the MPC distribution. We find consumption inequality not to have economically meaningful cyclical variation.

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. The flip side of this finding is that discretionary monetary and fiscal policies are less effective in aggregate demand management than many HANK-type models may suggest; nevertheless, other more direct tools such as targeted transfers to high-MPC individuals still have high AD effects even in such economies. But the positive takeaway is that automatic stabilizers do, in fact, stabilize: the tax and transfer system is working to mitigate the effects of aggregate shocks through distributional dynamics.

The theory, measurement, and results presented in this paper place it within several strands of the **literature**. The first strand we contribute to is the burgeoning literature on TANK and HANK models used to analyze fluctuations and the effects and design of policies. These include earlier TANK models (Galí et al. (2007), Bilbiie (2008)), rich-heterogeneity HANK models (Oh and Reis (2012), McKay and Reis (2016), Kaplan et al. (2018), Gornemann et al. (2016), Guerrieri and Lorenzoni (2017), Auclert (2019), Ravn and Sterk (2017), Den Haan et al. (2017), Challe et al. (2017), Luetticke (2021), Bayer et al. (2019), Auclert et al. (2024), Hagedorn et al. (2019)), and tractable-HANK models leveraging insights from the former class to capture key channels of the latter (Bilbiie (2020, 2024), Broer et al. (2020), Debortoli and Galí (2024a), Acharya and Dogra (2020), Ravn and Sterk (2020), Bilbiie et al. (2022b), Debortoli and Galí (2024b)). Our contribution to this literature is to provide model-consistent measurement using high-quality Norwegian data.

In the realm of pure measurement, we also contribute to the vast empirical literature on MPC estimation, by using 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 works of Souleles et al. (2006), Shapiro and Slemrod (1995, 2009a); Sahm et al. (2010), Gruber (1997); Jacobson et al. (1993), Blundell et al. (2008), as well as more recent studies 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, 2024), Orchard et al. (2025), Nielsson 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, 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. (2024), and Ganong et al. (2020). Finally, we follow a strand of recent work that shifts the focus from atypical, transitory income shocks (e.g., Fagereng et al. (2021); Golosov et al. (2024)) to more typical sources of income variation, such as unemployment (e.g.,Ganong and Noel (2019) Ganong et al. (2020) Ilut and Valchev (2021)), allowing us to study the full population. We also test whether our results are robust against estimating MPCs based on other, and potentially more transitory, income shocks by using the Blundell et al. (2008)method (extended by Commault (2022)).

We also contribute to the empirical literature on inequality over the cycle, and to the related notion of income risk. Leading 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), as well as Bergholt et al. (2024) who present a detailed study of income and consumption inequality over time in Norway. Estimated HANK models employing cross-sectional time series to elicit the role of heterogeneity channels include Bayer et al. (2024), Auclert et al. (2020), and Bilbiie et al. (2022a). Measurement exercises using micro data include the Patterson paper reviewed above and Pekanov (2024), who using EU data obtains evidence pointing to amplification through the earnings betas. Related studies 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 learn 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) and Townsend (1994). What our paper adds to these empirical strands is the use of transactions-level consumption data and matched administrative gross net income, as well as wealth data.

Finally, our results are also consistent with an important recent study by Orchard et al. (2025). Using a different approach complementary to ours—historic macro counterfactuals as diagnostic tools—and focusing on the Great Recession in the US, they argue that high micro MPCs can be made consistent with low macro amplification (low multipliers) through general-equilibrium dampening forces. Our exercise, in particular when focusing on the Great Recession in Norway, sheds complementary light using granular data from Norway and unveils via sufficient statistics different micro channels through which such general-equilibrium dampening may occur.

2 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 models 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 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, but in a model with non-trivial heterogeneity (no full insurance), the mode-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 employ often. In theoretical settings we use aggregated to denote model consistent aggregates rather than averages, similar to the usage of the correct price index in models where aggregate consumption consists of individual varieties. In empirical settings we use it to denote aggregates correctly built up from the micro data that we 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 classic "Keynesian Cross" representation of aggregate demand (AD) pioneered by Samuelson (1948). This starts by postulating a generic aggregate consumption function, or "Planned Expenditure":

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

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 apply to any AD shocks and policies.

The effect of an interest rate change dR is 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", planned 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 effect, we need to make extra assumptions and impose market clearing. In the simplest case where there is no saving in equilibrium, we only need to impose that consumption equals total disposable income $C = Y$. Adding this, we can see *general equilibrium* AD amplification, which 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 models adds the following layer of complexity: to find the aggregated, economy-wide MPC of an economy populated by many agents we need to aggregate them carefully. Start from an individual j's "MPC" out of *aggregate* income:

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

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 the "beta". In our terminology, the beta is the elasticity of an individual variable with respect to the aggregate one, in this case the elasticity of individual income to aggregate income:

$$\beta^j_y \equiv \frac{dy^j_t}{dy_t}, \tag{3}$$

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

The **aggregate(d) 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 empirically estimate this important object.

The aggregate MPC is a sufficient statistic for the general equilibrium multiplier. A useful formalization, due to Auclert and Patterson, of how amplification works through this channel in a many-agent economy is as follows. We can rewrite the aggregate MPC as

$$\mathbf{A}MPC = \bar{\beta} \times aMPC + \mathrm{Cov}(MPC^j, \beta^j) \tag{5}$$

where $\bar{\beta} = \sum_j \frac{Y^j}{Y} \beta^j$ and $aMPC \equiv \sum_j MPC^j \frac{Y^j}{Y}$ is the simple "average" MPC, obtained by merely aggregating individual MPCs using the population shares, i.e. assuming that changes in all individual incomes are proportional to the aggregate one, $\beta_y^j = 1$ for all j. The key object is thus

$$\mathrm{Cov}(MPC^j, \beta^j) \equiv \sum_j \frac{Y^j}{Y} (MPC^j - aMPC) \times (\beta^j - \bar{\beta}), \tag{6}$$

the (income share) weighted covariance. If $\bar{\beta} = 1$, which is true when we compute the elasticity of individual variables relative to the *aggregated* variable, the AMPC equals the aMPC plus the (weighted) covariance between MPCs and betas.The sufficient statistic to have an amplification effect due to heterogeneity is therefore:

$$Cov(MPC^j, \beta^j) > 0 \Longleftrightarrow \mathbf{A}MPC > aMPC|_{\beta_y^j = 1}. \tag{7}$$

Amplification occurs when the aggregate MPC just defined is larger than the simple "average" MPC. Naturally, when higher-MPC agents' income is more cyclical, an expansion is amplified as an aggregate income change is distributed more to higher-MPC agents, who shift their demand further, triggering 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:

$$GE \text{ Multiplier} = \frac{1}{1 - \mathbf{A}MPC}. \tag{8}$$

To study 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 simply need to compare the above multiplier with its counterpart in that case, given by $1/\left(1 - aMPC\right)$. The ratio of the two is thus a direct measure of the amplification driven by heterogeneity in the betas.[2]

2.1.1 MPC-Free Multiplier in TANK

It is instructive to note that the sufficient statistic 1 is intimately related to and a generalization of the channel of "cyclical *inequality*" (between hand-to-mouth and savers) emphasized in Bilbiie (2008, 2020)). To illustrate this, consider 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:

$$\mathbf{A}MPC = \tilde{\lambda}_y \times MPC^H \times \beta_y^H + \left(1 - \tilde{\lambda}_y\right) \times MPC^S \times \beta_y^S,$$

where $\tilde{\lambda}_y = \lambda \frac{Y^H}{Y}$ denotes the income share of the entire H group. Since in this model $MPC^H = 1$ and MPC^S is very low (S are permanent-income consumers with an MPC of $1 - e^{-\rho}$, where ρ is the discount rate), we infer that the general equilibrium, multiplier-like amplification is governed by an aggregate MPC of the form $\tilde{\lambda}_y \beta_y^H$. Indeed, it can be easily shown (see the above-cited references) that in this model the general equilibrium multiplier of an interest rate cut is:

$$\frac{dc_t}{d\left(-r_t\right)} = \sigma \frac{1 - \tilde{\lambda}_y}{1 - \tilde{\lambda}_y \beta_y^H}, \tag{9}$$

where small letters denote log deviations from steady state. 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 $1 - \tilde{\lambda}_y$ agents (weighed by income) 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 $\sigma\left(1 - \tilde{\lambda}_y\right)$. However, there are also indirect effects amplifying this initial direct effect through the (new) Keynesian-cross mechanism described above, manifested in the expression in the denominator, with $\tilde{\lambda}_y \beta_y^H$ as the aggregate MPC. There is amplification when these indirect effects governed by $1/\left(1 - \tilde{\lambda}_y \beta_y^H\right)$ dominate the direct effects $1 - \tilde{\lambda}_y$, i.e. whenever

[2] For the case $\bar{\beta}$ =1 and denoting the covariance by Cov for brevity, the ratio of the multipliers is $\frac{1 - aMPC}{1 - \mathbf{A}MPC} = 1 + \frac{Cov}{1 - \mathbf{A}MPC}$. Patterson (2023) instead uses the difference of the multipliers, which is a rescaling of this since $\frac{1}{1 - \mathbf{A}MPC} - \frac{1}{1 - \mathbf{a}MPC} = \frac{Cov}{1 - \mathbf{a}MPC}$.

$\beta^H_y > 1$, which trivially makes the aggregate MPC $\tilde{\lambda}_y \beta^H_y$ larger than the average MPC $\tilde{\lambda}_y$, because the high-MPC agent's income reacts disproportionately to aggregate income. The covariance condition (7) is a many-agent generalization of this simple logic. Indeed, it is straightforward to show that the weighted covariance between MPCs and betas in the analytical TANK case is:

$$\mathrm{Cov}(MPC^j, \beta^j) = \tilde{\lambda}_y e^{-\rho} \left(\beta^H_y - 1\right) = \mathbf{A}MPC - aMPC.$$

This, instead, is proportional (and has the same sign as) the cyclicality of income inequality between savers and hand-to-mouth, defined as:

$$\frac{d\left(y^S_t - y^H_t\right)}{dy_t} = \beta^S_y - \beta^H_y = \frac{1 - \beta^H_y}{1 - \tilde{\lambda}_y}.$$

In other words, a positive covariance ($\mathbf{A}MPC > aMPC$) is equivalent to countercyclical income inequality $\beta^H_y > 1$. We will use this in our measurement exercise below to directly gauge whether there is amplification or not directly, by studying the cyclicality of income inequality, or of the income betas—without needing to resort to the MPC estimates, insofar as we can use wealth data to classify a group of agents as "hand-to-mouth". This also tightens the link to our sufficient statistic 2, based on consumption betas.

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:

$$c^S_t = E_t c^S_{t+1} - \sigma r_t, \tag{10}$$

where c^j_t is the log consumption of type $j = H, S$, σ the elasticity of intertemporal substitution and r_t the real interest rate. Define consumption inequality as $\Gamma_t \equiv \frac{C^S_t}{C^H_t}$, and in log deviations:

$$\gamma_t = c^S_t - c^H_t. \tag{11}$$

Using this definition to rewrite the loglinearized aggregation equation as $c_t = \tilde{\lambda}_c c^H_t + \left(1 - \tilde{\lambda}_c\right) c^S_t$ (defining the *consumption* share of the whole H group as $\tilde{\lambda}_c \equiv \lambda \frac{C^H}{C}$.), we have:

$$c_t = c^S_t - \tilde{\lambda}_c \gamma_t. \tag{12}$$

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:

$$c_t = E_t c_{t+1} - \tilde{\lambda}_c \left(\gamma_t - E_t \gamma_{t+1}\right) - \sigma r_t \tag{13}$$

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 (2024a).

Proposition: There is AD *amplification*, i.e.the effect of an interest rate change on aggregate demand is larger under heterogeneity than in the $\lambda = 0$ representative agent economy, iff consumption inequality is *countercyclical, i.e.*

$$\gamma_c \equiv \frac{d\gamma_t}{dc_t} < 0 \iff \beta_c^H > \beta_c^S, \tag{14}$$

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

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:[3]

$$\frac{dc_t}{d\left(-r_t\right)} = \sigma \frac{1}{1 + \tilde{\lambda}_c \gamma_c} \tag{15}$$

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 $\left(1 + \tilde{\lambda}_c \gamma_c\right)^{-1}$ and assessing whether it is larger than 1.

It is important to emphasize that the key statistic 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.[4] Yet it is useful to notice that in the simplest case

[3]This follows by direct inspection of (13) and holds more generally under time-invariant cyclicality of consumption inequality (see Bilbiie (2020) for a derivation under persistent shocks).

[4]In particular, it depends on the labor market structure, tax-and-transfer system progressivity, what types of financial income accrue to which agents and how that maps into the wealth distribution. Our next section expands on this last element. In a recent contribution, Broer et al. (2024) show that a similar "multiplier" logic

with 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, since now $\tilde{\lambda}_c = \tilde{\lambda}_y$ and $\gamma_c = \frac{1-\beta_y^H}{1-\tilde{\lambda}_y}$, replacing this in the multiplier (15) shows directly that it is identical to the multiplier derived based on sufficient-statistic 1, found in (9). 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 time varying savings

The brief derivation in the previous subsection assumes no net assets as well as 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 necessarily 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:

$$C_Y c_t^S + \frac{X_Y}{1-\lambda} x_t = Y_Y^S y_t^S, \tag{16}$$

where $Z_Y \equiv Z/Y$ is the share of variable Z in total income, $Z \in \{C, X, Y^S\}$, and we assume without loss of generality equal steady-state consumption shares across agents, $C^H = C^S = C$. For the HtM agents we still have:

$$c_t^H = y_t^H. \tag{17}$$

Through some simple algebra outlined in 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:

$$\gamma_t = c_t^S - c_t^H = \frac{1-\beta_y^H C_Y}{(1-\lambda)C_Y} y_t - \frac{X_Y}{(1-\lambda)C_Y} x_t. \tag{18}$$

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.[5]

extends to a wide menu of fiscal instruments in a richer model encompassing search and matching. The point is that such details notwithstanding, the cyclicality of inequality is a sufficient statistic for gauging amplification.

[5] 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 (2024)) or investment in physical capital

3 Data and Descriptive Statistics

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 records provide near-complete population coverage, along with detailed demographic information and data on various income and wealth components. Since the balance sheet data are derived from tax records, the potential for measurement issues is minimal compared to survey data.[6]

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 zero observations (no labor earnings). 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,[7] (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)), we focus on the prime working age population and restrict our sample to individuals aged between 25 and 55 years.

(Auclert et al. (2020), Bilbiie et al. (2022b))

[6] Third party reported means that information to the tax authorities in large part is provided by the individual's bank and employer. Fagereng et al. (2020); Holm et al. (2021); Ring (2024) are some recent applications of these data on household saving and consumption behavior.

[7] 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.

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 wire transfers where the recipient is a bank.[8] 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 and 2018, and we see no reason to why the excluded 20 percent of payments are systematically different to the included payments along any of the dimensions relevant to this paper.

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).[9] Additionally, our dataset includes 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 consumption reported in the Norwegian national accounts. In Figure 1 we compare the nominal household consumption in the national accounts with aggregated, total electronic spending; the left panel compares levels, and the right panel quarterly growth rates. The national accounts series is constructed as domestic household consumption, excluding imputed owner-occupied housing. Appendix B presents details on the cleaning of the electronic transactions data.[10] In levels, our aggregated series is representative of the Norwegian economy as a whole, as illustrated by the left panel; our aggregate consumption series is on average 90% of the national accounts series. The correlation of the quarterly growth rates of the two time series is 0.83 over the whole sample period.

We merge the electronic transactions data with the administrative records described above.

[8] 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.

[9] In contrast, across the Euro area the average was 80 percent (Esselink and Hernández, 2017). By value, the corresponding averages were 3 and 53 percent in Norway and the Euro area, respectively.

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

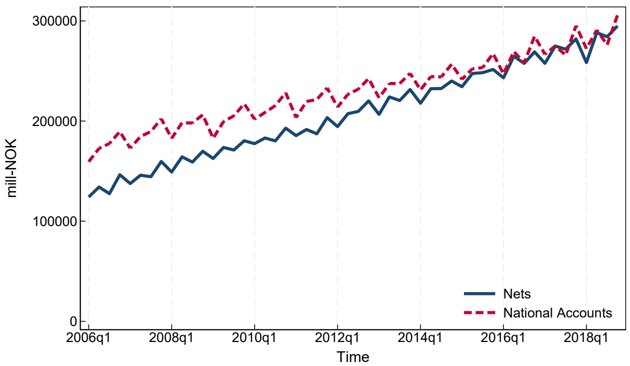

(a) Levels

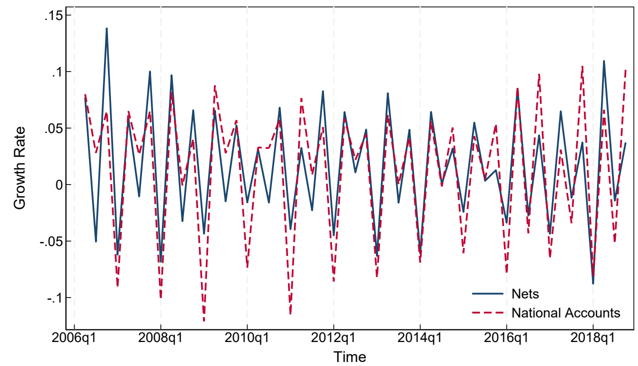

(b) Quarterly Growth Rates

Notes: Panel (a) displays quarterly levels of nominal household consumption from the national accounts, excluding imputed owner-occupied housing (dashed line), and aggregate household electronic expenditures (solid line). Panel (b) displays the corresponding quarterly growth rates. For the electronic expenditure measure we remove single wire transfers above 100,000 NOK (12,500 USD2015), person-to-person transfers, and all transfers to banks, except imputed payments of credit card bills

Figure 1: Consumption in Norway: Aggregated Transactions Data vs National Accounts

This gives us a dataset 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 discussed in detail in Appendix B.

3.3 Institutional details of the 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 rising average tax burdens at higher incomes.[11] By international standards, taxes are high, consistently above the OECD average. Since the 1970s, tax revenue has remained at 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 benefits. Figure 3.3 shows 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 earnings decile received, on average, half of their income from transfers. For most people, labor income remains the primary source of earnings. At higher income levels, transfers largely consist of pensions. As part of the welfare state, education and healthcare are almost completely free.

Despite Norway's extensive welfare system, unemployment leads to substantial income losses.

[11] A flat 22% tax applies to labor and capital income, with additional bracketed marginal taxes reaching 17.7%.

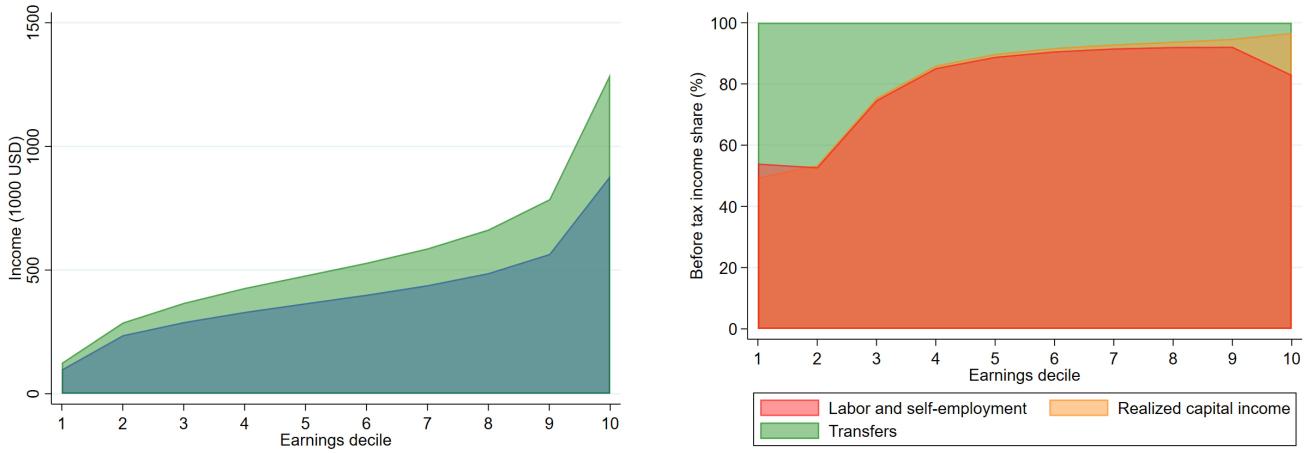

(a) Before and After Tax Income

(b) Components of Before Tax Income

Notes: Panel (a) shows household income before and after taxes across earnings deciles. Panel (b) presents the corresponding shares of before-tax income derived from labor and self-employment, realized capital income, and transfers

Figure 2: Norwegian Income, Taxation and Transfers: By Income Deciles

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. (2024) estimations). The reason for this is that UI is less generous than in many other Western European countries. The 62.4% replacement rate is the lowest in the Nordics, comparable to Germany, but Germany offers a higher benefit ceiling and, unlike in Germany, Norwegian benefits are taxed.[12] Eligibility is stricter than in many other Western European countries, requiring one year of prior employment and income above a threshold. However, Norway offers a longer benefit duration than most European countries, and UI payments are easy to calculate with an online tool. High earners face the largest income losses upon unemployment due a cap on UI benefits.[13] Other factors also shape disposable income changes during unemployment: wealth and debt levels, as wealth taxes and debt payments persist, and children, as child benefits help smooth income drops.

3.4 The Great Recession in Norway

As part of our empirical analysis in Section 6.3, we focus on the most severe contraction within our sample period: the Great Recession. This episode, marked by large adjustments in aggregate consumption (see Figure 3), offers a natural laboratory for studying consumption dynamics at the micro level. The downturn in aggregate consumption expenditure was both sharp and long lasting: it took around eight quarters for real consumption growth to return to its pre-recession peak of Q3 2007, while real GDP growth did not return to its pre-recession level within our sample

[12] 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.

[13] The UI cap is adjusted annually, in 2024 it was around 140% of the national average after tax income .

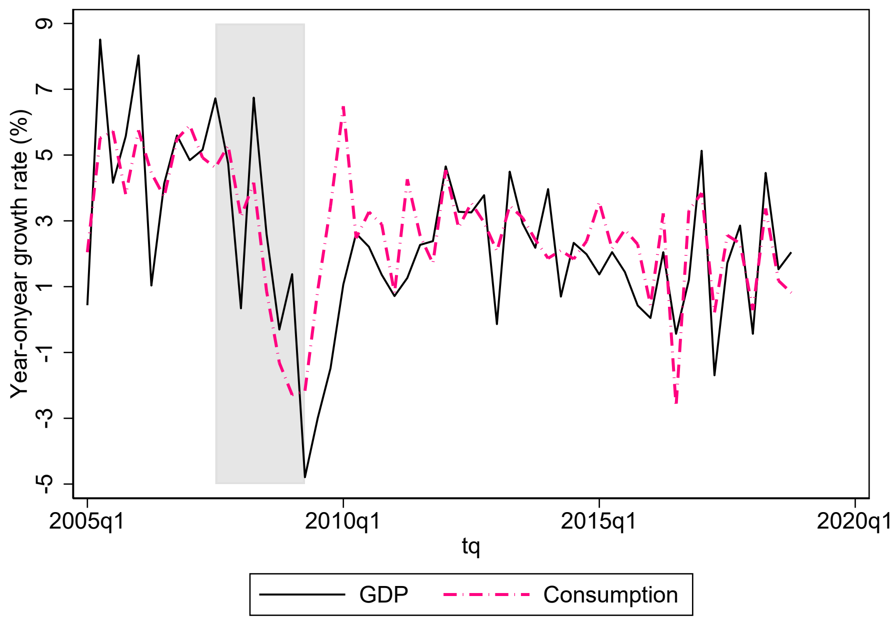

Notes: The figure shows the annual growth rate of real household consumption and GDP by quarter from the national accounts

Figure 3: Norwegian GDP and Real Consumption Expenditure of Households

window. Average annual GDP growth declined from 2.2 percent prior to the Great Recession to 1.8 percent in the years that followed.

3.5 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 similar to the ones documented for the U.S. by the aforementioned studies.

3.5.1 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 4 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 in both income

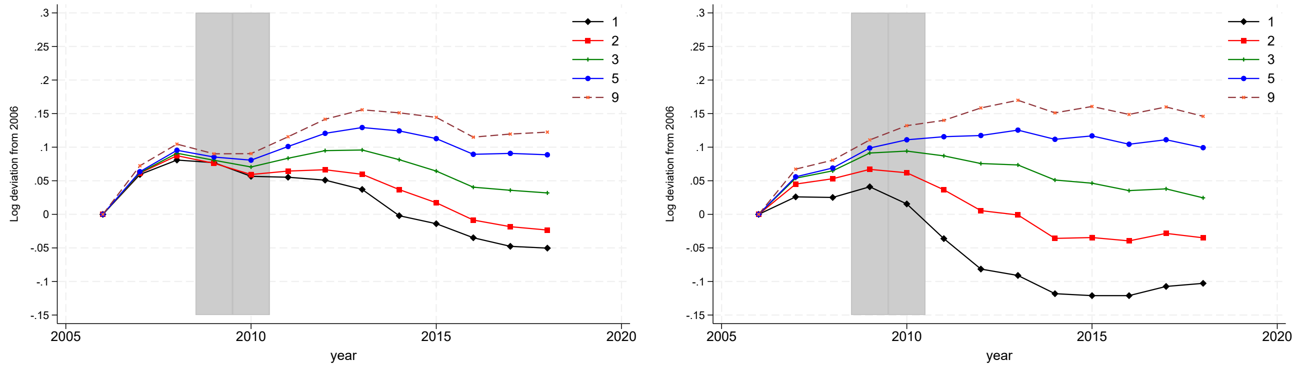

(a) Total Income After Tax

(b) Consumption

Notes: The figures show the evolution of income and consumption inequality. Each year, we sort households in 10 deciles based on their income after tax, and compute the median income after tax and consumption within each decile. The figures report the cumulative growth since 2006, by selected deciles. The shaded area corresponds to the years 2009-10.

Figure 4: Household Income and Consumption: by Income Level.

and consumption has been increasing over time, as shown by the distributions fanning out. More importantly for our purpose, the recession does seem to affect the bottom deciles relatively more. See Bergholt et al. (2024) for an in-depth study of income and consumption inequality dynamics in Norway that is complementary to ours.[14]

3.5.2 Norwegian "worker betas"

We now compute so-called worker betas as in Guvenen et al. (2017) for various population groups. These betas measure the elasticity of an 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 workers labor earnings betas for the same partitioning of the population as in their study.

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

$$\Delta y_{i,t} = \alpha_g + \beta_g \Delta Y_t + \varepsilon_{i,t} \tag{19}$$

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 5, 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.

In the rest of the paper, we will measure inequality in both disposable income and consumption, how inequality in these variables manifests itself in the MPC dimension, and how MPCs and betas are related in the data. We will do this analysis guided by the theory discussed in section 2.

[14] The authors find that inequality, defined as cross-sectional dispersion in the respective variable's distribution, has been procyclical in Norway for both consumption and income.

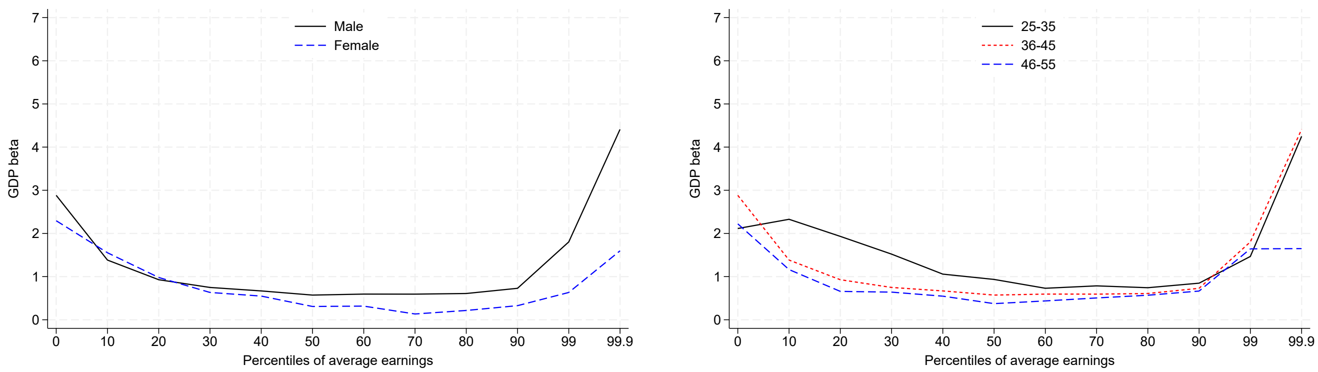

(a) Gender

(b) Age

Notes: Panel (a): GDP-beta by gender at age 36-45. Panel (b) Male GDP-beta by age-groups.

Figure 5: Worker Betas (Labor Earnings)

4 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 types of data from a variety of countries. Our approach is different from existing studies in that we use transactions-level consumption data with broad population coverage. Still,our estimates of both the average MPC and the dispersion are similar to estimates obtained using different estimation techniques on other Norwegian datasets as well as estimates from other countries using similar (or different) estimation techniques.

4.1 Estimating the MPC distribution of Norway

Our unit of analysis it the individual and the goal is to estimate MPCs at the individual level. To do so, 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 proceed by estimating the following relationship between consumption and income changes:

$$\Delta C_{i,t} = \sum_{x} (\mu_x \Delta Y_{i,t} \times x_{i,t-1} + \alpha_x x_{i,t-1}) + \delta_{t,s} + \varepsilon_{i,t}. \tag{20}$$

Here, $C_{i,t}$ is real total consumption of individual i at time t (measured as individual i's total household consumption), $Y_{i,t}$ the real after tax total income of individual i, and $x_{i,t}$ are individual characteristics.[15] The empirical specification in (20) allows these individual characteristics to affect con-

[15]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

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

Since both consumption and income are endogenous and jointly influenced by many factors, a naive OLS of (20) 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[16]

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

is instrumented with:

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

The advantage of this approach is that it captures a typical business-cycle shock experienced by workers which is large and therefore generates identifying variation in income. The disadvantage is that it may not always be completely exogenous, and it is often a persistent (rather than purely transitory) income shock. While commonly used and informative (see for example Ganong et al. (2020); Fagereng et al. (2024)), this method requires caution, especially when comparing groups with differing levels of unemployment persistence. To address this concern, Section 5.5 introduces a robustness check using using the estimator proposed by Commault (2022).This alternative method estimates MPCs out of transitory income shocks. .

We estimate an average unconditional MPC of 0.38, with substantial cross-sectional variation: 80% of MPCs fall between 0.21 and 0.52 (see 6). This degree of heterogeneity is in line with the literature, despite differences in data sources and identification strategies. Using a similar methodology in U.S. administrative data Patterson (2023) also documents a wide variation in household MPCs.Lewis et al. (2024), applying clustering methods that allow for latent factor-driven heterogeneity, reports similar dispersion in response to the 2008 U.S. stimulus. Survey-based evidence from Jappelli and Pistaferri (2014, 2020) likewise find substantial variation in Italian data. Experimental evidence from Boehm et al. (2025) using randomized cash transfers through a French bank also supports the presence of sizable MPC heterogeneity. While few papers estimate full MPC distributions, the consistent finding across methods and settings is that average MPCs mask a wide

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 having gross liquid wealth below two weeks net income, while low net wealth is defined as below two months net income. Since we only observe income annually, we proxy weekly income as annual income divided by 52. These wealth dummies are related to often used proxies for hand-to-mouth status used in the literature.

[16]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.

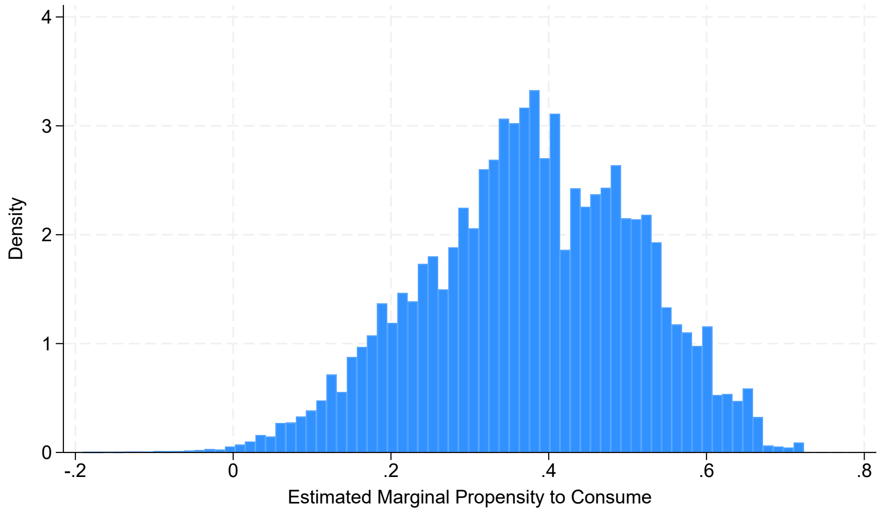

Notes: This figure reports the full distribution of individual MPC estimates $\widehat{MPC}_{i,t} = \sum \hat{\mu}_x x_{i,t-1}$ obtained from the IV regression of (20).

Figure 6: The MPC Distribution of Norway

range of household responses.

Our average MPC is in line with previous estimates.. Patterson (2023) reports that unemployed U.S. households reduce consumption by 33-55 cents per dollar lost from labor income. Ganong et al. (2020), using transaction-level data from the JPMorgan Chase Institute, also finds high MPCs between 0.27 and 0.42, depending on the identification strategy. One approach exploits delays in unemployment insurance (UI) payments due to administrative backlogs, comparing individuals who received benefits at different times; another compares unemployed individuals to matched employed counterparts. Both strategies consistently yield high MPC estimates. In the Norwegian context, using imputed total consumption, Fagereng et al. (2024) find an MPC of 0.41 from unemployment shocks comparing unemployed to matched employed workers.I.[17] These results support the view that high MPCs may arise from liquidity constraints, and that job loss is a major disruption to household finances.

4.2 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 7 we report conditional MPC estimates, from bi-variate regressions based on the set of characteristics captured by the covariates in equation (20): gender, partnership status, and age in the left panel; lagged income, 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

[17] Fagereng et al. (2021) reports an even higher 50% MPC from Norwegian lottery winnings, likely due to the nature of the shock or the characteristics of lottery players.

4.1. In Table 6 in appendix 4.1, we report the full results of the multivariate regression containing all determinants jointly.

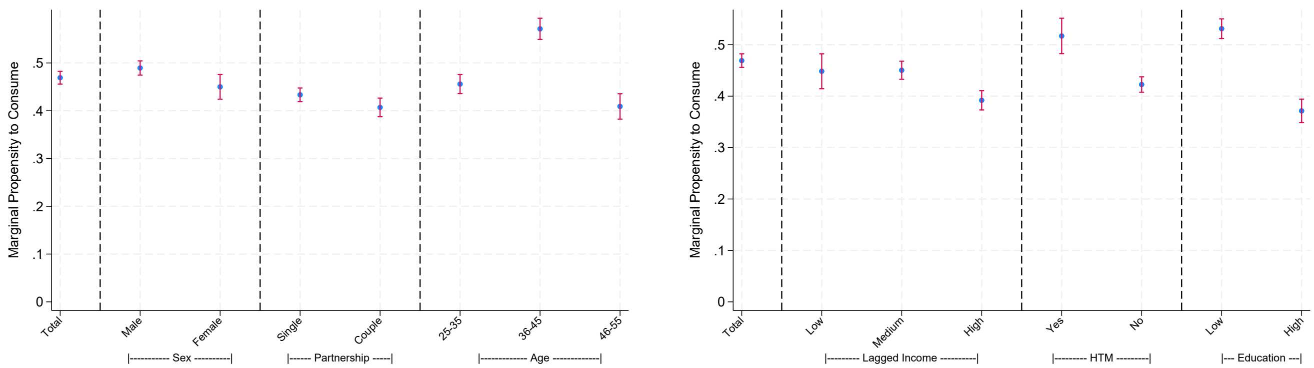

a. Sex, partnership status and age

b. Lagged income, HtM status and education

Notes: The figures report point estimates and 95% CI for various bivariate MPC regressions. The dots represent point estimates obtained from IV estimation of equation (20), excluding the vector **x**, on the sub-sample of individuals in the subgroup indicated below the horizontal axis. Total refers to the whole sample; sex to males or females, partnership to singles or individuals with a partner; age to three age groups; lagged income to the 2nd-3rd-4th quintile of the lagged income distribution; HtM to individuals with gross liquid wealth below two weeks net income; education to those with or without a college degree.

Figure 7: MPC Determinants

In doing analysis of this nature, while most categories are straightforward (gender, cohabitation status etc. are directly observable in the registry data), hand-to-mouth status is not. In many HANK-type models, being HtM and 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.

Since our dataset allows us to separate total wealth into liquid and illiquid components, we can investigate whether either total wealth or liquid wealth have any bearing on estimated MPCs. Following Aguiar et al. (2024) and Kaplan et al. (2014), respectively, we classify agents as HtM based on total wealth if their total wealth is less than two months' after-tax income and as HtM based on liquid wealth if their liquid assets are less than half of monthly after-tax income.[18]

We have two important findings to report. The first is that while income by itself is not an indicator of MPC level, education is. Specifically, individuals with a university education tend to have a lower MPC. Our findings correspond with previous estimates that suggest an inconclusive or nonlinear relationship between income and MPC but a strong correlation between education and MPC (see for example Shapiro and Slemrod (2009b), Misra and Surico (2014), Lewis et al. (2024), Jappelli and Pistaferri (2014), Commault (2024)).[19] Second, we find that HtM status is also an important predictor of MPC. However, this predictive power *only* holds for liquid wealth, not for total wealth. Figure 7 shows these findings, with the liquidity-based HtM measure shown in

[18] All variables are measured yearly, so monthly income is measured as one twelfth of yearly income.

[19] 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 is available upon request. As for age, we find a non-linear effect on the MPC; interpreting this through the lens of a life-cycle model is challenging, as other factors beyond the standard expected lifespan–such as cohort effects, bequest motives or health shocks–may also influence the relationship.

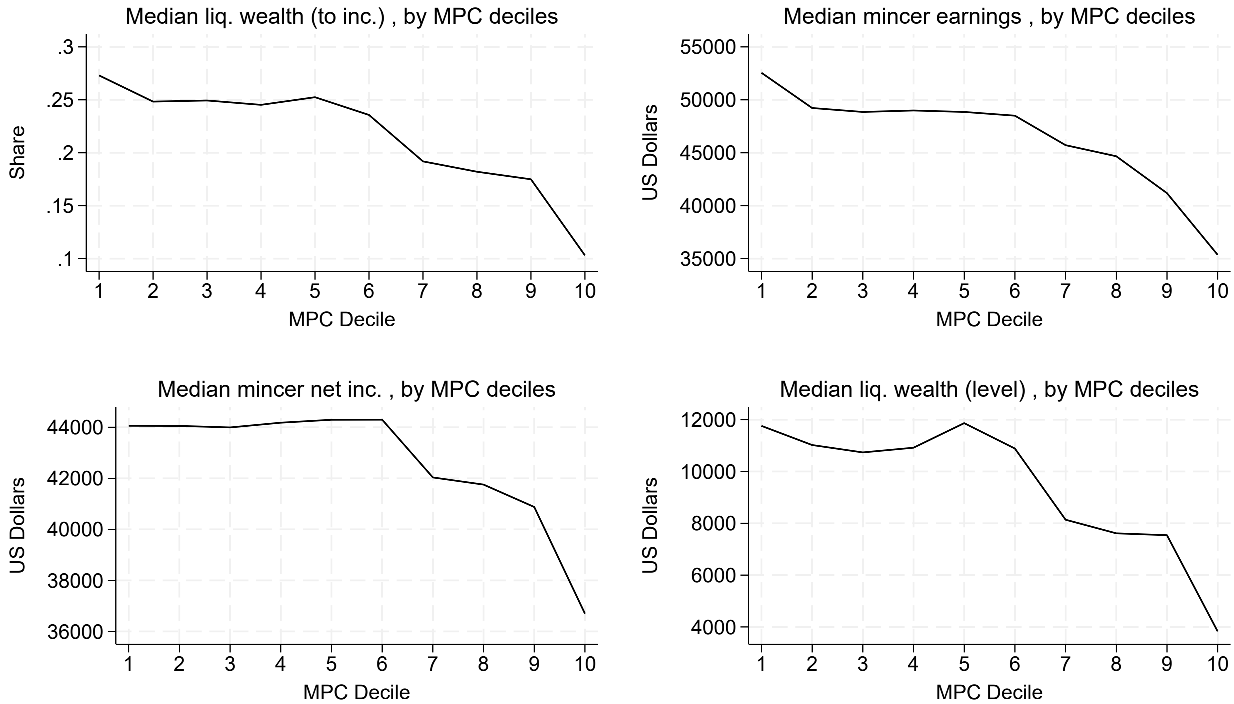

Notes: Measures of liquidity and permanent income across the MPC distribution. The upper-left panel shows the median liquid-wealth-to-income ratio within each MPC decile, while the bottom-right panel displays the median level of liquid wealth. The upper-right and bottom-left panels repeat the exercise for two alternative measures of permanent income derived from Mincerian regressions.

Figure 8: Liquid Wealth and Permanent Income in the MPC distribution

the right panel. These findings reinforce the importance of liquid assets in economic modeling, as proposed by Kaplan et al. (2014).[20]

Another way to examine determinants of MPCs is by analyzing how certain individual characteristics vary along the MPC distribution. In figure 8 we explore this by considering liquid wealth and life-cycle income. The latter is estimated 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.[21] 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 characterized by a markedly lower level of liquid assets and low Mincerian income. This provides further support for the often-used assumption in HANK models that high-MPC, hand-to-mouth

[20]Total wealth, on the other hand, is not informative about MPC status. We do not illustrate this in the figure for brevity but the results are available upon request.

[21]We run the regression $ln(y_{it}) = \beta_1 D_{it}^{age} + \beta_2 D_{it}^{edu} + \beta_3 D_i^{male} + \beta_4 D_t^{year} + \beta_5 D_{it}^{partner} + \beta_6 D_{it}^{age} D_i^{male} + \beta_7 D_{it}^{age} D_{it}^{edu} + \beta_8 D_i^{country}$, where D^{age}, D^{year}, D^{edu} and $D_i^{country}$ are vectors of age, year, education (4 levels) and country-of-birth dummies, while $D_{it}^{partner}$ and D_i^{male} are dummies for having a partner and for being male. The left-hand-side variable is either log of labor earnings or net income. The Mincerian income is the individual's predicted income from this regression.

agents are the liquid-asset poor. We leverage this below in our computation of sufficient statistics based directly on consumption.

5 Sufficient Statistic 1: The Aggregate MPC

We are now in a position to compute the first sufficient statistic from section 2, the aggregate MPC, by linking an individual's MPC to a measure of their cyclical income sensitivity. 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.

5.1 Income Betas Across the MPC Distribution

To move towards computing our object of interest, we estimate the betas as in section 3.5.2, but we make two changes. First, we use *disposable income* instead of earnings. And second, we group individuals 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 9 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).[22] 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

[22] In the figures that plot betas across the MPC distribution, we use a slightly different regression specification than in section 3.5.2. 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 Appendix C.3, where we also show that this specification also ensures that the group-level betas will approximately sum to one when weighted by the group income shares. The income betas with respect to GDP do not sum to one across the MPC deciles, however. The reason is that while the growth rate of aggregated income is the weighted sum of the growth rate of income within each decile in a given year, this is not the case for aggregate GDP growth. Nevertheless, what matters for our purposes is that we find a similar slope of the betas along the MPC distribution whether we use GDP or aggregated income on the right-hand-side of the regression.

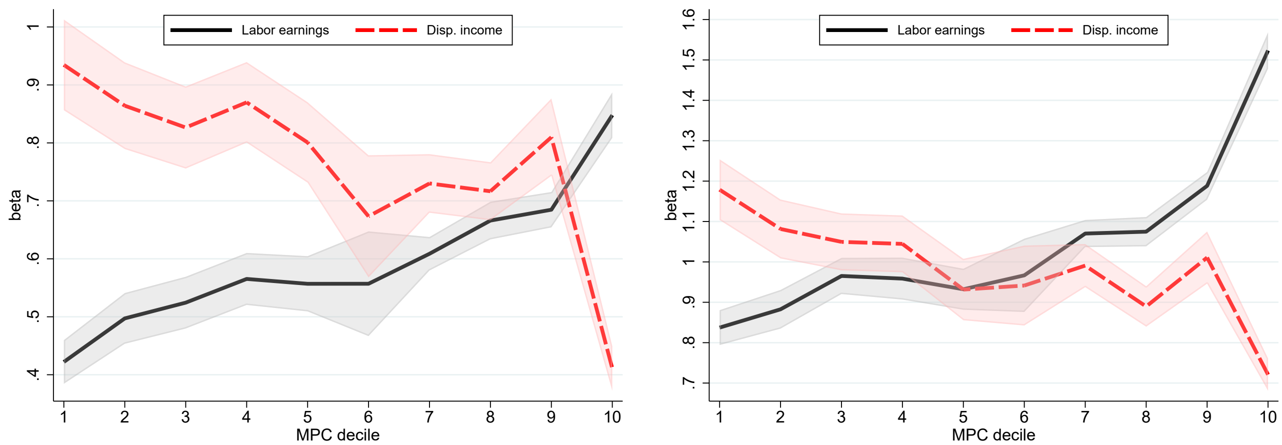

(a) GDP betas (b) Aggregated income betas

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 9: Betas across the MPC distribution: earnings vs. disposable income.

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.

5.2 A Decomposition: The Role of Capital Income, Taxes and Transfers

The income beta and MPC relationship changes sign and becomes negatively sloped as we move from using labor earnings to using disposable income as the income measure. We decompose disposable income into its constituent components to understand the reason. Figure 10 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.[23]

Adding capital income to labor earnings mildly flattens the beta curve, in particular for individuals with very high and very low MPCs. To understand why, in figure 11 we show both the decomposition of disposable income by MPC decile and the volatility of each (aggregate) income component over time. Since capital income is highly cyclical and low MPC individuals–who on

[23]Figure 23 in the appendix displays beta estimates for each income measure with confidence intervals.

average have higher income and net worth than high MPC ones–on 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 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 than low MPC individuals, as shown in figure 11. Notably, in the Great Recession of 2008-2009, the growth rate of both labor earnings and capital income dropped, while the growth rate of transfers increased. This counter-cyclicality of aggregate transfers contributes to the slightly negative beta for income inclusive of transfers for high-MPC individuals. Second, as detailed in section 5.3, 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 ones. For instance, while a low-income person is more likely to lose her job in a recession, contributing to the positive slope of the beta curve for labor earnings, unemployment benefits make up for a larger share of pre-employment income for the low income person than for a high-income one. 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.

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 10: Betas Decomposition: From Earnings to Net Income

The positive slope of the beta curve across the MPC deciles, known in the literature for other

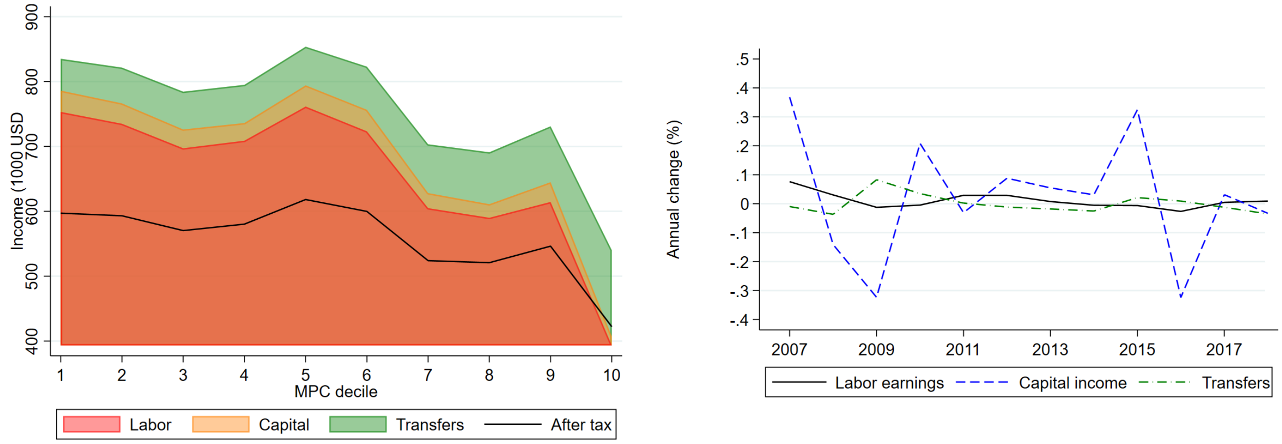

(a) Components of income

(b) Aggregate income growth

Notes: Panel (a) shows a decomposition of mean disposable income into its components (labor earnings, capital income, transfers and taxes), by decile of the MPC distribution. Panel (b) shows the annual growth rate, by year, of the aggregated income components for our sample.

Figure 11: Income components by MPC decile and volatility of income components.

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.

5.3 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 high-MPC individuals' incomes. 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 US to our Norwegian micro data, then re-estimate the beta curve for disposable income imputed from the US system.[24] We use the specification estimated by Ferriere et al. (2023) on Current Population Survey data, consisting of a flat capital income tax, a labor tax rate that varies with labor income, and a transfer that varies with total 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

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

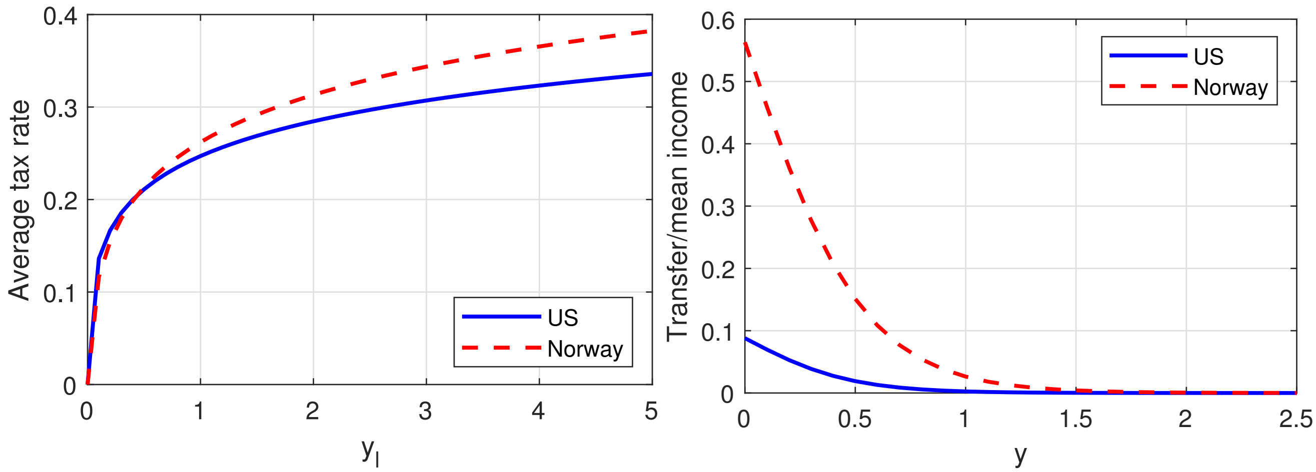

(a) Labor tax function

(b) Transfer function

Notes: The figures illustrate the estimated functions for average labor tax and transfer relative to mean population 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 39 and 40, and the parameter values are found in table 7.

Figure 12: Estimated tax and transfer functions for US and Norway.

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.[25]

Figure 12 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 ones. Of more consequence, low income individuals in Norway receive substantially more government transfers, both in levels and as a share of their income.[26] Figure 13 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 people 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 they receive. 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 14. 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

[25]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.

[26]The exact parameter values, as well as measures of fit, can be found in table 7.

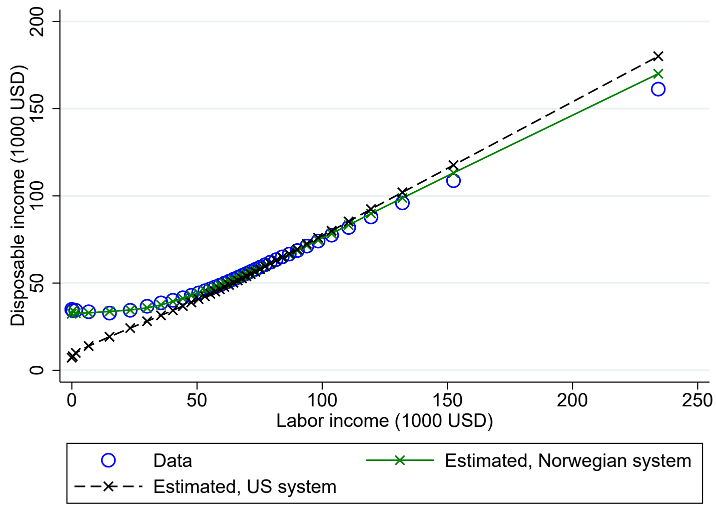

Notes: Disposable income by 2.5% quantiles of labor income, in data and based on estimated tax and transfer system for Norway and the United States. NOK values are converted to USD using the average exchange rate over the period 2006-2018.

Figure 13: Imputed disposable income by labor income, based on estimated tax and transfer system for Norway and the US.

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.[27]

5.4 Measuring Amplification: The Aggregate MPC and 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 and quantify the degree of amplification by estimating the aggregate MPC in equation 4. To that end, we employ three separate estimation methodologies that give very similar results. We start by calculating the aggregate MPC based on the 10 MPC decile groups used in section 5.1, employing equation 4 with estimates of MPCs and income betas at the group level.[28]

[27]The indirect implication that heterogeneity channels play a limited role in amplifying business cycles in the US too, despite large micro MPCs, corroborates and complements Orchard et al. (2025), who argue this compellingly based on analyzing fiscal stimulus in the US Great Recession, and emphasize different channels delivering this general-equilibrium dampening.

[28]Specifically, our AMPC estimate is given by $\frac{1}{T}\sum_t \sum_{d=1}^{10} \overline{\widehat{MPC}_{d,t}} \times \frac{Y_{d,t}}{Y_t}\hat{\beta}_d$, where d denotes the income beta decile, T is the number of years in the sample, $Y_{d,t}$ is the total income of individuals in income d at time t, and $\overline{\widehat{MPC}_{d,t}}$ is the average MPC of individuals in decile d at time t, weighted by their respective income shares within the decile. We calculate the covariance between the MPCs and the income betas as the difference between

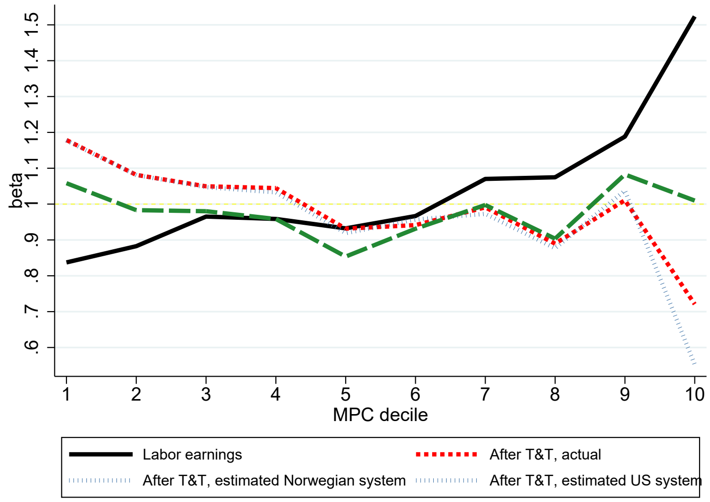

Notes: The figure shows the point estimates for income betas across 10 MPC deciles based on labor earnings, actual disposable income (after taxes and transfers), and disposable income imputed from the estimated taxes and transfer systems for Norway and the United States, respectively.

Figure 14: Income betas with US counterfactual tax and transfer system.

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. We also report the GE multiplier based on the AMPC (labeled Mult. A), and the multiplier based on the uniform-betas, average MPC (labeled Mult. a). Finally, the last column reports the ratio of the two multipliers, which is a measure of the amplification due to heterogeneity in betas.

Consistent with the group-level results reported in section 5.1, we find that the aggregate MPC is higher than the average MPC for labor income, reflecting a positive covariance between earnings betas and MPCs—and generating an amplification of 2.57 percent relative to the uniform-betas case.[29] 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 in a representative agent model where the agent has–for whatever reason–an MPC of 0.367; in other words, heterogeneity generates *dampening* by 1.84 percent.[30]

In appendix section C.5, we show that our estimates of amplification are virtually unchanged

the AMPC and the average MPC, consistent with equation 6.

[29] Part of the reason why our covariance is lower than that obtained by Patterson for the US is that the variance of the MPC distribution is lower: indeed, our standard deviation is 2.34 times lower. If we increase our MPC standard deviation ceteris paribus (keeping fixed the betas' standard deviation and the correlation coefficient), we get a covariance of 0.037, closer to the numbers obtained by Patterson (which range from 0.06 to 0.099).

[30] 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.

Table 1: Aggregate MPC of Norway:

Income measure	AMPC	avg. MPC	$Cov(MPC, \beta)$	Mult. A	Mult. a	Mult. A/a
Labor income	0.374	0.358	0.016	1.597	1.557	1.025
+ capital	0.365	0.358	0.071	1.574	1.556	1.011
+ transfers	0.356	0.363	-0.007	1.553	1.571	0.988
- taxes	0.355	0.367	-0.012	1.550	1.579	0.981
US counterfact.	0.358	0.360	-0.001	1.558	1.562	0.997

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, the general equilibrium multiplier based on the AMPC and equation 8 (Mult. A), the multiplier based on the aMPC (Mult. a), and finally the ratio of the two multipliers.

under two alternative ways of estimating the aggregate MPC and covariance. First, we show that the results reported in table 1 are not sensitive to the grouping of individuals into 10 MPC bins, as we calculate the AMPC by summing over all individuals in the sample, weighting the individual-level MPC estimates by individual-level beta estimates. Second, we employ the methodology used by Patterson (2023), regressing individual-level income growth on their MPC estimate and the interaction between MPC and aggregate income growth. Both of these methodologies result in amplification numbers that are very similar to those reported in table 1.

Lastly, we also gauge amplification directly by estimating betas when individuals are grouped according their hand-to-mouth status (according to several HtM categorizations), and applying the multiplier formula in equation (9). This has the advantage of not using the MPC estimates and thus circumventing any potential bias pertaining to their estimation. We describe the results in the next section, after discussing the respective HtM categorizations. The bottom line and common conclusion is that the results are broadly in line with the conclusion obtained from the AMPC estimation using the more disaggregated data: estimating the cyclicality of inequality (and betas) in earnings points to amplification, whereas estimates of the cyclicality of inequality in disposable income overturns that result and points to mild dampening. These numbers also provide a more direct benchmark of comparison for our sufficient statistic 2, which uses consumption betas for the same sample splits.

Our results have important policy implications. Policy analysis employing HANK models will lead to erroneous 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. At the same time, we also show that amplification numbers due to heterogeneity are quantitatively small under any measure of income.[31] A pos-

[31]Our findings thus provide support to the original Campbell and Mankiw (1989) assumption of different (by MPC) agents' incomes being proportional to aggregate income. Connecting back to the theory, this instead

itive interpretation of this result is that automatic stabilizers are doing their job as they should, smoothing out macroeconomic fluctuations and mitigating their incidence in the cross-section.

5.5 Dealing with shock persistence

Since our main result concerns the ordering of betas along the MPC distribution, it is important that our results are robust to the methodology used to estimate MPCs. To address robustness to both the estimation procedure and the persistence of income shocks, we group households into income beta quintiles and estimate an MPC for each quintile. For this we use the semi-structural approach developed by Commault (2022) (BPP-C), extending the Blundell et al. (2008) estimator by allowing temporary income shocks to exhibit short-term persistence through an MA component. This estimator isolates the consumption response to transitory (but potentially persistent) income fluctuations. Following BPP-C, we restrict the sample to stable married couples. Appendix C.7 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 15).[32] Second, the BPP-C point estimates are lower than the average MPC from our baseline method, but not significantly so. 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. Using US data, Commault (2022) estimates an MPC of 0.32 , remarkably close to our estimate on Norwegian data.

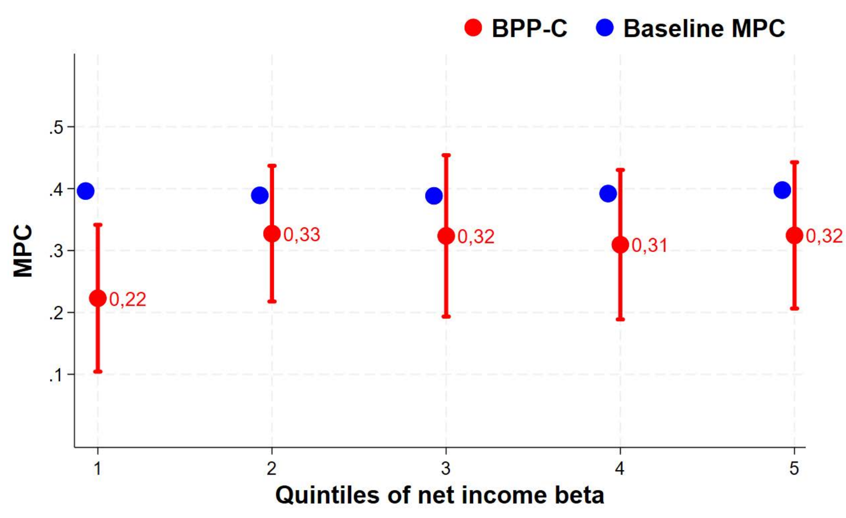

Notes: The figure shows (in red) the point estimates and standard errors for the MPCs estimated using the BPP-C method, alongside our baseline estimates (blue dots) by quintiles of disposable income betas

Figure 15: MPCs estimated by the BPP-C method as a function of net income betas.

implies no aggregate effects of heterogeneity (as long as consumption equals income), as shown by Bilbiie (2008, 2020) and Werning (2015) in TANK and HANK models, respectively.

[32] BPP-C estimates for smaller beta quintiles are noisier but also flat.

6 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.

As in the previous section, the analysis is conducted at the individual level.[33] In 6.1, we illustrate how the consumption betas vary across the distribution of MPCs and liquid assets. In Section 6.2, we measure the extent of amplification by considering various ways of splitting the population into two groups, where we assume that the consumption of one of the groups follows an Euler equation, while the other group consume their income every period (HtM).

6.1 Consumption betas: The extent of consumption risk sharing

In Figure 16, we report the betas computed directly using our consumption data, as a function of the MPC decile (panel a) and deciles of liquid assets normalized by income (panel b). As before, we compute the betas both with respect to a national income accounts measure of aggregate consumption and with respect to the aggregated consumption measure itself, following the methodology outlined in appendix C.3. The picture is the same regardless of the measure used. As the figure shows, the consumption betas are close to flat over both the MPC and liquid assets distributions, indicating a high degree of consumption risk sharing, or insurance, in a manner similar to Townsend (1994), Mace (1991) and Cochrane (1991). The second panel suggests that, while the distribution is still relatively flat, lower-liquidity individuals have slighty higher consumption betas. We next quantify precisely how much amplification this generates through the lens of the model and our second sufficient statistic.

6.2 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 categorizations, 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 our estimates 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 consumption cyclicality, the magnitude is negligible. In other words, regardless of the classification criterion, the consumption betas of HtM and non-HtM individuals are not that different. This is fully consistent with Figure 16, which showed that the consumption betas are very similar for all MPC deciles. Here, we learn that grouping the population in other ways does not change the consumption beta invariance result. Clearly,

[33]In this section, individual i's consumption is defined as the total consumption of its household. If the household consists of two adults (either married or cohabiting), the consumption measure is divided by two.

Table 2: Consumption Betas Norway:

HtM Categorization	Share in		Consumption betas		Multipliers	
	Pop.	C	Agg'D c_is	NA C	Agg'D c_is	NA C
(a) Liq.<2 wks perm. y						
HtM	0.16	0.126	1.06	1.11	1.009	1.012
Non-HtM	0.84	0.874	0.99	1.01		
(b) Liq.<2 wks of y_{-1}						
HtM	0.17	0.14	1.078	1.11	1.012	1.015
Non-HtM	0.83	0.86	0.989	0.999		
(c) Liq. to perm. y (bottom 30%)						
HtM	0.31	0.27	1.074	1.116	1.028	1.037
Non-HtM	0.69	0.73	0.971	0.982		
(d) Liq. to y_{-1} (bottom 30%)						
HtM	0.30	0.26	1.09	1.12	1.033	1.04
Non-HtM	0.70	0.74	0.965	0.97		
(e) Liq. level (bottom 30%)						
HtM	0.29	0.24	1.049	1.081	1.015	1.021
Non-HtM	0.71	0.77	0.985	0.994		
(f) Liq. level (bottom 10%)						
HtM	0.08	0.06	0.984	1.085	0.999	1.004
Non-HtM	0.92	0.94	1.001	1.011		
(g) Not owning stocks						
HtM	0.58	0.5	1.023	1.008	1.02	0.994
Non-HtM	0.42	0.5	0.984	1.02		
(h) Net wealth < 2 months y_{-1}						
HtM	0.32	0.28	1.082	1.078	1.029	1.024
Non-HtM	0.68	0.72	0.98	0.992		
(i) Net wealth level (bottom 30%)						
HtM	0.27	0.24	1.036	1.147	1.01	1.045
Non-HtM	0.73	0.76	0.994	0.966		

Notes: The table displays the consumption-betas and general equilibrium multipliers for various HtM categorizations and for two ways of constructing aggregate consumption growth: either GDP growth (NA C) or the growth rate of aggregated individual consumption in the sample (Agg'D c). The beta estimates are obtained from regressing equation (38) in appendix C.3, with income replaced by consumption. The unit of analysis the individual and consumption is measured as individual i's total household consumption (normalized by 2 if individual i has a partner, either married of cohabitant). When constructing the HtM splits, wealth and income are measured as individual i's total household wealth and income. The multiplier is constructed as the ratio of the multiplier based on the aggregate MPC (AMPC) and the average MPC (aMPC; when individuals are weighted by their income shares), defined in equation (eq. 8). Liq (liquid wealth) and net wealth are defined in appendix B.1. In panels (a) and (c) income, labeled y, is measured as the average net income of the past three years, and in panels (b), (d) and (h) as last year's net income. Bottom 30% and 10% refers to being below the 30th or 10th percentile of the within-year distribution of the variable. In panel (a) and (b) weekly income is proxied as annual income divided by 52, while in panel (h) monthly income is proxied by annual income divided by 12.

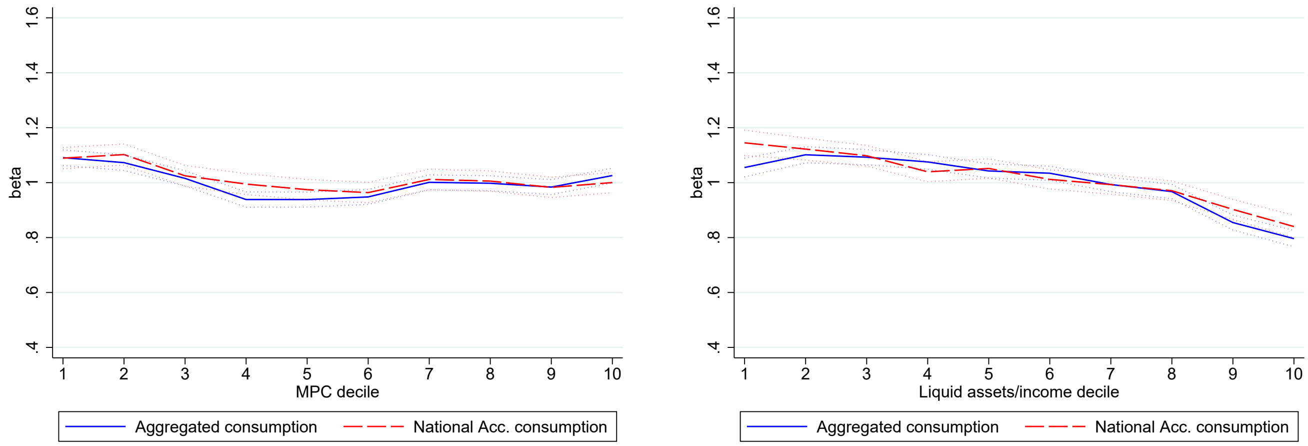

(a) by MPC decile

(b) by Liquid assets/income decile

Notes: The figure shows the estimated consumption beta for various subgroups of the population. The beta estimates are obtained from regressing equation (38) in appendix C.3, with income replaced by consumption. The dashed line represents consumption betas with respect to national accounts consumption, while the solid line with respect to the aggregated individual consumption in the sample. The unit of analysis is the individual, and consumption is measured as individual i's total household consumption (normalized by 2 if individual i has a partner, either married of cohabitant). In panel (a) we consider deciles of the mpc distribution obtained in section 4. In panel (b) we consider deciles of the distribution of liquid wealth to income, where both wealth and income is measured as individual i's total household wealth and income. Dotted lines show 95% confidence bands.

Figure 16: Consumption betas.

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 substantiate this, the last two columns report the multipliers (computed using the expression in Equation 15) for each of the two corresponding ways of estimating the betas. As can be seen by inspecting the numbers, the degree of amplification would be merely 1.009 or 1.012 for the first split—i.e. a mere 1% larger effect of demand shocks and policies. Similarly modest numbers apply for the other splits (some indeed imply dampening), up to at most 1.045 for the eighth and last split, i.e. a 4.5% larger effect through heterogeneity.[34] We provide a further summary calculation of the implied multiplier for the case study of the Great Recession below.

6.2.1 Consumption versus income betas

To have an even more direct benchmark for gauging the role of the distinction between income and consumption, we compute the income betas (for both earnings and disposable income) for the same HtM splits as in Table 2. We then draw on the theoretical insights in section 2, in particular on the multiplier based on income betas in Equation 9. In Table 3, we report the corresponding numbers for two main splits (the first and fourth in Table 2), estimating the betas with respect to

[34]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.

the aggregated income measure only. We outline the results from all the other HtM categorizations and for both measures of aggregate income for completion in Appendix C.8.

Table 3: Multipliers from Earnings and Income Betas by HtM Status

HtM Categorization	HtM share in		Betas		Multipliers	
	Earnings	Disp. Income	Earnings	Disp. Income	Earnings	Disp. Income
(a) Liq.<2 wks perm. y						
HtM	0.11	0.13	1.279	0.799	**1.036**	**0.97**
Non-HtM	0.89	0.87	0.962	1.031		
(b) Liq. to y_{-1} (bottom 30%)						
HtM	0.25	0.27	1.16	0.859	**1.056**	**0.95**
Non-HtM	0.75	0.73	0.945	1.053		

Notes: The table presents income betas and general equilibrium multipliers for various hand-to-mouth (HtM) categorizations and income measures. The columns labeled "HtM share in" refers to the share of total labor earnings or disposable income earned by each group. The Betas estimates are obtained by regressing equation (38) in appendix C.3 for each group and each income measure, using sample-aggregated income as the right-hand-side variable. The last two columns report the ratio of the general equilibrum multipliers: the multiplier based on AMPC and the multiplier based on the aMPC, defined in equation (8).

The results are remarkably in line with those obtained by calculating the AMPC by estimating MPCs and betas at the disaggregated, individual level. Based on the cyclicality of earnings inequality (the difference in betas between savers and HtM individuals), there is amplification, while based on disposable income there is dampening. These results are reassuring in that they circumvent the need to compute MPCs and are thus robust to any bias in the MPC estimation. Furthermore, they constitute a natural benchmark for comparison with the conclusion of the consumption betas-based estimation.

Note that consumption betas can imply amplification, despite disposable income betas indicating dampening. In an economy with saving, these two metrics do not overlap, with the consumption beta-based measure being the correct sufficient statistic. That the direct, consumption-based sufficient statistic does not point to significant multiplier amplification due to heterogeneity suggests the aggregate quasi-irrelevance of the interplay of the several heterogeneity mechanisms we emphasized, including the novel one that is captured by the sufficient statistic 2. Recall that HtM (or high MPC) agents' consumption can be relatively more cyclical than non-HtM (or low MPC) agents', even if the former group's 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). The overall joint contribution of these "cyclical savings" channels is gauged by comparing the consumption-betas multipliers in Table 2 with the corresponding, dampening-inducing ones based on disposable income in Table 3. For the two HtM categorizations we emphasize, the ratio of these is 1.04 (1.009/0.97) and 1.087, respectively, indicating that cyclical savings per se amplify fluctuations relative to a uniform-income hypothetical scenario by 4 and 8.7 percent, respectively.

Our findings of flat consumption betas suggest that overall, the lower cyclicality of disposable incomes of high MPC consumers is offset by the consumption smoothing ability and cyclical savings of low MPC agents. This relates to the cyclicality of income risk and its concentration; measuring these channels and teasing out their roles is the subject of a follow-up paper.

6.3 Case Study: The Great Recession in Norway

As an alternative to gauge the degree of 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? In this section, we calculate betas for the Great Recession to learn whether this large shock created a more pronounced heterogeneity in consumption betas compared to the period 2006-2018 as a whole. In Figure 17 we plot the growth rate of both consumption and the various measures of income for both the HtM and the non-HtM group, based on two classification methods, as well as the aggregate growth rate. As a benchmark, consider the perfect insurance case where both groups' consumption growth is the same as the aggregate growth rate.

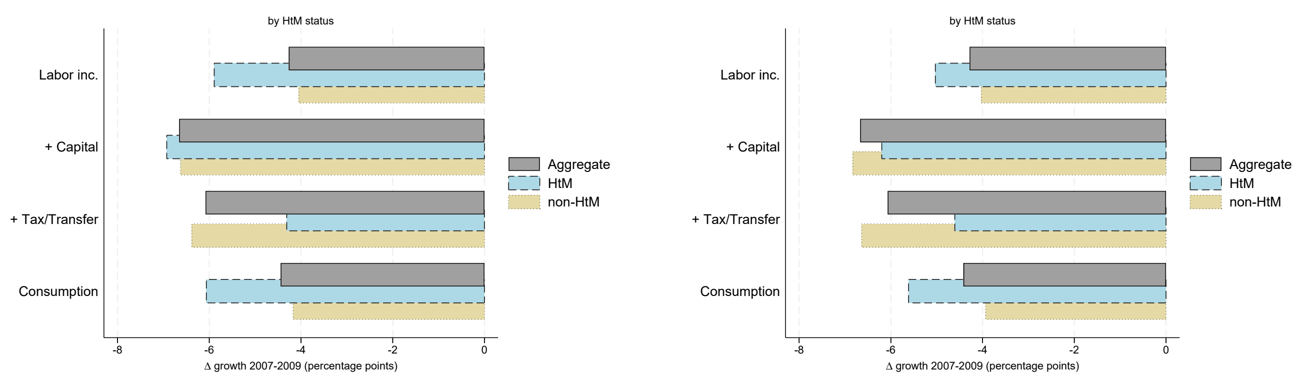

(a) Liq.<2 weeks perm. y

(b) Liq. (by y_{-1}) in bottom 30%

Notes: This figure depicts the two year growth rate (between 2007 and 2009) for successively more comprehensive income measures as well as consumption, both for the total sample as well as for two subgroups that together make up the total. In (a) the subgroups are defined based on low liquidity relative to permanent income (average of last three years of net income), while in (b) based on being below the 30th percentile of the distribution of liquid wealth relative to lagged income. The partitioning of people into subgroups is performed in 2007 and we compute both the total and group-specific growth rates for people who are observed in both years 2007 and 2009.

Figure 17: The Great Recession and Inequality in Norway: A Decomposition

Our findings indicate that at most, with Great Recession betas for the HtM and non-HtM groups of 1.355 and 0.948, respectively, we would observe a mere 1.054 multiplier based on the first HtM categorization in the left panel. To move towards obtaining an *upper bound* and give these mechanisms the best chance to deliver substantial amplification, we consider the second

HtM categorization based on splitting at the bottom 30% of the liquid assets distribution (normalized by previous-period income). Even in this case, the Great Recession betas (1.30 vs 0.89) deliver a multiplier of 1.12, similar to what was reported in Table 1.[35] Even in the Great Recession period we thus find scant evidence of amplification of shocks due to heterogeneity.[36]

6.4 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.[37] To address this, we re-estimate all our beta regressions instrumenting the dependent variable with the past aggregate variable (earnings, net income, and consumption, respectively) instead of using the current one directly, as in Flynn and Sastry (2024). We do this for both the national accounts variables and for the micro-aggregated variables. The results are summarized in Figure 18, 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.

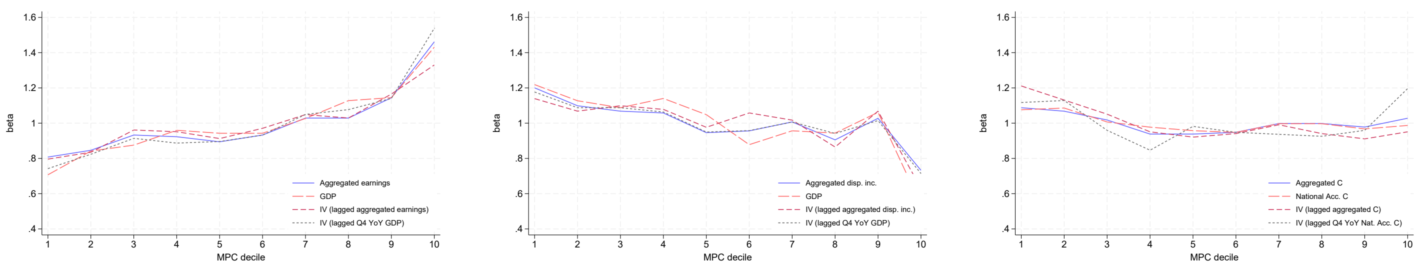

(a) Labor Earnings Betas (b) Income After Tax Betas (c) Consumption (transactions)

Notes: This figure depicts beta estimates for deciles of the mpc distribution. Aggregated earnings and GDP, refers to the baseline right-hand-side variables. In IV (lagged aggregated earnings) we instrument aggregated earnings with its lagged value. In IV (lagged Q4 YoY GDP) we instrument GDP with lagged 4th quarter year-on-year growth. Similar instrumentation is performed for after-tax income and consumption in (b) and (c).

Figure 18: Betas with current vs (instrumented) lagged aggregate variables.

We run a two-stage least squares regression where in the first stage we regress the contemporaneous aggregate growth on the instrument and in the second stage we regress the individual

[35]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).

[36]This conclusion is broadly consistent with and complements Orchard et al. (2025) who use different methods to study the Great Recession in the US and find that (different) general-equilibrium forces act in a dampening fashion, thus reconciling low amplification with high micro MPCs.

[37]We thank Joel Flynn for raising this point.

growth rates on the predicted contemporaneous growth, where the predicted value is the fitted value from the first stage. We use two IVs 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 accounts elasticities).

Figure 18 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.

7 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.

We first propose two sufficient statistics that determine whether dynamics due to heterogeneity will contribute to aggregate fluctuations. We show that these are equivalent in a world where assets are in net zero supply, but will differ when net aggregate savings are possible and time-varying. In that case, sufficient statistic 2, based directly 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 (betas) and MPCs are positively correlated. Based on this estimate, there is amplification of aggregate shocks through heterogeneity: while the model consistent aggregate MPC is 0.374, a simple weighted average of individual MPCs that does not correct for the individual exposures is 0.357. The finding of amplification based on earnings data is qualitatively similar to what Patterson (2023) finds based on US earnings data. However, our data allows us to work with disposable income as well; using that theory-consistent income measure, we find that heterogeneity if anything mildly *dampens* aggregate fluctuations, with an aggregate MPC that is slightly lower than the average MPC.

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 are mainly responsible for changing the cyclical dynamics of income across MPC groups. Furthermore, we show by means of a counterfactual exercise that even the much less progressive US tax and transfer system would have essentially offset any heterogeneity-induced amplification of shocks. These results have a positive connotation: the tax and transfer system is doing its job, cushioning the effects that business cycle shocks could have through inequality and distributional dynamics. In other words, automatic stabilizers do stabilize.

Turning to the second 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 are quite similar across groups, regardless of how these are defined. Hence, heterogeneity does not add much to aggregate fluctuations based on this sufficient statistic either. Even in response to the substantial income loss induced by the Great Recession, heterogeneity contributed *at most* to a 12% amplification.

This paper focuses on the effect of heterogeneity in amplifying shocks on impact and found limited evidence of such amplification. There are, however, other ways heterogeneity may be consequential, above and beyond the importance of distributional issues per se. Not all consumption responses 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. (2024), Hagedorn et al. (2019), and Bilbiie (2024) is a complementary exercise needed to gauge amplification and propagation more broadly. 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 a line we are currently pursuing. These are all questions central to the 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. (2024). Who are the hand-to-mouth? *The Review of Economic Studies*, page rdae056.

Ahn, S., Galaasen, S. M., and Mæhlum, M. (2024). The cash-flow channel of monetary policy - evidence from billions of transactions. Working Paper 20, Norges Bank.

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. (2020). Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model.

Auclert, A., Rognlie, M., and Straub, L. (2024). The intertemporal keynesian cross. *Journal of Political Economy*.

Bayer, C., Born, B., and Luetticke, R. (2024). Shocks, frictions, and inequality in us business cycles. Technical Report 5.

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.

Bergholt, D., Furlanetto, F., and Mori, L. (2024). A tale of procyclical inequality: Facts and implications. Mimeo.

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

Bilbiie, F. O. (2024). Monetary policy and heterogeneity: An analytical framework. *Review of Economic Studies*.

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., Druedahl, J., Harmenberg, K., and Ãberg, E. (2024). Stimulus effects of common fiscal policies. CEPR Discussion Paper 19823, Centre for Economic Policy Research.

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.

Commault, J. (2024). Heterogeneity in mpc beyond liquidity constraints: The role of permanent earnings. Working Paper hal-03870685, Sciences Po. Available at HAL open archive: https://sciencespo.hal.science/hal-03870685v2/document.

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. (2024a). Heterogeneity and aggregate fluctuations: Insights from tank models. *NBER Macroeconomics Annual*.

Debortoli, D. and Galí, J. (2024b). Idiosyncratic income risk and aggregate fluctuations. *American Economic Journal: Macroeconomics*, 16(4):279â310.

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. (2024). The consumption expenditure response to unemployment: Evidence from norwegian households. *Journal of Monetary Economics*, page 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.

Flynn, J. P. and Sastry, K. (2024). The macroeconomics of narratives. NBER Working Paper 32602, National Bureau of Economic Research.

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. (27552).

Ganong, P. and Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review*, 109(7):2383–2424.

Golosov, M., Graber, M., Mogstad, M., and Novgorodsky, D. (2024). How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income. *The Quarterly Journal of Economics*, 139(2):1321–1395.

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.

Guerrieri, V. and Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap*. *The Quarterly Journal of Economics*, 132(3):1427–1467.

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.

Halvorsen, E., Holter, H. A., Ozkan, S., and Storesletten, K. (2024). Dissecting idiosyncratic earnings risk. *Journal of the European Economic Association*, 22(2):617–668. Published online: 25 August 2023.

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.

Ilut, C. and Valchev, R. (2021). Imperfect problem solvers. *The Quarterly Journal of Economics*, 136(3):1399–1457.

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.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531–542.

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.

Nielsson, U., Rangvid, J., Saidi, F., Seyrich, F., and Streitz, D. (2025). Empirical monetary-fiscal equivalence. Mimeo.

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.

Samuelson, P. A. (1948). *Economics: An Introductory Analysis*. McGraw-Hill, first edition.

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. (2009a). Did the 2008 tax rebates stimulate spending? *American Economic Review*, 99(2):374–379.

Shapiro, M. D. and Slemrod, J. (2009b). 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.

Werning, I. (2015). Incomplete markets and aggregate demand.

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:

$$C_Y c_t^S + \frac{X_Y}{1-\lambda} x_t = Y_Y^S y_t^S, \tag{21}$$

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.[38]

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

$$c_t^H = y_t^H. \tag{22}$$

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

$$y_t = C_Y c_t + X_Y x_t. \tag{23}$$

Aggregate consumption and income are given by:

$$c_t = \lambda c_t^H + (1-\lambda) c_t^S \tag{24}$$

$$y_t = \lambda Y_Y^H y_t^H + (1-\lambda) Y_Y^S y_t^S. \tag{25}$$

Replacing in the budget constraints (21) and (22):

$$c_t^H = \beta_y^H y_t \tag{26}$$

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

We solve again for savers' consumption:

$$c_t^S = \frac{1-\lambda \beta_y^H \Omega}{1-\lambda} c_t, \tag{27}$$

where

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

that is when the share of saving (investment, or liquidity) in total income, aka the savings rate, is procyclical. This governs how much more volatile total income is relative to consumption, solved

[38] 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).

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:

$$c_t = E_t c_{t+1} - \frac{1 - \lambda}{1 - \lambda \beta_y^H \Omega} r_t. \tag{28}$$

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

$$\frac{dc_t}{d\left(-r_t\right)} = \frac{1 - \lambda}{1 - \lambda \beta_y^H \Omega}. \tag{29}$$

There can be amplification even if $\beta_y^H < 1$ (disposable-income betas inversely correlated with MPCs) when there is non-zero saving $X_Y > 0$ and the savings rate is procyclical $\eta > 1$.

A.1 Adding risk (for liquidity)

We use the simplest version of the THANK model in Bilbiie (2024) 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 $1 - s$:

$$c_t^S = s E_t c_{t+1}^S + \left(1 - s\right) E_t c_{t+1}^H - r_t. \tag{30}$$

Replacing individual consumptions (26) and (27) in (30) delivers the next Proposition.

The aggregate Euler equation with idiosyncratic risk and saving is:

$$c_t = \Theta_c E_t c_{t+1} - \Theta_r r_t,\ \Theta_c \equiv 1 + \left(1 - s\right) \frac{\beta_y^H \Omega - 1}{1 - \lambda \beta_y^H \Omega}; \tag{31}$$

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:

$$\beta_y^H \Omega > 1 \rightarrow \eta > 1 + \left(1 - \beta_y^H\right) \frac{1 - X_Y}{X_Y}, \tag{32}$$

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 (32) 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 (2024, 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 :

$$c_t^S - c_t^H = \frac{1-\beta_y^H C_Y}{(1-\lambda)C_Y} y_t - \frac{X_Y}{(1-\lambda)C_Y} x_t. \tag{33}$$

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:

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

$$= \frac{1-\lambda\beta_y^H Y_Y^H}{1-\lambda} y_t - C_Y \beta_y^H y_t \tag{35}$$

$$C_Y c_t^S - C_Y c_t^H + \frac{X_Y}{1-\lambda} x_t = \frac{1-\beta_y^H\left(\lambda Y_Y^H + (1-\lambda)\,C_Y\right)}{1-\lambda} y_t \tag{36}$$

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

$$C_Y\left(c_t^S - c_t^H\right) = \frac{1-\beta_y^H C_Y}{1-\lambda} y_t - \frac{X_Y}{1-\lambda} x_t \tag{37}$$

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

Our income and wealth variables is extracted from individuals' annual tax records. From here we observe the following annual income measures:

- Labor earnings: sum of cash wages and salaries, taxable in-kind benefits, sickness benefits and parental leave benefits
- Income from self-employment: Net income from operating sole proprietorships

- Capital income: sum of interest and dividends, realized net capital gains and property income
- Transfers: Sum of taxable and tax-free transfers (excluding sickness benefits and parental leave benefits)
- Gross income: sum the above.
- Income after tax. Gross income net of taxes

In addition to each of the main components listed above, for some of the categories we observe sub-components. In particular we observer unemployment benefits. From the tax records we also obtain detailed wealth variables measured as of December 31 each year. From this information we construct liquid and total wealth:

- Liquid wealth: sum of cash, deposits, securities (stocks, bonds and funds)
- Other wealth: sum of real estate, private business, cars etc.
- Total gross wealth: Liquid wealth + Other wealth
- Total net wealth: Total gross wealth - Total debt.

Using unique anonymized personal identifiers, we combine income and wealth with demographic information such as age, education, and family characteristics. Education is measured as the highest completed level of schooling. The family, from which we construct household-level variables, is defined as a group of individuals residing in the same dwelling and related to each other as spouses, registered partners, cohabitants, and/or parents and children (regardless of the child's age). Unique anonymized family identifiers enable us to link individuals within the same family to one another.

B.2 Electronic Transactions Data

This section covers details about the transactions data and draws heavily on Ahn et al. (2024). Our electronic transaction data covers all Norwegian residents over the period 2006-2018 and contains information on debit card purchases processed by the BankAxept system and online wire transfers processed by the Norwegian Interbank Clearing System (NICS). At the most granular level, our database provides information on number of transactions made and total payments, in a given week, consumption category (26 in total) and zip code (5-digit level) We observer debit card transactions and online wire transfers separately. The aggregation of transaction level data to week-category-zip level was performed by the data provider prior to sending the data to Norges Bank. For the debit card purchases, the mapping to consumption category and zip was based on meta-data about the terminal in which the card was used. For online wire transfers the mapping was based on meta-data about the firm receiving the payment (its NACE code and address). Transactions with no such meta-data are placed in separate categories for missing consumption and/or zip.

B.2.1 Coverage and pre-cleaning

Figure 19 compare the volume and value of total debit card transactions in our database with those from official statistics. While our data base tracks official statistics closely from 2012, prior to the coverage is somewhat lower in our database. The reason, as explained in Ahn et al. (2024) is that the data provider did not retain the link between cards and card owners for bank accounts that where closed prior to 2012. The likely source of bank account closure is either death or migration prior to 2012, implying that people with missing card-owner links are likely completely missing from the data. The fact that per-capita transactions displays no jump in 2012 is consistent with this conjecture (see Figure 20).

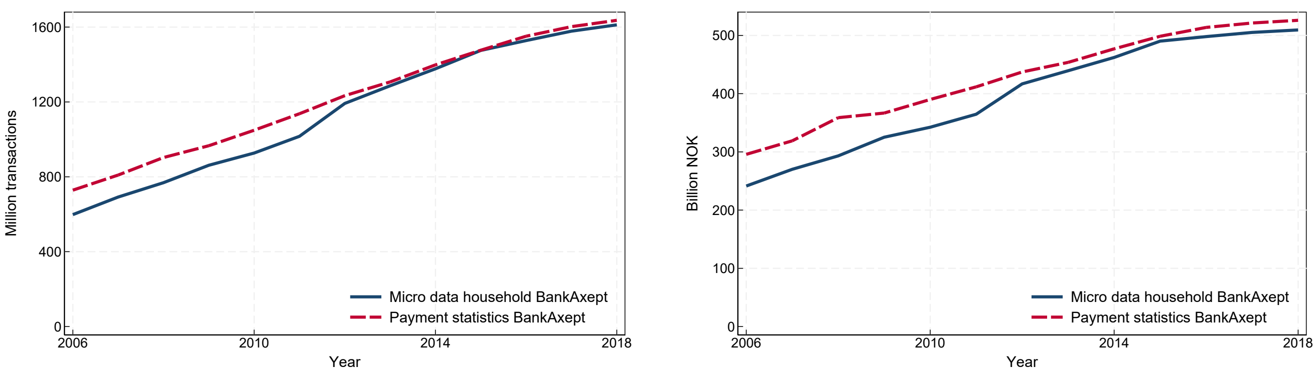

a: Transactions

b: Amounts

Notes: This figure shows aggregate number of transactions and the total transaction amount from two sources. The solid line represents the aggregated volumes over all households in our data base. The dashed line represents the official numbers reported in Norges Bank (2023). Panel (a) and (b) show debit card transactions processed using BankAxept.

Figure 19: BankAxept Transactions and Amounts

Figure 21 compare the volume and value of payments made through online bank wire transfers in our database with those from official statistics. On average our data captures 77 percent of all transactions (panel a). The reason for this gap is that not all wire transfers are processed by NICS. The jump in the official statistics between 2017-2018 reflects the introduction of peer-to-peer mobile payments. Panel (b) reveals that in terms of value, the coverage of wire transfers is slightly lower. The total value is however very sensitive to unusually large single transactions. For example, in 2009 and 2011, one single transfers accounts for more than 5 percent of the total transfer value in those years. We now present three pre-cleaning steps performed on the wire transfer data:

1. Large transactions: Since large transactions are unlikely to capture consumption, we pre-clean the data by removing large single transactions from our data base. We follow Ahn et al. (2024) and remove transactions exceeding 12,500 USD (2015).

2. Credit card imputation: Moreover, one of the consumption categories in our database is payments to banks (as identified by the 5 digit bank NACE codes). These payments are pre-

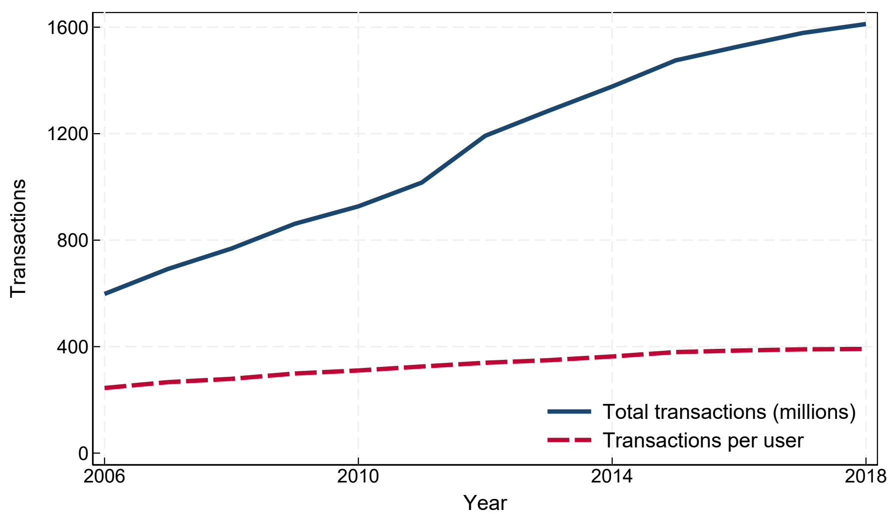

Notes: This figure plots the total transactions (in millons) and transactions per debit card user in our electronic transactions database.

Figure 20: BankAxept Transactions: Total and per users

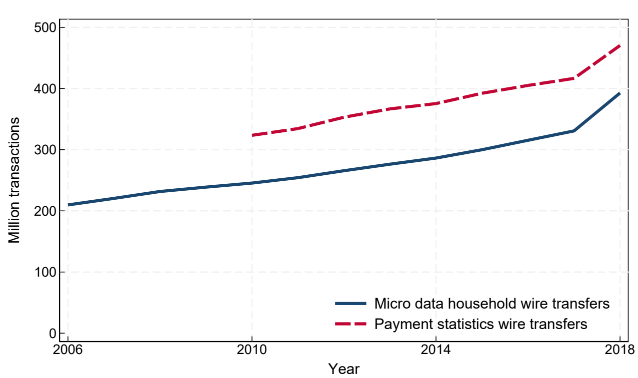

a: Transactions

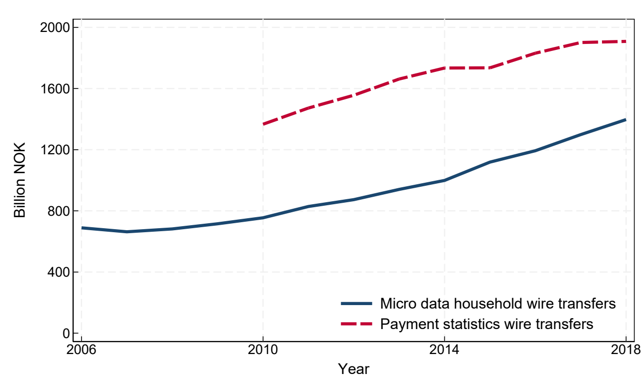

b: Amounts

Notes: This figure shows aggregate number of transactions and the total transaction amount from two sources. The solid line represents the aggregated volumes over all households in our data base. The dashed line represents the official numbers reported in Norges Bank (2023). Panel (a) and (b) show bank wire transfers. The split between firm and household wire transfers is only available from 2010 in the official statistics.

Figure 21: Wire Transfer Transactions and Amounts

dominantly combinations of mortgage payments and payments of credit card bills. Since the former to a large extent reflects savings rather than consumption expenditures, we clean this category by removing transactions imputed as being mortgage related following procedure explained in Ahn et al. (2024). We retain the transactions we are left with after this imputation, and these will to a large extent capture payments of credit card bills.

3. Intra-household transfers: Additionally, since we are measuring consumption at the household level we would ideally like to remove intra-household transfers from the data to avoid double counting consumption. We do not directly observe intra-household transfers, but person-to-person transfers are concentrated in combination of missing consumption category and zip code. We therefor remove all transfers with missing information on both these two dimensions.

After this pre-cleaning step we are left with a transaction based consumption measure that tracks household consumption in the national accounts closely, as depicted in Figure 1.

B.3 Sample Selection and Summary Statistics

When collapsing at the household level, our full household panel, after imposing the age restriction of age 25-66, contains 16.4 million household-year observations. Since the transaction data is noisy with a non-negligible fraction of households with exactly zero consumption, we need to impose certain sample restrictions before proceeding with the analysis. First, we drop households who are self-employed and who migrate in or out. Based on reported self-employment income in the tax records and migration data from the population records we check if household members ever migrate out of Norway or ever are self-employed, and remove all the individuals in the households entirely from our sample. This reduces the total sample size by 4.8 million household-year observations The restriction on non-migration and non-self-employed is related to only partially observing consumption for people who stay abroad for the former group and contamination from business-related expenditures for the latter group. The majority of the 4.8 million dropped comes from the latter group, accounting for 85 percent of the dropped households. Next we remove households with little or no activity in the consumption data, by imposing the condition that households must make at least one debit card transaction per week in at least 50% of the years. If not, we remove the household and all its members entirely from the sample. Finally, we trim extreme observations by removing household-year observation with a consumption-to-income ratio outside the 2.5-97.5 percentile range of the full sample, each year, and remove household-years with less than 12 card transactions. These three last conditions further remove an additional 1.4 million household-year observations.

Table 4 reports household-level summary statistics for both the full and final sample. Since most of our analyses are performed at the individual level following Patterson (2023), Table 5 reports the individual level summary statistics. Compared with income, the distinction between private and household level consumption and wealth is less clear. Hence for individuals with a partner, we measure consumption and wealth variables as total household level counterpart (divided by 2 if the individual has a partner, married or cohabitant).

Table 4: Income and Consumption Statistics Households 2006-2018

	N (mill.)	mean	sd	p10	p50	p90
Panel (A): Unrestricted sample 2006-2018:						
Labor earnings	16,4	54600	41312	0	54899	97477
Total after-tax income	16,4	50375	69387	24959	47101	76738
Cons	16,4	36057	34057	5577	31995	64802
pst. w Cons>0	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 (C): Restricted sample 2006-2018:						
Labor earnings	10,2	59803	38710	7881	58889	99789
Total after-tax income	10,2	51911	30774	29453	48157	75574
Cons	10,2	35056	19504	13559	32408	58537
pct. w Cons>0	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 household level 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 consiting of couples, values are divided by two. Panel (A) reports values for the unrestricted sample for the years 2006-2018 while in panel (C) we report summary statistic for the final analysis sample.

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

In Table: 6 we report the coefficient estimates from the full multivariate MPC regression.

Table 5: Income and Consumption Statistics Individuals 2006-2018

	N (mill.)	mean	sd	p10	p50	p90
Panel (A): Unrestricted sample 2006-2018:						
Labor earnings	24.7	54872	48232	0	54307	100237
Total after-tax income	24,7	49416	92009	23615	45210	75357
Cons	24,7	37864	34048	8510	33761	66134
pct. Cons>0	24,7	96	20	100	100	100
Deposits	24,7	21232	78986	516	8007	48562
Debt	24,7	126985	181049	1158	104322	264523
Securitites	24,7	31348	1241357	0	0	20292
Housing	24,7	153484	169258	0	127807	347923
Age	24,7	40	8	28	40	52
Panel (C): Restricted sample 2006-2018:						
Labor earnings	15,1	59104	44731	1424	57504	102030
Total after-tax income	15,1	49963	37815	26809	45967	73775
Cons	15,1	35866	19063	14703	33402	58783
pct. Cons>0	15,1	100	0	100	100	100
Deposits	15,1	19113	46711	607	7690	44292
Debt	15,1	123121	133869	2987	106054	253261
Securitites	15,1	19875	452674	0	0	16800
Housing	15,1	160355	157981	0	140380	348739
Age	15,1	40	8	28	39	51

Notes. This table shows individual level summary statistics for our sample of individuals 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 individuals living with a partner, consumption and wealth variables are total household counterpart divided by two. Panel (A) reports values for the unrestricted sample for the years 2006-2018 while in panel (C) we report summary statistic for the final analysis sample.

C.2 Appendix: Disposable-Income Betas

The results in Figure 22 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 3.5.2, we used specification

Table 6: 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. Coefficient estimates from estimation of equation (20). Standard errors, clustered at the individual level, in parenthesis.

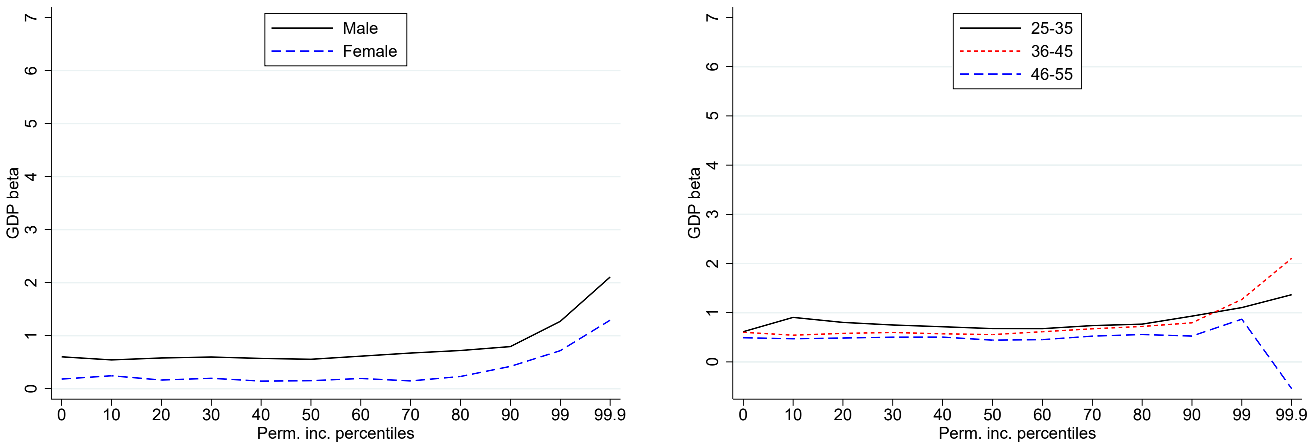

(a) Gender

(b) Age

Figure 22: Disposable Income Betas.

19 from Guvenen et al. (2017), 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 5.1, the relevant object is the group-level beta 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

$$\frac{Y_{i,t} - Y_{i,t-1}}{\bar{Y}_{g,t-1}} = \alpha_g + \beta_g \left(\frac{Y_t - Y_{t-1}}{Y_{t-1}} \right) + \varepsilon_{n,t}, \tag{38}$$

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 38 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 23 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

$$\tau_l(y_l) = \exp\left(\log(\lambda) \left(\frac{y_l}{\bar{y}} \right)^{-2\theta} \right), \tag{39}$$

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. 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

$$T(y) = m\bar{y} \frac{2\exp\left(-\xi\left(\frac{y}{\bar{y}}\right)\right)}{1 + \exp\left(-\xi\left(\frac{y}{\bar{y}}\right)\right)}. \tag{40}$$

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 7, which also shows the square root of the mean squared deviation between actual income

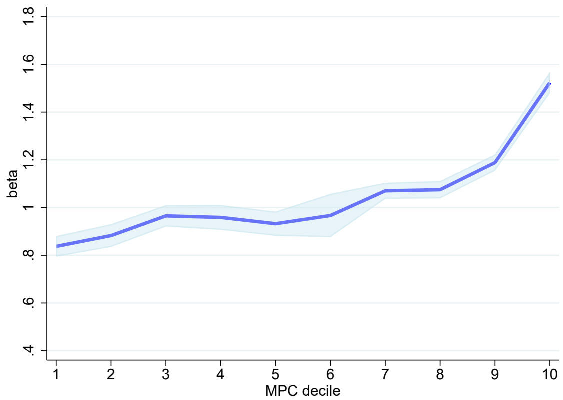

(a) Labor earnings

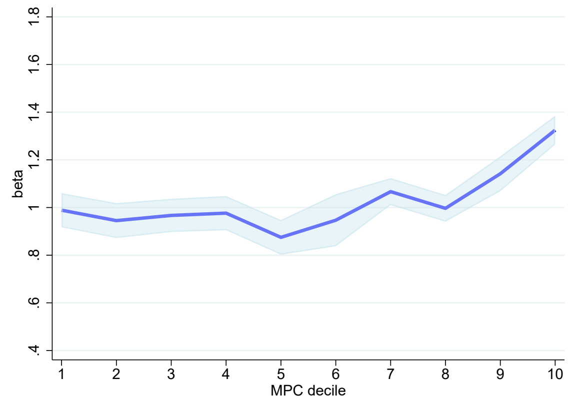

(b) Labor earnings and capital income

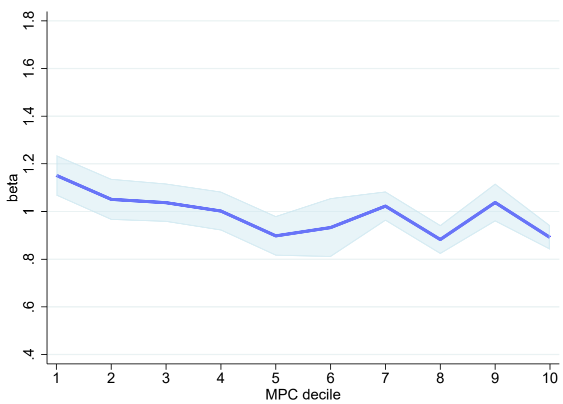

(c) Labor earnings, capital income and transfers

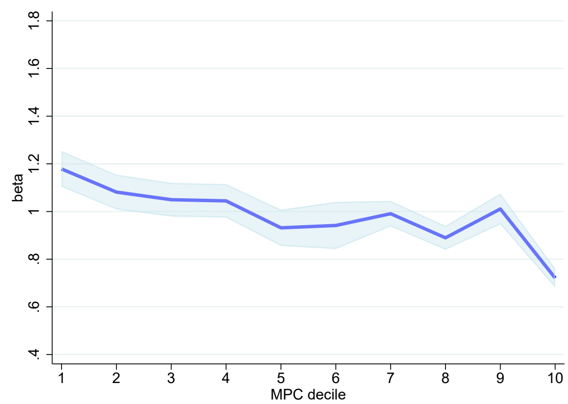

(d) Disposable income

Notes: Estimated income betas from regression 38, by MPC decile. Dotted lines indicate 95% confidence bands based on heteroskedasticity robust standard errors.

Figure 23: Estimated income betas by MPC decile.

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.[39] 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.[40] 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 \ref{eq:ussystem_taxfunc} by non-linear least squares. Third, we define y as labor earnings and capital income, and we estimate equation 40 using non-linear least squares.

The second line of table 7 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.

Table 7 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 24 shows the fit of the estimated average taxes and transfers along the distribution of labor income. [41]

C.4.3 Income betas at the household level

Figure 26shows 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 labor income, actual disposable income and imputed disposable income based on the US tax and transfer system.

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

[40] 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.

[41] 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.

Table 7: 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 39 and 40, 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.

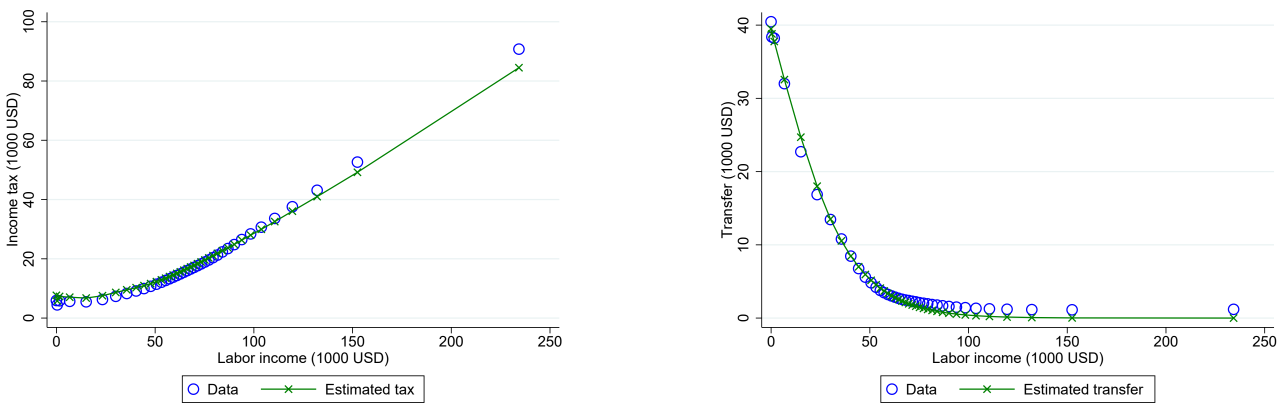

(a) Fit of labor tax function

(b) Fit of transfer function

Notes: The figures demonstrate the fit of the estimated tax and transfer functions in equations 39 and 40, by 2.5% quantiles of labor income. NOK values are converted to USD using the average exchange rate over the period 2006-2018.

Figure 24: Fit of tax and transfer functions for Norway.

C.5 Appendix: Robustness of Aggregate MPC Estimates

In this section we describe how we calculate the aggregate and average MPCs based on, respectively, income betas at the individual level and the method used by Patterson (2023). We also show the results of applying these methods to our data.

C.5.1 Income Betas at the Individual Level

We calculate the aggregate MPC directly as the sum over all individuals in our sample, based on equation 4. Since the estimation of regression 19 at the individual level is sensitive to outliers – due to individuals with income very close to zero in at least one year – we can instead estimate $\tilde{\beta}^j$ for each individual j, where $\tilde{\beta}^j \equiv \frac{dY_t^j}{dY_t}$. In other words, we estimate betas for income in levels

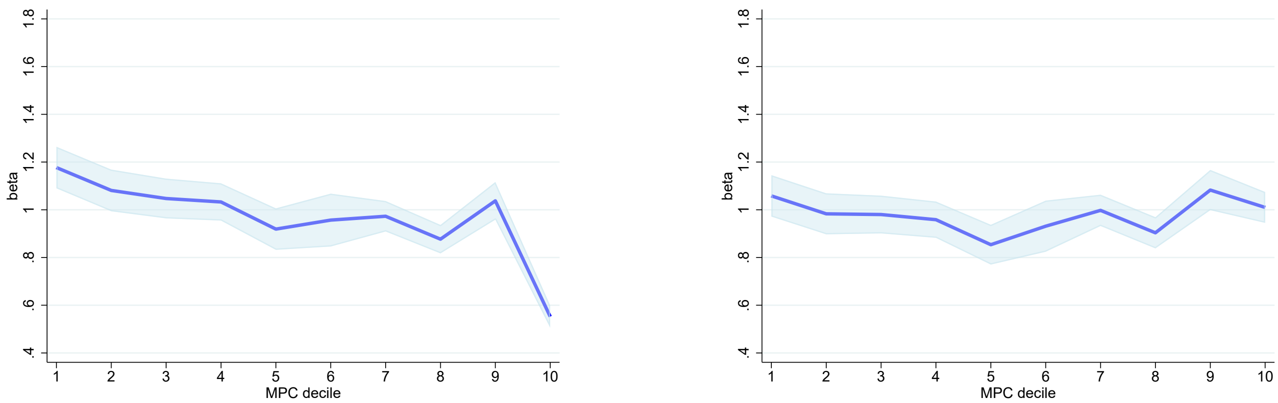

(a) Norwegian tax and transfer system
(b) US tax and transfer system

Notes: Estimated income betas from regression 38, 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 25: Estimated disposable income betas based on estimated tax and transfer system.

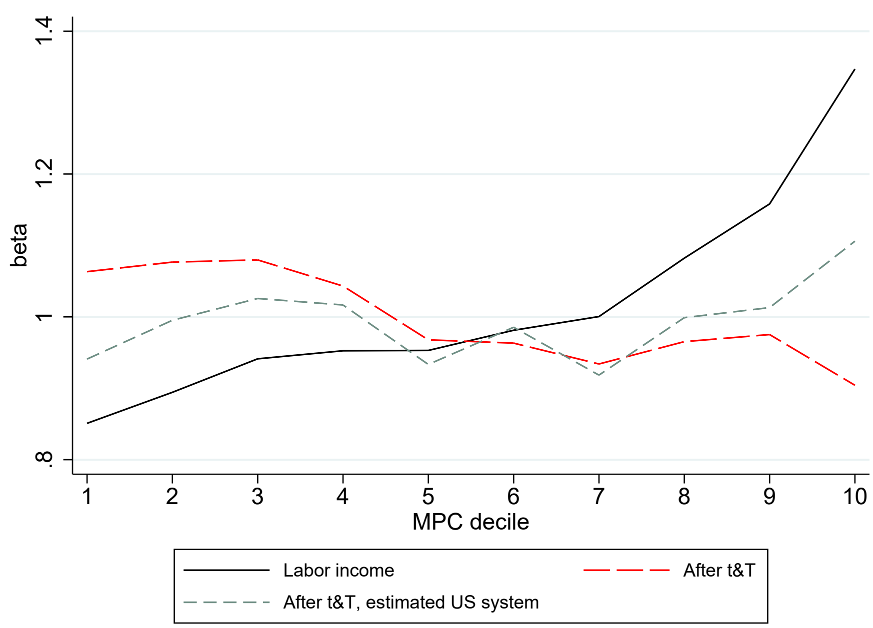

Notes: Estimated income betas at the household level for labor income, actual disposable income (after tax and transfers) and imputed disposable income based on the US tax and transfer system.

Figure 26: Estimated income betas at household level.

instead of in logs. Then the aggregate MPC is given by

$$AMPC = \sum_j MPC^j \tilde{\beta}^j. \tag{41}$$

We estimate the betas by running the regression

$$\Delta Y_{j,t} = \tilde{\alpha}^j + \tilde{\beta}^j \Delta Y_t + \varepsilon_{j,t} \tag{42}$$

for each individual j separately. Here Y_t is the aggregated value of income in year t for all individuals in the sample. The sample consists of all individuals with at least two consecutive years of observations, and we trim the sample at the top and bottom 1% of income growth. The aggregate MPC is calculated by assigning $\hat{\tilde{\beta}}^j$ to individual j for each year that individual appears in the sample, then summing over all individuals and years.

The results are reported in table 8. The degree of amplification is very similar to those reported in table 1, measured by either the covariance between MPC and beta, or the difference between the multiplier based on the aggregate MPC and that based on the average MPC. The covariance is 0.0144 for labor income, but drops to -0.00917 based on disposable income.

Table 8: Aggregate MPC of Norway: Individual-level Income Betas

Income measure	AMPC	avg. MPC	$Cov(MPC, \beta)$	Multiplier A	Multiplier a
Labor income	0.376	0.361	0.0144	1.601	1.565
+ capital	0.370	0.361	0.00924	1.588	1.565
+ transfers	0.363	0.367	-0.00386	1.571	1.580
- taxes	0.361	0.370	-0.00917	1.565	1.588
US counterfact.	0.363	0.363	-0.0000766	1.570	1.570

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 8. Aggregate MPC is based on estimated income betas at the individual level. Covariance calculated as difference between aggregate MPC and average MPC.

C.5.2 The Patterson (2023) Method

We now use the method preferred by Patterson (2023) on our data to calculate the degree of amplification. Specifically, we run the regression

$$\Delta_s Y_{j,t} = \alpha_1 M\hat{P}C_{j,t-1} + \alpha_2 M\hat{P}C_{j,t-1} \times \Delta \log Y_t + \delta_t + \varepsilon_{j,t}, \tag{43}$$

where $\Delta_s Y_{j,t} \equiv \frac{Y_{j,t} - Y_{j,t-1}}{0.5Y_{j,t} + 0.5Y_{j,t-1}}$ is the midpoint symmetric growth rate of individual j's income growth, and δ_t is year-fixed effects. Each individual in the regression is weighted by their share of total income Y_t. As in Patterson (2023), we calculate the covariance between MPCs and income betas as the estimate $\hat{\alpha}_2$ multiplied by the income-share weighted variance of the MPCs. Our estimate of the aggregate MPC is the sum of the average MPC and the covariance term, following equation (5). The results are reported in table 9. As in section C.5.1, we find estimates of amplification that are very similar to those reported in the main body of the paper. The covariance is 0.0122 for labor income, but drops to -0.00578 based on disposable income.

Table 9: Aggregate MPC of Norway: Patterson (2023) Methodology

Income measure	AMPC	avg. MPC	$Cov(MPC, \beta)$	Multiplier A	Multiplier a
Labor income	0.374	0.362	0.0122	1.599	1.568
+ capital	0.372	0.362	0.00972	1.592	1.568
+ transfers	0.370	0.368	0.000936	1.586	1.583
- taxes	0.373	0.379	-0.00578	1.595	1.610
US counterfact.	0.371	0.364	0.00651	1.589	1.573

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 8. Covariance is based on the linear regression in Patterson (2023). Aggregate MPC is the sum of the average MPC and the covariance term.

C.6 Appendix: External Validation/Sanity Check: Beta-Implied MPCs

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

$$\frac{\beta_C^j}{\beta_Y^j} = \frac{dc_t^j}{dc_t} / \frac{dy_t^j}{dy_t} \rightarrow \tilde{MPC}^j = \frac{\beta_C^j}{\beta_Y^j} \frac{dc_t}{dy_t} \frac{C_t^j}{Y_t^j}$$

C.7 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), Commault (2022) 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

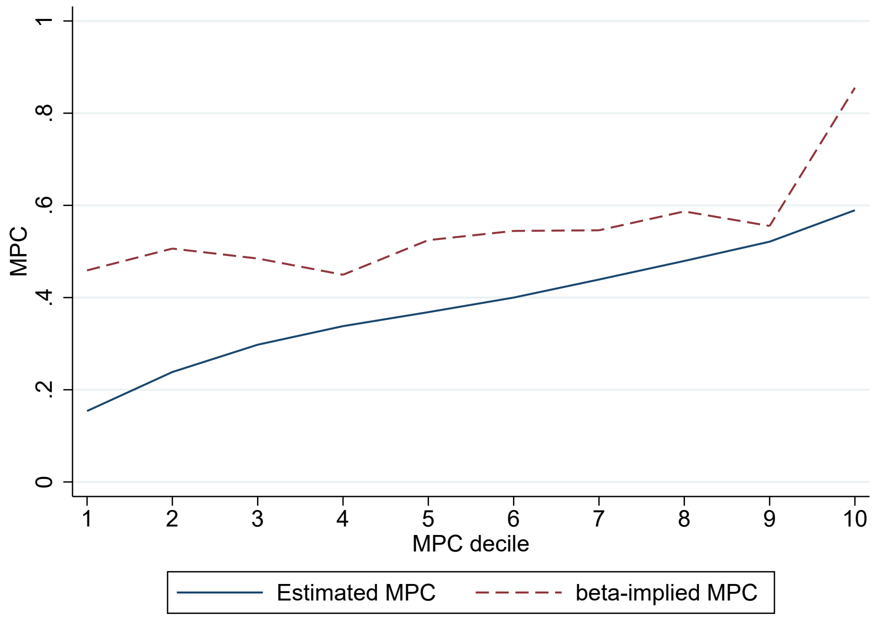

Figure 27: Betas-Implied MPCs.

component

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

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

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

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.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 10: BPP-C MPC estimate for the whole sample

MPC	0.300^{***}
	(0.0273)
Observations	859321

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C.8 Multipliers from Earnings and Income Betas by Hand-to-Mouth Status

This section reports the betas and implied multipliers calculated by splitting the sample according to the various HtM categorizations. We report these in Tables 11 and 12 for both earnings and disposable income, respectively, in a manner that parallels the consumption betas and implied multipliers reported in Table 2. Table 3 in text presents a distilled version of this analysis (the first and fourth HtM categorizations).

Table 11: Multipliers and Earnings Betas by HtM Status

HtM Categorization	Share in		Earnings betas		Multipliers	
	Pop.	Y	Agg'D y_is	NA Y	Agg'D y_is	NA Y
Liq.<2 wks perm. y						
HtM	0.16	0.11	1.279	0.962	1.036	1.019
Non-HtM	0.84	0.89	0.723	0.545		
Liq.<2 wks of y_{-1}						
HtM	0.17	0.12	1.256	0.717	1.036	1.02
Non-HtM	0.83	0.88	0.961	0.546		
Liq. bottom 30% perm. y						
HtM	0.31	0.25	1.171	0.657	1.06	1.031
Non-HtM	0.69	0.75	0.942	0.536		
Liq. bottom 30% y_{-1}						
HtM	0.31	0.25	1.16	0.66	1.056	1.031
Non-HtM	0.69	0.75	0.945	0.538		
Liq. bottom 30% (level)						
HtM	0.29	0.21	1.211	0.693	1.06	1.034
Non-HtM	0.71	0.79	0.94	0.535		
Liq. bottom 10% (level)						
HtM	0.08	0.05	1.501	0.923	1.027	1.019
Non-HtM	0.92	0.95	0.974	0.55		
Not owning stocks						
HtM	0.58	0.49	1.03	0.563	1.027	0.995
Non-HtM	0.42	0.51	0.976	0.573		
Net wealth < 2 months y_{-1}						
HtM	0.32	0.27	1.121	0.64	1.058	1.029
Non-HtM	0.68	0.73	0.915	0.535		
Net wealth bottom 30%						
HtM	0.27	0.23	1.073	0.608	1.026	1.012
Non-HtM	0.73	0.77	0.962	0.556		

Table 12: Multipliers and Disposable-Income Betas by HtM Status

HtM Categorization	Share in		Disp. income betas		Multipliers	
	Pop.	Y	Agg'D y_is	NA Y	Agg'D y_is	NA Y
Liq.<2 wks perm. y						
HtM	0.16	0.13	0.799	0.524	0.97	0.961
Non-HtM	0.84	0.87	1.031	0.828		
Liq.<2 wks of y_{-1}						
HtM	0.17	0.15	0.818	0.529	0.969	0.956
Non-HtM	0.83	0.85	1.031	0.829		
Liq. bottom 30% perm. y						
HtM	0.31	0.28	0.856	0.584	0.947	0.927
Non-HtM	0.69	0.72	1.054	0.862		
Liq. bottom 30% y_{-1}						
HtM	0.30	0.27	0.859	0.587	0.95	0.931
Non-HtM	0.70	0.73	1.053	0.858		
Liq. bottom 30% (level)						
HtM	0.29	0.24	0.727	0.483	0.921	0.913
Non-HtM	0.71	0.76	1.083	0.878		
Liq. bottom 10% (level)						
HtM	0.08	0.06	0.617	0.375	0.976	0.974
Non-HtM	0.92	0.94	1.024	0.81		
Not owning stocks						
HtM	0.58	0.51	0.761	0.541	0.807	0.804
Non-HtM	0.42	0.49	1.227	1.018		
Net wealth < 2 months y_{-1}						
HtM	0.32	0.29	0.904	0.682	0.894	0.917
Non-HtM	0.68	0.71	1.039	0.848		
Net wealth bottom 30%						
HtM	0.27	0.24	0.957	0.749	0.962	0.954
Non-HtM	0.73	0.76	1.003	0.796		