CONSUMPTION INEQUALITY, HOUSEHOLD RISKS, AND THE BUSINESS CYCLE*

Drago Bergholt†, Francesco Furlanetto‡ and Lorenzo Mori§

March 2026

Abstract

Conventional economic theory holds that consumption is smoother and less disperse than income, and that inequality rises in downturns. Using electronic transaction data linked to administrative income records for the universe of Norwegian households, we document the opposite. Consumption is more unequally distributed than disposable income, more volatile within households, and both income and consumption inequality move procyclically. We show that excess consumption dispersion primarily reflects idiosyncratic expenditure shocks rather than income risk or life-cycle factors. Procyclical inequality is driven by heightened business-cycle sensitivity in the upper tail of the consumption and income distributions. These findings call for a reassessment of precautionary saving motives and highlight a disconnect between household risk and the aggregate business cycle.

Keywords: *Consumption and income inequality, precautionary saving, expenditure shocks, business cycles.*

JEL Classification: *D3, E2, E3, H3.*

*First draft: June 2024. This draft should not be reported as representing the views of Norges Bank, Bank of Italy, or the Eurosystem. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank, Bank of Italy, or the Eurosystem. A special thank to Jeff Campbell and Edmund Crawley for detailed written comments on the paper, to Sigurd Galaasen and Ragnar Juelsrud for patient help with the data, and to our discussant Elisa Guglielminetti for useful comments on an earlier draft. Moreover, we are grateful for the valuable feedback by Knut Are Aastveit, Erich Battistin, Jeff Campbell, Fabio Canova, Cristiano Cantore, Efrem Castelnuovo, Edouard Challe, Olivier Coibion, Davide Debortoli, Michael Ehrmann, Andreas Fagereng, Luca Fosso, Luisa Fuster, Jordi Galí, Alessandro Guerriero, Jonathon Hazell, Jonathan Heathcote, David Hémous, Luca Gambetti, Yuriy Gorodnichenko, Luigi Guiso, Martin Holm, Nicolò Maffei Faccioli, John Kramer, Gisle Natvik, Pablo Ottonello, Serdar Ozkan, Fabrizio Perri, Giorgio Primiceri, Ørjan Robstad, Tiziano Ropele, Kasper Roszbach, Juan Rubio-Ramirez, Kjetil Storesletten, Mathias Trabandt, Nicolas Werquin, Giordano Zevi and Fabrizio Zilibotti, as well as seminar participants at Federal Reserve Board, University of Copenhagen, Federal Reserve Bank of St. Louis, Sveriges Riksbank, Norges Bank Mini Workshops on Monetary Policy in 2022 and 2024, Norges Bank 2024 Spring Institute, 2024 Workshop on Household Heterogeneity at Sapienza University, 2024 IFO Conference on Survey data, 2025 SNDE conference in San Antonio, Barcelona Summer Forum 2025, Sailing the Macro 2025 in Siracusa, 2024 and 2025 Dolomiti Macro Meetings, CEPR MEF Symposium 2025 in Gerzensee, 2025 joint conference Bundesbank, Danmarks Nationalbank and Norges Bank, SMN conference in Madrid, and 2026 NBRE Spring Institute in Venastul.

†Norges Bank. Email: drago.bergholt@norges-bank.no

‡Norges Bank. Email: francesco.furlanetto@norges-bank.no

§Bank of Italy. Email: lorenzo.mori@esterni.bancaditalia.it

1 INTRODUCTION

Standard macroeconomic theory assumes that households adjust saving to make consumption smoother than income. Friedman's (1957) permanent income hypothesis, for example, posits that consumption decisions are guided by expected lifetime resources rather than by current income, implying that temporary income shocks have little effect on spending. Modern household finance and intertemporal macroeconomics build on this principle. One natural implication is that household-level consumption should be less volatile and disperse than household-level income. Another is the widely held view that consumption inequality tends to be countercyclical, as poorer households are less able to hedge against income shocks. Indeed, contemporary macroeconomic models often incorporate frictions—such as financial constraints (Kaplan, Moll, and Violante, 2018)—that are more likely to bind for poorer households. In these models, countercyclical consumption inequality amplifies aggregate shocks and plays an important role in shaping business cycle dynamics.

Despite its theoretical importance, there remains considerable uncertainty about both the level and cyclicality of consumption inequality. A major challenge is data: household consumption is difficult to measure, and the main sources—surveys and imputation methods—are prone to significant measurement error. For example, Attanasio and Pistaferri (2016) review prominent measures of U.S. consumption inequality and document wide variation in both their levels and cyclical properties. Moreover, because surveys are typically repeated cross sections, there is limited evidence on the joint evolution of income and consumption at the household level—evidence that is essential for understanding consumption choices and precautionary saving behavior.

We address these measurement challenges by leveraging unique, high-quality panel data on electronic transactions to track household-level consumption in Norway. Our dataset covers approximately 80% of all electronic payments and bank transfers made by the universe of Norwegian households from 2006 to 2018. In a predominantly cashless economy like Norway, these transactions offer an exceptionally accurate and granular view of both the level and cyclical behavior of consumption – making the data ideal for tracking consumption inequality over time with limited measurement error. By linking these transactions data to administrative income and tax records from the Norwegian Tax Authority, we are able to analyze the joint dynamics of consumption and income inequality across the entire population.

We use these data to document that consumption inequality in Norway *does not* behave as conventional wisdom would imply. Our first key—and somewhat unexpected—finding is that consumption inequality at annual frequency is substantially larger than disposable income inequality. This result holds even if we purge away between-group inequality due to life-cycle trends and demographic characteristics, and if we control for time-invariant between-household differences, for example arising from potentially unobserved but permanent differences in productivity or wealth. A second result follows: within-household consumption volatility exceeds that of within-household income volatility. Our analysis indicates that the two observations are closely related, as within-household consumption volatility accounts for a sizable share of the cross-sectional consumption differences in our data. The excess dispersion and volatility in household consumption are only in part explained by the lumpy nature of durable good spending. In

fact, the results are confirmed across different consumption categories, different demographic groups among households and different datasets for consumption expenditures. Furthermore, we show that our results do not violate lifetime budget constraints, as consumption inequality converges toward disposable income inequality when measured over longer horizons (multiple years).

What can explain these empirical findings? We document that idiosyncratic expenditure shocks play a central role. Examples include sudden health expenses, car repairs, a funeral, an expensive holiday, a newborn kid, kitchen renovations, or a wedding banquet. In all cases, income remains relatively unchanged while consumption fluctuates. We point out that excess consumption volatility and household-level consumption smoothing can coexist only if households experience large expenditure shocks independent of income. Put differently, households' consumption dynamics are shaped not only by income shocks (which induce smoothing), but also by transitory expenditure shocks.

To assess the role of expenditure shocks in data, we estimate a set of models designed to quantify their contribution after controlling for income. In the spirit of Hall (1978), we begin with simple regressions projecting household consumption on lagged consumption and contemporaneous income, as well as various controls for life-cycle trends and demographic characteristics. We then consider a more structural model linking consumption to permanent and transitory income shocks, as in Blundell, Pistaferri, and Preston (2008). Importantly, we also identify a consumption expenditure shock that is orthogonal to income changes. Two results emerge. First, income shocks do transmit to consumption. However, households partially smooth them out, especially when the income shocks are transitory. Second and most importantly, after controlling for life-cycle trends, demographics and macroeconomic factors, expenditure shocks account for roughly 85-90% of household-level consumption volatility. Because our data are based on actual transaction records rather than surveys or imputations, we argue that this volatility cannot be attributed solely to measurement error, but instead reflects the pervasive role of expenditure shocks. They account for a major share of the observed dispersion in consumption across households, as well as a major share of consumption volatility within households.

We next turn to the cyclical properties of income and consumption inequality. The conventional view is that economic inequality tends to be countercyclical—inequality rises during downturns and falls during booms—mainly because households in the left tails of the distributions are less able to insure themselves against adverse macroeconomic events. Our analysis challenges this view, as we find that both disposable income and consumption inequality move procyclically in our data.

These results are documented in two ways: first, we use Bayesian time series techniques to estimate a structural vector autoregressive (SVAR) model. This exercise allows us to quantify how inequality responds to a main business cycle shock, following the identification strategy in Angeletos, Collard, and Dellas (2020). Interestingly, we find that procyclical inequality largely arises due to excess sensitivity in the upper tails: following a contractionary business cycle shock, both consumption and income drop significantly among households located at the 90-percentiles and above, while households below the median of both distributions are relatively unaffected. The second way we investigate the cyclicality of inequality is by means of an event study of the Great Recession. Following Guvenen, Ozkan, and Song (2014), we employ a difference-in-differences approach and calculate heterogeneous household exposures *across* the respective distributions. How-

ever, while their analysis is limited to male earnings, we instead consider household-level labor income, market income (defined as the sum of labor income and net capital income), disposable income, and consumption. This exercise not only confirms our SVAR results but adds further granularity: earnings inequality was relatively unaffected during the Great Recession, as labor earnings fell by similar amounts for most households except those at the very top (where bonuses likely played a major role). By contrast, both consumption and disposable income inequality declined because households higher up in the distributions were disproportionately affected by the negative consequences of the recession. The pronounced cyclicality of consumption at the top of the distribution may appear surprising at first, but is consistent with event studies of the Great Recession in the U.S. based on survey data (Parker and Vissing-Jørgensen, 2010; Attanasio and Pistaferri, 2014). Overall, our results highlight how economic environments and policies can protect households with relatively little income and consumption from bad economic events, effectively making income and consumption inequality procyclical.

The empirical results presented in this paper have several implications for macroeconomic theory and policy making. First, they challenge the perceived wisdom that households' consumption variance is driven primarily by income variance —either in the form of innovations to permanent income as in the Friedman (1957) tradition, or from shocks to current income as in the Keynesian tradition. Similarly, income risk typically plays a key role in precautionary savings behavior in modern equilibrium models with heterogeneous agents (Kaplan et al., 2018, Auclert, 2019, Bayer, Born, and Luetticke, 2024, Debortoli and Galí, 2024). Yet, our results point to a major role for idiosyncratic expenditure shocks that are orthogonal to income, suggesting that fundamental consumption risk may represent a separate but key motivation for precautionary savings and insurance. Second, our results not only imply a disconnect between income and consumption at the household level, but also a disconnect between the sources of microeconomic risk and macroeconomic volatility. The reason is that the key drivers of household risk—such as expenditure shocks—appear idiosyncratic and largely average out when aggregated across households and over time. Likewise, aggregate business cycle shocks only account for a minor share of households' consumption allocations. Thus, drivers of aggregate, distributional summary measures such as the 90-10 percentile ratio matter far less for the individuals populating these distributions. Third, despite the apparent micro-macro disconnect, we still argue that macroeconomic stabilization policy can play an important role. In fact, targeted insurance mechanisms provided by government policies may, if successful, cause the disconnect in the first place. This view is supported by our finding that procyclical inequality arises because of redistribution through taxes and transfers, rather than from cyclical differences in market income, or from the wage compression policies emphasized by Mogstad, Salvanes, and Torsvik (2025). In any case, we conclude that the procyclical nature of income and consumption inequality helps to dampen aggregate business cycle fluctuations, making them less relevant for most households except those at the very top. Standard calibrations of heterogeneous agent models are at odds with this conclusion.

This paper contributes to different strands of the literature on household-level consumption inequality and insurance. First, we complement existing evidence on the cross-sectional dispersion in income and consumption—much of which relies on survey data (especially for consumption). We build on the seminal papers of Storesletten, Telmer, and

Yaron (2004), Heathcote, Perri, and Violante (2010) and Guvenen et al. (2014) for the US and we make progresses on the measurement of consumption inequality and volatility. In this sense, our work connects to a large literature that relies on survey data or imputation methods to study the level (Krueger and Perri, 2006; Attanasio, Battistin, and Ichimura, 2007; Primiceri and van Rens, 2009; Aguiar and Bils, 2015; Coibion, Gorodnichenko, and Koustas, 2021; Meyer and Sullivan, 2023) and the cyclicality (Coibion, Gorodnichenko, Kueng, and Silvia, 2017; De Giorgi and Gambetti, 2017; Chang and Schorfheide, 2024; Mangiante and Meichtry, 2025) of consumption inequality. As emphasized by Attanasio and Pistaferri (2016), however, no clear consensus has emerged.

Second, while ample literature has studied the role of income risk for consumption, we contribute to an emerging strand focusing on the role of expenditure shocks. Fulford and Low (2024) document that these shocks are at least as large and frequent as income shocks in US survey data, while Miranda-Pinto, Murphy, Walsh, and Young (2025) provide a theoretical framework where stochastic consumption thresholds explain why households may prioritize saving over consumption. Further highlighting their impact on financial vulnerability, Briglia (2025) identifies expenditure risk as the leading driver of transitions into hand-to-mouth status, and Guiso and Jappelli (2024) find that such shocks are major drivers of precautionary savings in Italian survey data. Relatedly, Blundell, Borella, Commault, and De Nardi (2024) study US health shocks while Campbell and Hercowitz (2019) document widespread saving in anticipation of major expenditures like home purchases and college education. We contribute to this literature by quantifying the importance of expenditure shocks in large-scale transactions data, where we can track a panel of households over time, directly observing both their consumption expenditures, income, asset positions, as well as an array of other household-level characteristics.

Third, we contribute to the growing literature surveyed in Baker and Kueng (2022), which construct measures of individual spending from card transactions or financial applications data. Of particular interest is the paper of Buda, Hansen, Rodrigo, Carvalho, Ortiz, and Rodríguez Mora (2023), which builds a measure of aggregate consumption based on the universe of transactions made by retail clients of a large bank in Spain. They find a higher level of consumption inequality than in survey data. We do not rely on data from a single bank but on the universe of transactions within the Norwegian banking system. A related paper is also Bilbiie, Galaasen, Gurkaynak, Maehlum, and Molnar (2025), which examines whether heterogeneity in individuals' marginal propensities to consume (MPCs) amplifies the aggregate effects of demand shocks and finds a limited role for such heterogeneity in Norway.

The rest of the paper is organized as follows: Section 2 describes our household-level panel dataset. Section 3 documents that consumption inequality is substantially larger than disposable income inequality in the cross-section of households, and substantially more volatile within households over time. Section 4 quantifies the importance of expenditure shocks for consumption dynamics at the household level. Section 5 examines the cyclical properties of our inequality measures using structural vector autoregressions and event-study methods. Section 6 discusses some of the key implications of our results while Section 7 concludes.

2 THE DATA

In this section, we describe how the dataset is constructed as well as some of its key features. As a sanity check, we also compare its aggregates with relevant counterparts from the national accounts.

2.1 INDIVIDUAL-LEVEL ECONOMIC ACCOUNTS

To measure consumption, we use a panel dataset with electronic transactions covering the universe of Norwegian residents over the period 2006–2018. These data are provided by the Norwegian retail clearing institution Nets Branch Norway. Our consumption data are constructed using two main sources of consumer expenditures: (1) debit card transactions processed through the BankAxept network, that is, the national payment system owned by Norwegian banks; and (2) online bank wire transfers cleared through the Norwegian Interbank Clearing System (NICS). Individual transactions are at the person-level and aggregated by the data provider by week and consumer category. A total of 26 distinct consumer categories are available—24 COICOP groups plus wire transfers to the government and to banks. The groups include broad consumption measures such as food, housing, water and electricity, purchase of vehicles, various recreation services, personal care, and financial services. A complete list of the 26 consumer categories is provided in Table A1 in the appendix. Although credit card transactions are not directly observed, we infer them by identifying wire transfers made to banking institutions, following the approach of Ahn, Galaasen, and Maelhum (2024). We cannot, however, disentangle different consumer categories when credit cards have been used. Remaining transfers to banks are excluded, as they most likely correspond to mortgage payments. Our records do not include debit-card payments processed by VISA or Mastercard (e.g., online or international transactions), nor do they include wire transfers settled outside NICS. Nevertheless, Ahn et al. (2024) show that BankAxept and NICS together account for roughly 80% of all household electronic payments from 2006 to 2018 and provide more details about the electronic transactions data. Since transactions are reported at high frequency, the dataset is routinely used to nowcast aggregate consumption in Norway (Aastveit, Fastbø, Granziera, Paulsen, and Torstensen (2024)).

Our data offer substantial advantages over alternative popular consumption measurements, particularly for studying cross-sectional dynamics as we do. Survey data – the primary sources of consumption data – are well known to suffer from measurement error, most notably systematic under-reporting among high-income and high-consumption individuals (Aguiar and Bils, 2015), which biases estimates of consumption inequality downward. Attanasio, Battistin, and Leicester (2006) document a pronounced decline over time in the extent to which survey-based consumption data are able to replicate the patterns observed in official statistics. In addition, surveys typically follow individuals for only a short period—often just a single wave—whereas our data allow us to track their consumption comprehensively, alongside income, wealth, and other covariates (as described below), over many years. Imputation methods, in turn, define consumption residually as disposable income minus saving, so that measurement error in wealth translates directly into error in imputed consumption. Because wealth under-reporting is especially pronounced among the wealthy (Baker, Kueng, Meyer, and Pagel, 2022), this approach

risks overstating consumption inequality. Moreover, Eika, Mogstad, and Vestad (2020) show that standard imputation techniques require discarding nearly half of the sample to mitigate error, a prohibitive limitation in our context.

Our registry data on income and wealth are provided by the Norwegian Tax Authorities, who levy income and wealth taxes on all Norwegian residents. Measurement issues are significantly mitigated by the extensive reliance on third-party reporting. Employers and financial institutions supply information on labor income and financial assets directly to the tax authority, thus making tax evasion very difficult. The data are annual and cover the universe of Norwegian residents since 1993. They include detailed, individual-level accounts of income, wealth, and personal taxes and transfers. The richness of the income accounts allows us to separate between labor income and various forms of capital income, including returns from interest-bearing financial assets, stock returns, capital gains, and other sources of capital income. Moreover, we have a complete overview of registered transfers and taxes, including unemployment benefits, pensions, child support, and social welfare benefits. Finally, the individual-level registry data allow us to separately observe liquid assets (e.g. bank deposits, etc), and to distinguish between gross and net wealth. We also observe individual characteristics such as gender, age, and education level. Finally, while the registry data on income and wealth are widely regarded as being of exceptionally high quality, we nevertheless acknowledge that even these data have limitations, in particular when it comes to the measurement of retained capital gains and income from housing services. We refer to Aaberge, Mogstad, Vestad, and Vestre (2021) for a detailed analysis of these issues.

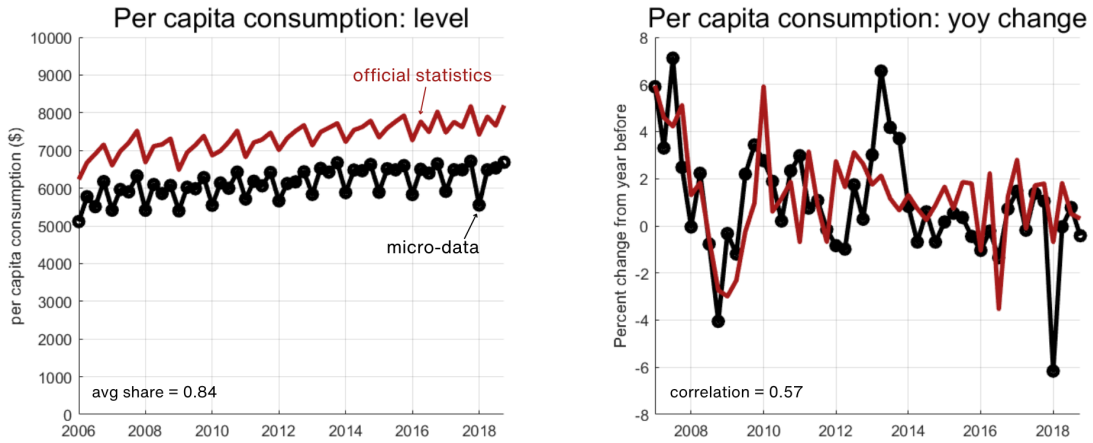
2.2 Household-level aggregation and sample selection

We consider the household as our baseline unit of observation in the paper. Consumption and saving decisions are typically made at the household level, implying that individual-level expenditure transactions may be disconnected from actual individual consumption when people live together. We aggregate all individual accounts to the household-level, and then compute equivalized values by dividing the household accounts by the square root of the number of people in the household. This is a widely used approach to measure variables at the household level, see OECD (2013).

Following standard practice in the literature, we focus on a sample of household-year observations with some (but not necessarily strong) labor-market attachment. We also require regular use of electronic payments. This seems natural given the nature of our consumption data. For each year t we remove household–year observations with labor earnings below $Y_{min,t} = 5\%$ of median equivalized earnings, or consumption below $C_{min,t} = 5\%$ of median equivalized consumption (see Halvorsen, Holter, Ozkan, and Storesletten, 2023 for a similar restriction). We further trim away outliers by excluding household–year observations with consumption-to-income ratios below the 2.5$^{\text{th}}$ percentile of the full sample. A few additional steps are taken to clean the data: first, we remove unusually large single transactions that are unlikely to represent consumption expenditures.[1] Second, we remove observations with equivalized self-employment

[1] We set this threshold at 60,000$ for all categories except car-related transactions, where it is 600,000$. To be conservative, we also drop transactions in the "Insurance," "Financial services," and "Miscellaneous: Other services" categories.

Figure 1: Consumption in micro data vs. official national accounts



Notes. Left panel: level of (quarterly) per capita consumption in the micro-data vs. official statistics. All values are deflated by Norwegian CPI (2015 =1) and converted into US dollars using the exchange rate in 2015. Right panel: year-on-year (yoy) growth rates of the series. To aggregate the micro-data, we consider our baseline sample selection described below (we sum consumption across retained households and divide by the total number of individuals in those households).

earnings above $Y_{min,t}$, given that it may be difficult to separate actual consumption from business-related expenditures in electronic transactions for self-employed workers. Finally, following Heathcote et al. (2010), we drop households if no member is of working age (25–60 years).

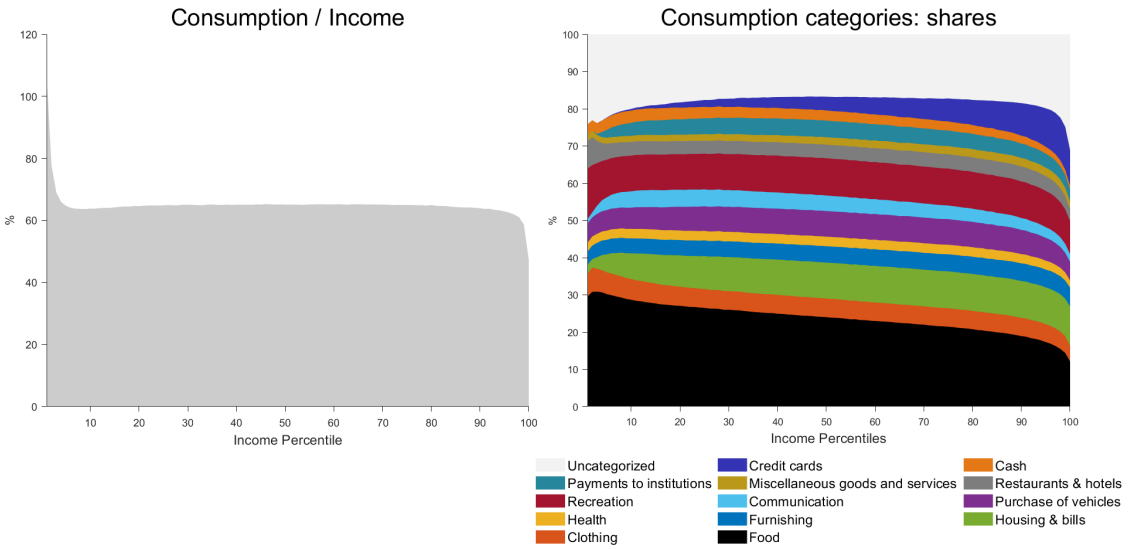
After imposing the age bound—which yields an average sample of 1.6 million households per year—the application of our restrictions gives a final average sample of 1.2 million households per year. Table A3 in the Appendix presents summary statistics for the main variables considered in the estimation sample, along with the role of the sample selection procedures. Overall, we view our sample restrictions as relatively mild compared with much of the existing literature.[2] This is deliberate and motivated by our aim to retain a significant share of the households at the bottom of the distributions.

2.3 A FIRST DESCRIPTIVE LOOK AT THE DATA

Figure 1 provides a smell-test of how electronic transactions fare as a measure of consumption. The left panel compares quarterly levels of per capita consumption in the national accounts with the counterpart when we aggregate up our household-specific consumption series. The right panel compares the quarterly year-on-year growth rates in both series. Overall, the micro data closely mirror the official consumption series, capturing both the levels——accounting for over 80% of total consumption——and the growth rates, with a correlation of nearly 0.6. The match with national accounts in our case is comparatively better than, for example, the widely used US Consumer Expenditure Survey (CES).

[2]For example, an alternative to our cutoff restrictions would be to exclude households with income and consumption below the Norwegian social-security minimum, as in Holm, Paul, and Tischbirek (2021). This minimum was approximately 11,000$ in 2015. Our baseline threshold, in contrast, implies a cutoff of 1,800$ for labor earnings and 1,300$ for consumption in 2015. We consider this alternative cutting in the Appendix.

Figure 2: Consumption and income shares



Notes. First panel: share of consumption as fraction of income, along the income distribution. Second panel: share of categories' consumption as a fraction of total consumption along the income distribution. Median shares within income and year, averaged across all years. We consider our baseline household level sample described below.

In the latter case, the (annual) correlation with the official statistics is 0.41 over the period 1994-2019, with a coverage of 72% of total per capita consumption (see Heathcote, Perri, Violante, and Zhang, 2023).

The left panel in Figure 2 reports sample averages in various consumption shares across the disposable income distribution. The first panel shows total consumption in our transactions data relative to disposable income. Consumption accounts for nearly all of disposable income among households in the bottom 10%. At the very top of the income distribution, in contrast, the consumption share drops sharply as income increases further. This is consistent with the conventional view that high-income households save a larger portion of their income. Except for these observations in the tails, the consumption-income ratio remains relatively constant throughout most income distribution. This suggests a fairly flat savings rate across most Norwegian households, consistent with the findings of Fagereng, Holm, Moll, and Natvik (2019). It is important to note, however, that the saving rate implied in our Figure 2 overestimates the true level. This is because (i) our data do not capture the universe of consumption transactions, and (ii) we exclude all mortgage-related transactions.

The right panel in Figure 2 reports the consumption shares in our data across the income distribution (see Figure A1 in the Appendix for each category shown separately). Food constitutes a substantial share of total consumption for lower-income households—around 30% at the very bottom. But this share declines with income, reaching about 12% at the top. Conversely, the shares for other categories are relatively stable across the income distribution. Finally, the mode of payment also differs systematically across the distribution: cash withdrawals at points of sale, while playing a very limited role overall (about 2% of total consumption), are more common for income-poor households. Credit

card usage in contrast increases substantially with income. Taken together, our preliminary inspections support the view that electronic transactions data capture consumption patterns relatively accurately.

3 INCOME AND CONSUMPTION INEQUALITY

This section documents a set of empirical facts about income and consumption inequality in our panel data. We first examine inequality in the raw data, and then quantify the contributions of life-cycle trends, demographic characteristics, and within-household volatility to overall cross-sectional dispersion. Finally, we conduct a series of robustness exercises that help to place our empirical findings in a broader context.

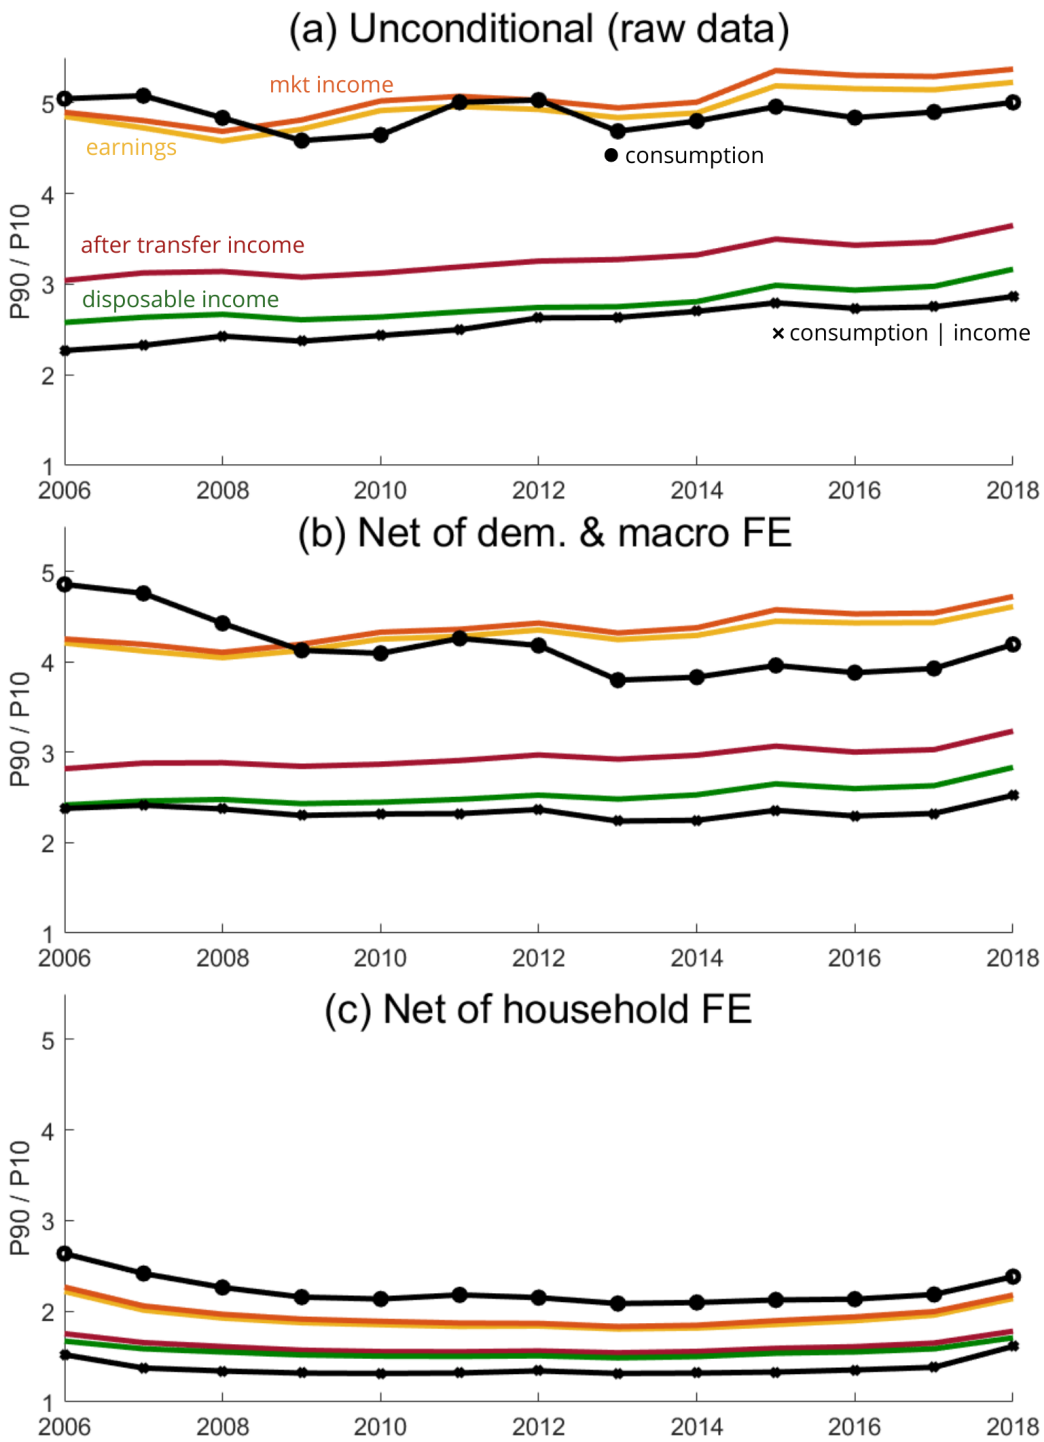
3.1 INEQUALITY, THE LIFE-CYCLE, AND HOUSEHOLD VOLATILITY

Figure 3 reports the levels of income and consumption inequality over time in our panel data. All inequality measures are computed in equivalized units at the household level. We separately report inequality in labor income, market income, income after transfers, disposable (after tax) income, as well as inequality in consumption expenditures from our transactions data (see table for A2 for detailed definitions). As a baseline, we define inequality as the ratio between the 90th and the 10th percentiles in the cross-sectional distribution at annual frequency. However, alternative definitions are provided in later sections and in the appendix, and none of the results are qualitatively affected. Panel (a) in Figure 3 considers measures of inequality in the observed data. Panel (b) reports inequality after we purge the data for between-group drivers such as demographic characteristics and time fixed effects. Panel (c) also controls for household-level fixed effects.

We first consider the *observed, unconditional* levels of inequality in Figure 3(a). The ranking across alternative income measures aligns with expectations: inequality in market income is slightly higher than in labor income, while income inequality after taxes and transfers—that is, inequality in disposable income—is substantially lower.[3] The first observation reflects a more unequal distribution of capital income relative to labor earnings,while the second results from redistribution towards relatively income-poor households. More strikingly, Figure 3(a) shows that *consumption inequality is substantially higher than disposable income inequality* in our data—around 60-100% higher when measured by the P90/P10-ratios. In fact, households in the top 90th percentile of the annual consumption distribution spend nearly 5 times more than those in the bottom 10th percentile. This implies that consumption is almost as unequally distributed across households as income *before* taxes and transfers. Importantly, this observation does not seem to be driven by income: if we rank households by their disposable income, the consumption gap between the 90th and 10th percentiles (black line with crosses) is even smaller than the corresponding income gap. Thus, consumption inequality *conditional on income* is indeed lower than income inequality. However, in the unconditional, observed data, many

[3]None of the inequality measures in our data display pronounced trends. There is a slight increase in these measures over the sample period, but this largely reflects that our sample begins in 2006—a historical low point for income inequality in Norway (Aaberge et al., 2021, Bennett and Salvanes, 2024, and Bjørnland, Chang, and Onshuus (2025)). Hence, unlike in the U.S., the trend dynamics of these measures are not especially informative.

Figure 3: Levels of income and consumption inequality in Norway



Notes. All measures of inequality are computed at the household level and in equivalized units. Income definitions: (1) labor earnings (yellow) = wages and salaries; (2) market income (red) = labor earnings + net capital income; (3) pre tax income (purple) = market income + public transfers; (4) after tax income (green) = pre tax income - taxes. Left panel: raw data in levels. Mid panel: data purged for variation attributed to demographic characteristics and aggregate (common) dynamics. Right panel: data purged for variation in demographic characteristics, aggregate dynamics, and household-level fixed effects.

income-poor households appear to spend relatively much in a given year, while many income-rich households spend relatively little. This latter observation provides an early indication of a disconnect between (current) income and consumption at the household level, a topic we examine in detail in the following sections.

How much of the excess cross-sectional dispersion in consumption reflects systematic household differences arising from life-cycle trends and other demographic characteristics? Panels (b) and (c) in Figure 3 address this question. There we consider the residual variation in income and consumption after removing an array of commonly considered predictable components. Specifically, panel (b) reports alternative inequality measures when we regress household-level income and consumption on a full set of interacting year–age fixed effects. These fixed effects allow us to control non-parametrically for all the common variation due to aggregate fluctuations and business cycle shocks, as well as for cohort or age-specific determinants that may vary by year. In addition, we include fixed effects for household size, the presence of children, and immigrant status. The inequality levels in panel (b) are computed from the residuals of those regressions, see the appendix for further details. While most of the inequality measures in panel (b) decline modestly relative to their counterparts in panel (a), and while part of the upward drift in income inequality disappears (and, if anything, a slight decline in consumption inequality emerges), both the ranking between the different income measures and the excess dispersion in consumption relative to disposable income remain salient features of the data. Consistent with this observation, we also find that life-cycle trends, demographics and macroeconomic factors jointly explain about 19% of the observed variation in household-level consumption expenditures. Thus, commonly considered consumption predictors cannot account for the elevated levels of consumption inequality observed in our panel data.

Figure 3(c) presents income and consumption inequality when we add household-level fixed effects to the set of controls. This latter step serves two purposes: first, it purges any contribution to observed inequality stemming from time-invariant household characteristics—such as permanent productivity differences or permanent differences in wealth. Second, by removing all between-household variation in the data, the household-level fixed effects enable us to isolate and quantify *within-household volatility* in income and consumption over time. Two important observations emerge from the inspection of panel (c): first, the ranking of the different inequality measures remains, including the elevated level of consumption inequality relative to income inequality. In fact, the contrast becomes even sharper, with consumption inequality now exceeding that of market income inequality. Second, although all inequality measures decline markedly relative to their counterparts in panels (a) and (b)—with the predictors now explaining about 68% of the variation in household consumption—ample cross-sectional dispersion remains. For example, households in the 90th percentile still spend more than twice as much as those in the 10th percentile, even after between-household variation has been purged away. Put differently, around 50% of the total observed consumption inequality remains even when we remove the permanent differences across households. Importantly, these observations reveal that within-household consumption volatility (i) exceeds within-household income volatility, and (ii) plays a major role for the overall level of consumption inequality in our data. They also indicate (iii) more mobility within the consumption distribution, as idiosyncratic innovations that move individual households around within each distribution

play a greater role for consumption than for income. In Appendix C we estimate so-called rank regressions, and confirm that this is indeed the case.

INEQUALITY AT DIFFERENT HORIZONS. Our focus so far has been on household inequality at annual frequencies. From an accounting point of view, the heightened level of consumption inequality relative to disposable income inequality that we document raises the issue of whether household-level budget constraints limit consumption choices in our data. To fix ideas, recall the accounting identity $Y_{i,\tau} = C_{i,\tau} + S_{i,\tau}$ where, for each household i over the period τ (τ can be a year, a decade, or the lifetime of the household), $Y_{i,\tau}$ denotes disposable income, $C_{i,\tau}$ consumption, and $S_{i,\tau}$ savings. It follows that cross-sectional household variation in consumption can only exceed the cross-sectional household variation in income if consumption and savings are sufficiently negatively correlated.[4] While this is indeed a possibility over short to medium horizons, we expect a closer link between income and consumption in the very long run (when τ becomes large). In Figure 4 we plot disposable income and consumption inequality at different horizons, ranging from one year (as in Figure 3) up to ten years. For each horizon, we take the sum of households' consumption and income over the entire period, before computing the ratio of percentiles. Two clear patterns emerge: first, as the aggregation period τ increases, the levels of income and consumption inequality decrease. This is intuitive: the role of within-household variation is strictly decreasing in τ, implying that only between-household variation can explain overall inequality levels at sufficiently long horizons. Second, the decline in consumption inequality at lower frequencies is significantly more pronounced: consumption inequality falls by about 40% when we go from annual to ten-year averages, while the corresponding decline in income inequality is only about 18%. The right panel in Figure 4 illustrates the consequence of this last observation: as the time horizon expands, the level of consumption inequality converges towards that of income inequality, suggesting that lifetime budget constraints start to kick in.[5]

3.2 HOW GENERAL ARE THESE RESULTS?

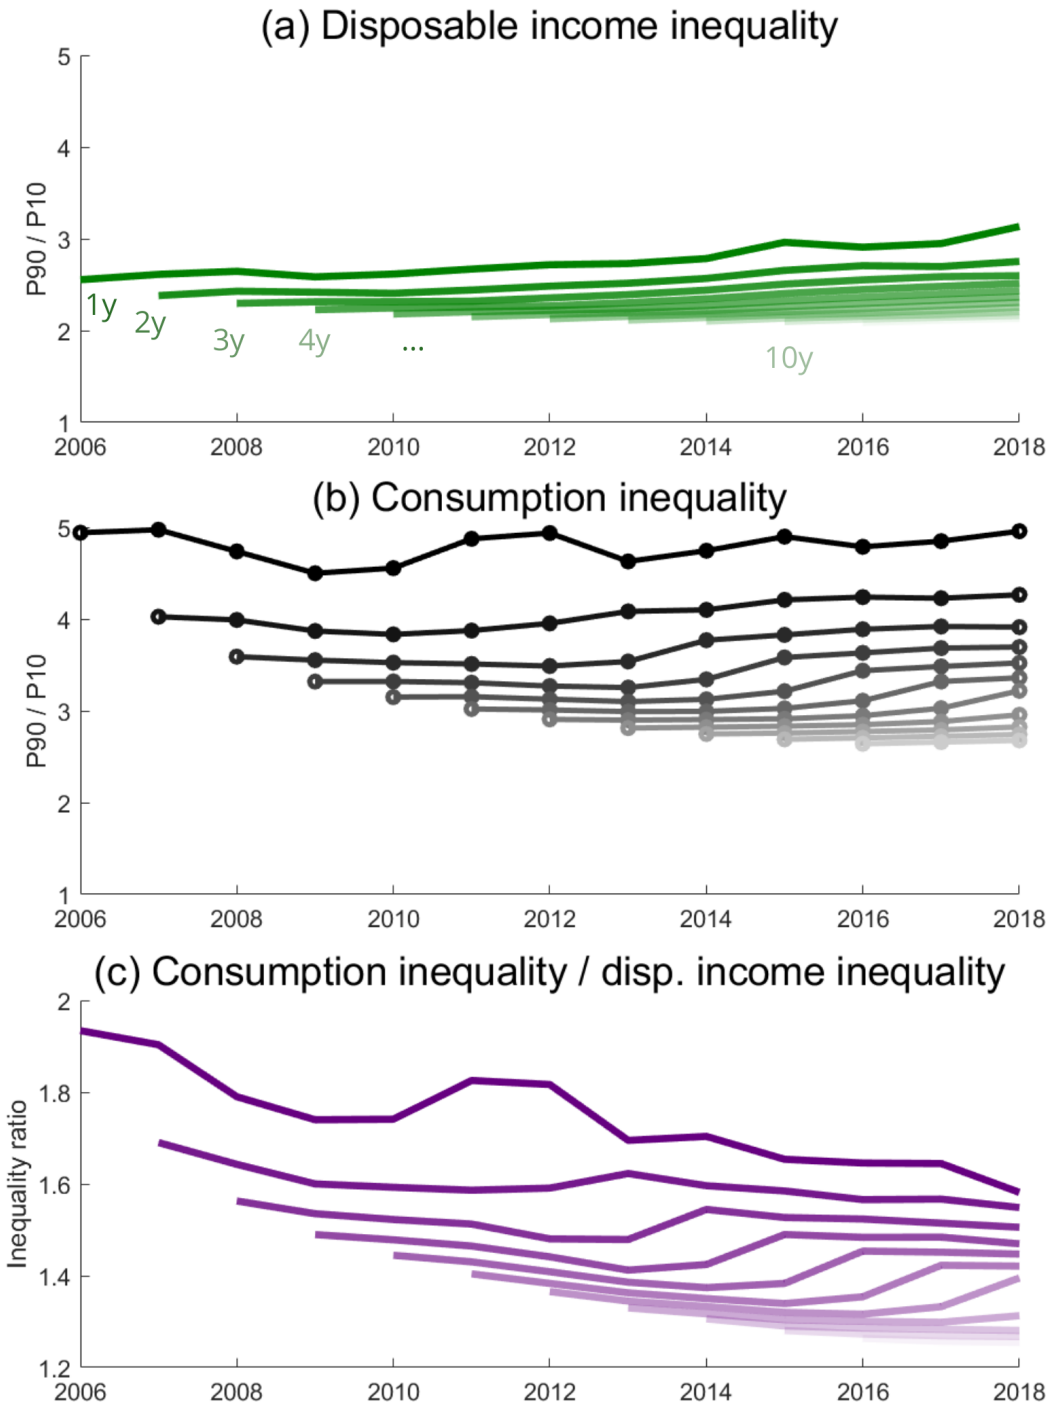
Next, we inspect whether the high level of consumption inequality in Figure 3 is confirmed when i) we isolate the role of lumpy expenditures, ii) we focus on various demographic groups, iii) we consider other data sets for Norway, and iv) we use alternative measures of inequality.

THE ROLE OF LUMPINESS. The level of consumption inequality is not driven by specific expenditure categories. In fact, a natural concern is that cross-sectional dispersion only reflects lumpy purchases of durable goods, such as vehicles. The left panel of Figure 5 shows that this channel accounts for part—but not all—of the observed inequality: excluding durable categories (vehicles; furnishings and household equipment; and credit-card payments) leads to a sizable decline in consumption inequality. However, even after removing these categories, consumption inequality remains above disposable income

[4] Note that $var\left(Y_{i,\tau}\right) = var\left(C_{i,\tau}\right) + var\left(S_{i,\tau}\right) + 2cov\left(C_{i,\tau}, S_{i,\tau}\right)$, so that excess consumption variability requires $var\left(S_{i,\tau}\right) + 2cov\left(C_{i,\tau}, S_{i,\tau}\right) < 0$.

[5] Quantitatively, the ratio of consumption to income inequality remains above one even at the ten-year horizon. A natural interpretation is that substantial wealth inequality—well documented for Norway (see the World Income Inequality Database)—continues to sustain dispersion in consumption levels even at longer horizons.

Figure 4: Income and consumption inequality at longer time horizons



Notes. First two panels report the inequality measures for consumption and disposable income when measured at different frequencies of time aggregation ranging one (darkest lines) to ten years (most transparent lines). Third plot: ratio of consumption / income inequality at different frequencies.

inequality. In Figure A5, we report inequality measures at the category level. Inequality is particularly large in categories typically associated with infrequent or lumpy purchases—such as vehicles, furniture, and housing-related expenditures—where spending differences across households are pronounced. By contrast, inequality in essential categories, such as food, while economically meaningful, is substantially flatter and broadly comparable in magnitude to disposable income inequality. These patterns suggest that spending on durable and discretionary goods varies much more strongly with overall consumption than spending on necessities, which changes more gradually across households.

INEQUALITY WITHIN HOUSEHOLDS GROUPS. Figure A6 documents that the high level of consumption inequality is widespread across various demographic groups (single households, households with kids, native households, young, middle-aged and older households). Consumption inequality is particularly elevated among singles (households composed by only one member) and younger households when compared with (aggregate) disposable income inequality. In addition, the dashed black line in Figure A6 shows that consumption inequality remains high when applying a higher minimum consumption threshold for sample selection (following Holm et al., 2021).

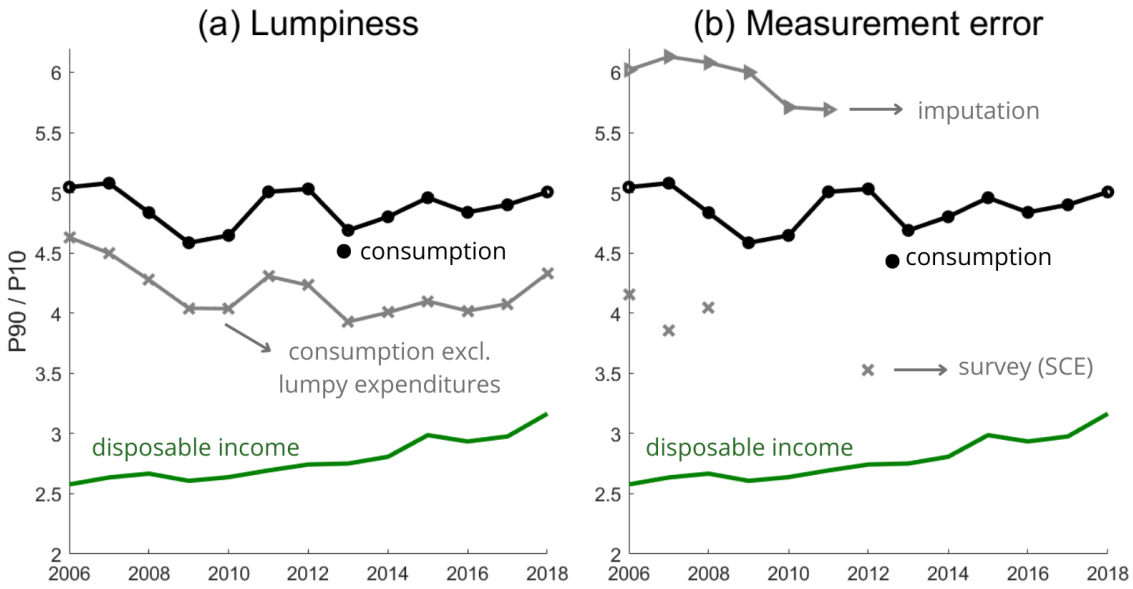
MEASUREMENT. The right panel in Figure 5 compares our baseline inequality measures using transactions data with counterparts using two alternative data sources: First, we plot the P90/P10 ratio over time when household-level consumption is imputed from registry data on income and net wealth, as in Fagereng and Halvorsen (2017). The idea is to infer consumption for household i in year t from the accounting identity $C_{i,t} = Y_{i,t} - S_{i,t}$, using observations on income $Y_{i,t}$ and changes in net wealth, $S_{i,t} = W_{i,t} - W_{i,t-1}$. The left panel in Figure 5 reveals that consumption inequality is even greater in the imputed consumption data than in our transactions data. Does this imply that the transactions data systematically under-report actual consumption differences? Not necessarily: one limitation with imputed consumption is that measurement errors in the wealth data—which we suspect can be relatively large—translate one-to-one into measurement errors in imputed consumption. Moreover, these measurement errors may be correlated at the household level with the household's position in income or wealth distributions (Baker et al., 2022). To the extent that registry data disproportionately understate wealth at the very top of the distribution, imputations of consumption based on wealth can naturally amplify inferred levels of consumption inequality.

Second, we compare our baseline with household-level consumption in the Norwegian Survey of Consumer Expenditures (SCE).[6] Consumption inequality in the survey data is lower than in our baseline, albeit still significantly higher than inequality in disposable income. A lower level of inequality in the survey data is consistent with the view that consumption expenditures among income-rich and wealthy households tend to be under-reported in survey data, thus, leading to a downward bias in the estimated level of aggregate consumption inequality (Aguiar and Bils, 2015).[7] In any case, all three data

[6]The SCE is managed by Statistics Norway and provides a detailed overview of household consumption by categories of goods and services. Data for the survey have been collected on an irregular basis since 1974. We use the variable "Total consumption expenditure" ("Samlet forbruksutgift"). No surveys were conducted after 2012 until the end of our sample period.

[7]The use of survey data also allows us to place Norwegian consumption inequality in an international perspective, as consumption inequality is typically measured using household surveys. Krueger, Perri, Pistaferri, and Violante (2010) summarize evidence for nine countries around the year 2000—well before the marked rise in inequality observed over the past two decades—and report an average cross-sectional

Figure 5: Alternative measures of consumption inequality



Note: Evolution of consumption inequality. First panel: our measure of inequality vs. those obtained from survey data and imputation techniques. Second panel: our measure of inequality vs. the one obtained by excluding lumpy categories of consumption. As a benchmark, we also report the evolution of disposable income inequality in both panels.

sources that provide us with annual observations of household-level consumption in Norway signal that consumption inequality is higher than disposable income inequality.

INEQUALITY IN THE FULL DISTRIBUTIONS. Finally, we stress that the high level of consumption inequality is a general feature of the data and not restricted to the P90 and P10 percentiles. Figure A2 in the appendix plots the full unconditional distributions of the log-levels of income and consumption, and the excess cross-sectional dispersion in household-level consumption is visually clear even there. It is also apparent in Figure A3, where we plot the full distributions of residualized income and consumption—that is, after we have purged away the variation driven by the full set of fixed effects described earlier. Finally, Figure A4 in the appendix reproduces Figure 3(a) but with the cross-sectional variances of log-levels instead of the P90/P10 ratios. Consumption inequality significantly exceeds that of income inequality in all cases.

4 INCOME AND EXPENDITURE SHOCKS

To better understand the high level of consumption inequality in our data, consider a variant of the canonical, permanent–transitory income decomposition of Hall and Mishkin (1982), which has been widely used to analyze the joint dynamics of income and consumption heterogeneity. In particular, for every household i in year t, the log income level $y_{i,t}$ is decomposed into a permanent component $\bar{y}_{i,t}$ and a mean-reverting transitory

consumption inequality (P90/P10) of approximately 5.2. For disposable income, the corresponding average is about 7. Taken together, these comparisons suggest that consumption inequality in Norway is relatively low by international standards.

component $\tilde{y}_{i,t}$. Income growth

$$\Delta y_{i,t} = \Delta \bar{y}_{i,t} + \Delta \tilde{y}_{i,t}$$

is stationary. In the spirit of Blundell et al. (2008) and subsequent literature, we suppose that consumption growth $\Delta c_{i,t}$ can be decomposed as follows:

$$\Delta c_{i,t} = \phi \Delta \bar{y}_{i,t} + \psi \Delta \tilde{y}_{i,t} + \gamma_{i,t}$$

The parameters ϕ and ψ govern the pass-through from permanent and transitory income to consumption, while the residual term $\gamma_{i,t}$ captures the variation in households' consumption orthogonal to income.

This formulation describes the relationship between household-level consumption and income in a flexible, reduced-form way: as noted by Blundell et al. (2008), the parametrization $\phi = \psi = 0$ is implied by complete market models with full insurance, $\phi = 1$ and ψ close to zero follows from standard models of the permanent income hypothesis, while $\phi = \psi = 1$ is implied by hand-to-mouth consumption, as under autarky. More generally, $0 < \phi < 1$ and $0 < \psi < 1$ represent intermediate scenarios with partial insurance against income risk.

The standard setting is one where permanent and transitory income components are orthogonal to each other, and we make the same assumption for $\gamma_{i,t}$. Thus, the variance in total income growth is simply the sum of the variances in permanent and transitory components, while the variance in consumption growth is a weighted sum of the two income components and the variance in $\gamma_{i,t}$. Given the income and consumption processes described above, it is straightforward to show that the *cross-sectional variance of consumption growth exceeds that of income growth* if and only if the following inequality holds:

$$var(\gamma_{i,t}) > \left(1 - \phi^2\right) var(\Delta \bar{y}_{i,t}) + \left(1 - \psi^2\right) var(\Delta \tilde{y}_{i,t})$$

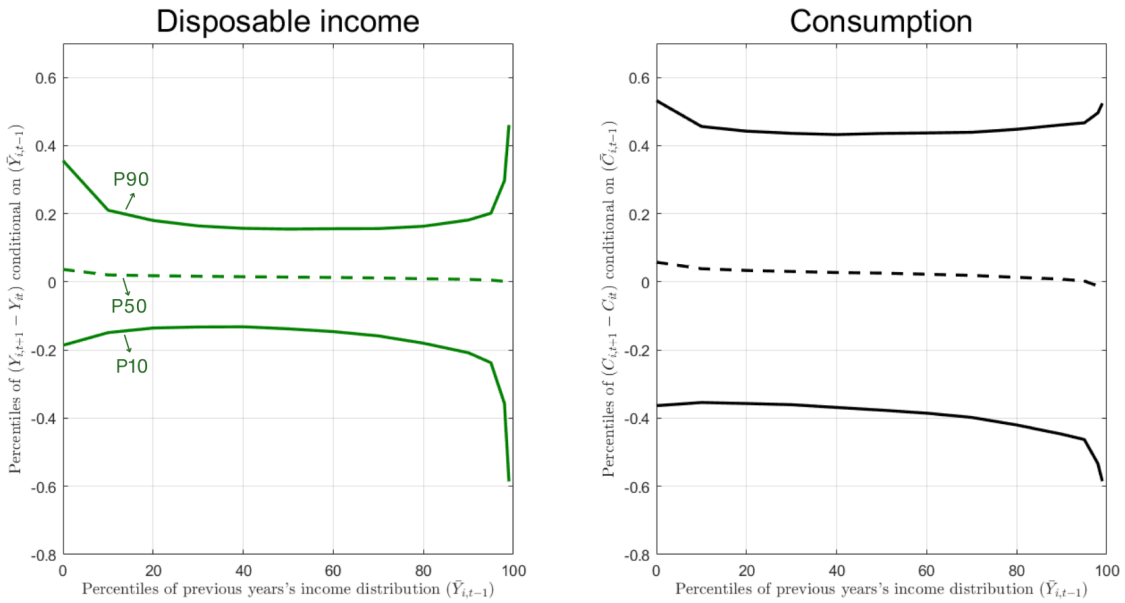
The literature typically focuses exclusively on income shocks as the sole drivers of consumption, thereby implicitly imposing the restriction that $var(\gamma_{i,t}) = 0$.[8] In that case, the presence of at least some insurance (i.e. not full autarky) implies that consumption must be less volatile—and hence its distribution more compressed—than income. By how much naturally depends on the values of ψ and ϕ, and on the composition of the income components $\Delta \bar{y}_{i,t}$ and $\Delta \tilde{y}_{i,t}$.[9]

This stark prediction, while being standard in the literature, abstracts from the role of $\gamma_{i,t}$. In this paper, we instead interpret fluctuations in $\gamma_{i,t}$ as consumption *expenditure shocks* and incorporate them in the analysis. The inequality constraint above illustrates that household-level dispersion in consumption growth exceeds household-level dispersion in income growth when (i) expenditure shocks are sufficiently volatile relative to

[8] More precisely, the common approach is to ignore any variation in $\gamma_{i,t}$ under the assumption that this variation is solely due to measurement errors.

[9] Consider the permanent income hypothesis as an example, i.e. the parametrization with $\phi = 1$ and ψ close to zero: in that case, income volatility will be substantially higher than consumption volatility if the transitory component $\Delta \tilde{y}_{i,t}$ dominates the variation in overall income. Variation in the permanent component $\Delta \bar{y}_{i,t}$, by contrast, results in similar levels of income and consumption volatility with this parametrization.

Figure 6: Income and consumption change volatility along the distribution



Notes. Percentiles of the disposable income, and consumption growth distribution, along the income distribution.

income shocks (high $var(\gamma_{i,t})$ relative to $var(\Delta\bar{y}_{i,t})$ and $var(\Delta\tilde{y}_{i,t})$), and (ii) insurance against income shocks is sufficiently muted (ϕ and ψ take large enough values). In either case, the presence of expenditure shocks provides a natural (and in simple income-consumption models such as the one above necessary) explanation for the high levels of consumption inequality observed in the data.

These expenditure shocks naturally operate as household-level liquidity shocks in the sense emphasized by Diamond and Dybvig (1983). When households face unanticipated timing and size of liquidity needs (medical bills, repairs, funerals, etc.), they demand liquid deposits or short-term safe instruments rather than illiquid, higher-return assets; banks then provide liquidity insurance by pooling idiosyncratic withdrawal risk.

4.1 A BIRD'S-EYE VIEW OF VOLATILITY

To better understand the sources of consumption dispersion, we group households by past income realizations, and then compute the distribution of *future outcomes within* each group. Taking disposable income as an example, let the income level in household i in year t be denoted by $Y_{i,t}$, and $\bar{Y}_{i,t-1} = \frac{1}{\tau}\sum_{s=1}^{\tau} Y_{i,t-s}$ represent average past income over some horizon τ. Conditional on each percentile of $\bar{Y}_{i,t-1}$, we build a distribution of future income and consumption growth, $\Delta y_{i,t+1} = \ln(Y_{i,t+1}) - \ln(Y_{i,t})$ and $\Delta c_{i,t+1} = \ln(C_{i,t+1}) - \ln(C_{i,t})$. We refer to future outcomes for a given percentile group as uncertain if the corresponding distribution for that group is disperse.

Figure 6 reports the results of this exercise. We condition on past averages over $\tau = 2$ years and plot the 90th, 50th, and 10th percentiles of each distribution. For disposable income, future income growth displays systematically greater uncertainty at the two ends of the distribution, compared with those in the middle 60%. Notably, households in the upper

tail face the largest income uncertainty.[10] Turning to consumption, the uncertainty about future consumption outcomes is substantially larger than for income, consistent with the theme in section 3 of excess dispersion in consumption relative to income. Moreover, unlike income, consumption uncertainty is distributed relatively homogeneously across the distribution and inherits only partially the shape of the income change distribution (both display higher uncertainty in the tails, even if this is more pronounced for income). Taken together, these facts suggest that changes in consumption cannot be fully accounted for by changes in income, particularly since households—at least partially—should be able to smooth income fluctuations. This points to the presence of additional forces shaping consumption dynamics, with fundamental shocks to expenditures emerging as a natural candidate in this respect.

4.2 Quantifying idiosyncratic expenditure shocks

We quantify the role of idiosyncratic expenditure shocks in three complementary ways: first, by estimating a linear regression model in the spirit of Hall (1978). Second, by extending the first specification to allow for permanent income shocks. And third, by estimating second moments in income and consumption using GMM, as in Blundell et al. (2008). In all three exercises we work with residualized income and consumption to purge out the variation due to life-cycle trends, demographics, and common macroeconomic factors. We refer to the appendix for details about the construction of the residualized panel data.[11]

The first empirical specification we consider is a simple panel regression linking consumption to income changes. The idea is to gauge how important household-level income fluctuations are for household-level consumption. In the spirit of Hall (1978), we project household consumption on the idiosyncratic, residual variation in contemporaneous income, as well as lagged, residualized consumption:

$$c_{i,t} = \mu + \rho c_{i,t-1} + \beta y_{i,t} + \gamma_{i,t}$$

The inclusion of lagged consumption implicitly controls for omitted, past drivers of consumption, such as the level of expected permanent income. We stress that since income and consumption data have been purged for systematic predictors such as life-cycle trends and other demographic characteristics, we base our inference on residual income variation that resembles income shocks, even if there is no attempt in this exercise to distinguish between permanent and transitory components. Nevertheless, if consumption largely depends on innovations to income, the regression should display strong predictive power for household-level consumption. Column (i) in Table 1 summarizes the results. Note, first, that the regression suggests significant pass-through from income to consumption; a one percentage point increase in income is associated with a 0.46 percentage point increase

[10] In Figure A7, we repeat this exercise for earnings (Guvenen et al., 2014) and find that the relative uncertainty at the tails flips: households at the bottom of the earnings distribution experience greater uncertainty than those at the top. This pattern suggests that redistribution—moving from earnings to disposable income—largely shields low-income households from extreme shocks, whereas high-income households remain exposed to additional risks stemming from capital income.

[11] Life-cycle trends, demographics, and common macroeconomic factors jointly explain about 68% of the observed variation in the *level* of household consumption expenditures, as reported in section 3.

Table 1: Decompositions of consumption growth

	(i) Hall regression simple	(ii) Hall regression extended	(iii) Blundell et al. model
	Dependent variable: consumption expenditures		
Passthrough y_t	0.46	x	x
Passthrough $\bar{y}_t$	x	0.51	0.79
Passthrough $\tilde{y}_t$	x	0.38	0.27
$var(\Delta c)$|income	14%	14%	9%
$var(\Delta c)$|exp. shocks	86%	86%	91%

Notes. Specifications: (i) $c_{i,t} = \mu + \rho c_{i,t-1} + \beta y_{i,t} + \gamma_{i,t}$. (ii) $c_{i,t} = \mu + \rho c_{i,t-1} + \bar{\beta}\bar{y}_{i,t} + \tilde{\beta}\tilde{y}_{i,t} + \gamma_{i,t}$. (iii) Moments from model with $\Delta c_{i,t} = \phi\eta_{i,t} + \psi\varepsilon_{i,t} + \gamma_{i,t}$, as described in the text.

in consumption. The remaining half is naturally absorbed by increased savings. Second, despite such a large pass-through elasticity, income explains only about 14% of the variation in consumption. The rest, about 86%, is attributed to idiosyncratic fluctuations in residualized consumption, i.e. to something that may resemble expenditure shocks.

However, a key limitation of the simple Hall regression above is that it ignores differences between permanent and transitory income changes. This may lead to biased pass-through estimates, as some of the income variation risks ending up in the consumption residual. Next we follow Friedman's proposal to proxy households' permanent income ($\bar{y}_{i,t}$) with their average income over a three-year period. Transitory income ($\tilde{y}_{i,t}$) is then defined as the deviation in current income from that average (the underlying idea is that liquidity constraints and uncertainty may prevent distant future income from affecting current choices, see Parker, Souleles, and Carroll, 2014). We then project consumption on its lag and on the contemporaneous values of permanent and transitory income:

$$c_{i,t} = \mu + \rho c_{i,t-1} + \bar{\beta}\bar{y}_{i,t} + \tilde{\beta}\tilde{y}_{i,t} + \gamma_{i,t}$$

Column (ii) in Table 1 reports pass-through coefficients and the variance decomposition in this case. Note first that we obtain a somewhat higher pass-through elasticity for average past income than for transitory current income—with point estimates of 0.51 versus 0.38—in line with the idea that a greater share of transitory income fluctuations is absorbed by savings rather than by consumption. Second, we still find that most of the variation in residualized consumption is attributed to its own shock.

Lastly, we estimate the second moments of income and consumption shocks, and quantify their relative importance in data, as in Blundell et al. (2008). In short, the model separately identifies: (i) a permanent income shock, modeled as a random walk; (ii) a transitory income shock, assumed to be an MA(1); and (iii) a transitory consumption (expenditure) shock, modeled as white noise. The income processes are summarized as follows:

$$\begin{aligned} y_{i,t} &= \bar{y}_{i,t} + \tilde{y}_{i,t} \\ \bar{y}_{i,t} &= \bar{y}_{i,t-1} + \eta_{i,t} \\ \tilde{y}_{i,t} &= \varepsilon_{i,t} + \theta\varepsilon_{i,t-1} \end{aligned}$$

Here, $\eta_{i,t}$ and $\varepsilon_{i,t}$ represent permanent and temporary income shocks, respectively. They

are assumed to be household-specific and drawn from distributions common across households, with time-varying variances denoted by $\sigma^2_{\eta,t}$ and $\sigma^2_{\varepsilon,t}$. The MA coefficient is assumed fixed. Collecting terms, we can write income growth as a function of the shocks:

$$\Delta y_{i,t} = \eta_{i,t} + \varepsilon_{i,t} - (1-\theta)\,\varepsilon_{i,t-1} - \theta\varepsilon_{i,t-2}$$

It is assumed that $cov(\eta_{i,t}, \eta_{i,s}) = cov(\varepsilon_{i,t}, \varepsilon_{i,s}) = cov(\eta_{i,t}, \varepsilon_{i,s}) = 0 \;\forall\; t$ and s. Thus, persistence in income growth over time comes solely from the MA structure. Turning to consumption, the following consumption equation is conjectured:

$$\Delta c_{i,t} = \phi\eta_{i,t} + \psi\varepsilon_{i,t} + \gamma_{i,t}$$

Here, we refer to $\gamma_{i,t}$ as a household-year specific consumption expenditure shock. Moreover, it is orthogonal to income changes, as we assume that $cov(\gamma_{i,t}, \gamma_{i,s}) = cov(\gamma_{i,t}, \eta_{i,s}) = cov(\gamma_{i,t}, \varepsilon_{i,s}) = 0 \;\forall\; t$ and s. The model is estimated by GMM, using the moment conditions described in the Appendix. Finally, equipped with the estimated parameters, we can compute the cross-sectional variance in consumption growth in year t as

$$\Omega_{\Delta c} = \phi^2\sigma^2_\eta + \psi^2\sigma^2_\varepsilon + \sigma^2_\gamma,$$

with the variance share attributed to expenditure shocks given by $\frac{\sigma^2_\gamma}{\Omega_{\Delta c}}$, for example. A variance decomposition of household-level income can be computed similarly.[12]

Column (iii) in Table 1 reports estimates when they are assumed fixed over time. Two main results stand out: first, the pass-through to consumption is substantially higher for permanent than for transitory income shocks—the estimated coefficients are $\hat{\phi} = 0.79$ and $\hat{\psi} = 0.27$ respectively —in line with the previous Hall-type regression.[13] Second, also in line with our previous exercises, consumption volatility is predominantly driven by expenditure shocks. They account for 91% of the dispersion in household-level consumption growth, while permanent income shocks account for about 8%. Transitory income shocks play a minimal role for consumption: they explain only 1%, despite accounting for more than 50% of income growth.[14] Such an important role for expenditure shocks may seem implausible at first sight. However, we reiterate here that we are decomposing the variation in *residualized* consumption, which accounts for only about 32% of the observed variation in the level of household consumption expenditures.

4.3 Interpretation of expenditure shocks

The robustness of the results presented above is analyzed Appendix F. It is shown that expenditure shocks play a major role also if we remove durable consumption goods, control for the use of credit cards, or restrict the analysis to food expenditures. The remainder

[12] Our baseline estimates are obtained from a constant-parameter model. In Appendix E, we allow coefficients to vary by year following Blundell et al. (2008), letting the pass-through elasticities (ϕ_t, ψ_t) and the cross-sectional variance of expenditure shocks ($\sigma^2_{\gamma,t}$) evolve over time. The results are robust along both quantitative and qualitative dimensions.

[13] Our estimates for Norway are in the ballpark of those in Eika et al. (2020) (using imputed consumption) and Fagereng, Holm, and Natvik (2021) (who exploit lottery prize data to identify transitory shocks). For the US, Blundell et al. (2008) estimate $\hat{\phi} = 0.64$ and $\hat{\psi} = 0.05$. However, Crawley (2020) shows that $\hat{\psi}$ increases to 0.24 when correcting for time aggregation.

[14] Interestingly, we find that the variance of expenditure shocks is significantly larger than that of income shocks. Specifically, we obtain $\sigma^2_\eta = 0.02$, $\sigma^2_\varepsilon = 0.02$, and $\sigma^2_\gamma = 0.16$ in the estimation with constant parameters.

of this section, instead, discusses possible interpretations of expenditure shocks. Given that measurement errors are likely to play a relatively limited role in our data, the residual "expenditure shock" term may arise from at least three distinct sources, each of which can be either unexpected or anticipated.[15] First, true consumption-expenditure shocks: unplanned, involuntary outlays such as emergency medical bills, urgent car or home repairs, or a funeral that requires spending even without a contemporaneous rise in income; or the planned timing of large expenditures (anticipated) such as a booked wedding or a scheduled renovation for which saving and timing decisions preceded the outlay. Second, preference or demand shocks: voluntary shifts in tastes or spending desires — for example an impulsive luxury purchase or a deliberate decision to take an expensive vacation — which may arrive without warning or may be premeditated and funded through past savings. Third, wealth shocks: changes in household wealth that affect consumption independently of current income — unexpected transfers or inheritances, or announced/planned transfers and realizations that are known in advance.

Disentangling the three interpretations goes beyond the scope of this paper. However, one simple exercise can isolate at least in part the role of planned expenditures. We revisit the household's period budget constraint:

$$W_{i,t} + C_{i,t} = E_{i,t} + (1 + r_{i,t}) W_{i,t-1} - \mathcal{T}_{i,t}$$

As before we let $C_{i,t}$ denote consumption, $E_{i,t}$ represent labor income $(1 + r_{i,t}) W_{i,t-1}$ represent the gross returns on the stock of wealth, and $\mathcal{T}_{i,t}$ be taxes net of transfers (we abstract from the role of capital gains for simplicity). Denoting disposable income by $Y_{i,t} = E_{i,t} + r_{i,t}W_{i,t-1} - \mathcal{T}_{i,t}$, the identify $Y_{i,t} = C_{i,t} + S_{i,t}$ follows, where $S_{i,t} = \Delta W_{i,t}$ represent household-level savings. Equivalently, the law of motion for wealth is

$$W_{i,t} = W_{i,t-1} + Y_{i,t} - C_{i,t}.$$

This simple identity highlights that any consumption shock *orthogonal to income* must, by construction, also appear as a shock to savings or wealth. A discretionary increase in consumption implies a one-for-one decline in savings, either through the decumulation of current assets or as a consequence of prior asset accumulation. It is the latter that we seek to assess as the main source of households' observed consumption volatility.

Panel D in Table F5 reports results when we redo the exercises in Table 1, taking past wealth into account. In specifications (i) and (ii) we proceed as follows: first, we residualize household-level consumption, income *and liquid wealth* following the procedures described earlier. We include the full set of controls, including household-level fixed effects. Second, we add lagged, residualized wealth to the Hall-style regressions (i) and (ii). In the Blundell et al. (2008) setup, i.e. specification (iii), we simply include lagged liquid wealth in the factor structure $F^c_{i,t}$ before proceeding with the GMM estimation. Thus, inference during the GMM estimation is based on consumption variability orthogonal to the variation in lagged wealth. Across all three specifications we make two observations that speak against a major role for planned expenditures: first, both the pass-through coefficients and the variance decompositions remain very similar to those in Table 1. Second,

[15] Since we cannot rule out an anticipated component, the label "expenditure shocks" used so far implies a slight abuse of language.

the pass-through coefficients governing the conditional correlation between lagged liquid wealth and consumption are low, about 0.03 in specifications (i) and (ii). A word of caution is needed at this point. These results do not imply that wealth does not matter for consumption. First, permanent differences in wealth across households do matter but are controlled for when we compute residualized consumption. Second, contemporaneous shocks to wealth (orthogonal to income) may affect consumption and be a component of the estimated expenditure shocks.

All-in-all, we believe that the robustness exercises presented here jointly help to build confidence about our residualized consumption variation as being closely related to unanticipated disturbances broadly defined as above. First, the major role of these shocks is found across consumption categories, e.g. both when we exclude durables and only consider food. This significantly limits the likelihood that volatility can be arise from lumpy expenditures, or from measurement errors. It seems unlikely that either of these explanations would occur systematically across consumption categories. Second, the last exercises above also provide evidence against the idea that residual variation mainly arises from expenditures planned in the past.

5 Inequality and the Business Cycle

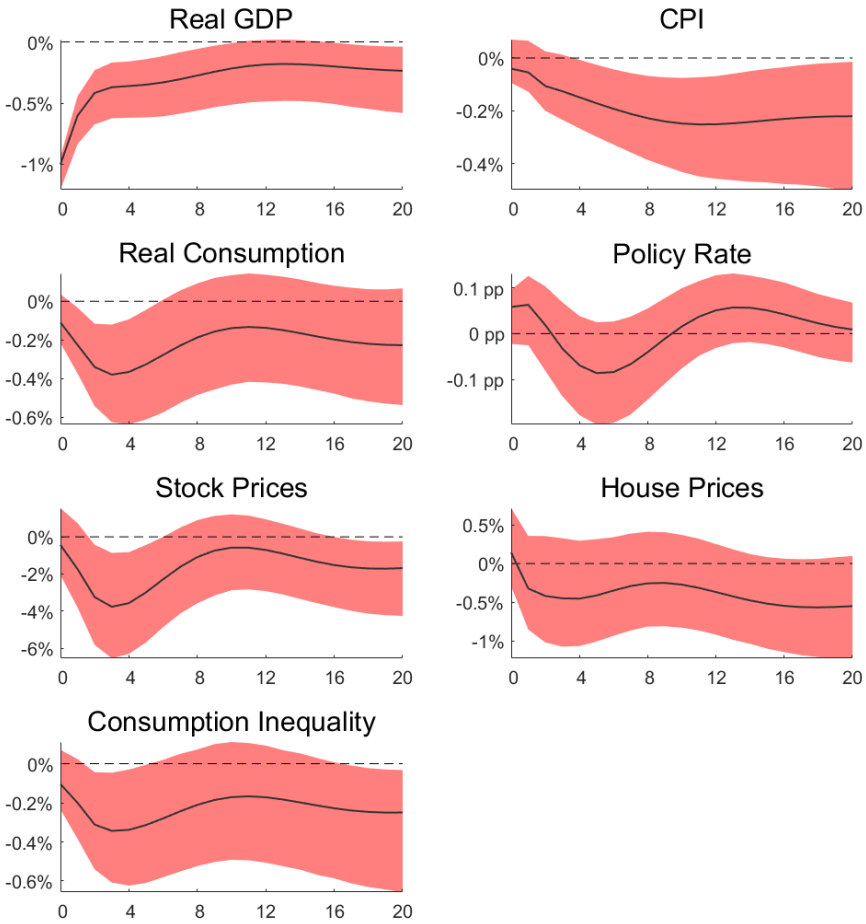
The analysis so far has documented two key forms of disconnection in household-level consumption. First, household consumption appears only weakly related to household income, despite sizeable marginal pass-through elasticities. Second, aggregate macroeconomic fluctuations account for a relatively small share of the variation in consumption across households—less than 20% according to the results in section 3. Since these findings pertain to allocations *within* the cross-sectional distribution at a given point in time, they naturally raise the question of how *the distributions themselves* evolve over the business cycle. Do distributional summary measures of income and consumption—such as P90-P10 ratios—systematically respond to aggregate shocks? And if they do, are these responses consistent with the conventional view of countercyclical inequality, as documented, for example, by Heathcote et al. (2010) for the U.S.? These are the questions we address in this section.

We proceed in two steps: first, we use Bayesian time series techniques to estimate a structural vector autoregressive (SVAR) model. This allows us to quantify the dynamic responses of selected distributional moments to identified business cycle shocks. Second, we perform a reduced-form event study of the Great Recession and calculate the heterogeneous exposure of households to the recession *along* the consumption and disposable income distributions, using a difference-in-differences approach.

5.1 Time Series Evidence

Our baseline SVAR model is estimated on quarterly time series data that include real GDP, the CPI, real consumption *as measured in the national accounts*, the policy interest rate, stock prices, and house prices. In addition to these key macro aggregates, we add selected distributional moments one-by-one and estimate the joint system. In the spirit of Angeletos et al. (2020), we identify a "main business cycle shock" as the disturbance

Figure 7: Impulse responses conditional on a recessionary business cycle shock



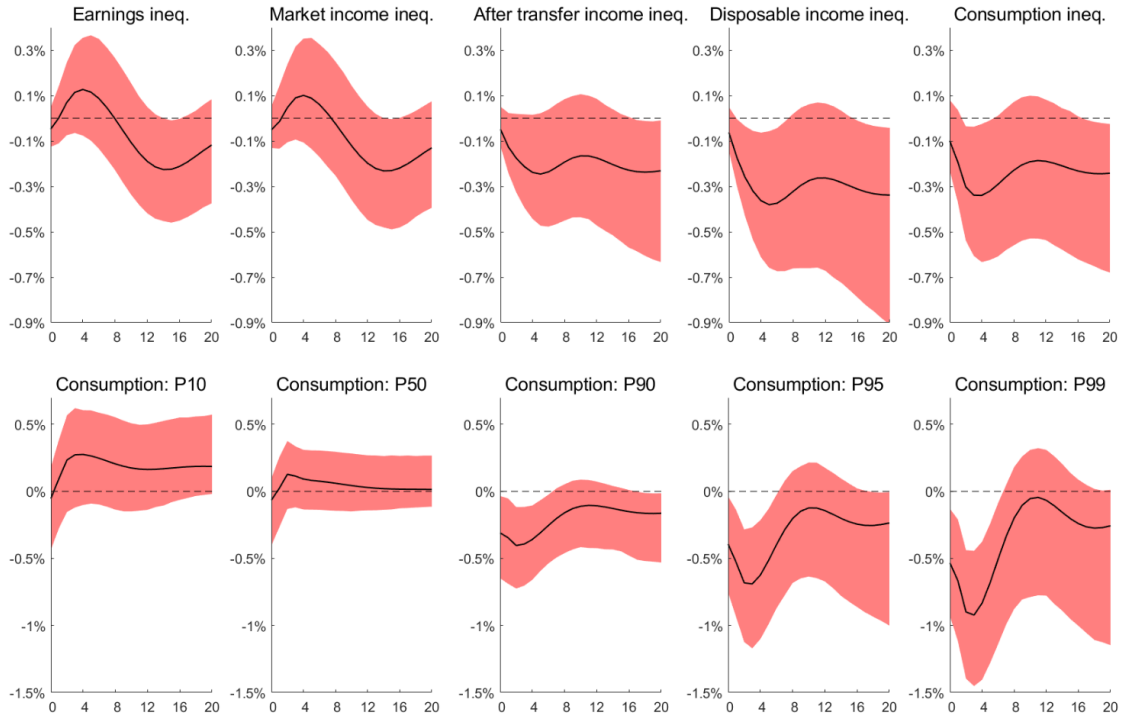
Notes. Impulse response functions for a "main business cycle shock" generating a 1% fall of GDP on impact. We report 80% posterior coverage bands. Horizons: quarters.

that accounts for the bulk of unexpected fluctuations in real GDP at business cycle frequencies.[16] We operationalize this by setting up a recursive VAR with real GDP ordered first among the observables. In this sense, the VAR exercises serve as a tool to extract the conditional cyclical correlation between macroeconomic aggregates and inequality. All variables are in log-levels, and the estimation sample covers the period 2006–2018. Given the short time span, we use two lags and rely on Bayesian estimation techniques to mitigate overfitting. Specifically, we adopt Normal-Inverse Wishart and Minnesota-type priors, with the tightness of the priors optimally chosen following Giannone, Lenza, and Primiceri (2015).

In the first exercise, we include the quarterly version of the P90–P10 ratio from the consumption distribution in the SVAR model. Impulse responses to a contractionary main business cycle shock are shown in Figure 7. Both aggregate quantities and prices decline—resembling a negative demand disturbance—while the interest rate response re-

[16] An alternative would be to rely on a more structural identification scheme, for example, by estimating the effects of monetary policy shocks (Coibion et al., 2017). However, such shocks are typically found to explain only a limited share of the business cycle, thus being less informative for our purposes.

Figure 8: Impulse responses across income components and consumption percentiles



Notes. Impulse response functions for a "main business cycle shock" generating a 1% fall of GDP on impact. We report 80% posterior coverage bands. Horizons: quarters.

mains muted (we are not concerned with zero lower bound issues, as the Norwegian policy rate never went below 0.5 percentage points during the sample period). Quantitatively, a 1% decline in output is accompanied by a 0.4% reduction in aggregate consumption from the national accounts. But most importantly for our purposes, the P90-P10 consumption ratio from our transactions data also decreases during the downturn, about 0.3% after two years, indicating that consumption inequality is procyclical.

To shed light on the sources of procyclical consumption inequality, we replace the P90-P10 consumption ratio with selected counterparts from different income distributions of interest (all captured by P90-P10), one at the time.[17] The top row in Figure 8 summarizes the results and compares with the response of consumption inequality. Several important findings stand out: first, earnings and market inequality are relatively acyclical or even mildly countercyclical in the short run. Both measures decline slightly after 2-3 years, but the responses are barely significant. The relatively similar responses of earnings and market income inequality suggest that disproportionate capital losses among wealthy households play a minor role for aggregate inequality dynamics. Second, inequality in income after taxes and transfers—disposable income inequality—shows an economically significant and quite persistent decline. About three years after the shock, the posterior median fall is about 0.35% per 1%-decline in GDP. Thus, disposable income inequality is indeed procyclical according to our estimates. Moreover, the response of disposable income inequality is quantitatively similar to the response of consumption inequality.

[17] To obtain quarterly measures of income inequality based on the annual registry data, we interpolate the annual data using a shape-preserving, piecewise cubic interpolation.

Finally, to understand *where* in the consumption distribution procyclicality comes from, we replace the P90-P10 consumption ratio with selected consumption percentiles, one at the time. The bottom row in Figure 8 reports the results: consumption responses at the bottom (10th percentile) and median (50th percentile) of the distribution are essentially acyclical or even mildly countercyclical. However, the consumption responses become stronger in magnitude and increasingly procyclical as we move towards the top of the distribution. The 90th percentile, for example, responds negatively and tracks the aggregate consumption response fairly well. Moreover, the 95th and 99th percentiles decline by about 0.7% and 1%—almost 2-3 times the fall in aggregate consumption from the national accounts. Thus, the procyclicality of consumption inequality arises from excess business cycle sensitivity at the top of the consumption distribution, while the bottom half remains relatively insulated. We stress that the large consumption responses at the top of the distribution can be quantitatively important also in the aggregate, as households at the top account for a disproportionately large share of aggregate consumption. In our data, the top 5% (1%) of the households account, on average, for 17% (6%) of aggregate consumption.

5.2 The 2008-2009 recession as an event study

While the SVAR analysis tracked business cycle dynamics in selected distributional moments, it did not account for compositional changes within distributions. Thus, it implicitly ignored the idea that any given percentile might be populated by different households over time. Next, we exploit the panel dimension of our data by tracing individual households throughout a business cycle downturn, the 2008-2009 Great Recession. We adopt primarily non-parametric methods, in the spirit of Guvenen et al. (2014) who studied earnings among males in the US. We extend their analysis to household-level disposable income and consumption, an area not yet explored in the literature. For comparability, the appendix replicates all exercises for the income variables, restricting the sample to prime-age males, and shows that the results are similar.

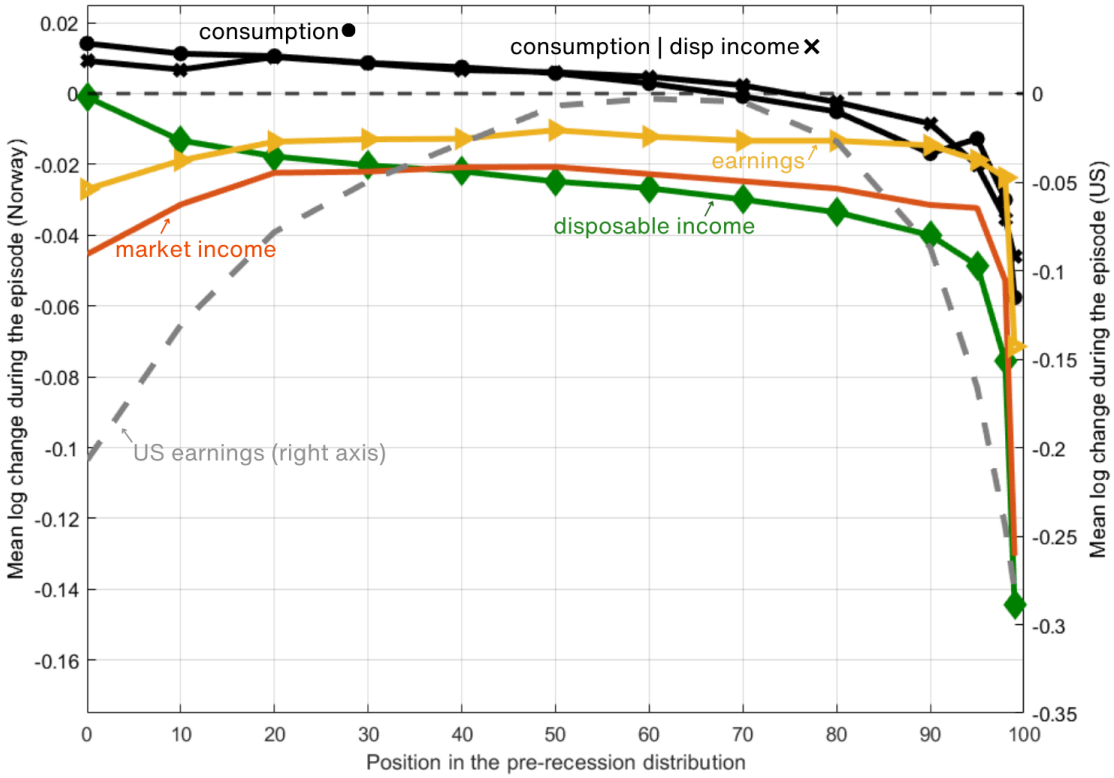
We proceed as follows: first, for each household we compute the average levels of earnings, market income, disposable income, and consumption over the pre-recession years 2006–2007. We then assign households to percentiles of each respective pre-recession distribution.[18] Based on these initial rankings, we follow households over time and compute the mean log change in income and consumption between 2008 and 2009 at each percentile level. This gives us a non-parametric reduced-form measure of how different segments of the distribution were affected by the recession. To account for possible pre-trend differences across the distributions (e.g., due to life-cycle effects), we benchmark our results against "normal times" and adopt a difference-in-differences strategy: we first compute the growth rates by percentile during the 2008–2009 recession, and then subtract the average percentile-specific growth rates observed in subsequent non-recession years.[19]

Figure 9 illustrates the results of this exercise (in Figure A8, we show that the re-

[18] We focus on the following percentile groups: 1-10, 10-20, ..., 80-90, 90-95, 95-98, 98-99, more than 99. See Guvenen, Schulhofer-Wohl, Song, and Yogo (2017) for a similar classification.

[19] To define "normal times", we use the years 2010–2015 and 2017–2018, excluding 2016 due to a temporary GDP contraction and a peak in unemployment.

Figure 9: Great Recession losses along the income and consumption distributions



Notes. Average log change along the distribution in labor earnings, income and consumption during recessions: Norway vs. US. For Norway: labor earnings (yellow line); market income (red); income after tax (green); consumption (black circled); consumption given income (black starred). Percentile of the pre-recession distribution calculated during 2006-2007. Change in log average earnings / income / consumption calculated over the years 2008-2009 (which is the only period during which Norway experienced a negative real annual GDP growth in our sample). For US: Change in log average earnings during all historical recession (average). Source: Guvenen et al. (2014) (Figure 15).

sults remain virtually unchanged when controlling for life-cycle trends and demographic characteristics). As a baseline reference we also include the percentile-specific earnings losses during an "average" recession in the US, as documented by Guvenen et al. (2014). They find that, relative to normal times, the earnings growth losses during recessions are relatively large both at the bottom and the top of the earnings distribution for prime-age males. The growth rates in the 10th and 90th percentiles, for example, fall by 10-13 percentage points, while males located in between the 50th and 70th percentiles experience only minor earnings losses.

By contrast, the household-level earnings losses in our Norwegian microdata (yellow) are remarkably stable across most of the pre-recession distribution, with only 1-3 percentage points declines among the bottom 80%.[20] The top 2% households experienced larger earnings losses, albeit even here we observe average losses less than 10 percentage points. Thus, overall there was little change in household-level earnings inequality during

[20]The contraction of earnings at the bottom of the distribution is more pronounced at the individual level (Figure A9). However, as we will see for household income, even individual disposable income at the bottom is largely shielded from the recession.

the recession.

Turning to market income (red), we note that the losses were somewhat larger across the distribution than for earnings, reflecting excess business cycle sensitivity in capital returns. However, the vast majority of households rely on labor income as the main income source, implying a gradient for market income that is relatively similar to that for earnings. The exception is found at the top 90th-95th percentiles and above, where capital income becomes a major income source.

Percentile-specific losses in disposable income (green) tell a different story: these losses were strictly increasing in households' pre-crisis income, ranging from nearly zero at the bottom to 12-13 percentage points at the very top. Thus, unlike earnings and market income, the disposable income gradient was negative throughout the distribution, implying that disposable income inequality was distinctly procyclical during the recession. Figure 9 is also informative about the role of redistribution, illustrated as the vertical differences between the green and red lines. Our general observation is a clear negative relationship between households' pre-recession income level and their net gain from redistribution during the recession.

Finally, turning to consumption (black dotted), we note that the consumption losses were substantially smaller than the disposable income losses in all percentiles, reflecting significant household-level consumption smoothing across the entire pre-recession distribution. Moreover, the consumption gradient is strictly negative and relatively similar to that of disposable income: while the bottom 80% households did not experience any decline in consumption, those at the top went through a 6 percentage points decline. Interestingly, these numbers are very similar if we instead consider consumption losses conditional on income (black crosses).

Overall, we conclude that consumption inequality in Norway is distinctly procyclical, both in an estimated time series model, and when we track individual households through the Great Recession. Importantly, the procyclicality in consumption inequality mirrors that of disposable income inequality, while market income is relatively acyclical. Moreover, the bulk of the income and consumption losses during a recession take place at the top of the distribution, while households in lower segments are relatively unaffected. Thus, procyclical inequality seems to be driven by systematic redistribution, which effectively insures households outside the top segments against bad macroeconomic events.

6 TAKING STOCK

Using high-quality panel data for Norwegian households, this paper documents a set of key empirical facts that are at odds with common views in the literature: we find more dispersion and higher volatility in household-level consumption than income, and distinct procyclical dynamics in both income and consumption inequality.

Are these results relevant beyond the institutional context of Norway and other Scandinavian countries? Based on the results presented here, one cannot conclude that consumption is generally more dispersed and volatile than income in most other countries, nor that income or consumption inequality is typically procyclical. For instance, the Scandinavian wage bargaining system is known to compress wage differences (Mogstad et al., 2025), while comprehensive social safety nets shield households from significant income risks. However, despite these institutional differences, we still believe that the mechanisms we

identify—such as substantial idiosyncratic expenditure volatility and the role of cyclical redistribution and insurance—should operate in other countries as well. Indeed, certain types of expenditure shocks may play an even more prominent role in countries like the U.S., where households can be exposed to relatively large and abrupt expenses in health or education. Moreover, redistributive policies have been shown to shape procyclical inequality dynamics in other countries and, more recently, even in the U.S., as emphasized by Heathcote et al. (2023) in the context of the COVID-19 recession. We therefore find it useful to outline and discuss broader implications of our findings that likely apply in a wider international context.

Motives for precautionary savings and insurance. The first conclusion we want to put forward concerns households' motives for precautionary savings and insurance. Most economic models of household behavior feature income risk—often linked to the increased likelihood of unemployment during downturns—as the key source of consumption volatility. The assumed goal of insurance in these models is therefore to insulate desired consumption plans from unwanted income shocks. We instead show that most household-level consumption or expenditure volatility in our data is largely uncorrelated with income, pointing to fundamental consumption or expenditure risk as a separate motivation for precautionary saving behavior (for additional sources of precautionary savings, see also Benhabib, Bisin, and Zhu (2011) and Savoia (2026)). In contemporaneous work, Guiso and Jappelli (2024) use a consumption expectations survey in Italy to elicit the probability distribution of future consumption growth and construct a measure of expected consumption risk at the individual level. They find that only one-third of precautionary savings act as a buffer against labor income shocks, while the remainder is driven by expenditure shocks. This result is qualitatively consistent with our findings, and we conclude that models aiming for empirical relevance should incorporate fundamental shocks to consumption as an explicit source of household-level uncertainty, as done in Briglia (2025).

A disconnect between micro risk and aggregate business cycles. A second observation from our analysis is that the key drivers of consumption volatility at the household level play only a minor role in shaping aggregate consumption dynamics. This is because the expenditure shocks we consider are idiosyncratic and largely average out when aggregated across households and time. Likewise, the main drivers of the aggregate business cycle explain relatively little of the consumption variation within households. This is evident, for example, from the panel regressions in section 4, where the common time fixed effects explain only a small share of household-level consumption variation. Thus, while macroeconomic forces are important for distributional *summary moments*—for instance, the P90–P10 ratio analyzed in section 5—they capture relatively little of the micro-level variation. In this sense, we document a disconnect between microeconomic risk and macroeconomic volatility, suggesting that precautionary saving motives are largely detached from the aggregate business cycle.

The role of policy. This brings us to the question of whether stabilization policy, too, is disconnected from allocations at the household level. At first glance, one might suspect that the micro-macro disconnect in the sources of volatility makes macroeconomic stabilization policy largely irrelevant for most households. We believe such a conclusion is premature. On the contrary, our findings in section 5—that inequality in labor and capital income before taxes and transfers is relatively acyclical, while inequality in

disposable income and consumption expenditures is distinctly procyclical—suggest that macroeconomic policy plays an important role through redistribution in favor of income- and consumption-poor households. While our analysis stops short of quantifying the importance of aggregate risk in counterfactual redistribution scenarios, it seems likely that the redistributive nature of taxes and transfers helps insulate households in the lower part of the distribution from macroeconomic risk. This may be a key reason why we observe a disconnect between micro risk and aggregate business-cycle volatility in our data.

BUSINESS CYCLE VOLATILITY. A related question is whether business cycle fluctuations are dampened or amplified by the pattern of income redistribution over the cycle. The cyclical distribution of income is a central ingredient in models with household heterogeneity. It was first featured in Galí, López-Salido, and Vallés (2007) and derived analytically in Bilbiie (2008) in models with two agents. It plays a key role also in HANK frameworks (see Auclert, Rognlie, and Straub (2025) for a review and Debortoli and Galí (2025) for a comparison).

The standard calibration of these models typically implies amplification of business cycles. In general, the degree of amplification of demand shocks—such as monetary or fiscal expansions—increases with the share of constrained households and depends on the redistribution mechanisms in place. Amplification is particularly strong when constrained agents—who tend to be concentrated at the lower end of the income distribution—are disproportionately exposed to aggregate fluctuations (see Patterson, 2023 and Auclert and Rognlie, 2018). In that case, inequality moves countercyclically. Yet, the empirical results presented in section 5 strongly suggest that redistribution in Norway leads to procyclical income and consumption inequality. This pattern implies dampened rather than amplified business cycle volatility: poorer households are relatively better off in recessions, and since these households typically hold fewer liquid assets and may have higher marginal propensities to consume (Fagereng et al., 2021), the cyclical nature of inequality documented in this paper acts as an automatic stabilizer. Our results also complement those of Bilbiie et al. (2025), who estimate the degree of amplification through MPC heterogeneity in Norway. Interestingly, they find no evidence that individuals with high MPCs have more cyclical disposable income or consumption. We, instead, find that both disposable income and consumption are more cyclical among income- or consumption-rich households, implying procyclical inequality. Moreover, as illustrated in Figure 8, the cyclical relationship between disposable income and consumption inequality is nearly one-to-one, suggesting that procyclical consumption inequality primarily reflects procyclical income inequality. Taken together, these results indicate that income levels—rather than MPCs—are associated with business cycle exposure. However, the inverse relationship between disposable income and exposure points to cyclical, income-based redistribution as an important mechanism for stabilizing the aggregate business cycle in Norway.[21]

[21] Bilbiie (2025) further distinguishes the cyclical inequality channel from a separate cyclical risk channel while Bilbiie, Primiceri, and Tambalotti (2023) quantify empirically the importance of the two channels in an estimated model for the US. We find that consumption risk is acyclical while income risk is only mildy countercyclical (see Figure in the Appendix). This suggests that cyclical risk is not an important source of business cycle amplification in Norway.

7 CONCLUDING REMARKS

This paper documents and analyzes several new empirical facts about the dispersion and volatility of household-level consumption dynamics. Our motivation is found in two fundamental premises underlying much of modern economics: (i) households rely on precautionary savings and other insurance mechanisms to smooth consumption in response to income shocks, and (ii) poorer households—who have access to fewer savings and face tighter borrowing constraints—are more exposed to abrupt economic events. These premises imply that consumption should be more equally distributed than income across households, less volatile than income within households over time, and more equally distributed in good times than in bad times.

The key contribution of this paper is to document that none of these patterns hold in high-quality registry and transaction data covering the universe of Norwegian households. On the contrary, we find greater cross-sectional dispersion and higher within-household volatility in consumption than in income, as well as distinctly procyclical dynamics in both income and consumption inequality. Importantly, these results remain even after controlling for life-cycle effects, household characteristics including household-level fixed effects, as well as the consumption dynamics that are correlated with past wealth. We also show that procyclical inequality arises only after accounting for taxes and transfers, driven by excess business-cycle sensitivity among households at the top.

Thus, to shed light on our findings, we quantify the role of permanent and transitory income shocks for household-level consumption dynamics. After controlling for life-cycle trends, demographics, and common macroeconomic drivers, our estimates indicate that these two income shocks together explain only around 10% of the cross-sectional variance in consumption growth. Such a minor role for income shocks is established despite significant marginal pass-through from income to consumption. Our findings stand in stark contrast to the common assumption that income risk is the dominant source of consumption volatility. The remaining 90% of the dispersion is attributed to idiosyncratic expenditure innovations, which we interpret as consumption or expenditure shocks. These shocks likely represent a major motive for precautionary saving and insurance, largely orthogonal to traditional income risk.

That said, the nature of these expenditure shocks remains an open question. On one hand, they may reflect forced or involuntary spending needs—such as unexpected medical costs, car repairs, or family-related expenses. They may even arise indirectly from frictions that are not observable. On the other hand, they may stem from fundamental preference heterogeneity, or "preference shocks", which lead to voluntary and desired spending choices such as vacations or durable purchases. While our finding that expenditure shocks are uncorrelated with past wealth helps to rule out many anticipated expenses, we do not observe whether these shocks are chosen by or imposed on households. The role of shocks to wealth should also be investigated. It is therefore essential to emphasize that we do not take a stance on their precise origin or welfare implications. Nor do we interpret high household-level consumption volatility as evidence of market inefficiency or insufficient risk sharing. These open questions are left for future research.

REFERENCES

Aaberge, R., M. Mogstad, O. L. Vestad, and A. Vestre (2021). Economic inequality in Norway in the 21st century. Statistics Norway Report 2021/33.

Aastveit, K. A., T. M. Fastbø, E. Granziera, K. S. Paulsen, and K. N. Torstensen (2024). Nowcasting Norwegian household consumption with debit card transaction data. *Journal of Applied Econometrics 14*(2), 207–242.

Aguiar, M. and M. Bils (2015). Has consumption inequality mirrored income inequality? *American Economic Review 105*(9), 2725–2756.

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

Angeletos, G.-M., F. Collard, and H. Dellas (2020). Business-cycle anatomy. *American Economic Review 110*(10), 3030–3070.

Attanasio, O., E. Battistin, and H. Ichimura (2007, 11). What really happened to consumption inequality in the united states? In *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*. University of Chicago Press.

Attanasio, O. and L. Pistaferri (2014). Consumption inequality over the last half century: Some evidence using the new PSID consumption measure. *American Economic Review 104*(5), 122–26.

Attanasio, O. P., E. Battistin, and A. Leicester (2006). From micro to macro, from poor to rich: Consumption and income in the UK and the US. Working Paper University College London.

Attanasio, O. P. and L. Pistaferri (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. and M. Rognlie (2018). Inequality and aggregate demand. NBER Working Paper 24280.

Auclert, A., M. Rognlie, and L. Straub (2025). Fiscal and monetary policy with heterogeneous agents. *Annual Review of Economics 17*.

Baker, S. R. and L. Kueng (2022). Household financial transaction data. *Annual Review of Economics 14*(1), 47–67.

Baker, S. R., L. Kueng, S. Meyer, and M. Pagel (2022). Consumption imputation errors in administrative data. *Review of Financial Studies 35*, 3021–3059.

Bayer, C., B. Born, and R. Luetticke (2024). Shocks, frictions, and inequality in US business cycles. *American Economic Review 114*(5), 1211–1247.

Benhabib, J., A. Bisin, and S. Zhu (2011). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica 79*(1), 123–157.

Bennett, P. and K. Salvanes (2024). Changing patterns of inequality in Norway: the roles of gender, education, immigration and unions. *Fiscal Studies 45*(2), 205–224.

Bilbiie, F., S. Galaasen, R. Gurkaynak, M. Maehlum, and K. Molnar (2025). HANKSSON. CEPR Discussion Paper 20090.

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

Bilbiie, F. O. (2025). Monetary policy and heterogeneity: An analytical framework. *Review of Economic Studies 92*(4), 2398–2436.

Bilbiie, F. O., G. Primiceri, and A. Tambalotti (2023). Inequality and business cycles. NBER Working Paper 31729.

Bjørnland, H., Y. Chang, and H. Onshuus (2025). Five facts about income inequality in a resource rich economy. Manuscript.

Blundell, R., M. Borella, J. Commault, and M. De Nardi (2024). Old age risks, consumption, and insurance. *American Economic Review 114*(2), 575–613.

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

Briglia, L. (2025). Idiosyncratic expenditure risk as a driver of hand-to-mouth transition dynamics. *Working paper available at: https://www.luigibriglia.com/*.

Buda, G., S. Hansen, T. Rodrigo, V. M. Carvalho, Á. Ortiz, and J. V. Rodríguez Mora (2023). National accounts in a world of naturally occurring data: A proof of concept for consumption. CEPR Discussion Paper 17519.

Campbell, J. R. and Z. Hercowitz (2019). Liquidity constraints of the middle class. *American Economic Journal: Economic Policy 11*(3), 130–155.

Chang, M. and F. Schorfheide (2024). On the effects of monetary policy shocks on earnings and consumption heterogeneity. NBER Working Paper 32166.

Coibion, O., Y. Gorodnichenko, and D. Koustas (2021). Consumption inequality and the frequency of purchases. *American Economic Journal: Macroeconomics 13*(4), 449–82.

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

Crawley, E. (2020). In search of lost time aggregation. *Economics Letters 189*.

Crawley, E., M. Holm, and H. Tretvoll (2025). A parsimonious model of idiosyncratic income. *International Economic Review, forthcoming*.

De Giorgi, G. and L. Gambetti (2017). Business cycle fluctuations and the distribution of consumption. *Review of Economic Dynamics 23*, 19–41.

Debortoli, D. and J. Galí (2024). Idiosyncratic income risk and aggregate fluctuations. *American Economic Journal: Macroeconomics 16*(4), 279–310.

Debortoli, D. and J. Galí (2025). Heterogeneity and aggregate fluctuations: Insights from tank models. *NBER Macroeconomics Annual 39*(1), 307–357.

Diamond, D. W. and P. H. Dybvig (1983). Bank runs, deposit insurance, and liquidity. *Journal of political economy 91*(3), 401–419.

Eika, L., M. Mogstad, and O. L. Vestad (2020). What can we learn about household consumption expenditure from data on income and assets? *Journal of Public Economics 189*, 104163.

Fagereng, A. and E. Halvorsen (2017). Imputing consumption from Norwegian income and wealth registry data. *Journal of Economic and Social Measurement 42*(1), 67–100.

Fagereng, A., M. B. Holm, B. Moll, and G. Natvik (2019). Saving behavior across the wealth distribution: The importance of capital gains. NBER Working Paper 26588.

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

Friedman, M. (1957). *Theory of the Consumption Function*. Princeton University Press.

Fulford, S. and D. Low (2024). Expense shocks matter. *Consumer Financial Protection Bureau Office of Research Working Paper No. 24-08*.

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

Giannone, D., M. Lenza, and G. Primiceri (2015). Prior selection for vector autoregressions. *Review of Economics and Statistics 97*(2), 412–435.

Guiso, L. and T. Jappelli (2024). Anatomy of consumption risk. CSEF Working Paper 732.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

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

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy 86*(6), 971–987.

Hall, R. E. and F. S. Mishkin (1982). The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica 50*, 461–481.

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

Heathcote, J., F. Perri, and G. L. Violante (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., F. Perri, G. L. Violante, and L. Zhang (2023). More unequal we stand? Inequality dynamics in the United States, 1967–2021. *Review of Economic Dynamics 50*, 235–266.

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

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

Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? Evidence and theory. *Review of Economic Studies 73*(1), 163–193.

Krueger, D., F. Perri, L. Pistaferri, and G. L. Violante (2010). Cross-sectional facts for macroeconomists. *Review of Economic Dynamics 13*(1), 1–14.

Mangiante, G. and P. Meichtry (2025). On the distributional effects of conventional monetary policy and forward guidance. Banque de France Working Paper 996.

Meyer, B. D. and J. X. Sullivan (2023). Consumption and income inequality in the US since the 1960s. *Journal of Political Economy 131*(2), 247–284.

Miranda-Pinto, J., D. Murphy, K. J. Walsh, and E. R. Young (2025). A model of expenditure shocks. *Journal of Monetary Economics 154*, 103807.

Mogstad, M., K. G. Salvanes, and G. Torsvik (2025). Income equality in the Nordic countries: Myths, facts, and lessons. *Journal of Economic Literature 63*(3), 791–839.

OECD (2013). OECD framework for statistics on the distribution of household income.

Parker, J. A., N. S. Souleles, and C. D. Carroll (2014). *The Benefits of Panel Data in Consumer Expenditure Surveys*, pp. 75–99. University of Chicago Press.

Parker, J. A. and A. Vissing-Jørgensen (2010). The increase in income cyclicality of high-income households and its relation to the rise in top income shares. *Brookings Papers on Economic Activity 41*(2), 1–70.

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

Primiceri, G. E. and T. van Rens (2009). Heterogeneous life-cycle profiles, income risk and consumption inequality. *Journal of Monetary Economics 56*(1), 20–39.

Savoia, E. (2026). Intertemporal self-insurance and excess sensitivity of affluent households. Sveriges Riksbank Working paper series 464.

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

Appendix

CONSUMPTION INEQUALITY, HOUSEHOLD RISKS, AND THE BUSINESS CYCLE

Drago Bergholt
Norges Bank

Francesco Furlanetto
Norges Bank

Lorenzo Mori
Bank of Italy

November 2025

A ADDITIONAL TABLES AND FIGURES

Table A1: Consumption categories in the dataset

Category	Description
01	Food and non-alcoholic beverages
02	Alcoholic beverages, tobacco and narcotics
03	Clothing and footwear
04	Housing, water, electricity, gas and other fuels
05	Furnishings, household equipment, and routine household maintenance
06	Health
07	Transport
071	Purchase of vehicles
072	Operation of personal transport equipment
073	Transport services
08	Communications
09	Recreation and culture
091	Audio-visual, photographic, and information processing equipment
092	Major durables for outdoor recreation
093	Other recreational items and equipment, gardens and pets
094	Recreational and cultural services
095	Newspapers, books, and stationery
10	Education
11	Restaurants and hotels
111	Restaurants
112	Hotels
12	Miscellaneous goods and services
121	Personal care
123	Personal effects
124	Social protection
125	Insurance
126	Financial services
127	Other services
13	Payments to banks[a]
14	Payments to public institutions[a]

Notes: Aggregation level for consumption category. The category numbers corresponds to the 1999 COICOP version. Category 13 and 14 are not part of the COICOP classification, and apply only to bank wire transfers made via NICS. We refer to Ahn et al. (2024) for details.

Table A2: Definition of income variables

Category	**Description**
Labor earnings	Employee income is the sum of cash wages and salaries, taxable payments in kind, sickness benefits, and parental benefits received during the calendar year.
Market income	Labor earnings plus capital income (the sum of interest received, share dividends received, realized capital gains or losses, and other property income received during the calendar year).
After-transfer income	Market income plus transfers (taxable and tax-free transfers received during the calendar year).
Disposable income	Sum of wages and salaries, income from self-employment, property income, and transfers received, minus total assessed taxes and negative transfers.

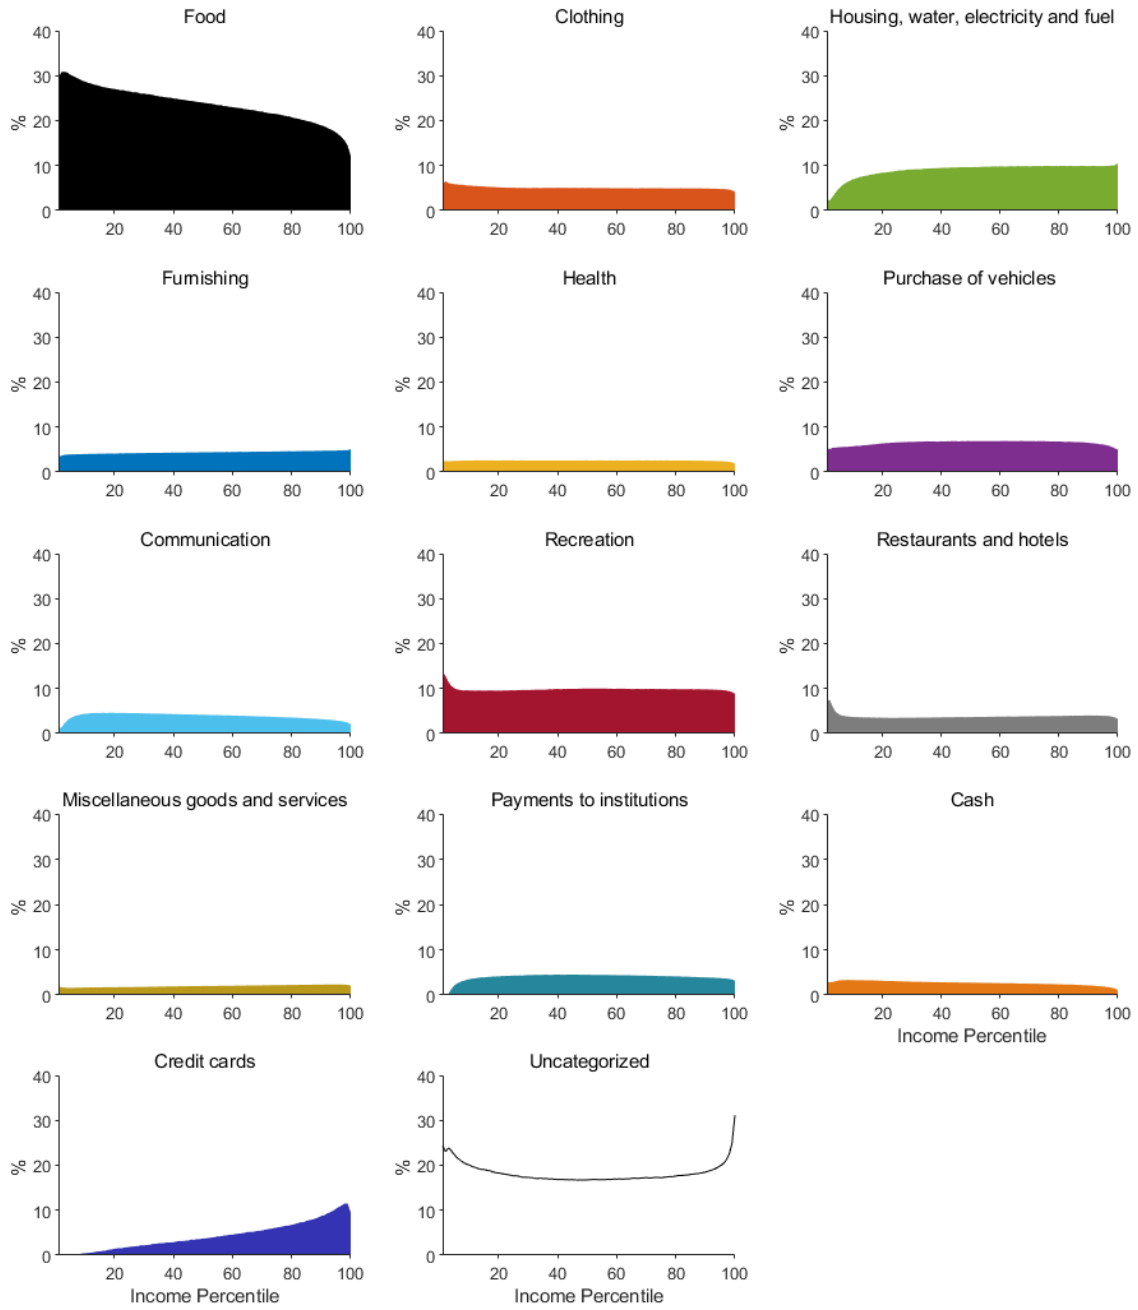
Notes: Income variables are defined by Statistics Norway.

Table A3: Summary Statistics

	Panel A: Full Sample				
Variable	Mean	SD	P10	P50	P90
Labor earnings	50 124	40 008	98	48 583	93 079
Market income	53 161	73 191	989	49 569	97 368
Pre tax income	61 333	71 382	22 884	55 475	101 005
Disposable income	47 448	62 475	22 427	43 865	73 410
Consumption	34 746	40 247	6 184	28 364	63 672
Age (head)	45	12	29	44	60
	Panel B: Restricted Sample				
Variable	Mean	SD	P10	P50	P90
Labor income	59 313	37 889	20 927	55 338	98 392
Market income	62 579	73 223	21 493	56 332	103 102
Pre tax income	68 803	72 217	33 195	61 305	106 599
Disposable income	50 370	54 080	27 141	46 203	74 438
Consumption	34 796	32 729	12 206	29 893	59 759
Age (head)	44	12	29	44	60

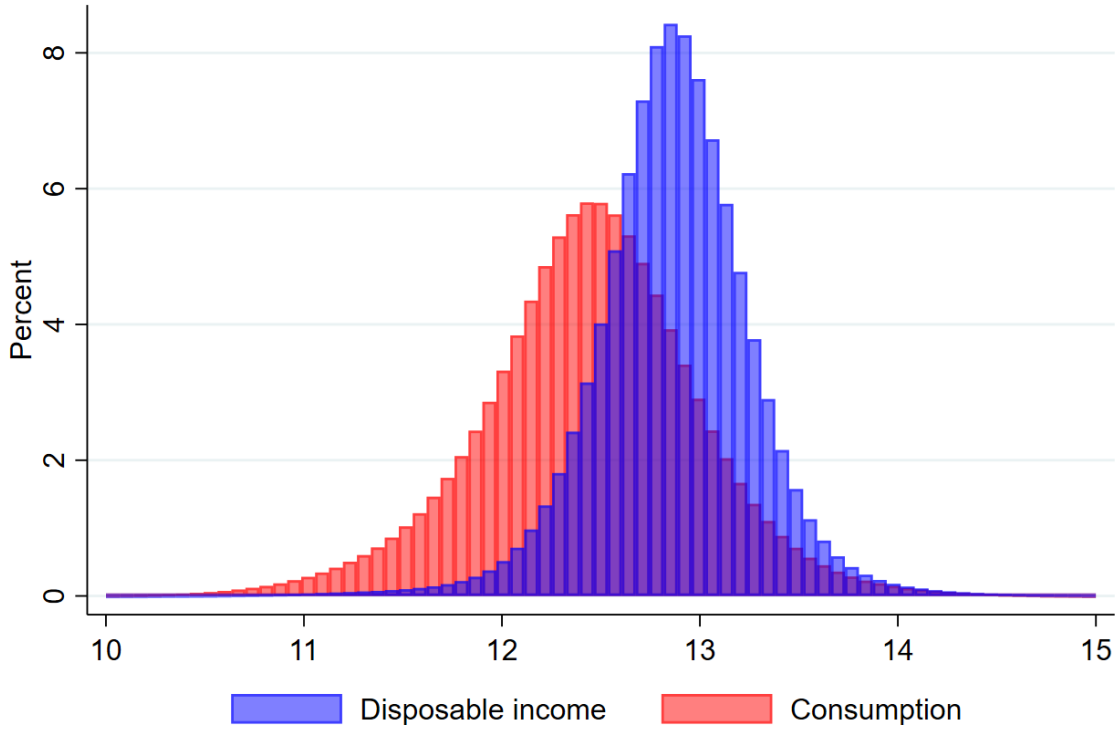
Notes. The table shows summary statistics for the estimation sample (2006-2018), households satisfying the age restriction. Panel A: unrestricted sample. Panel B: restriced sample. All values except age are inUSdollars (using the 2015 exchange rate), 2015 prices. Income and consumption variables are equivalized taking into account the dimension of the household.

Figure A1



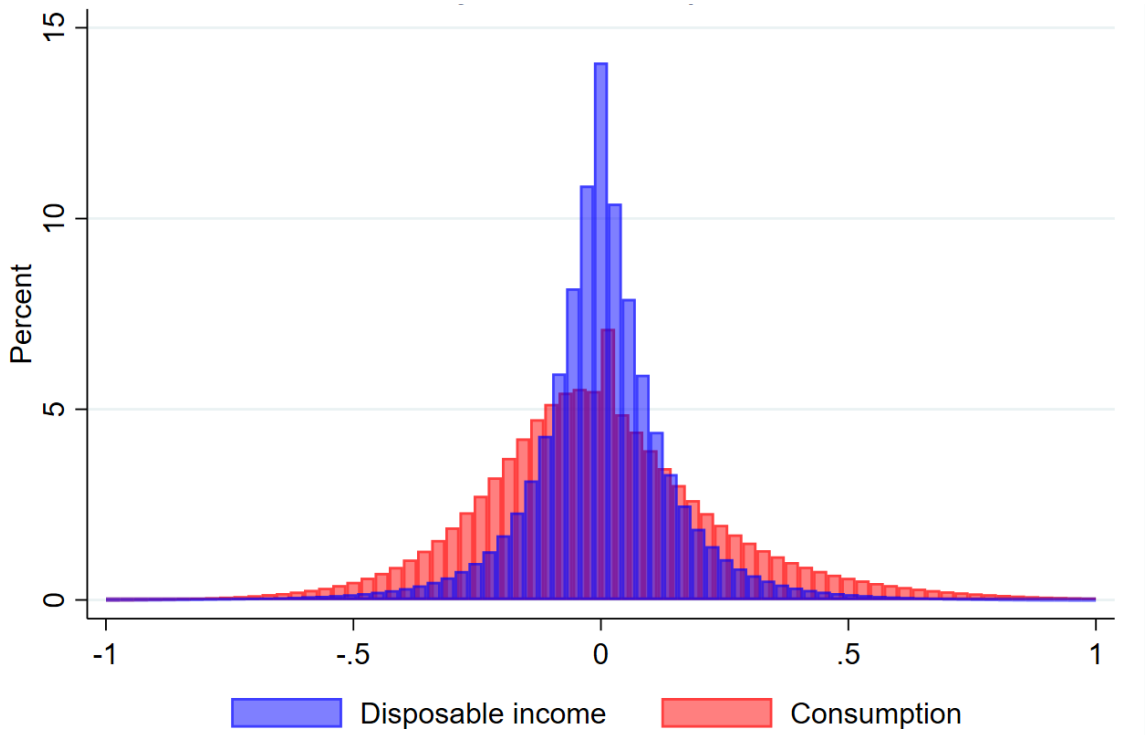
Notes. Share of categories' consumption as a fraction of total consumption along the income distribution. Median shares within income and year, averaged across all years. We consider our baseline household level sample described in the main text.

Figure A2: Unconditional income and consumption distributions



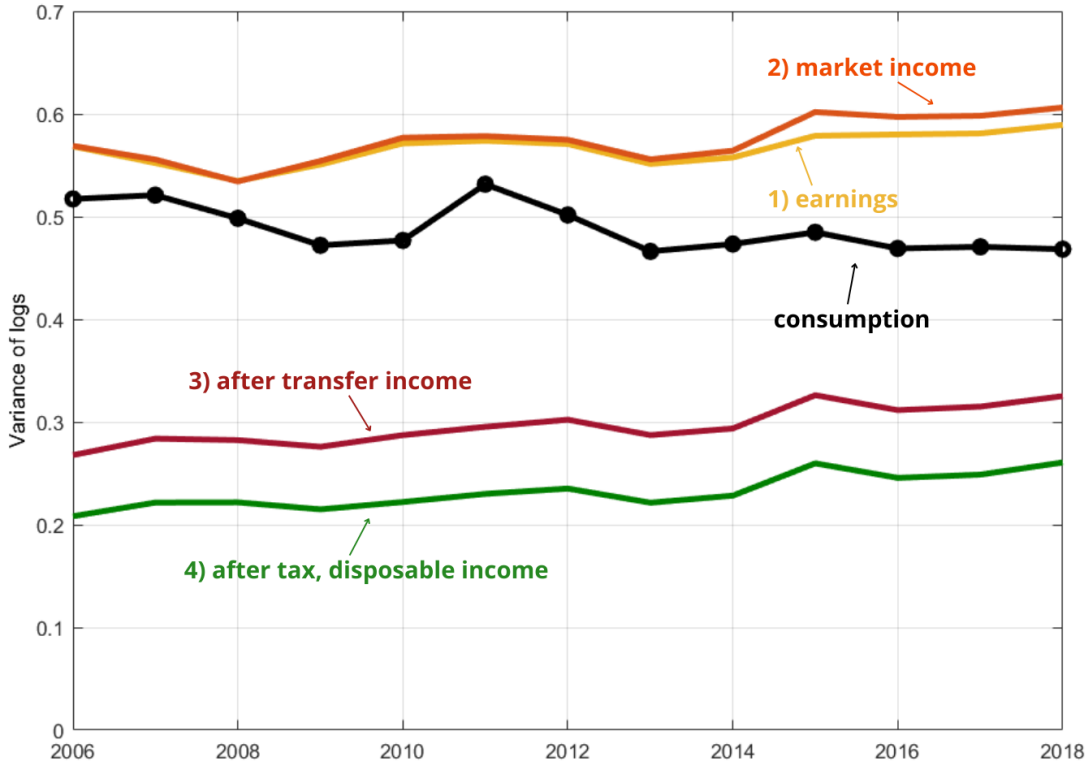
Note: Equivalized household level distributions of disposable income and consumption. Both measures are expressed in real terms (deflated by the CPI).

Figure A3: Residualized income and consumption distributions



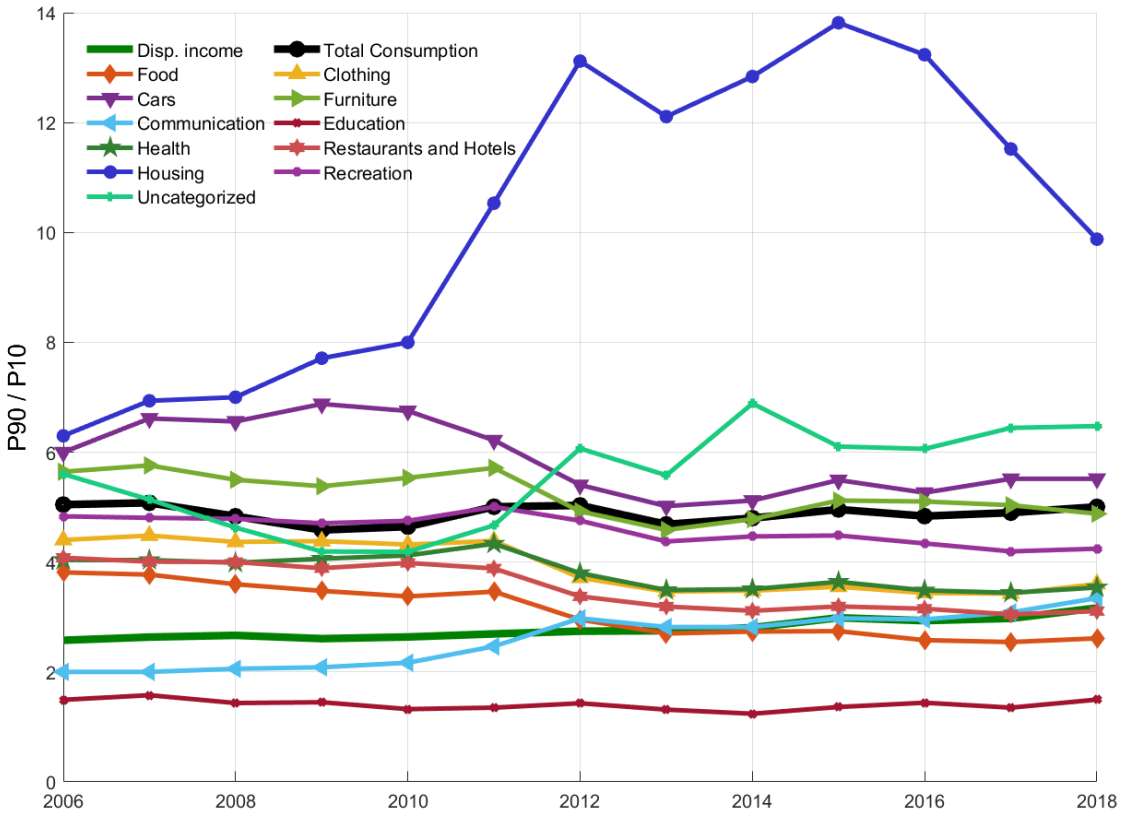
Note: Equivalized household level distributions of residualized disposable income and consumption. Both measures are expressed in real terms (deflated by the CPI).

Figure A4: Levels of income and consumption inequality in Norway



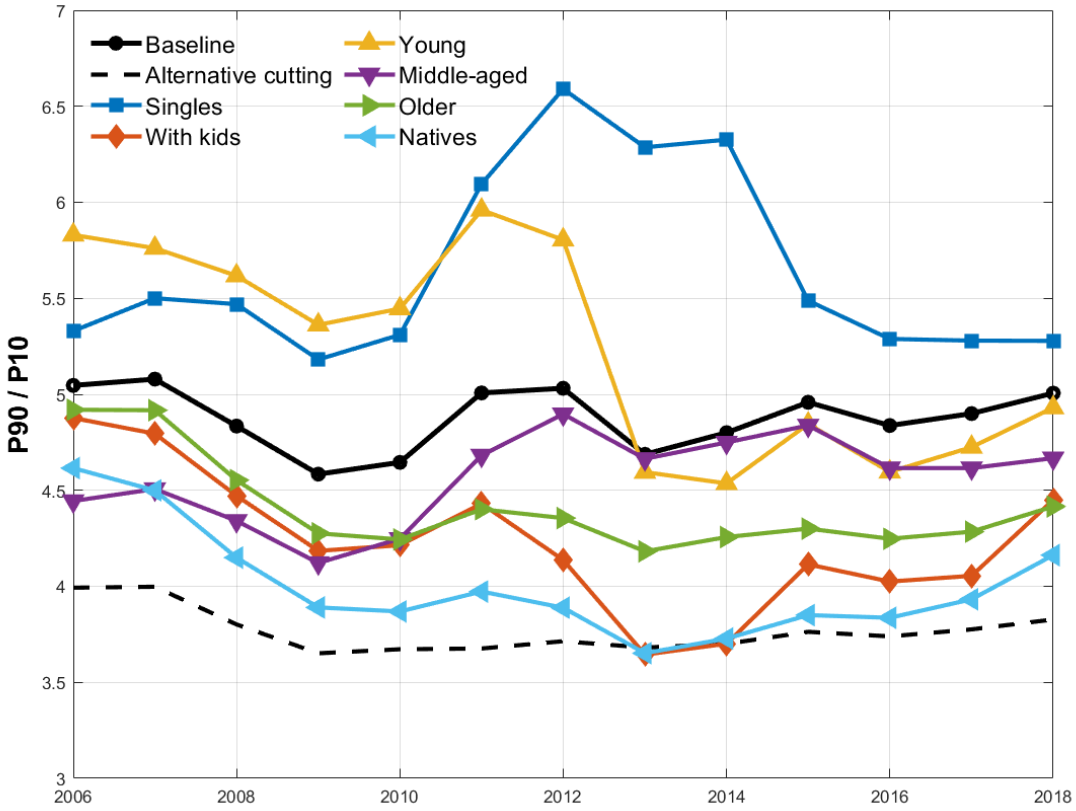
Notes. We consider the the variance of the log of the distribution as a measure of inequality. Definitions: labor earnings = wages and salaries; market income = labor earnings + net capital income; pre tax income = market income + public transfers; after tax income = pre tax income - taxes.

Figure A5: Consumption inequality across expenditure categories



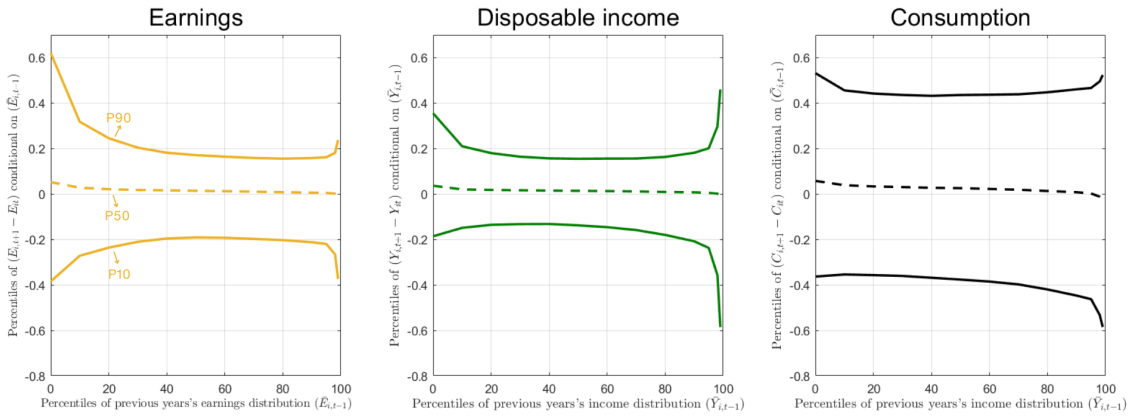
Notes. Consumption inequality in different categories of expenditures. For each year, households are ranked by total consumption. We define two narrow windows around the 10th and 90th percentiles of the total consumption distribution and compute, for each category, the ratio of median expenditure in the upper window relative to the lower window. Conditioning on the rank in total consumption allows us to mitigate distortions arising from zero expenditures, lumpy purchases, and measurement error that can mechanically inflate unconditional inequality statistics at the category level.

Figure A6: Consumption inequality within groups



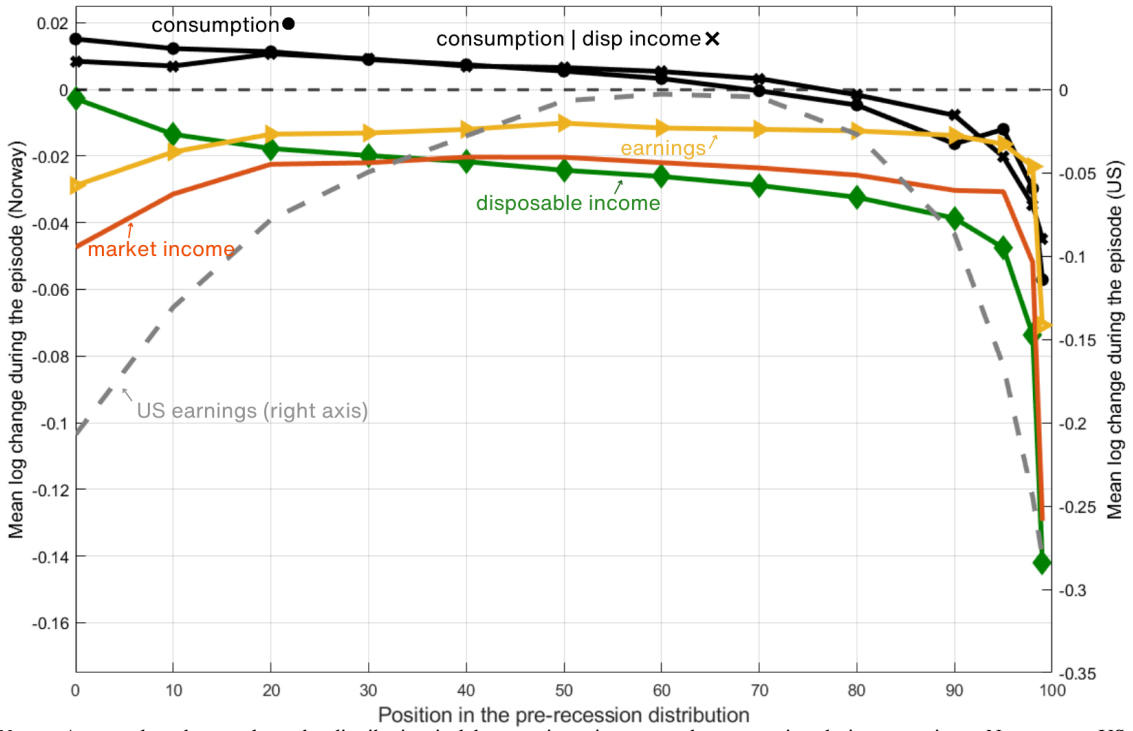
Notes. Consumption inequality under an alternative cutting and across population subgroups. Alternative cutting: rather than using the baseline cutoff for minimum annual household consumption, $C_{min,t} = 5\%$ (equal to \$1,800 for labor earnings and \$1,300 for consumption in 2015), we adopt the Norwegian social-security minimum (as in Holm et al., 2021) – approximately \$11,000 in 2015. Age groups are based on the age of the household head: young (below 30), middle-aged (30–40), and older (above 40). We classify as "natives" households those with all members born in Norway.

Figure A7: Income and consumption change volatility along the distribution



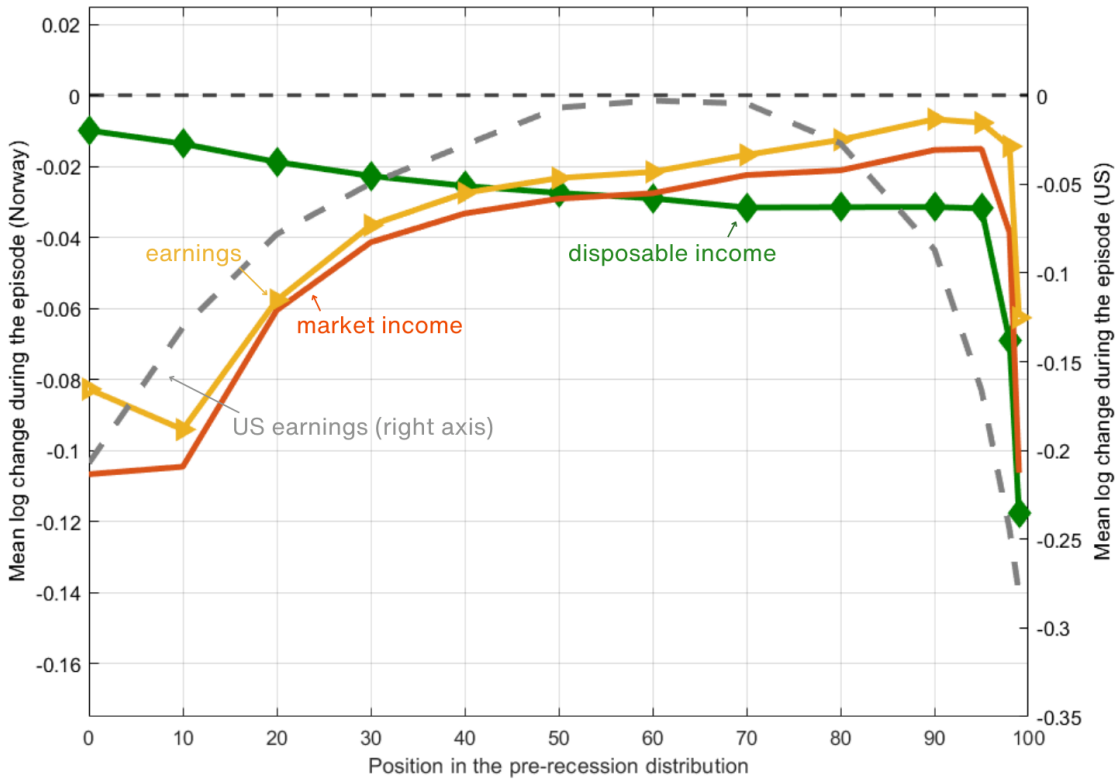
Notes. Percentiles of the earnings, disposable income, and consumption growth distribution, along the distribution.

Figure A8: Great Recession losses along the income and consumption distributions: residualized variables



Notes. Average log change along the distribution in labor earnings, income and consumption during recessions: Norway vs. US. For Norway: labor earnings (yellow line); income after tax (green); consumption (black circled); consumption given income (black starred). Percentile of the pre-recession distribution calculated during 2006-2007. Change in log average earnings / income / consumption calculated over the years 2008-2009 (which is the only period during which Norway experienced a negative real annual GDP growth in our sample). Variables are first cleaned of deterministic and demographic components, following the procedure described in Section 4.2 and Appendix B. Since we focus here on common shocks, variables are not residualized with respect to time fixed effects. For US: Change in log average earnings during all historical recession (average). Source: Guvenen et al. (2014) (Figure 15).

Figure A9: Great Recession losses along the income distribution: sample of prime age males



Notes. Individual level sample as in Guvenen et al. (2014): prime-age (35-54 years) males only. Average log change along the distribution in labor earnings, income and consumption during recessions: Norway vs. US. For Norway: labor earnings (yellow line); income after tax (green). Percentile of the pre-recession distribution calculated during 2006-2007. Change in log average earnings / income calculated over the years 2008-2009 (which is the only period during which Norway experienced a negative real annual GDP growth in our sample). For US: Change in log average earnings during all historical recession (average). Source: Guvenen et al. (2014) (Figure 15).

B RESIDUALIZED INCOME AND CONSUMPTION

Here we provide more details about the construction of residualized income and consumption at the household-year level. Let residualized income and consumption be given by:

$$y_{i,t} = \ln(Y_{i,t}) - F^y_{i,t} \qquad\qquad c_{i,t} = \ln(C_{i,t}) - F^c_{i,t}$$

$F^y_{i,t}$ and $F^y_{i,t}$ are non-parametric factor structures that purge out the systematic variation in data. They may capture life-cycle trends or cohort effects, family characteristics, etc.:

$$F^y_{i,t} = \sum_{j=1}^{J} \xi^y_{j,i,t} \qquad\qquad F^c_{i,t} = \sum_{j=1}^{J} \xi^c_{j,i,t}$$

We include J factors, modeled as a flexible system of fixed effects: first, we include individual fixed effects to control for all time-invariant drivers of income and consumption, such as permanent productivity differences at the individual and household level. Second, a full set of interacting year-age specific fixed effects allows us to control for (i) all the common variation in household-level income due to macroeconomic fluctuations and business cycle shocks, and (ii) cohort-specific determinants of income that may vary over time. Third, we include fixed effects for households with kids, as these may affect both the composition and the per capita level of household's consumption basket. For similar reasons, we include household size fixed effects. Finally, we purge away the fixed effect for households with immigrants.

C RANK REGRESSIONS

To quantify the degree of household-level mobility within a distribution over time, we estimate rank mobility as captured by changes in households' *rank position* in the income and consumption distributions. For both income and consumption, we run the following panel regression:

$$r_{i,t} = \alpha_i + \rho r_{i,t-1} + \sum_{k=1}^{K} \beta_k x_{ik,t} + \varepsilon_{i,t}$$

The variable $r_{i,t}$ represents *the distribution rank* for household i in year t, while $\{x_{ik,t}\}$ represents various control variables which may be systematically correlated with rank. We consider two statistics of interest: ρ and R^2. The parameter ρ is a direct measure of rank mobility. At one extreme, if households' rank in the respective distribution is fixed

across time, then $\rho = 1$. In contrast, if households' rank is purely idiosyncratic over time, then $\rho = 0$. Thus, higher mobility implies lower value of ρ, and higher mobility in the consumption distribution than the income distribution implies that our estimate of ρ in the rank regression for consumption should be lower than its estimated counterpart for income. The second statistic, R^2, summarizes the extent to which households' rank position is predictable by the included regressors.

Results are shown in Table C4. Columns (1A) and (1B) suggest that, before controlling for household fixed effects and life-cycle characteristics, the household rank is more persistent in income than in consumption even if last year's rank has substantial explanatory power in both cases. The autoregressive coefficient is 0.86 for income and 0.76 for consumption, implying that shocks to a household's position within a distribution have a half-life of about 4.6 years for income, but only 2.5 years for consumption.[1] Obviously, these numbers are heavily influenced by household characteristics such as age, kids, etc. Columns (2A)-(2B) control for the household fixed effects described previously. Now both autoregressive coefficients drop substantially—the one for consumption declines by almost two thirds—but the main observation that movements in the consumption distribution are less persistent remains. The same applies when we also control for the life-cycle trends and demographic characteristics described previously, see columns (3A)-(3B). Finally, the computed values for R^2 further strengthen the view that past positions in the consumption distribution are less informative about current positions, compared with the counterparts for income.

Table C4: Income and consumption rank regressions

	(1A) Inc.	(1B) Cons.	(2A) Inc.	(2B) Cons.	(3A) Inc.	(3B) Cons.
ρ	0.86	0.76	0.46	0.28	0.42	0.24
R^2 total	0.76	0.59	0.76	0.59	0.54	0.38
HH-level FE	No	No	Yes	Yes	Yes	Yes
Life-cycle trends	No	No	No	No	Yes	Yes

D GMM ESTIMATION Á LA BLUNDELL ET AL. (2008)

After having obtained the estimates of residualized income and consumption, $y_{i,t}$ and $c_{i,t}$, we proceed with the estimation of income and consumption shocks. We separately identify (i) a permanent income shock assumed to be a random walk, (ii) a transitory income

[1] The half-life is computed as $h = ln(0.5)/ln(\rho)$.

shock assumed to be MA(1), and (iii) a transitory consumption expenditure shock which is white noise. We allow for year-specific second moments of all shock distributions, as well as year-specific pass-through elasticities from income shocks to consumption (see Blundell et al., 2008, Crawley, Holm, and Tretvoll, 2025). The income processes are summarized below:

$$\begin{aligned} y_{i,t} &= \bar{y}_{i,t} + \tilde{y}_{i,t} \\ \bar{y}_{i,t} &= \bar{y}_{i,t-1} + \eta_{i,t} \\ \tilde{y}_{i,t} &= \varepsilon_{i,t} + \theta\varepsilon_{i,t-1} \end{aligned}$$

Household-level income growth follows:

$$\Delta y_{i,t} = \eta_{i,t} + \varepsilon_{i,t} - (1-\theta)\,\varepsilon_{i,t-1} - \theta\varepsilon_{i,t-2}$$

Turning to consumption, we follow common practice in the literature and conjecture the following consumption growth equation:

$$\Delta c_{i,t} = \phi_t \eta_{i,t} + \psi_t \varepsilon_{i,t} + \gamma_{i,t}$$

Thus, the reduced-form elasticities ϕ and ψ capture pass-through to consumption from the permanent and transitory income shocks, respectively. Finally, $\gamma_{i,t}$ represents an expenditure shock that is orthogonal to changes in income.

The model is estimated by GMM. We let $\sigma^2_{\eta,t}$ and $\sigma^2_{\varepsilon,t}$ represent the time-varying variances of permanent and transitory income shocks, and $\sigma^2_{\gamma,t}$ represent the time-varying volatility of expenditure shocks. These three shock volatilities, as well as the parameters ϕ, ψ and θ, are confronted with empirical moments in data. Identification arises from the autocovariances in income and consumption growth in data.

We follow standard procedure in the literature to estimate the model: first, we compute the relevant sample moments for income and consumption growth. Second, we form a system of moment conditions that allows us to identify the parameters of interest. Third, we estimate the model by choosing parameters to match selected theoretical moments to the sample moment counterparts. Let $\Omega_{\Delta x_t, \Delta z_{t+j}}$ denote the covariance between variables Δx in period t and Δz in period $t+j$. For $j = [0, 1, 2]$ we derive the following autocovariance matrices between consumption and output:

$$\Omega_{\Delta y_t, \Delta y_t} = \sigma^2_{\eta,t} + \sigma^2_{\varepsilon,t} + (1-\theta)^2\,\sigma^2_{\varepsilon,t-1} + \theta^2\sigma^2_{\varepsilon,t-2} \tag{D.1}$$

$$\Omega_{\Delta y_t, \Delta y_{t+1}} = -(1-\theta)\,\sigma^2_{\varepsilon,t} + \theta\,(1-\theta)\,\sigma^2_{\varepsilon,t-1} \tag{D.2}$$

$$\Omega_{\Delta y_t, \Delta y_{t+2}} = -\theta\sigma^2_{\varepsilon,t} \tag{D.3}$$

$$\Omega_{\Delta c_t, \Delta y_t} = \phi_t \sigma^2_{\eta,t} + \psi_t \sigma^2_{\varepsilon,t} \tag{D.4}$$

$$\Omega_{\Delta c_t, \Delta y_{t+1}} = -\psi_t (1-\theta) \sigma^2_{\varepsilon,t} \tag{D.5}$$

$$\Omega_{\Delta c_t, \Delta y_{t+2}} = -\psi_t \theta \sigma^2_{\varepsilon,t} \tag{D.6}$$

$$\Omega_{\Delta c_t, \Delta c_t} = \phi_t^2 \sigma^2_{\eta,t} + \psi_t^2 \sigma^2_{\varepsilon,t} + \sigma^2_{\gamma,t} \tag{D.7}$$

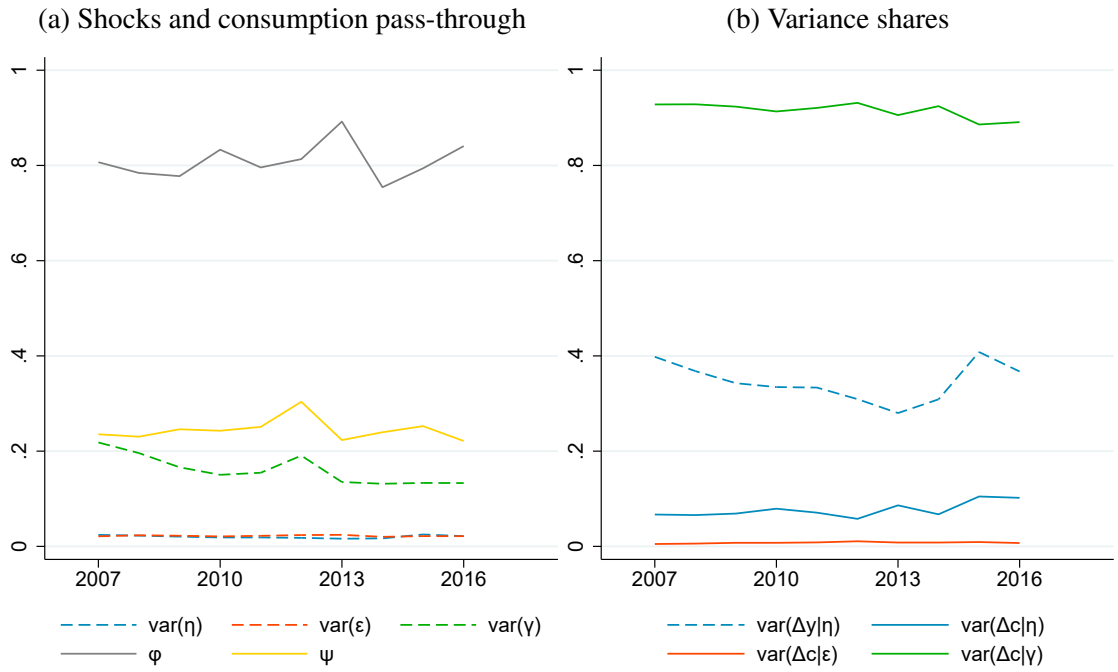
The autocovariances for $j > 2$ are zero. Note that the system is over identified—at each point in time we have 7 equations to identify the 6 parameters $\sigma^2_{\eta,t}$, $\sigma^2_{\varepsilon,t}$, $\sigma^2_{\gamma,t}$, ϕ_t, ψ_t and θ. The first equation governs the overall variance of the two income shocks. The second and third equations exploit the auto-covariance structure in income to identify the relative volatility of transitory income disturbances. The fourth equation governs the overall pass-through from income shocks to consumption while the fifth and sixth equations pin down the pass-through from the transitory components. Finally, the variance of the expenditure shock, $\sigma^2_{\gamma,t}$, is determined residually in the seventh equation to match the overall volatility of consumption.

Overall, identification is achieved in two ways: first, we disentangle permanent and transitory income shocks from each other because only the latter gives non-zero *auto*-covariances (predictability) in income and consumption growth. Second, we disentangle income shocks from consumption shocks because only the former lead to non-zero *cross*-covariances between income and consumption growth, while the latter soaks up any residual consumption variation that is ortogonal to income.

E ROBUSTNESS: TIME VARIATION

We assess the stability of the estimated income and expenditure shocks over time. To this end, we follow Blundell et al. (2008) and explicitly estimate year-specific pass-through elasticites ϕ_t and ψ_t, as well as year-specific cross-sectional shock variances $\sigma^2_{\eta,t}$, $\sigma^2_{\varepsilon,t}$ and $\sigma^2_{\gamma,t}$. Results are documented in Figure E10 while the details about moment conditions are relegated to the appendix. The time-varying estimates in Figure E10a largely confirm our results so far: first, the estimated time series for ϕ_t fluctuates between 0.75 and 0.89, while ψ_t takes values between 0.22 and 0.30. Thus, while permanent income shocks display much higher pass-through to consumption than temporary income shocks, we still find that a quarter of the temporary income shocks pass on to consumption. Reassuringly, our estimates are in the ballpark of those based on U.S. survey evidence obtained by Blundell et al. (2008), as well as the results with Norwegian microdata in Eika et al. (2020) (using imputed consumption) and Fagereng et al. (2021) (who exploit lottery prize data to identify transitory shocks). Second, the the annual estimates of $\sigma^2_{\eta,t}$ and $\sigma^2_{\varepsilon,t}$ vary between 1.6 and 2.5, while $\sigma^2_{\gamma,t}$ varies between 13 and 22.Thus, consumption shocks are

Figure E10: Year-specific estimates of income and expenditure shocks



Notes. Panel (a): annual estimates of $\sigma^2_{\eta,t}$, $\sigma^2_{\varepsilon,t}$, $\sigma^2_{\gamma,t}$, ϕ_t and ψ_t. Panel (b): annual variance decompositions of household-level income and consumption growth.

substantially more volatile than income shocks. Moreover, none of the parameters display any apparent trend, the sample averages for ϕ_t and ψ_t are about 0.8 and 0.25, respectively, while the sample averages for $\sigma^2_{\eta,t}$, $\sigma^2_{\varepsilon,t}$ and $\sigma^2_{\gamma,t}$ are 0.02, 0.02 and 0.16.

Figure E10b reports annual decompositions of the household-level, cross-sectional dispersions in income and consumption growth into the three shocks. While permanent income shocks explain about 28-41% of income—and transitory income shocks the rest—the two income shocks combined can account for only 7-11% of the overall dispersion in consumption growth. On average, the permanent (transitory) income shocks account for 8% (1%) of the volatility in consumption. The rest, about 91% on average, is attributed to the dispersion across households in realizations of $\gamma_{i,t}$. Thus, our estimated expenditure shocks remain key drivers of consumption volatility in our data even when we allow for time variation in all of the relevant second moments.

F Robustness: Expenditure Innovations

To further gauge the importance of idiosyncratic expenditure shocks, we assess three useful perturbations to the baseline specification presented in the main text. First, we ask whether expenditure shocks remain key if we restrict attention to non-durable consumption expenditures. To this end, we remove all expenditures on durable goods from data

and redo the three exercises summarized in Table 1. Results are reported in Table F5 in the appendix, Panel A. As can be seen from the table, both the estimated pass-through coefficients and the share of consumption variance attributed to expenditure shocks remain relatively similar to the baseline estimates. The largest difference is found in model (iii), where we impose moment conditions using GMM. But even in this case, we only see a small decline in the role of expenditure shocks for (non-durable) consumption, from 91% in the baseline to 87% when durable goods are excluded.

Second, we exclude the data on credit cards as these may be more exposed to measurement errors, and also used for particularly large transactions. Results are reported in Panel B in Table F5. Again both the coefficient estimates and the implied variance decompositions of consumption remain very similar to the baseline counterparts in Table 1.

Third, we restrict attention to expenditures on food consumption. Food expenditures tend to be well measured also in survey data, and the majority of food expenditures are spent on non-durable items. Panel C in Table F5 summarizes the estimates when we only consider food consumption. Estimates once again remain similar to their baseline counterparts. Still, we note that the pass-through elasticities from income to consumption fall somewhat in all three specifications. This is especially true for the pass-through from permanent income $\bar{y}_{i,t}$ to consumption $c_{i,t}$ in specification (ii). One possible explanation is that food consumption to largely consists of necessary goods (with Engel curve slopes substantially lower than 1), so that a given increase in income implies only a small rise in food consumption. However, when we do a more structural separation of permanent and transitory income shocks, as in specification (iii), the pass-through coefficient from permanent income to food expenditures remains very similar to that in the baseline when all consumption goods are considered.

Table F5: Robustness exercises: perturbations to residualized consumption

	(i) Hall regression simple	(ii) Hall regression extended	(iii) Blundell et al. model
	A: Consumption excluding durables		
Passthrough y_t	0.43	x	x
Passthrough $\bar{y}_t$	x	0.48	0.77
Passthrough $\tilde{y}_t$	x	0.37	0.26
$var(\Delta c)$|income	16%	15%	13%
$var(\Delta c)$|exp. shocks	84%	85%	87%
	B: Consumption excluding credit cards		
Passthrough y_t	0.45	x	x
Passthrough $\bar{y}_t$	x	0.50	0.77
Passthrough $\tilde{y}_t$	x	0.38	0.27
$var(\Delta c)$|income	13%	13%	12%
$var(\Delta c)$|exp. shocks	87%	87%	88%
	C: Only food consumption		
Passthrough y_t	0.34	x	x
Passthrough $\bar{y}_t$	x	0.28	0.76
Passthrough $\tilde{y}_t$	x	0.34	0.21
$var(\Delta c)$|income	22%	22%	11%
$var(\Delta c)$|exp. shocks	78%	78%	89%
	D: Consumption variation purged for past liquid wealth		
Passthrough y_t	0.45	x	x
Passthrough $\bar{y}_t$	x	0.49	0.85
Passthrough $\tilde{y}_t$	x	0.39	0.28
Passthrough wealth $t-1$	0.03	0.03	x
$var(\Delta c)$|income	15%	14%	12%
$var(\Delta c)$|exp. shocks	85%	86%	88%

Notes. Panels A, B, C: Coefficients of the three empirical specifications considered in the main text, with alternative definitions of consumption expenditures as the dependent variable. Panel D: baseline regression specifications where we also include (lagged) liquid wealth.