A TALE OF PROCYCLICAL INEQUALITY: FACTS AND IMPLICATIONS*

PRELIMINARY DRAFT. WORK IN PROGRESS

DRAGO BERGHOLT†, FRANCESCO FURLANETTO‡ AND LORENZO MORI§

February 28, 2025

Abstract: We analyze the level, the trend, and the cyclical patterns of consumption and income inequality in Norway, a modern welfare state. To achieve this, we utilize a panel dataset of card transactions and administrative income data, encompassing the entire population of Norwegian households from 2006 to 2018. Contrary to the increased inequality observed in the U.S., we find that both consumption and income inequality in Norway have remained relatively stable over time, albeit with considerable mobility within the distributions. Moreover, while pre-tax and transfer inequality appears largely disconnected from the business cycle in our data, we identify a pronounced pro-cyclical pattern in the inequality across households related to consumption and disposable income. This pro-cyclicality is evident both unconditionally and in response to major business cycle shocks. Finally, we show that the right tails of the income and consumption distributions exhibit excess cyclical sensitivity, while households in the lower percentiles remain relatively insulated from macroeconomic fluctuations. Interpreted through the lenses of heterogeneous agent models, the pro-cyclicality of inequality may explain why the Norwegian economy has been rather unaffected by economic recessions in recent decades: redistribution to low-income households with high MPC dampens the effects of recessionary shocks, thereby limiting their transmission to the aggregate economy. Consequently, our findings offer new evidence on the role of distributional policies, as well as important implications for models with household heterogeneity.

Keywords: *Consumption inequality, micro-data, earnings inequality, structural VAR models.*
JEL Classification: *C3, E3, E5.*

---

*First draft: June 2024. This draft should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. A special thank to Sigurd Galaasen and Ragnar Juelsrud for patient help on the data. We would like to thank for useful comments our discussant Elisa Guglielminetti, and Thor Andreas Aursland, Erich Battistin, Fabio Canova, Cristiano Cantore, Efrem Castelnuovo, Olivier Coibion, Davide Debortoli, Michael Ehrmann, Andreas Fagereng, Luca Fosso, Luisa Fuster, Martin Holm, Nicolò Maffei Faccioli, Gisle Natvik, Erlend Njølstad, Serdar Ozkan, Giorgio Primiceri, Ørjan Robstad, Juan Rubio-Ramirez, Nicolas Werquin, as well as seminar participants at SNDE 2024 conference in Padova, Norges Bank Mini Workshops on Monetary Policy in 2022 and 2024, Norges Bank 2024 Spring Institute in Venastul, Workshop on Household Heterogeneity at Sapienza University, and 2024 Dolomiti Macro Meeting in Castelrotto.
†Norges Bank. Email: drago.bergholt@norges-bank.no
‡Norges Bank. Email: francesco.furlanetto@norges-bank.no
§University of Padova. Email: lorenzo.mori.1@phd.unipd.it

1 INTRODUCTION

While high and rising levels of income inequality have been center stage in the policy debate over the last couple of decades (Piketty and Saez (2003) and Piketty, Saez, and Zucman (2018)), the level and the dynamics of consumption inequality are still highly uncertain, as discussed in Attanasio and Pistaferri (2016). Measuring well consumption inequality is crucial because consumption is arguably the best indicator of economic well-being and constitutes a key component of welfare in any macroeconomic model. Yet, such measurement is notoriously difficult because the main data source consists of surveys plagued by various forms of measurement error (Meyer and Sullivan (2023)). Our key contribution consists in measuring consumption inequality using a data set tracing electronic payments in Norway over the period 2006-2018. We focus on its level, trend and cyclicality, we compare it to various measures of income inequality and we evaluate some important theoretical implications of our results.

The data cover around 80 percent of electronics payments made by Norwegian households over the sample 2006-2018 and has been used first by Ahn, Galaasen, and Maelhum (2024) to study the cash-flow effects of monetary policy in Norway. It includes debit card payments at weekly frequency. Each transaction is associated to a consumer category, location and some demographic characteristics of the person who makes the payment. At the end of the sample, we detect more than two billions of transactions per year.[1] In addition, we have access to income and tax statements for the universe of Norwegian residents from the Norwegian Tax authority in order to measure various definitions of income inequality (Halvorsen, Ozkan, and Salgado (2022)).

We find that the level of consumption inequality is high in Norway. It is comparable to the level of earnings inequality and substantially higher than the level of disposable income inequality. This result is at first surprising and is confirmed also when using measures of imputed consumption (Fagereng and Halvorsen (2017)). When it comes to its evolution over time, consumption inequality is very stable and does not feature any visible trend. The same is true for various measures of income inequality, as already documented in Halvorsen et al. (2022). The evidence for Norway is substantially different from the case of the US. In fact, US data feature a clear upward trend in income inequality while the evidence for consumption is less clear, with most studies finding a rather stable level of inequality with the notable exception of Aguiar and Bils (2015). The evidence for Germany is somewhere in between, as documented in Fuchs-Schündeln, Krueger, and Sommer (2010).

Our second set of results is related to the cyclicality of inequality. While our sample period does not feature large recessions in Norway, we observe a stark reduction in income and consumption inequality during the Great Recession (which was not so great in Norway). In good times, all measures tend to increase, thus implying that inequality seems to be pro-cyclical in Norway. Crucially, the procyclicality of consumption inequality is driven by the top of the distribution: percentiles from 95 to 99 are by far the most cyclical while no clear cyclicality emerges for percentiles in the middle of the distribution. The lowest percentiles are even countercyclical. The same is true for various measures of income inequality. This is in stark contrast with results presented by Guvenen, Ozkan,

[1] An aggregate version of the same dataset is used to nowcast Norwegian consumption in Aastveit, Fastbø, Granziera, Paulsen, and Torstensen (2024).

and Song (2014) who document a strong cyclicality at the bottom of the distribution in the US. Previous evidence for Brazil is closer to the case of Norway (see Sonnervig (2023)).

Since our sample features only one clear contractionary episode, standard structural vector autoregression (SVAR) models can be more useful than event studies to evaluate the cyclicality of inequality. We use SVAR models to identify the so- called main business cycle shock (as in Angeletos, Collard, and Dellas (2020)), which has not a structural interpretation but captures a large share of variation at business cycle frequencies.[2] All measures of inequality appears to be pro-cyclical also conditional on such shocks. The procyclicality of earnings is milder while the one of disposable income and consumption is substantially larger. Once again, the cyclicality is mainly driven by the top percentiles. All in all, we offer a simple message: various measures of inequality seem to be pro-cyclical (conditionally and unconditionally), this cyclicality is driven by the top of the distribution and the results are particularly strong for consumption inequality (which is also high in level).

The estimated pro-cyclicality of consumption and income inequality have important implications for macroeconomic models with heterogeneous agents (Kaplan, Moll, and Violante (2018)) or models with simple forms of heterogeneity (Bilbiie (2024)), if applied to Norwegian data. The amplification properties of these models rely on counter-cyclical inequality (proxied by consumption inequality) and on countercyclical risk (proxied by earnings inequality).[3] If both are instead pro-cyclical, as it seems to be the case in Norway, demand shocks are expected to be dampened and leave limited traces on the macroeconomy. In such an economy, inequality and risk act as automatic stabilizers, thus potentially explaining why the Norwegian economy has been so stable since the banking crisis at the beginning of the 1990s.

We contribute to two strands of literature. First, we complement existing evidence on income and consumption inequality, much of which is primarily based on survey data (especially for consumption). For the US, Heathcote, Perri, and Violante (2010) provide a comprehensive overview across various measures, updated in Heathcote, Perri, Violante, and Zhang (2023). We rely on Guvenen et al. (2014) and we show how their main result, the counterciclicality of left-skewness in the earnings growth distribution, is not confirmed in Norwegian data.

We also contribute to a recent literature looking at the effects of macroeconomic shocks on inequality. Coibion, Gorodnichenko, Kueng, and Silvia (2017) find that expansionary monetary policy shocks reduce both earnings and consumption inequality in the US. Chang and Schorfheide (2022) build on the functional SVAR approach developed in Chang, Chen, and Schorfheide (2024) and confirm the evidence of countercyclical earnings inequality but find procyclical consumption inequality, although subject to a substantial amount of uncertainty. De Giorgi and Gambetti (2017) document the response of various deciles in the consumption distribution obtained from CEX survey data to TFP and uncertainty shocks and find that the highest deciles are more cyclical. Cantore, Ferroni, Mumtaz, and Theophilopoulou (2022) on the effects of monetary policy shocks on labor supply at the bottom of the income distribution, Gaudio, Petrella, and Santoro

[2] We plan to extend out analysis to more granular shocks like demand and supply shocks identified with sign restrictions or monetary policy shocks as in Holm, Paul, and Tischbirek (2021).

[3] Evidence on the relative importance of these channels for the US is provided in Bilbiie, Primiceri, and Tambalotti (2023).

(2021) on the effects of supply shocks on inequality and Furlanetto, Robstad, and Sarferaz (2024) on the link between immigration and inequality in Norway. As already mentioned, a key reference on the Norwegian consumption data and on their use for monetary policy is Ahn et al. (2024).

2 THE DATA

We use a comprehensive dataset of debit card transactions provided by the Norwegian retail clearing institution, Nets Branch Norway, spanning from 2006 to 2018. The dataset covers transactions processed through BankAxept, the national payment system owned by Norwegian banks, covering 80% of domestic card transactions for Norwegian residents. A detailed description of the data is provided in Ahn et al. (2024). For a comparison with the official statistics, see also Aastveit et al. (2024) The data are aggregated by week, postal code, consumer category (food, cars, furnitures, etc; for a total of 26 distinct categories), and individual level, and linked with annual administrative records detailing individuals' earnings, market income (labor, financial), and disposable income (after transfers and taxes), alongside demographic information (gender, age, education). Given that consumption decisions are very noisy at high frequencies and are typically made at the household level, we aggregate the data at quarterly frequency and household level.[4] We then apply an equivalization method to the household data, summing all incomes and consumptions and dividing by the square root of the household size, producing a panel dataset with household-level quarterly consumption data and annual income data from 2006 to 2018.

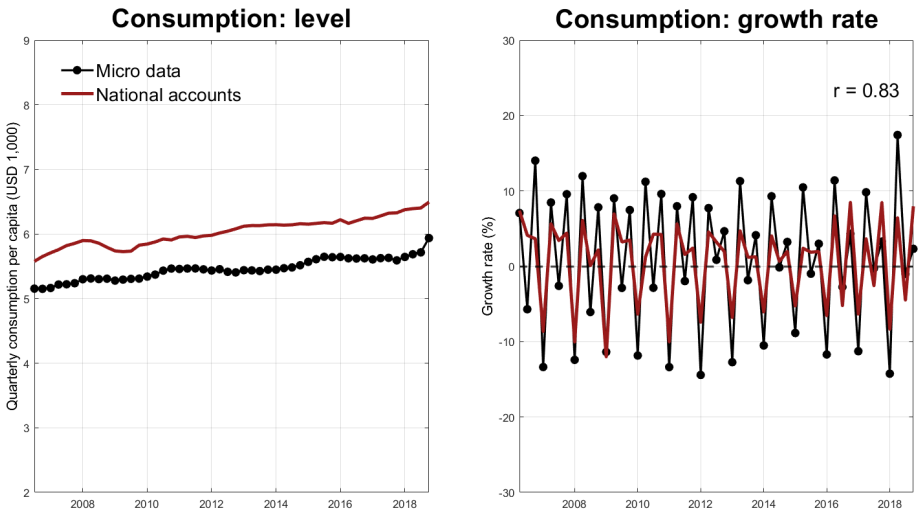
As a standard practice in the literature, we trim observations below a (time-varying) annual threshold of labor earnings and consumption to ensure the sample represents households with significant labor market participation and regular debit card use. Specifically, we drop household-year observations where equivalized yearly labor earnings and consumption fall below the minimum level set by the Norwegian social security scheme, approximately USD 11,000 in 2015 (see Holm et al. (2021) for a similar restriction). In addition, given that debit card transactions could be a poor measure of consumption for self-employed workers, we remove observations with equivalized self-employment earnings above half of the above mentioned threshold. Following Heathcote et al. (2023), we further restrict the sample by dropping a household if no household member is of working age, which is defined as between the ages of 25 and 60.[5]

To assess the effectiveness of using debit card transaction data as a proxy for consumption, Figure 1 contrasts the evolution of our per capita aggregate consumption with that of the measure taken from the national accounts. This comparison shows that the micro data closely mirrors the official consumption measures, capturing both the levels and growth rates effectively, as evidenced by a high correlation of 0.83. In addition, Fig-

[4] In Norway, the definition of household adopted by Statistics Norway is: "Household consists of persons who are permanent residents of the same dwelling (housing unit) or institution." We follow this metrics in our analysis.

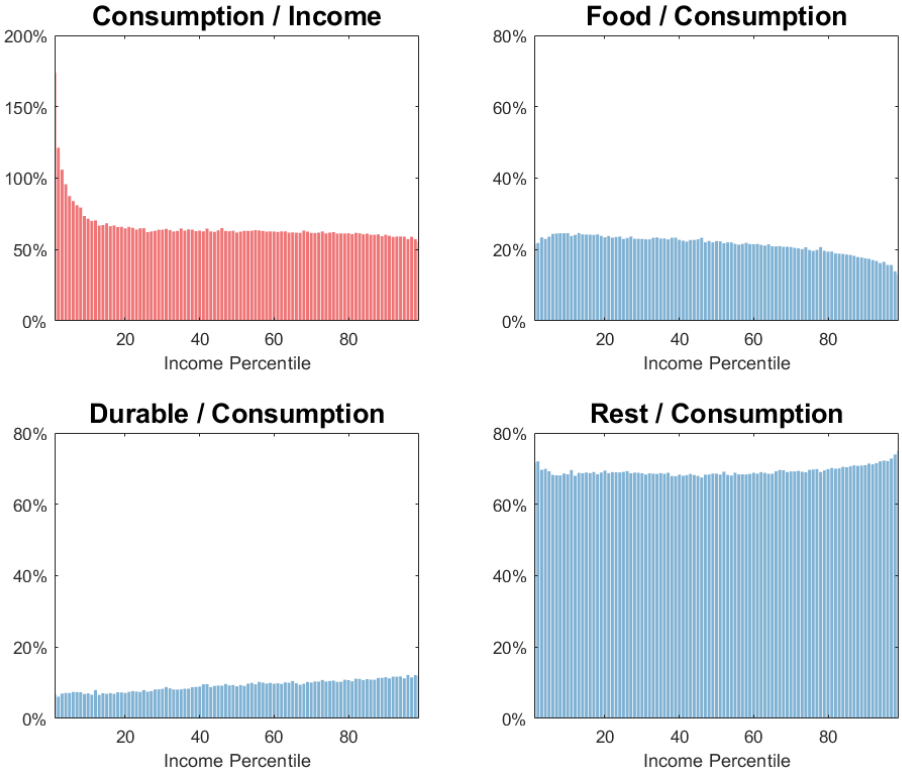
[5] After we impose the bound on age, the sample restrictions eliminate the following fraction of observations (recursively). Labor earnings: 13 percent - more than half of those are out of the labor force and receive zero labor earnings (see also Halvorsen et al. (2022)); consumption: 7 percent; self-employment income: 9 percent.

Figure 1



Notes. Left panel: level of per capita consumption in the micro data vs. national accounts (four quarter moving average). To make the series comparable, we exclude imputed rents for owner-occupied housing from the national accounts. All values are deflated by Norwegian CPI (2015 =1) and converted into US dollars using the exchange rate in 2015. Right panel: quarter-on quarter growth rates of the raw series.

Figure 2



Notes. In red: share of consumption as fraction of income, along the income distribution. In blue: share of categories' consumption as a fraction of total consumption along the income distribution. Food = spending for food and beverage; durable = spending for cars and furniture. Reference year: 2014.

ure 2 (first panel; in red) reports the share of consumption as fraction of income, along the income distribution. As expected, consumption represents a dominant share of income for the bottom tail of the distribution (people spend all or most of their disposable income), while such shares slowly decline over time. Digging deeper, we look at the share of different types of consumption (food, durables, and the residual component) as a share of households' consumption, once again along the income distribution. Figure 2 (blue panels) clearly illustrates that food consumption constitutes a substantial portion of total consumption, particularly for lower-income households, and diminishes in significance as income increases. Conversely, durable consumption is more prevalent among higher-income groups. These dynamics are consistent with conventional beliefs. It hence appears reasonable to employ our debit card data as a proxy for consumption patterns to analyze inequality dynamics. Although our method does not cover the universe of transactions, it offers a viable alternative to traditional consumption data, typically derived from survey information or imputation techniques. Crucially, leveraging micro-data enables more high-frequency consumption measurements, such as quarterly (and in principle even weekly) instead of annually, and reduce the potential mis-measurement.

3 Inequality trends and levels

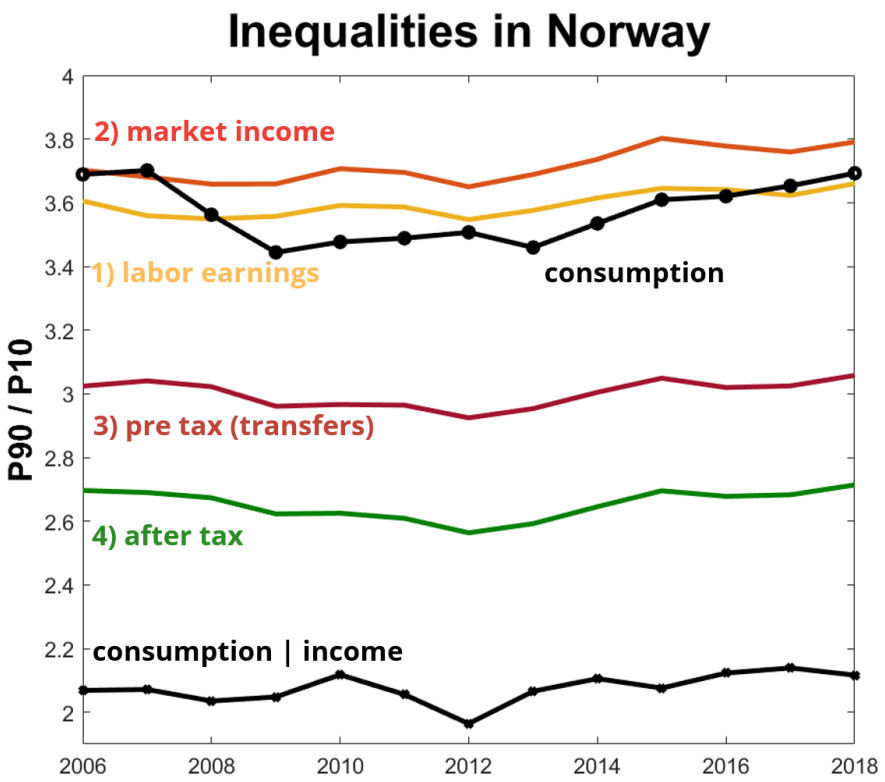
In Figure 3 we report a summary measure of inequality, P90/P10, as implied in our data set for various measures of income inequality and for consumption inequality. The starting point is the yellow line which refers to labor earnings inequality, the orange line considers market income, i.e. the combination of earnings with capital income, the red line refers to pre-tax income which sums or subtracts transfers to market income, the green line plots after tax income (or disposable income) inequality which takes into account the redistributive role of the tax system. Finally, the black line refers to consumption inequality as implied by our data set on electronic payments.

A first fact emerging from the figure is that all measures are relatively stable over time. No clear trend emerges neither for consumption, nor (and perhaps more surprisingly) for the various measures of income inequality, including earnings, in stark contrast with the evidence for the US (Heathcote et al. (2023)).

A second fact is that the level of earnings inequality is relatively low. Such a low level of inequality is further reduced when taxes and transfers are taken into account.

A third, and unexpected yet intriguing, fact is that consumption inequality is substantially higher than disposable income inequality and comparable even to earnings inequality. While at first puzzling, this turns out to be a robust results for Norway even when consumption inequality is estimated through imputation techniques, as documented by Fagereng and Halvorsen (2017). We will investigate what could explain this fact. At the moment, we can put forward only some conjectures. One possibility is that disposable income inequality is artificially low in the data because financial income is not properly measured at the top distribution (Aaberge, Mogstad, Vestad, and Vestre (2021)). Relatedly, there is evidence that wealth inequality is rather large in Norway (and more generally in Scandinavian countries), as discussed in Epland and Kirkeberg (2012). A relatively high level of wealth inequality could rationalize our finding on consumption inequality. Finally, it is important to stress that when we compute consumption inequality, we rank households by their level of consumption, as it is standard in the literature. If

Figure 3



Notes. We consider the P90/P10 ratio 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. For similar definitions, see Heathcote et al. (2023). Results are robust to considering variance of the log distribution or the Gini index.

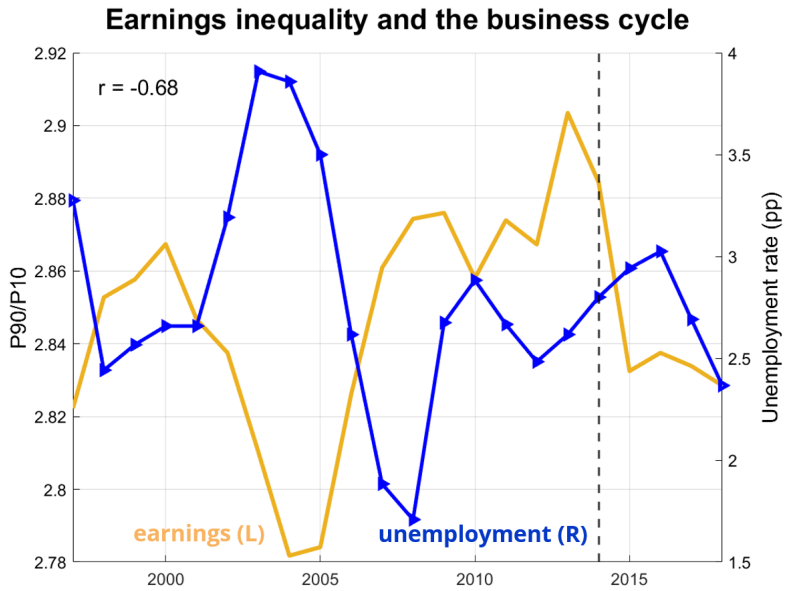
we rank households by income, as done in the black bottom line in Figure 3, the level of consumption inequality shrinks considerably.

In our data, the average level of the P90/P10 ratio is around 3.6 in a country where disposable income inequality is particularly low (on average 2.7, using the same statistics). While this result will need to be investigated in much more detail, results on Spanish data seem to indicate that consumption inequality may be much higher when measured with transactions rather than with surveys. Buda, Hansen, Rodrigo, Carvalho, Ortiz, and Rodríguez Mora (2023) use data from BBVA, a large spanish bank, and report a value of consumption inequality equal to 3.6 for P90/P10 in 2017, exactly as in our data for Norway. In addition, the top 1 percent in the consumption distribution accounts for 4.1 percent of total consumption in Spain. The same statistics for Norway is equal to 6 percent. The similarities between the two datasets are striking and indicate rather clearly that the evidence from billions of transactions is rather different from what we know from surveys, as discussed in detail in Buda et al. (2023).

4 THE CYCLICALITY OF INEQUALITY

We now investigate the cyclicality of inequality in our datasets. The standard result in the literature is that inequality is countercyclical. While this is more evident in US data than

Figure 4



Notes. Earnings inequality vs. unemployment rate in Norway. Earnings inequality is now based on an extended dataset from employer-employee registry records (individual level). A vertical dotted line indicates a break in the way earnings data are collected, complicating comparisons across periods.

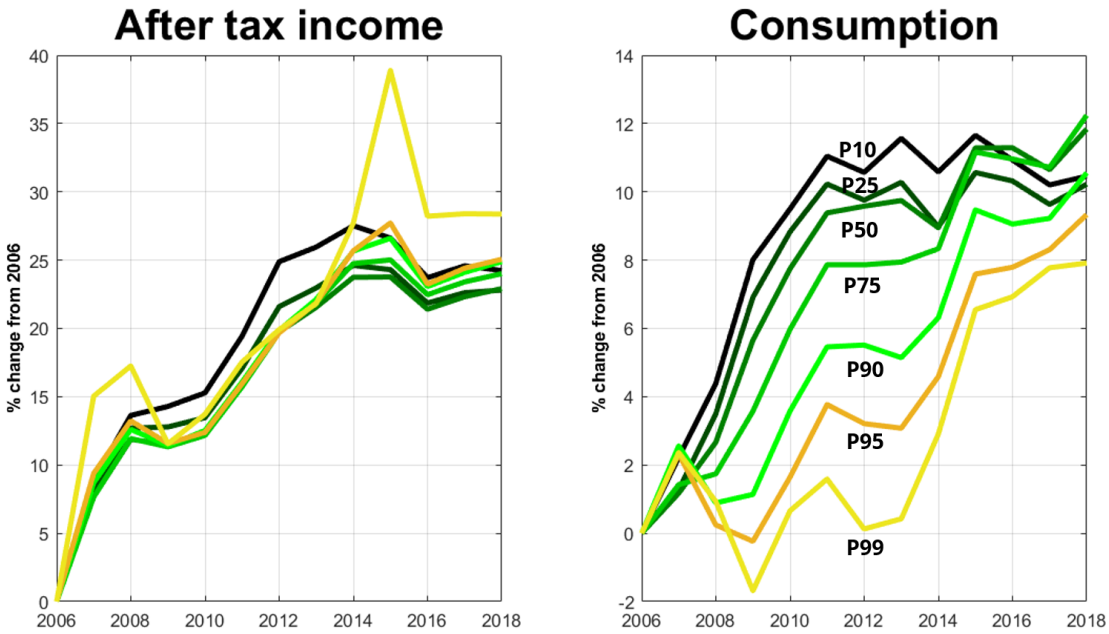
in european data and more evident for earnings than for consumption, it is fair to say that this is the conventional wisdom in the literature. Our main result is that all measures of inequality are instead procyclical in Norway, both unconditionally and conditionally on shocks.

4.1 Unconditional analysis

So far, we relied on a short sample starting in 2006, constrained by the availability of consumption data. However, earnings data are available since the end of the 1990s in an extended dataset at the individual level from the employer-employee registry. In order to set up the stage, in Figure 4 we plot earnings inequality (measured as P90-P10 in this auxiliary data set) against the registered unemployment rate (NAV). This allows us to include in the sample a period in which unemployment reaches 4 percent in 2003 (which is a high level given the very peculiar Norwegian standards) and the subsequent decline (reaching a record low of 1.7 percent) interrupted by the Great Recession. Despite fluctuating relatively little, the unemployment rate is considered as the best business cycle indicator in Norway. It is therefore striking to see that the two series are strongly negatively correlated: when unemployment is low, inequality is high and viceversa. Put differently, earnings inequality looks pro-cyclical (and quite strongly so). In the rest of this section, we check whether this result is confirmed for disposable income and consumption. In addition, we study whether the result survives when we condition on specific shocks in a SVAR model.

Figure 5 shows the evolution of selected percentiles of the annual real income and consumption distributions relative to 2006. Overall, we document an upward trend across

Figure 5



Notes. Percentiles of the log real after tax income and consumption distribution relative to 2006.

all percentiles, with cumulative growth approximately reaching +25% for income and +10% for consumption. However, the dynamics vary across the distribution. Specifically, income growth was stable and relatively uniform until 2008, but the Great Recession disrupted this pattern. The downturn resulted in slowed income growth for the lower percentiles, a slight decline for the middle range, and a significant decrease at the upper end (especially for the top 1%). Notably, similar patterns emerged in 2016, coinciding with the peak unemployment rate in our sample and a temporary contraction in GDP growth due to a sharp decline of oil prices. This raises the question: Does the cyclicality of income inequality between the top and bottom percentiles align with similar dynamics in consumption? A careful look of the right panel of Figure 5 suggests that this is indeed the case. In particular, during the 2008-09 recession, it is the top part of the consumption distribution that bore the consequences of the downturn.

The pronounced cyclicality at the top of Norway's distribution mirrors trends previously documented in the US, as discussed by Parker and Vissing-Jørgensen (2010); however, the relative resilience of the lower end of Norway's distribution during economic cycles starkly contrasts with the US experience. We dig deeper on this aspect. Guvenen et al. (2014) compare the change of US individuals' earnings during recessionary episodes along the distribution. We follow their leads for earnings and further do the same for income and consumption – given our focus on consumption, we keep our analysis at the household level. The primary focus of our analysis is the 2008-09 recession, which marks the only year of negative GDP growth in Norway within our dataset. We proceed as follows: first, we calculate the average household earnings, income, and consumption for the pre-recession years of 2006-07. We then assign households to percentiles within these distributions. Following this classification, we monitor these households over time, tracking their initial percentile positions, and compute the mean log change in earnings, income

and consumption at each percentile level. This allows us to assess the (reduced-form) responses of different segments of the distribution to the recession. A notable improvement of this approach is that it addresses an issue present in Figure 5, where individuals could shift between percentiles over time, introducing compositional effects. We mitigate this by consistently tracking individuals according to their original percentile positions prior to the recession. This provides a clearer picture of the recession's impact on different economic groups without the confounding variable of percentile shifts. Figure 6 shows the results of this exercise, where we also report the changes for earnings in an "average" recession in the US (from Guvenen et al. (2014)).[6] The Figure presents two key findings. First, the behavior of the income distribution's upper end in Norway compared to the US is strikingly similar during recessions: individuals in the highest percentiles (above the 95th) experience substantial income declines. For example, the top 1% in both countries suffers the most, with earnings in Norway and the US dropping by 13.1% and 12.7% respectively—far greater than the losses seen by those in the 90th percentile, which are only 2% and 0.1%. This sharp contraction at the top is further evidenced in Norway by significant decreases in disposable income and consumption, which plunge by 20% and 16%.

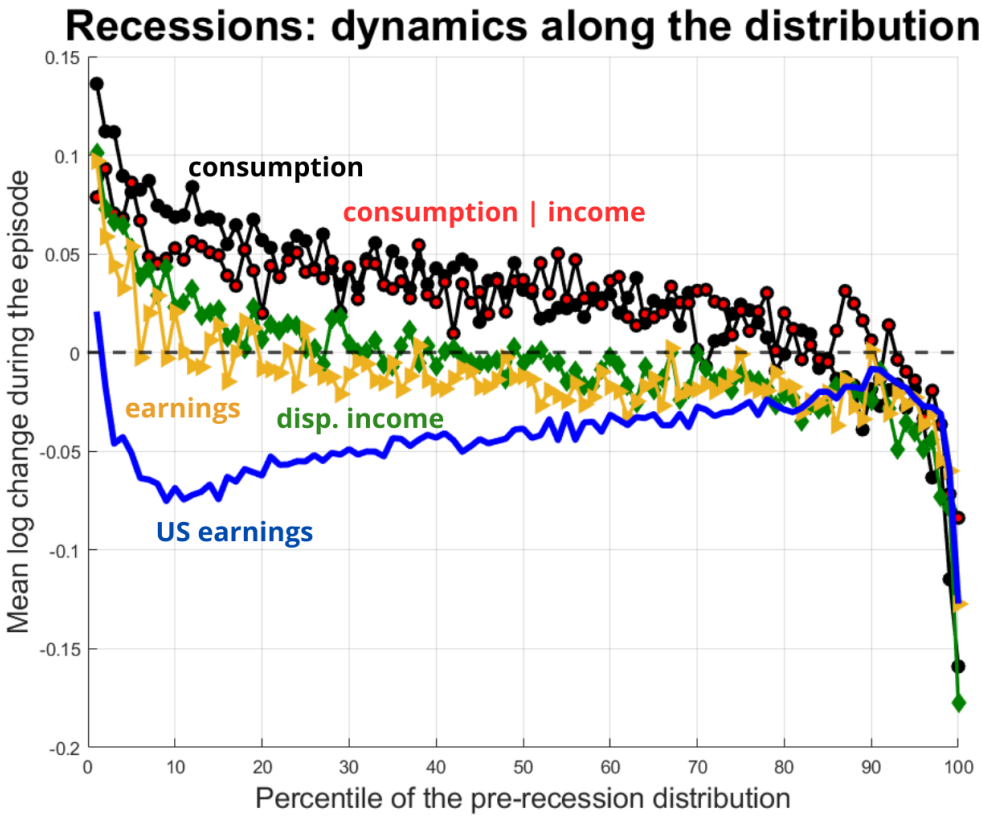
However, the analysis reveals marked differences when examining the central portion and lower end of the distribution. In the US, average earnings during the recession show a nearly linear decrease from the 10th to the 90th percentile, with lower-income individuals disproportionately affected—aside from the extremes. Conversely, in Norway, the lower end of the distribution appears to be largely shielded from the downturn, even when considering compositional effects. One natural explanation for this fact to the welfare system that shields a large share of households. Another institutional feature of relevance is that more than 30 percent of Norway's workforce is employed in general government, the highest value in the OECD.

Another important aspect of Norway that could explain the behavior at the bottom is the presence of frequent unconditional mobility *across* percentiles of the income and consumption distribution (say, for instance, because of life-cycle effects). In Figure 7, we present a graphical representation of such mobility. Using 2006 as the baseline year, households are grouped into a few percentile groups based on their rankings in the income and consumption distributions. Each group is assigned a specific color that remains the same over time. From 2007 to 2018 (denoted as year i), the colored areas represent the proportion of households in each group in year i that originated from the initial categories in 2006. As evident in the figure, household persistence in the bottom and top tails of the distribution is notably low, indicating substantial mobility over time. For example, about half of the households in the bottom decile in 2006 *traveled up* the distribution by 2007 and continued to move across ranks in subsequent years. Similarly, a substantial share of households in the top 2% in 2006 moved down over time. For example, more than 50% of those in the top 2% in 2006 had dropped out of this group by 2007.

All in all, our results for Norway are radically different from previous results for the US documented by Guvenen et al. (2014). Their main result, the countercyclicality of

[6] A key difference from Guvenen et al. (2014) is that they use individual-level earnings data from the US, whereas we focus on household-level data. With their individual-level data, they can also directly control for age (and the related life-cycle effects). We plan to extend our analysis in this direction. In Figure 1 of the Appendix, we report the 25th and 75th percentile of the log change distribution on top of the mean.

Figure 6



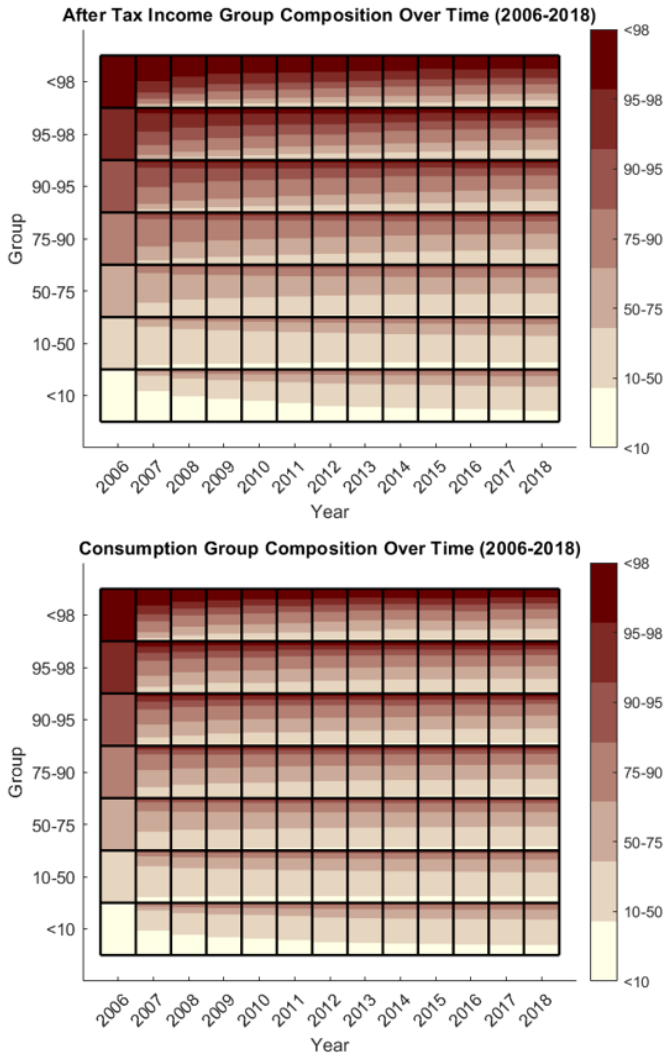
Notes. Change in log average labor earnings, income and consumption during recessions: Norway vs. US. For Norway: labor earnings (yellow line); income after tax (green); consumption (black); consumption given income (black and red). Percentile of the pre-recession distribution calculated during 200-2007. 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).

left-skewness in the earnings growth distribution, is not confirmed in Norwegian data, as shown in Figure 6.

4.2 Conditional analysis

Since our sample features only one clear contractionary episode, standard structural vector autoregression (SVAR) models can be more useful than event studies to evaluate the cyclicality of inequality. We are interested in shocks that drive a large share of economic fluctuations. Therefore, monetary policy shocks would not be useful for our purposes given that they explain routinely a minor share of aggregate fluctuations. We follow Angeletos et al. (2020) and identify a main business cycle shock defined as the shock that maximizes the explanatory power for GDP at business cycle frequencies. The baseline SVAR is estimated using data on real GDP, the CPI price index, real consumption, the policy rate, stock prices, house prices and P90-P10 as a measure of consumption inequality. The variables are in log-levels at quarterly frequency and the sample period is 2006-2018. Given the short sample, we consider two lags and employ Bayesian techniques to avoid

Figure 7



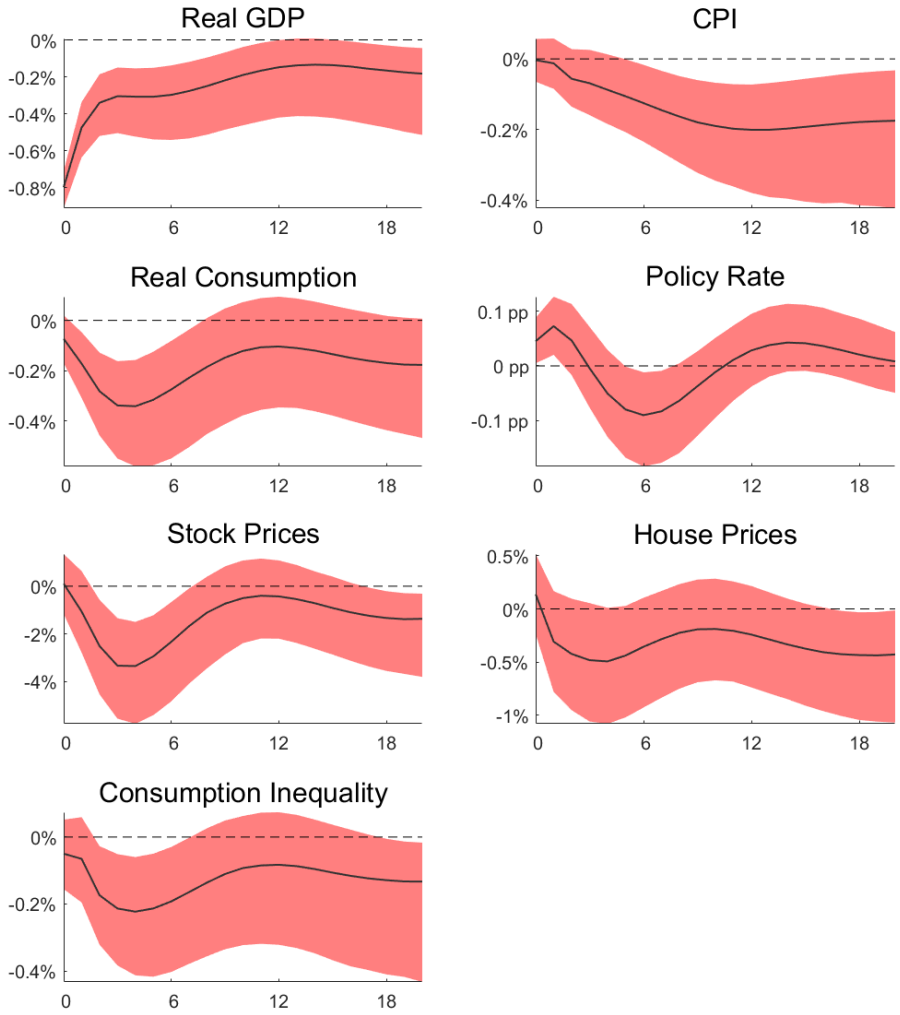
Notes. Mobility over time along the after tax and consumption distribution. We use the year 2006 as the starting point and categorize households into percentile groups based on their ranking in the income and consumption distribution. Each group is assigned a specific color, which remains consistent over time. From 2007 to 2018 (referred to as year i), the colored areas indicate the proportion of individuals in each group for year i who originated from these initial categories.

over-fitting. We use Normal Inverse-Wishart and Minnesota-type priors, whose tightness is calibrated following Giannone, Lenza, and Primiceri (2015).

Impulse responses to a contractionary main business cycle shock are plotted in Figure 8. The shock looks like a negative demand shock with a delayed monetary policy response (notice that Norway did not reach the zero lower bound during the sample). Most importantly, consumption inequality declines and is distinctively pro-cyclical, thus confirming the unconditional evidence provided above.

In a series of auxiliary exercises, we substitute consumption inequality with other variables, one at a time. Results are presented in Figure 9. We remark that consumption decreases only for the highest percentiles, which however account for a disproportion-

Figure 8



Notes. Impulse response functions for a unit-variance recessionary "main business cycle shock". We report 80% posterior coverage bands. Horizons: quarters.

ately large share of aggregate consumption. Intriguingly, consumption increases for P10. Both P90/P50 (referred as top consumption inequality) and P50/P10 (referred as bottom consumption inequality) are procyclical but the former more then the latter.

Finally, in Figure 10 we compare the response of the various measures of income inequality (always measured as P90/P10) against the response of consumption inequality showed earlier. Clearly, all measures are pro-cyclical. However, while earnings (labor income) and market income inequality are mildly pro-cyclical, the measures of income that account for transfers and taxes are substantially more pro-cyclical. Consumption inequality responds more than the other measures.

5 Implications for Theory

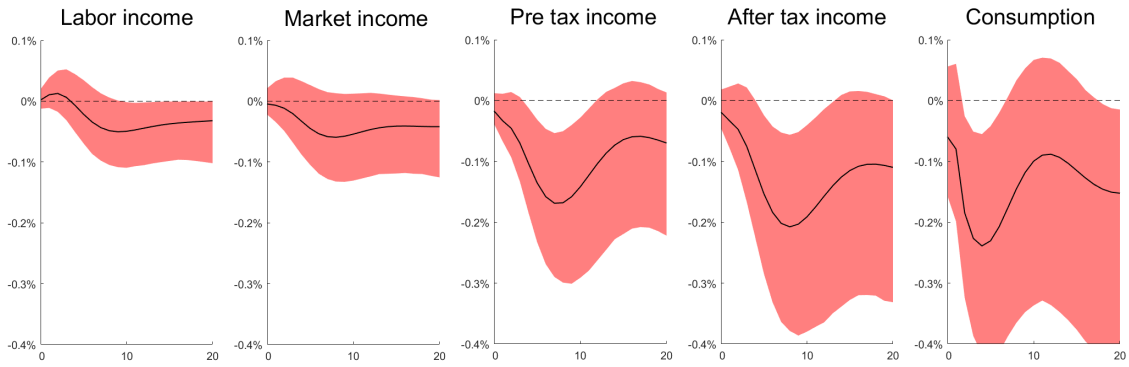
Our results have important theoretical implications for models with heterogeneous agents. For simplicity, we restrict our discussion here to models featuring simple forms of hetero-

Figure 9



Notes. Impulse response functions for a unit-variance recessionary "main business cycle shock". We report 80% posterior coverage bands. Horizons: quarters.

Figure 10



Notes. Impulse response functions for a unit-variance recessionary "main business cycle shock". We report 80% posterior coverage bands. Horizons: quarters.

geneity. It is well-known that demand shocks are amplified in models with consumption heterogeneity, often modeled by allowing for the presence of hand-to-mouth consumers

since the seminal papers of Galí, López-Salido, and Vallés (2007) and Bilbiie (2008). Under general conditions, monetary, government spending and investment shocks are amplified and the amplification is proportional to the share of constrained agents in the economy and depends on the redistributive scheme in place. Another source of amplification consists in the presence of idiosyncratic risk (especially if cyclical) that induces agents to engage in precautionary saving. Bilbiie (2024) shows analytically how the two channels interact while Bilbiie et al. (2023) quantifies their relative importance. Crucially, amplification relies on both consumption inequality (among the two agents) and risk (proxied by earnings inequality) being countercyclical. In Norway, according to our results both consumption and earnings inequality are procyclical leading to a double dampening of business cycle fluctuations. The two mechanisms behave as automatic stabilizers. According to this logic, government spending multipliers should be low in Norway and the cycle should be driven mainly by supply shocks, not because demand shocks are not present but because their propagation is dampened given the peculiar features of the Norwegian economy. In that sense, this analysis provide an ex-post rationalization of why recessions (which are traditionally driven by demand shocks) have so limited macroeconomic effects in Norway.

6 Tails

The goal of this section will be to investigate the tails of the earnings and consumption distributions. According to standard heterogeneous agents models, the Pareto coefficients of the upper tails of the consumption, capital income, and wealth distributions should theoretically all coin- cide . However, Gaillard, Hellwig, Wangner, and Werquin (2023) evaluate empirically these measures of top tail inequality and conclude that they systematically differ—thus rejecting the canonical model. Our data, and in particular the consumption data, can be useful to investigate the predictions of the standard model and of the extensions proposed by Gaillard et al. (2023) in the context of Norway.

7 Conclusion

At this stage, we have three important results. First, consumption inequality is high despite a relatively low level of disposable income inequality. Second, there is no pronounced trend neither in earnings inequality, nor in consumption inequality. Third, both earnings inequality and consumption inequality are procyclical in Norway both unconditionally and conditionally on shocks. All these features of the Norwegian economy seem to be in stark contrast with the available stylized facts for the US economy. These results have important implications: demand shocks are dampened in an economy in which heterogeneity is modeled explicitly with respect to a representative agent economy.

Since the results for Norway are so different from the ones for the US, one may wonder about the external validity of our analysis. It turns out that the results for Norway may not be so exotic. In fact, Ettmeier, Kim, and Schorfheide (2024) find some evidence in favor of procyclical earnings inequality also in Germany using functional VARs and cross-sectional units VARs (csuVAR) that combine aggregate variables with unit-level outcomes. One very tentative conclusion is that perhaps the results for Norway can be

applicable to other modern welfare states and the US constitute a more exceptional case. Substantial more work needs to be done to confirm these conjectures.

REFERENCES

Aaberge, R., M. Mogstad, O. L. Vestad, and A. Vestre (2021). Økonomisk ulikhet i norge i det 21. århundre. In *Økonomisk ulikhet i Norge i det 21. århundre: Aaberge, Rolf— uMogstad, Magne— uVestad, Ola Lotherington— uVestre, Arnstein*. Oslo: Statistisk sentralbyrå.

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*. forthcoming.

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. Manuscript available at https://sites.google.com/site/sigurdgalaasen/research.

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

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

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. (2024). Monetary policy and heterogeneity: An analytical framework. *Review of Economic Studies*. forthcoming.

Bilbiie, F. O., G. Primiceri, and A. Tambalotti (2023). Inequality and business cycles. Technical report, National Bureau of Economic Research.

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. Cambridge Working Papers in Economics CWPE2244.

Cantore, C., F. Ferroni, H. Mumtaz, and A. Theophilopoulou (2022). A tail of labor supply and a tale of monetary policy.

Chang, M., X. Chen, and F. Schorfheide (2024). Heterogeneity and aggregate fluctuations. *Journal of Political Economy*. forthcoming.

Chang, M. and F. Schorfheide (2022). On the effects of monetary policy shocks on earnings and consumption heterogeneity.

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

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

Epland, J. and M. I. Kirkeberg (2012). Wealth distribution in norway: evidence from a new register-based data source.

Ettmeier, S., C. H. Kim, and F. Schorfheide (2024). Measuring the effects of aggregate shocks on unit-level outcomes and their distribution. Manuscript available at https://www.stephanieettmeier.com/research.

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.

Fuchs-Schündeln, N., D. Krueger, and M. Sommer (2010). Inequality trends for germany in the last two decades: A tale of two countries. *Review of Economic Dynamics 13*(1), 103–132.

Furlanetto, F., Ø. Robstad, and S. Sarferaz (2024). Immigration and inequality: New macroeconomic evidence. manuscript.

Gaillard, A., C. Hellwig, P. Wangner, and N. Werquin (2023). Consumption, wealth, and income inequality: A tale of tails. Federal Reserve Bank of Chicago Working Paper.

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.

Gaudio, F. S., I. Petrella, and E. Santoro (2021). Supply shocks and asset market participation. Technical report, Working paper.

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

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

Halvorsen, E., S. Ozkan, and S. Salgado (2022). Earnings dynamics and its intergenerational transmission: Evidence from norway. *Quantitative Economics 13*(4), 1707–1746.

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.

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.

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.

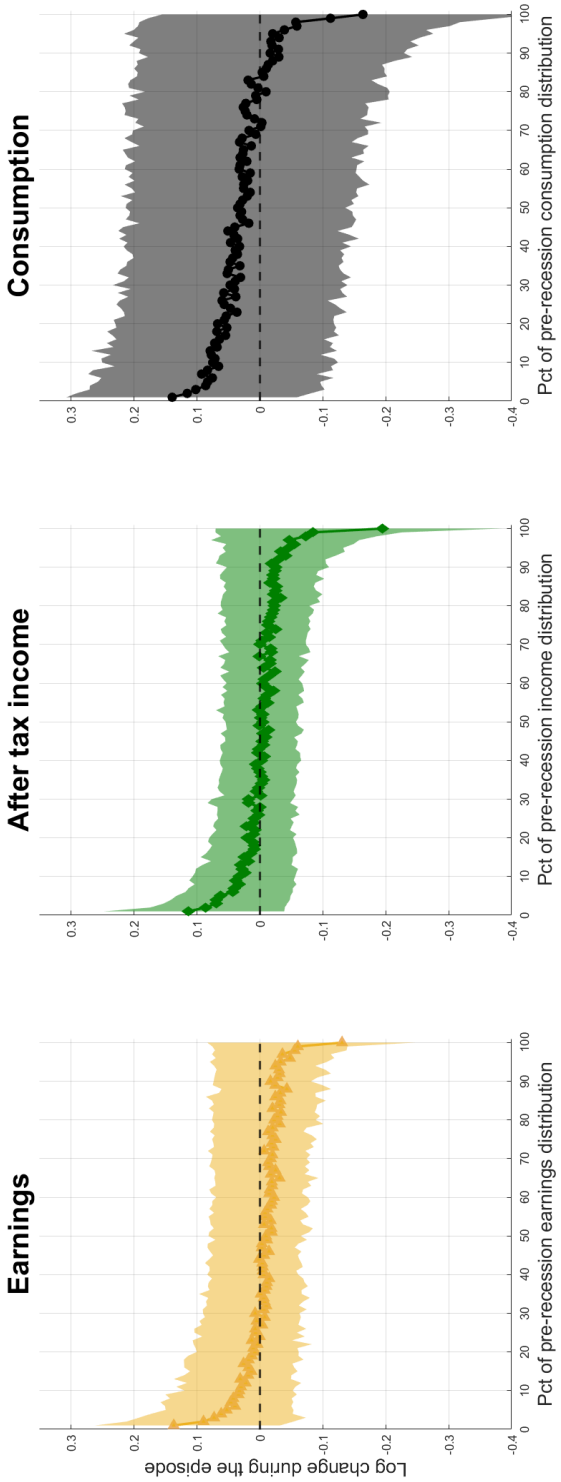
Piketty, T. and E. Saez (2003). Income inequality in the united states, 1913–1998. *The Quarterly Journal of Economics 118*(1), 1–41.

Piketty, T., E. Saez, and G. Zucman (2018). Distributional national accounts: methods and estimates for the united states. *The Quarterly Journal of Economics 133*(2), 553–609.

Sonnervig, M. (2023). Unequal business cycles. Manuscript.

8 APPENDIX

Figure 1



Notes. Change in log labor earnings, income and consumption during recessions: Norway. Average change together with $25th$ and $75th$ percentiles.