ARTICLE IN PRESS

Journal of Monetary Economics xxx (xxxx) xxx

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

The consumption expenditure response to unemployment: Evidence from Norwegian households☆

Andreas Fagereng[a,b,*], Helene Onshuus[a], Kjersti N. Torstensen[c]

[a] *BI Norwegian Business School, Norway*
[b] *CEPR, United Kingdom*
[c] *Norwegian Ministry of Finance, Norway*

ARTICLE INFO

JEL classification:
D12
E21
E24

Keywords:
Unemployment
Household finance
Consumption expenditure
Consumption smoothing
Household heterogeneity

ABSTRACT

This paper examines heterogeneity in household income and consumption responses to unemployment, using granular administrative tax data from Norway. On average, unemployment results in a significant, lasting income reduction, accompanied by a decrease in consumption expenditures of between one-third to one-half of the income loss. We find that households with greater liquid assets at the outset experience less of a decline in consumption, whereas those with higher levels of debt encounter a more substantial decrease. Notably, also the interaction of liquid assets and debt holdings matters for the consumption response. While households with larger initial liquid asset holdings on average respond less, the analyses show that this is not the case among households that simultaneously hold substantial amounts of debt, thus adding to a more nuanced view of the importance of household heterogeneity for economic outcomes. Furthermore, our investigation into heterogeneity across family composition and child age uncovers distinct patterns in consumption responses, highlighting the varied impacts of unemployment. Lastly, we find that spending patterns, as indicated by the marginal propensity to consume (MPC), become more pronounced during recessions.

1. Introduction

Job displacement and subsequent unemployment represent a significant financial disruption for households. Beyond the direct loss of income, job displacement introduces uncertainty regarding both the length of the unemployment spell and future wage prospects. For unemployed households, an important margin of adjustment is their consumption level. Changes in household consumption levels are important not only for the welfare of the individual household but, when aggregated, also for the broader economy. This dynamic has received longstanding attention from policymakers and economic researchers alike (Hall and Mishkin, 1982; Dynarski and Sheffrin, 1987). While the examination of average and aggregate consumption responses to unemployment

☆ This paper should not be reported as representing the views of Norges Bank or the Norwegian Ministry of Finance. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank or the Norwegian Ministry of Finance. Significant parts of this research project were carried out while Torstensen was working at Norges Bank. We thank Boragan Aruoba (Editor), Adam Guren (Associate Editor), and the anonymous referee for their valuable comments and suggestions. We are grateful to Kasper Roszbach for his help in facilitating the project and his continued support, and to Mette Ejrnæs, Moritz Kuhn, Gisle Natvik, and Kjetil Storesletten for their detailed feedback, guidance, and discussions. Fagereng acknowledges funding from the ERC under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 851891). This paper is part of the research activities at the Centre for Household Finance and Macroeconomic Research (HOFIMAR) at the BI Norwegian Business School.

* Corresponding author at: BI Norwegian Business School, Norway.
E-mail addresses: andreas.fagereng@bi.no (A. Fagereng), helene.onshuus@bi.no (H. Onshuus), kjersti-nass.torstensen@fin.dep.no (K.N. Torstensen).

https://doi.org/10.1016/j.jmoneco.2024.103578
Received 30 December 2022; Received in revised form 21 March 2024; Accepted 21 March 2024
Available online 27 March 2024

Please cite this article as: Andreas Fagereng et al., *Journal of Monetary Economics* , https://doi.org/10.1016/j.jmoneco.2024.103578

remains vital, recent advancements, particularly in theoretical work, have shed light on the importance of also understanding the heterogeneity in household responses (Krueger et al., 2016).

Yet, detailed empirical evidence in this setting is scarce, mainly due to limited data on unemployed households' economic behavior. To this end, recent studies such as the work of Ganong and Noel (2019) and Landais and Spinnewijn (2021) have made significant progress. These contributions provide more robust empirical evidence on the average household consumption response, and some along the margins of household heterogeneity, particularly in terms of the liquid assets and debt positions of unemployed households.

This paper extends the existing body of literature by further detailing how heterogeneity in household balance sheet components relates to consumption responses during unemployment. Using granular administrative tax data, the analyses reveal that both debt and liquidity are related to consumption responses, affirming suggestions in the prior literature. Our findings also reveal that the intersection of household balance sheet components matters, which is novel. Although households with larger initial liquid asset holdings on average respond less strongly to unemployment, this is not the case among those households that simultaneously hold substantial debt positions going into the unemployment spell. Here, debt plays a more dominant role than liquidity in shaping consumption patterns.

Our study is conducted in Norway, using administrative records that span over two decades to provide comprehensive annual income and balance sheet data for all Norwegian households. These records, intended for wealth and income tax assessments, are detailed and reliable, with financial institutions reporting directly to tax authorities regarding each household's assets and liabilities. The data, coupled with UI benefit spell data, enables tracking of household financial positions—liquid assets and debt—before, during and after unemployment periods. Additionally, it includes employer–employee records (detailing salaries, employment history, and employer characteristics), demographic data, educational information, and family relationships. Finally, and crucially for our analysis, the panel dimension and level of detail in the wealth and income data also allow us to construct a comprehensive measure of household consumption expenditure via the household budget constraint.

To select our sample, we align with the established literature on the persistent effects on income from job displacement (Jacobson et al., 1993). An important sample requirement in this regard is that workers are in stable job relationships and experience no unemployment spell in a defined period before job loss. This is crucial for the validity of our results, and mitigates concerns regarding unobserved heterogeneity and selection into unemployment, ensuring that unemployment outcomes are compared to a prior, or counter-factual state of employment (Jarosch, 2023). However, it also implies that the sample mainly consists of high-tenured workers for whom a job loss represents a more persistent negative income shock than the average in the universe of displaced workers. This aspect is important to keep in mind when interpreting our findings, and we also undertake robustness analyses where we relax some of these selection criteria.

We employ two methods to quantify the post-unemployment consumption response: First, we use an event study design, deploying a control group to construct counterfactual outcomes for the high-tenure households becoming unemployed. When constructing the control group, we take advantage of the granularity of the administrative data, matching on both balance-sheet components, employment history, employer information, and background characteristics to capture both the ability to smooth consumption, unemployment risk, and re-employment probabilities. Second, we use the control group to estimate the marginal propensity to consume (MPC) within the job-loss year, which not only enables a flexible interaction with covariates but also facilitates comparison to the existing literature on MPC heterogeneity.

We find that unemployment leads to a pronounced, persistent income drop, with earnings decreasing 20%–30%, and only starting to recover two years from the initial job loss. Post-tax labor income decreases 10%–15% instantly, maintaining a level below the control group for the following four years. Interestingly, households with low levels of liquid assets or high debt on average see smaller long-term decreases in earnings and after-tax income than high liquid assets or low debt households, even though the initial earnings drop is similar across debt and asset levels. The income decline is accompanied by a notable decrease in consumption expenditures, amounting to between one-third and one-half of the after-tax income drop. The expenditure decrease is less drastic among households with higher liquidity, while more indebted households see a more severe drop. A considerable fraction of high-liquidity households also bears substantial debt, perhaps out of liquidity concerns, as there are no tax benefits or institutional reasons to do so. Notably, even with substantial holdings of liquid assets at hand, these high-debt households markedly reduce their consumption upon unemployment. We furthermore document how these high-debt households spend a substantial fraction of their disposable household income on mortgage amortization and interest payments, suggesting an important role in preserving homeownership during periods of economic hardship.

Within the event study framework, we also undertake heterogeneity analyses along other important household margins, such as family composition. While families with children tend to recover in terms of income from unemployment somewhat faster, we notice a particular divergence in the consumption response among households with younger children (below 5 years). These households' consumption levels bounce back faster than for families with older children, which complements previous findings on investment in children in a partial insurance framework.

When measuring the household response as an MPC out of unemployment, we find that, on average, a one-dollar income loss leads to a spending decline of about 40 cents. We observe significant MPC heterogeneity across both debt and liquid assets. Further, we uncover a U-shaped relationship between debt-to-income (DTI) and MPC—middle DTI-tertile households have lower MPCs than those in both low and high DTI tertiles. Finally, we find that the U-shaped debt and MPC relationship persists across the liquid asset distribution.

Lastly, we investigate how the household responses differ over the business cycle. On average, we find that income drops appear somewhat less severe during recessions. For the MPC out of unemployment, we find a modest increase during recessions, and results furthermore point in the direction of the U-shaped relationship we observe between the DTI ratio and the MPC as being slightly more pronounced during recessions.

Our analyses point towards a relationship between household balance sheet components, and household debt in particular, and the consumption response to job loss. These findings may be relevant for researchers using estimates of the marginal propensity to consume to refine macroeconomic models or as benchmarks, and they offer insights into modeling choices in structural models where unemployment risk impacts household behavior. Furthermore, our estimates of heterogeneity in consumption responses may help specify policies aimed at buffering economic shocks, such as unemployment insurance schemes or targeted stimulus measures to sustain aggregate demand or serve as an input for policymakers focused on financial stability and macroprudential regulation.

Related literature

Our investigation leans on a well-established body of research on job displacement and its persistent effects on income. Seminal studies such as Jacobson et al. (1993) have documented the long-term income losses experienced by displaced workers in the U.S., a phenomenon consistent with findings in similar contexts, including this study and those reported by e.g. Bertheau et al. (2023) in other European countries.[1]

In exploring consumption responses to job loss, we draw upon a vast literature examining income changes and spending behavior, including foundational works by Hall and Mishkin (1982), Dynarski and Sheffrin (1987) and Blundell et al. (2008).[2] Other researchers have delved into how unemployment insurance (UI) mitigates the impact of income shocks, with Gruber (1997) and Browning and Crossley (2001) providing early insights into the relationship between UI benefits and consumption smoothing and regarding households with limited liquidity.[3]

Following in this line of research, Ganong and Noel (2019) examine the impact of unemployment on consumer spending using de-identified bank data in the US. They show that spending of the unemployed is highly responsive to the level of UI benefits, and drops sharply at both the onset of unemployment and at benefit exhaustion. While they do not observe liquid assets or liabilities directly, they apply an estimate of these positions and relate them to spending drops. In contrast, our work directly observes the household's complete balance sheet position, thereby enabling a more detailed examination of and focus on, for instance, their interactions. Gerard and Naritomi (2021) investigate the degree of consumption smoothing among Brazilian households who receive a sizable severance payment and find excess sensitivity with regards to the lump-sum transfer at unemployment onset. They do not explore the role of heterogeneity in initial balance-sheet positions. Landais and Spinnewijn (2021) explore different approaches to estimating the value of unemployment insurance using data from Sweden, in a setting comparable to ours. In some of the supplementary analyses, they do detect a larger spending drop among households with more leverage. However, the main aim of their article is to improve the understanding of the average valuation of UI, and while they undertake some analyses regarding heterogeneity, they largely leave the door open to further studies in this domain. Andersen et al. (2023) aim to quantify how unemployed households smooth consumption in the Danish setting, finding that using liquid assets is the primary method to mitigate the impact of income loss on spending. Building on this insight, they show that households with high liquid asset holdings have a smaller reduction in spending relative to households with low liquid asset holdings. They do not, however, delve into further analyses of household heterogeneity.

This paper also contributes to research on estimating the marginal propensity to consume (MPC) through quasi-experimental designs that identify exogenous income shocks. Studies have analyzed various sources of such shocks, including unemployment, firm-level shocks (Baker, 2018), lottery winnings (Fagereng et al., 2021), and stimulus payments (Jappelli and Pistaferri, 2014; Kaplan and Violante, 2014; Johnson et al., 2006; Broda and Parker, 2014; Misra and Surico, 2014).[4] Several of these studies explore the role of liquid assets, showing empirically that low levels of liquid assets are associated with higher MPCs. Relatedly, other studies explore the impact of household debt on spending during the Great Recession (Dynan et al., 2012; Mian et al., 2013; Andersen et al., 2016). Baker (2018) find that households with higher debt levels show a greater consumption response to income changes, attributed to credit and liquidity constraints.

More broadly, our work contributes to a burgeoning field within macroeconomics that integrates micro-level heterogeneity. Household consumption behavior is a key component of many macroeconomic models, and understanding how individual household decisions influence broader economic outcomes is key for leveraging macroeconomic models in policymaking. Empirical estimates of structural parameters, such as the MPC, can be used to evaluate model performance. Nakamura and Steinsson (2018) highlight the value of MPC estimates in distinguishing between competing models, a sentiment reinforced by Kaplan and Violante (2018) who emphasize the critical role of heterogeneity. Kaplan et al. (2014) illustrate this by identifying consumption patterns among wealthy hand-to-mouth households based on MPC estimates. In the context of our study, the observed relationship between debt and the consumption response to job loss should be informative for researchers evaluating consumption behavior in structural models where unemployment risk is a pivotal driver of consumption behavior.

1 See also other studies documenting the impact of job loss on workers earnings across time and space: Couch and Placzek (2010), Davis et al. (2011), Huttunen et al. (2011), Kawano and LaLumia (2017), Krolikowski (2017), Flaaen et al. (2019), Lachowska et al. (2020), Jarosch (2023) and Schmieder et al. (2023).

2 For comprehensive literature reviews, refer to Browning and Lusardi (1996) and Jappelli and Pistaferri (2010).

3 See also some of the further literature on evaluating costs and benefits of social insurance and the design of optimal welfare policies (Baily, 1978; Chetty, 2006) studying the value of unemployment insurance (Engen and Gruber, 2001; Chetty, 2008; Hendren, 2017; Kolsrud et al., 2018).

4 See also work by Kueng (2018), Olafsson and Pagel (2018), Bunn et al. (2018).

Estimates of structural parameters are useful for both model calibration and as causal indicators in model estimation (Carroll et al., 2017). Several studies derive MPC estimates from positive shocks (lottery wins, tax rebates). Our findings in this study are derived from negative income shocks (unemployment), thus adding nuance to the current understanding of these parameters, and to the application of such estimates. For instance, the MPC from unemployment is key for examining unemployment insurance (UI) as an economic stabilizer, affecting aggregate demand considerations (McKay and Reis, 2021; Kekre, 2022) and informing macroprudential policies for financial stability (Farhi and Werning, 2016; Korinek and Simsek, 2016).

From here, the paper proceeds in the following way: Section 2 describes the institutional setting, the data used in the analysis, and the sample selection. Section 3 presents the empirical framework, including the selection of a control group. In Section 4 we describe how income and spending develop after the onset of unemployment and investigate possible sources of heterogeneity in the spending responses, with particular attention paid to debt and liquid assets. In Section 5 we portray our results in terms of the marginal propensity to spend and investigate how it varies across the distribution of debt and liquid wealth and the business cycle. Section 6 concludes.

2. Context and data

2.1. Institutional setting

In Norway, participation in its welfare system is mandatory. While the welfare system is among the OECD's most generous, its unemployment benefits, at 62.4% of previous pre-tax labor income, are more limited and are also capped and subject to taxation. Eligibility demands a minimum of one year's prior employment and exceeding a specified income threshold. These benefits, available for a maximum of 104 weeks, apply to all job leavers, with an 8-week waiting period for voluntary leavers and three days for involuntary ones. Severance pay recipients face a benefits waiting period determined by the severance amount.

Norway's job termination procedures, regulated by the Working Environment Act, necessitate employers to establish valid reasons for termination, with downsizing recognized as a principal rationale (Addison and Teixeira, 2003). Other performance-related grounds require evidence of continuous underperformance or misconduct. When layoffs of ten or more employees are anticipated, firms must engage with employee representatives, usually from local labor unions. Layoff notification periods in Norway hinge on an employee's tenure and sometimes age. Typically, under 5 years of service warrants a month's notice, 5 to 10 years equals two months, and over 10 years secures three months. Employees aged 50 or 55 with a decade of service receive 4 or 5 months, respectively. Trial period workers receive a 14-day notice. Public employees, a minority in our dataset, have longer notification periods per their service years. For further details on the institutional setting see Appendix A.

2.2. Data sources

We use Norwegian administrative data from 1994 to 2015, which covers the universe of Norwegian individuals. The primary data sources are income and wealth data from the Norwegian Tax Authority ("Skatteetaten"), unemployment benefit registers from the Norwegian Labor and Welfare Administration, an archive of employer–employee relationships, housing market data from the Norwegian Mapping Authority ("Kartverket"), and extensive demographic data (including characteristics such as education, age, and marital status), which stem from various administrative archives provided by Statistics Norway. Each individual in Norway is assigned a unique identification number at birth or at immigration. This identification number, consistently de-identified for research purposes, facilitates the linkage of multiple data sources. Importantly, the data also include spousal identifiers, enabling us to aggregate household wealth and income metrics.

The income and balance sheet data, used primarily for wealth and income tax assessments, boast high reliability as financial institutions directly report each household's asset and liability positions to the tax authorities. Each source of income is reported and measured as the cumulative total throughout the calendar year. This data encompasses labor income, capital income, business income, pensions, and all other government transfers, along with taxes paid. Our key outcome variables include pre-tax labor earnings and a post-tax income measure, which also factors in unemployment insurance benefits.

Balance sheet components are reported by asset class and are measured as of December 31 each year. This data provides insights into deposits, mutual funds, both listed and non-listed stocks, cash, real estate, and other real assets. In subsequent analyses, "liquid assets" refer to the total of deposits, cash, mutual funds, and stocks, with the dominant component for most households being bank deposits (see Table 1). "Debt" includes the household's total debt amount, including student and car loans, consumer debt, and the primary component, mortgage debt.

In studies of household consumption dynamics, securing high-quality consumption expenditure data is crucial yet often challenging. While household surveys, used by studies such as Johnson et al. (2006) and Jappelli and Pistaferri (2014), offer direct measures of self-reported consumption, they are not without limitations like small sample sizes and potential measurement error (Meyer et al., 2015). We follow Fagereng and Halvorsen (2017) and Eika et al. (2020) in imputing consumption expenditure from Norwegian administrative data on income and wealth.[5] We calculate consumption expenditure as income net of savings, where the key challenge is to construct an "active" saving measure each period. Taking advantage of the panel data dimension, we

[5] For further details on the imputation procedure described here, see Appendix C. Other examples implementing this procedure include Browning and Leth-Petersen (2003) and Kreiner et al. (2014) using Danish data, and Koijen et al. (2015) and Kolsrud et al. (2020) using Swedish data.

Table 1
Summary statistics for targeted and non-targeted variables.

	Mean			Median		
	Unemployed	Control group	Possible controls	Unemployed	Control group	Possible controls
Targeted variables						
Demographics						
Age	44.7	44.7	46.1	45.0	45.0	46.0
Balance-sheet						
Male labor income	77,040	75,819	82,384	70,535	69,577	73,564
Debt	198,963	191,221	197,839	174,344	168,440	149,615
Liquid assets	35,828	34,543	124,112	16,882	16,460	27,824
Share homeowner (%)	0.94	0.95	0.89			
Share w/ risky assets (%)	53.71	53.70	59.32			
Education						
Low education	34.67	34.60	27.60			
High School Education	43.09	43.18	36.56			
Higher Education	22.24	22.22	35.39			
Industry composition						
Agriculture	0.60	0.44	0.82			
Education	2.74	2.29	4.99			
Health and social services	2.92	2.05	4.22			
Manufacturing and construction	37.46	39.36	27.74			
Other services	1.36	0.97	2.04			
Public admin. and defence	2.24	1.95	6.34			
Retail and services	47.19	48.67	31.93			
Unknown	5.49	4.27	21.91			
Employment						
Firm tenure	6.9	7.1	7.1	7.0	7.0	6.0
Non-targeted variables						
Demographics						
Share with children (%)	66.94	67.88	55.08			
Number of children	1.4	1.5	1.2	1.0	1.0	1.0
Balance-sheet						
Consumption	101,157	100,622	111,452	93,098	93,344	95,650
Spouse's labor income	47,198	46,885	45,644	48,020	47,724	45,932
Safe assets	25,576	24,098	49,389	12,616	12,145	19,098
Risky assets	10,252	10,445	74,723	192	220	1,117
Share receiving sickness benefits (%)	10.74	10.48	9.63			
Share receiving disability benefits (%)	0.22	0.33	3.41			
Employment						
Share public employer	0.08	0.06	0.22			
Estimated probability (%)	1.0	0.9	0.6	0.8	0.7	0.4
Days Unemployed	234			91		
Year unemployed	2007			2006		
N	11,497	147,027	251,618	11,497	147,027	251,618

Notes: Monetary values are CPI-adjusted with 2014 as base year, and measured in USD (NOK/USD = 6.3019). Possible controls include the full set of households that satisfy all sample selection criteria except actual job loss. Mean and median of the control group are weighted using CEM-weights, see Appendix E. All variables are measured two years prior to the year of job loss, except tenure and the probability of unemployment which is measured one year prior to job loss, and age which is measured in the year of job loss. The estimated probability of job loss is based on a probit regression with the following controls: public employer, tenure, education, industry, firm age, and firm size (see Appendix D for details).

impute active saving from the annual wealth change, assuming the return on households' risky assets (stocks and mutual funds) follows the general stock market. Interest earned on deposits and paid on debt is directly observable. We also observe housing market transactions and include net housing purchases in the imputation equation. To limit measurement error, we adhere to steps from previous literature: first, excluding households in the year of formation or dissolution, due to significant intra-year wealth movements, and second, excluding household-year observations where households own private businesses or farms due to poor measurement of both balance sheet components and income streams from these activities.

2.3. Identification of unemployment spells and sample selection

To identify unemployment spells, we use the registry of unemployment insurance benefit (UIB) applicants and recipients from the Norwegian Labor and Welfare Administration. We identify an individual as unemployed if they are registered as full-time unemployed for more than seven days (excluding both part-time unemployed and temporary furloughs), and require that the employer–employee relationship has ended within the last six months before registration as unemployed.[6] We exclude everyone who returns to the previous employer after an unemployment spell to remove seasonal workers who register as unemployed during the off-season from our sample.

For our analytical sample, we select married or cohabiting couples where the husband is identified as becoming unemployed between 1999 and 2014. Since consumption expenditure is not imputed in years of household formation or dissolution, and since the validity of the analysis will rely on the pre-trends of our measures, we require that the household has been married or cohabiting for four years leading up to the unemployment spell and that they remain married or cohabiting until one year after. We restrict attention to male unemployment to obtain a more homogeneous sample and to better be able to match the unemployed sample to a control group.[7]

We make two sample selection decisions that ensure that the households in the sample have stable labor market attachments. First, we restrict the sample to job losers who are eligible for unemployment benefits, and second, we require that neither adult in the household has had any type of unemployment spell in the four years leading up to the unemployment spell.[8] As discussed in the introduction, this requirement implies that our sample is one of high-tenured workers, who are likely to have a more persistent income loss associated with unemployment, which is important to keep in mind when interpreting our results. Since early retirement is widely available from 62 years of age, we restrict attention to workers who become unemployed at age 58 or younger in order to avoid selection bias where some workers choose early retirement instead of registering as unemployed (see Kyyrä and Wilke, 2007). Finally, we include only households for whom it is possible to impute consumption expenditure in the four years leading up to unemployment and in the year of unemployment, which implies the exclusion of households with business income or non-listed stocks.

In total, this yields a sample of 11,497 households (with male unemployment) that are eligible for our sample.[9] Summary statistics for this group can be viewed in columns 1 and 4 of Table 1. The median unemployment duration is 91 days, but the mean is higher, at 234 days, suggesting some individuals in our sample take a long time to find work or never return. Most unemployed workers come from retail and services (47%) or manufacturing (37.5%). The share of unemployed workers with high school education is 43%, while 22% have higher education. The average pre-tax labor income of the unemployed males is 77,000 USD, while the spouse on average earns 47,000 USD before taxes and transfers.[10]

3. Main empirical framework

3.1. A control group for the unemployed

The ideal experiment for investigating consumption responses after job loss would involve random, unexpected employee terminations. Since such settings are impractical, one alternative is using natural experiments where unemployment is quasi-randomly assigned. By conditioning on appropriate covariates, unemployment can be treated *as if* random. The literature using plant closures or mass layoffs to assess unemployment effects, bases its methodology on this notion, managing potential endogeneity issues up to a certain extent (Schwerdt, 2011). This approach minimizes selection issues in unemployment, yet it does not necessarily account for potential anticipatory (consumption) adjustments by workers foreseeing such layoff events, leading to possible bias in income and consumption correlations directly related to unemployment events. Moreover, using mass-layoff scenarios often leads to small sample sizes, making detailed analysis, such as relating portfolio composition to consumption responses, challenging.

Rather than adopting a mass-layoff approach, our strategy involves building a counterfactual through a control group method. Every worker who becomes unemployed is paired with a group of workers who, while similar in observables (taking advantage of extensive data to control for numerous variables influencing income growth and consumption), do not experience job loss during that period.[11] However, a latent concern persists regarding potential unobserved disparities, such as variations in ability or diligence, between displaced workers and their employed counterparts. To navigate this potential influence of unobserved variables on affecting income growth and consumption, we scrutinize the pre-trend of key variables. This aids in refining strategies to identify a control group that authentically mirrors the pre-trend of critical household income and balance-sheet components.[12]

---

6 If the job loser is registered with a new unemployment spell within 90 days after the initial spell ended, we consider this to be one unemployment spell.

7 In the data, the husband is the main income earner in about 9 out of 10 families.

8 The second restriction implies that if the husband has more than one subsequent unemployment spell within four years, only the first spell is included in our sample.

9 The number of observations is fairly small given that we use administrative data. A large fraction of the unemployment spells during the period are excluded when we impose the criteria of four years without unemployment prior to the observed unemployment spell, and a substantial share of these high-tenured workers are excluded when we further impose the sample selection criteria necessary for the imputation of consumption expenditure. In particular, the requirements of having no business income and being in a stable two-adult household reduce the number of observations substantially. In the Appendix (Figures A3 and A4) we undertake robustness analyses regarding these sample criteria, both regarding the job stability requirement (relaxing it from 4 to 2 years of job stability), and in terms of the selection criteria for imputation of consumption expenditures. The income path around unemployment is similar in shape, and size when we relax these criteria. However, the level of initial income is lower when we include single households, unstable relationships and individuals with business income.

10 According to Statistics Norway, the average (median) annual salary of workers in Norway in 2015, was 510,000 (460,200) NOK or about 81,000 (73,000) USD. For males, the corresponding numbers were 547,000 (470,000) NOK or about 87,000 (76,000) USD. See https://www.ssb.no/arbeid-og-lonn/faktaside/arslonn.

11 To sidestep the potential pitfall of conditioning on post-treatment outcomes during the selection of treatment and control groups, as cautioned by Krolikowski (2018), we do not preclude workers in the control group from encountering unemployment in *subsequent* years.

12 A related discussion and strategy for control group selection is found in Borusyak et al. (2017). See also Flaaen et al. (2019) for an alternative control group construction methodology.

3.2. High-dimensional near-neighbor matching

To identify the suitable control group we return to the detailed administrative data sources. The starting point of the matching procedure is a sample of households where we have imposed the same sample selection criteria as in Section 2 on the full population of households, the only difference being that these households do not become unemployed in that given year. This set of eligible households is termed "Possible controls".

The matching procedure employs both exact matching for discrete variables like education and home ownership, and interval-based matching ($\pm\alpha$) for continuous variables such as age and income. It incorporates age, education, and job tenure to comprehensively address labor market risks. To genuinely capture the likelihood of re-employment post-job loss, we select control group households that inhabit a labor market region similarly sized to their unemployed counterparts. Financial comparability between control and treatment groups is crucial, particularly when we want to study the spending response out of unemployment. Households are therefore selected for the control group based on similarity in liquid asset and debt levels on December 31st, two years before the job separation year, ensuring both groups are financially analogous. The households are also matched based on ownership of risky assets and housing, and income, measured at specific intervals before unemployment, sidestepping potential bias from strategic asset accumulation preceding unemployment episodes. We allow each control to be matched to multiple unemployed households and vice versa (n-to-m matching).[13]

3.3. Comparing samples on observables and pre-trends

From Section 2.3 we retrieve the sample of 11,497 households, which is our treatment group. In total our matching procedure selects 147,027 households from the set of Possible controls to be matched to these households. Table 1 allows a comparison of key characteristics among samples of displaced workers, the control group, and the set of possible controls from which the control group is chosen. The first panel of the table presents variables involved in the matching procedure. The set of possible controls is, on average (and at the median), slightly older, with higher income, and with different compositions when it comes to industry of occupation and education length. The balance sheet components of debt and liquid assets are more dispersed, both in terms of averages and median. The table shows that the matching procedure yields a sample closely resembling the sample of the unemployed in terms of observables, in most cases bringing both the average and median closer.

We also observe that the samples are aligned with respect to education and industry of employment, particularly within the sectors of manufacturing and construction, and retail and services where the discrepancies were most noticeable. Although the differences in labor income, debt, and financial assets means are still statistically significant between the unemployed sample and the chosen control group (see Appendix Table A3), they are now economically negligible.[14]

To further mitigate concerns regarding other unobserved differences between the samples that could confound our results, all our plots that follow later will include periods prior to unemployment. In addition, we here study the pre-trends of other key characteristics in the years leading up to unemployment, specifically. We use a simple event study specification:

$$Y_{i,t} = \sum_{k=-4}^{-1} \beta_k \mathcal{U}_{i,t}^k \times T_i + \varepsilon_{i,t}, \tag{1}$$

where i denotes household, calendar year is denoted by t, and time relative to the onset of unemployment is denoted by k. The dummy variable T_i indicates the individual belongs to the treatment group and does not change over time (i.e. the variable does not denote "treatment status"). Year relative to the onset of unemployment is indicated by $\mathcal{U}_{i,t}^k$ which takes the value one when k periods have passed since the year of unemployment, and zero otherwise. The error term is $\varepsilon_{i,t}$. We run this regression on the sample of matched treatment and control households.

Fig. 1 plots the development of income after tax, deposits, debt, the share owning risky assets, and housing in the pre-unemployment years. There are only minor observable differences between the treatment group and the control group. We see that there are virtually no visible differences in the development of male labor income, spousal labor income, debt, and deposits. Table A4 in Appendix F shows that although there are some statistically significant differences in the development year-by-year, the numbers are vanishingly small and economically insignificant.

Overall, the similarity of the unemployed and their matched group makes us confident in interpreting the chosen control group as a good proxy for the true counterfactual and we proceed using standard econometric techniques.

[13] This also implies that all of our regressions and statistics presented below are weighted, using CEM-weights, following Iacus et al. (2012). Details on how the weights are constructed, as well as the full breadth of details on the matching procedure and α-values for each variable chosen are available in Appendix E.

[14] In Appendix D we develop a statistical measure for job loss risk, based on observables such as tenure, firm age, education, and sector. This prediction is available for all households, and we include it in Table 1 "Estimated probability of job loss" to further assess the sample selection procedure. As we would expect, also this likelihood is aligned through the matching procedure. Appendix D also compares the distribution of probabilities between the different samples.

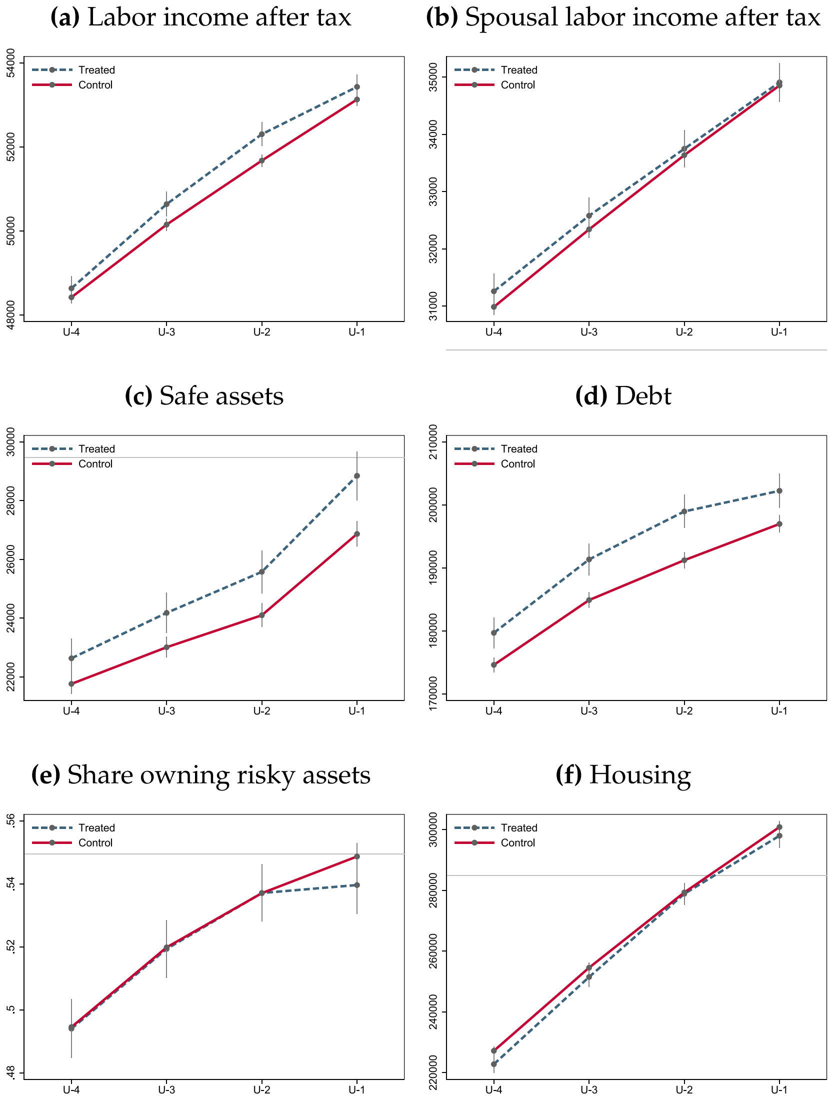

Fig. 1. Development of key observables before job loss.
Notes: Vertical lines show 95% confidence intervals, where standard errors are clustered at the matching group level. Observations are weighted using CEM-weight, described in Appendix E. Monetary values are CPI-adjusted with 2014 as base year, and measured in USD (NOK/USD = 6.3019).

4. Event study of job loss

In this section, we estimate the dynamic path of income, wealth, and expenditure in a four-year period after job loss. The regression specification is a simple difference in difference with staggered implementation:

$$Y_{i,j,t} = \alpha_j + \sum_{k\in\{-4:4\}} \beta_k U_{i,t}^k + \sum_{k\in\{-4,\dots,-2,0,\dots,4\}} \phi_k U_{i,t}^k \times T_i^k + \varepsilon_{i,t}, \tag{2}$$

where $Y_{i,j,t}$ denotes outcome variable in year t for household i belonging to matching group j. The matching group includes one household in the treatment group and their chosen set of control households. The outcome variable is regressed on a matching group fixed effect, α_j, a set of dummy variables $U_{i,t}^k$ indicating year relative to registering as unemployed, with $k \in \{-4, \dots, 0, \dots, 4\}$ denoting years passed since the onset of unemployment, a set of binary variables T_i^k indicating that the household is in the treatment group, and $\varepsilon_{i,t}$ is the error term which is assumed to be i.i.d-normally distributed. All standard errors are clustered at the matching groups, and the equation is estimated using weighted OLS.[15]

[15] See Appendix E for more details on the weights.

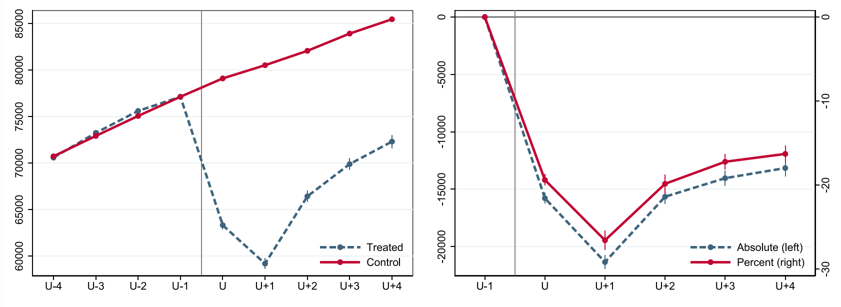

(a) Earnings, pre-tax

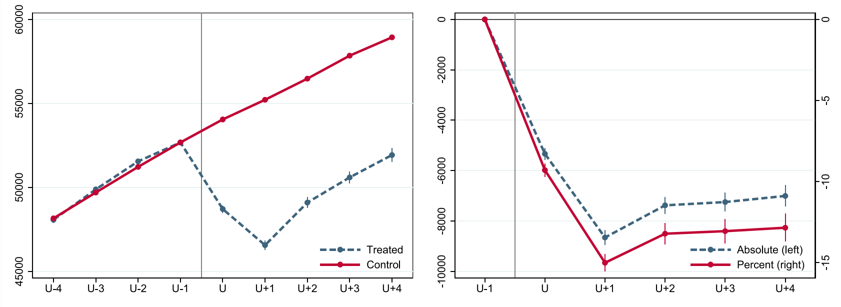

(b) After-tax labor income

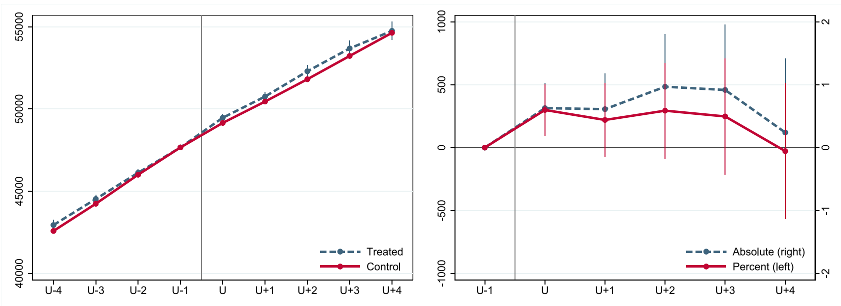

(c) Spousal earnings, pre-tax

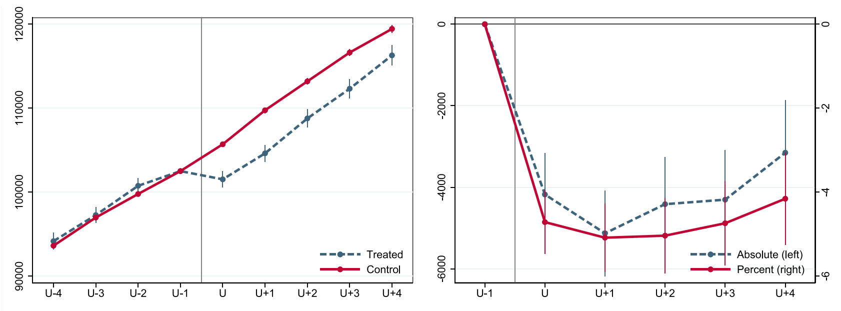

(d) Consumption expenditure

Fig. 2. Income loss and consumption expenditure responses after job loss.
Notes The left column shows the dynamic path of each outcome variable ($\beta_k + \phi_k \times T_i$), whereas the right column show the difference between the treated and control group (ϕ_k), measured both in absolute terms (left-hand side axes) and percentage change relative to pre-job loss average (right-hand side axes). Top and bottom 1% of observations are censored when estimating percentage change. Vertical lines show 95% confidence intervals, and standard errors are clustered at the level of the matching group. Observations are weighted using CEM-weight, described in Appendix E. Monetary values are CPI-adjusted with 2014 as base year, and measured in USD (NOK/USD = 6.3019).

Fig. 2 displays the results for key variables. The left column depicts estimated relative time dummies for treated and control groups, while the right shows the average treatment effect in 2014-USD and percentage deviation from pre-job loss averages.[16] Earnings drop post-job loss, averaging a decrease of over 15,000 USD upon unemployment registration and an additional 5,000 USD the following year. This second-year decline may stem from extended unemployment into the subsequent year, given the typical 24-month duration of UIB. With the income tax scheme's progressivity, the after-tax labor income (including unemployment benefits) impact is mitigated, totaling just over a 5,000 USD drop, or 13% of pre-job loss income, in the registration year.

Fig. 2 reveals that after-tax labor income does not recover four years post-unemployment onset, remaining more than 10% lower than the control group, mirroring findings from Jacobson et al. (1993) and similar studies. We identify a small statistically significant, but negligible in terms of economic significance, increase in spousal wage income in the year of job loss (third panel of Fig. 2). This is in line with the findings by Hardoy and Schøne (2014), Andersen et al. (2023), and Halla et al. (2020). Figures A5 and A6 in the Appendix show how other household income measures and income sources develop over the event window.

Although household income falls significantly both in the year of job loss and the year after, the last panel of Fig. 2 shows that the bulk of the fall in consumption expenditures (around 5%) happens in the first year of unemployment. They decline somewhat further in the subsequent year, but the bulk of the adjustment happens on impact, which is consistent with standard economic theory when unemployment is a mostly unforeseen and permanent shock to income. The size of the drop in consumption around unemployment is in the range with other estimates in the literature varying between 5 and 13% depending on the setting (e.g. Gruber, 1997; Browning and Crossley, 2001; Ganong and Noel, 2019; Landais and Spinnewijn, 2021; Andersen et al., 2023). Andersen et al. (2023) for instance find that spending drops on average around 5% during the first 24 months after the onset of unemployment using monthly data from Denmark, whereas (Landais and Spinnewijn, 2021), using data from Sweden, find a drop of more than 12% during the first year, and somewhat less than 10% in the consecutive 2–4 years.[17]

4.1. The importance of liquidity and debt

This section examines how income and spending responses to unemployment are associated with key household balance sheet components, centering on debt-to-income (DTI) and liquid-assets-to-income (LTI) ratios, indicative of financial constraints. The sample is divided into tertiles using DTI and LTI, derived from a two-year pre-unemployment household average, assessing respective income and consumption responses.[18] Ultimately, households are grouped into four distinct DTI and LTI categories for further analysis.

4.1.1. The separate importance of liquidity and debt

Panel (a) of Fig. 3 shows findings from the LTI split, revealing a stable initial pre-tax income drop in excess of 20% across groups but a slightly stronger recovery for the low-liquidity group 3–4 years after displacement. After-tax income trajectories underscore disparities, with the low-asset groups seeing a stronger recovery, possibly reflecting tax system progressivity. In terms of consumption, the most liquid group (with assets of on average 77,600 USD (Table 2)) encounters only a minor drop. The low and medium liquid asset groups exhibit larger and persistent drops in consumption up to 7 % through the sample period (Deaton, 1991; Carroll, 1997).

Panel (b) shows responses stratified by the debt-to-income distribution, with the more indebted households seeing a faster recovery of their income. But they also experience a larger and more lasting drop in consumption of up to 8%, whereas the least indebted households see a decline of 2%–3%. The group of the most indebted of households on average spend almost one-third of their income on mortgage commitments prior to job loss (Table 2), emphasizing these commitments' key role in household budget constraints.

These findings regarding the role of household heterogeneity are in line with the evidence from other studies, though a direct comparison is less straightforward. In the Danish setting, Andersen et al. (2023) define two groups of liquid assets (low and high), and find that those in the high asset group see a drop of up to 12% of their previous income, and the low asset group 9% over the first 24 months after the onset of unemployment. In our analyses using annual data, after-tax income falls between 8 and 15% across time and liquid asset groups. In terms of spending response, the high liquidity Danish households more or less perfectly smooth their consumption (which is close to our findings), whereas the other groups see larger than average declines.

This result is also mirrored in Landais and Spinnewijn (2021) finding that the 10% households with the most liquid assets prior to unemployment largely smooth consumption, whereas the remainders with some or none of liquid assets are not doing so. They also split the sample into 4 categories of initial debt holdings, and find that the group of the 25% most indebted households significantly reduces their consumption more than the average.

Ganong and Noel (2019) in their analyses find few differences in income responses between liquid asset categories based on a commercial bank estimate of liquid assets. Keeping in mind that they use a sample that stay unemployed, they find that in the initial phases, the high liquid asset group largely maintains consumption, whereas after the exhaustion of UI benefits at 6 months also the high liquid group suffers a more than 15% reduction in consumption. At the same time, the low liquid asset group faces an average reduction in consumption of more than 30% at the expiration of UI benefits.

[16] The pre-job loss average is household-specific and is an average of the two years prior to registering as unemployed.

[17] This persistence in the spending drop is consistent with our findings here, where we also follow workers up to 4 years after the onset of unemployment. The other studies discussed do not follow the workers beyond the horizons of 12 or 24 months.

[18] Given that the control group is selected (among other variables) along debt, liquid assets, and income margins, it inherently exhibits similar distributions in DTI and LTI as the unemployed sample. Various splits of these distributions have been tested, including division into quartiles, yielding consistent results. Refer to Appendix F for details.

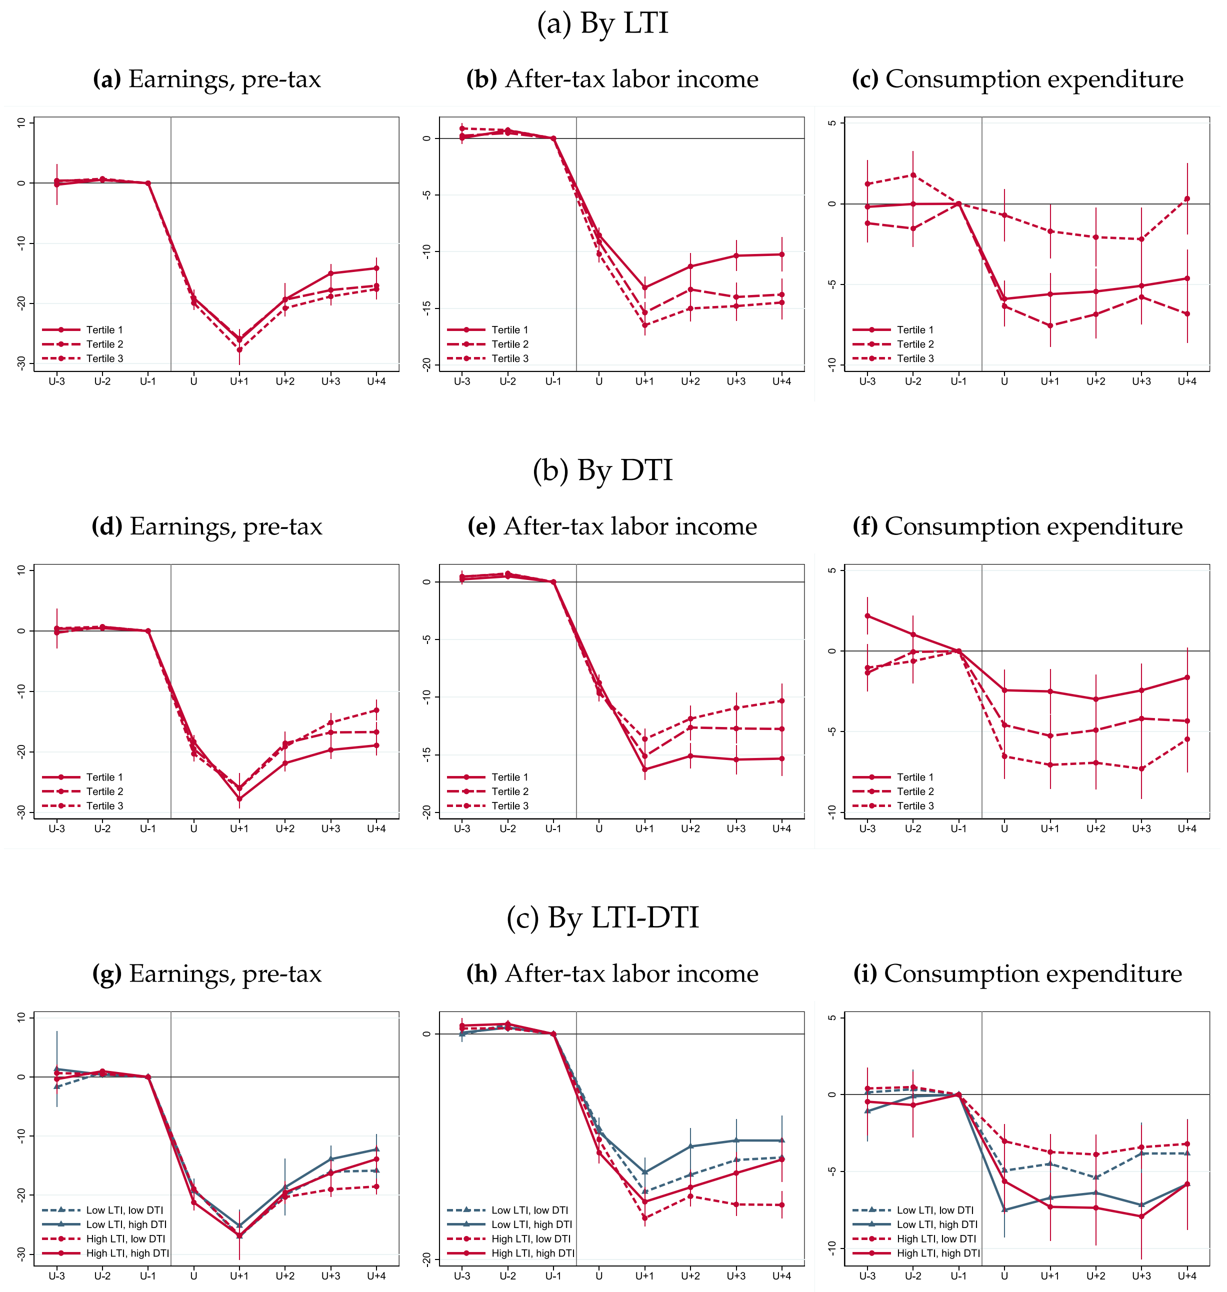

Fig. 3. Income loss and consumption expenditure responses across LTI and DTI groups.
Notes: Low (high) LTI refers to tertile 1 (2,3) of the distribution of liquid-assets-to-income. A high (low) DTI refers to tertile three (1,2) of the distribution of debt-to-income. Tertiles are computed for each calendar-year cohort of job losers using the mean of LTI or DTI in the two years preceding job loss. All variables are measured as percentage change relative to the pre-job loss average of the treatment group (year $U-1$). Top and bottom 1% of observations are censored. 95% confidence intervals. Standard errors are clustered at the matching group level.

4.1.2. The joint importance of liquidity and debt

We next examine the consumption responses arising from the interaction between the distributions of these two balance sheet components, an area where existing literature offers limited evidence. The bottom tertile of the LTI distribution, deemed closest to liquidity constraints, is defined as "low LTI", while the middle and top tertiles are grouped as "high LTI". Analogously, households on the upper spectrum of the DTI distribution are likely closer to credit constraints, categorizing the top tertile as "high DTI" and the bottom two as "low DTI".

Table 2
Summary statistics (averages) by LTI-DTI-groups.

Panel (a): By the separate distributions of LTI and DTI

	Low LTI	Mid LTI	High LTI	Low DTI	Mid DTI	High DTI
Demographics and unemployment						
Year	2006	2006	2006	2006	2006	2005
Age	42.5	43.2	45.4	47.9	43.4	39.9
Days unemployed	233	234	235	273	219	211
Balance sheet						
Debt	232,015	208,970	160,592	79,369	197,611	322,557
Debt-to-income	2.38	2.05	1.50	0.77	1.93	3.21
Liquid assets	5,372	19,467	77,639	54,210	27,487	20,838
Liquid assets-to-income	0.06	0.20	0.76	0.54	0.27	0.20
Interest paid	13,954	11,395	8,227	4,380	10,983	18,104
Interest paid /hh. income a.t.	0.15	0.11	0.08	0.04	0.11	0.18
Mortgage amortization /hh. income a.t.*	0.27	0.24	0.21	0.15	0.23	0.32
N	3,848	3,827	3,822	3,804	3,818	3,875

Panel (b): By the joint distribution of LTI and DTI

	Low LTI, low DTI	Low LTI, high DTI	High LTI, low DTI	High LTI, high DTI
Demographics and unemployment				
Year	2006	2005	2005	2006
Age	44.7	40.1	46.0	39.8
Days unemployed	254	211	243	211
Balance sheet				
Debt	158,861	312,515	131,312	331,561
Debt-to-income	1.63	3.21	1.25	3.21
Liquid assets	5,628	5,091	53,481	34,960
Liquid assets-to-income	0.06	0.05	0.53	0.34
Interest paid	9,609	18,736	6,996	17,538
Interest paid /hh. income a.t.	0.10	0.20	0.07	0.17
Mortgage amortization /hh. income a.t.*	0.22	0.32	0.19	0.32
N	2,016	1,832	5,606	2,043

Notes: Monetary values are CPI-adjusted with 2014 as base year, and measured in USD (NOK/USD = 6.3019). Low (high) LTI refers to tertile 1 (2,3) of the distribution of liquid-assets-to-income. A high (low) DTI refers to tertile three (1,2) of the distribution of debt-to-income. Tertiles are computed for each calendar-year cohort of job losers using the mean of LTI or DTI in the two years preceeding job loss. All variables are measured one year before job loss, except age and year of job loss, which is measured in the year of job loss. Mortgage amortization is a back-of-the-envelope calculation assuming 25 years downpayment plans starting at 30 years old and 5% interest rate, using group-averages for age and debt to calculate monthly payments: $amortization = \frac{interest\ rate \cdot debt}{1-(1+interest\ rate)^{55-age}}$.

The first two figures of Panel (c) reveal that high-LTI households generally experience a more gradual income recovery than their low-LTI counterparts, with the high-LTI low-DTI category revealing the most negative trajectory. Intriguingly, a reversal is noticeable in consumption response patterns. Despite experiencing a similar (or even slightly less negative) income decline, the high-LTI high-DTI group curtails consumption much more than the high-LTI low-DTI household. Again Table 2 illuminates the importance of mortgage commitments for the high-DTI groups. Still, it might be somewhat surprising that the high-LTI high-DTI group exhibits a relatively strong reduction in consumption. Potential explanations might derive from a precautionary savings motive if households now view the future as generally more uncertain, or perhaps from a reassessment of their permanent income. They might also conserve their liquid buffer for other investment opportunities or potentially for entrepreneurship. Exploring behavioral explanations, particularly those related to mental accounting biases in preserving home ownership, could also serve to explain the findings.

Given the complex interaction between liquid assets and debt, the observed income and consumption responses pose intriguing questions about household financial behavior and its further implications. Although it is beyond our scope to decipher all the nuances, it is compelling to ponder some implications of the findings for models and policy.

4.1.3. Implications for models and policy

The findings highlight the need to integrate household balance sheet dynamics into macroeconomic models, addressing shortcomings in standard Bewley and other heterogeneous agent models, especially concerning liquidity and debt. In contrast to these models, which tend to simplify the interplay between liquidity, consumption, and income shocks, it is crucial to adopt mechanisms, such as those proposed by Kaplan et al. (2014), that more accurately reflect complex financial decision-making.

The nexus between assets and debt in household behavior necessitates careful policy crafting. Policymakers must formulate policies that are attuned to and effective in diverse financial contexts. Our findings can inform policy design, such as unemployment insurance or post-crisis stimulus packages to stimulate aggregate demand, though one must apply them with caution to circumvent moral hazard when policy is linked to financial ratios like DTI or LTI. Rather than being applied off the shelf, our findings should serve as input into a wider policy consideration.

Another economic policy domain pertinent to our findings is macroprudential regulation. At first glance, our results suggest that DTI caps could soften consumption downturns for the unemployed. While such policies could temper the impacts identified in our study, it is crucial to thoroughly examine their broader implications, if implemented. For instance, a potential adverse effect would be if it led households to empty their liquid buffers to overcome a strict DTI cap when purchasing a home. Therefore, understanding both the immediate and the indirect impacts is essential when constraining households' financial leeway.

4.2. Heterogeneity along other margins

Navigating through further dimensions of heterogeneity, we explore the intersection of demographic variables and economic responses, specifically focusing on household compositions relative to child presence and age, as visualized in Fig. 4. The upper panels of this figure categorize households into subsets—those with and without children and further into varying child age groups. One observation surfaces: households without children experience a notably sharper income decline and a subsequently sluggish recovery compared to those with children. Arguably however it is hard to draw comparisons between these groups as households with and without children may be vastly different. Particularly, families with children under 17 years exhibit a rapid income recuperation, hinting towards a potential correlation with either the higher employment quality or an elevated income necessity driven by parental responsibilities. A notable divergence in consumption response is evident among households with younger children (below 5 years of age), showcasing a robust resurgence in consumption, indicative of a parental perspective that prioritizes essential consumption during these initial formative years. This pattern echoes (Carneiro et al., 2021), emphasizing the pivotal role of parental income in the early childhood phase (0–5 years) relative to later years (6–17).[19]

We also look into the difference between couples that are married vs cohabiting or by age (Figure A8), but find no significant differences in terms of the consumption responses.

5. The MPC out of unemployment

Every estimation of the marginal propensity to consume presupposes a level of income and consumption expenditure absent the income shock. Considering a framework of unobserved counterfactuals, define the outcome for individual (or household) i in year t as:

$$y_{i,t} = T_{i,t} \cdot y_{i,t}(T_{i,t} = 1) + (1 - T_{i,t}) \cdot y_{i,t}(T_{i,t} = 0) \tag{3}$$

Where T_i is 1 if the individual is unemployed, the observed outcome is $y_{i,t}(T_i = 1)$, and the unobserved counterfactual is $y_{i,t}(T_i = 0)$. The genuine treatment effect—i.e the job loss effect on income or spending (not the marginal propensity to spend)—is, $\tau^y_{i,t} = y_{i,t}(T_i = 1) - y_{i,t}(T_i = 0)$.

To estimate the unobserved counterfactual, we employ the household-specific control groups from Section 3.1, estimating income and spending growth using the control group's average growth rates, assuming that without job loss, a household would parallel the control group's average income and spending growth.

The income shock and the consumption response is constructed in the following way:

$$IncomeShock_{i,t} = INC^T_{i,t} - (1 + g^{INC}_{i,t}) * INC^T_{i,t-1}, \qquad \widetilde{\Delta C}_{i,t} = C^T_{i,t} - (1 + g^C_{i,t}) * C^T_{i,t-1} \tag{4}$$

where $INC^T_{i,t}$ is the observed after-tax labor income of the unemployed in household i in year t, and $C_{i,t}$ is observed consumption expenditure. Further, g^{INC} and g^C are the estimated growth rates of income and consumption expenditure in the control groups: $g^{INC}_{i,t} = \sum_{j=1}^{J_i} \frac{1}{J_i} \left(\frac{INC^C_{j,i,t}}{INC^C_{j,i,t-1}} \right)$, where $INC^C_{j,i,t}$ is the income of the male in household j in year t in the control group of household i, and J_i is the number of controls chosen for a treated household i. The growth rate of consumption is constructed in the same manner. The fundamental (although untestable) assumption is that had the treated household not encountered unemployment, their income and spending would have evolved akin to the control group's average, a presumption substantiated by observed pre-trend similarities between treatment and control groups (refer to Section 3).

5.1. The average marginal propensity to consume

To estimate the average MPC in the sample, we regress the consumption response $\widetilde{\Delta C}_{it}$ on the income shock, a constant, and a set of controls. Our baseline estimation equation is

$$\widetilde{\Delta C}_i = \beta_0 + \beta_1 IncomeShock_i + \mathbf{X}'_i \boldsymbol{\beta} + \varepsilon_i, \tag{5}$$

where $\mathbf{X}_i$ is a vector of control variables including a fourth-order polynomial in age, a second-order polynomial in the number of children below age 18, a dummy for whether the household is married or cohabiting, and a set of calendar year dummy variables.

The first panel of Table 3 reports the MPC out of an additional dollar lost due to unemployment within the year of job loss. The table presents results from six distinct specifications of the estimation equation, progressively and variably augmenting the control

[19] Carneiro and Ginja (2016) find that when they allow the parents' reaction to vary with the age of the child (in a partial insurance framework) the permanent income shocks have statistically significant effects only on inputs of children between ages 0 and 7.

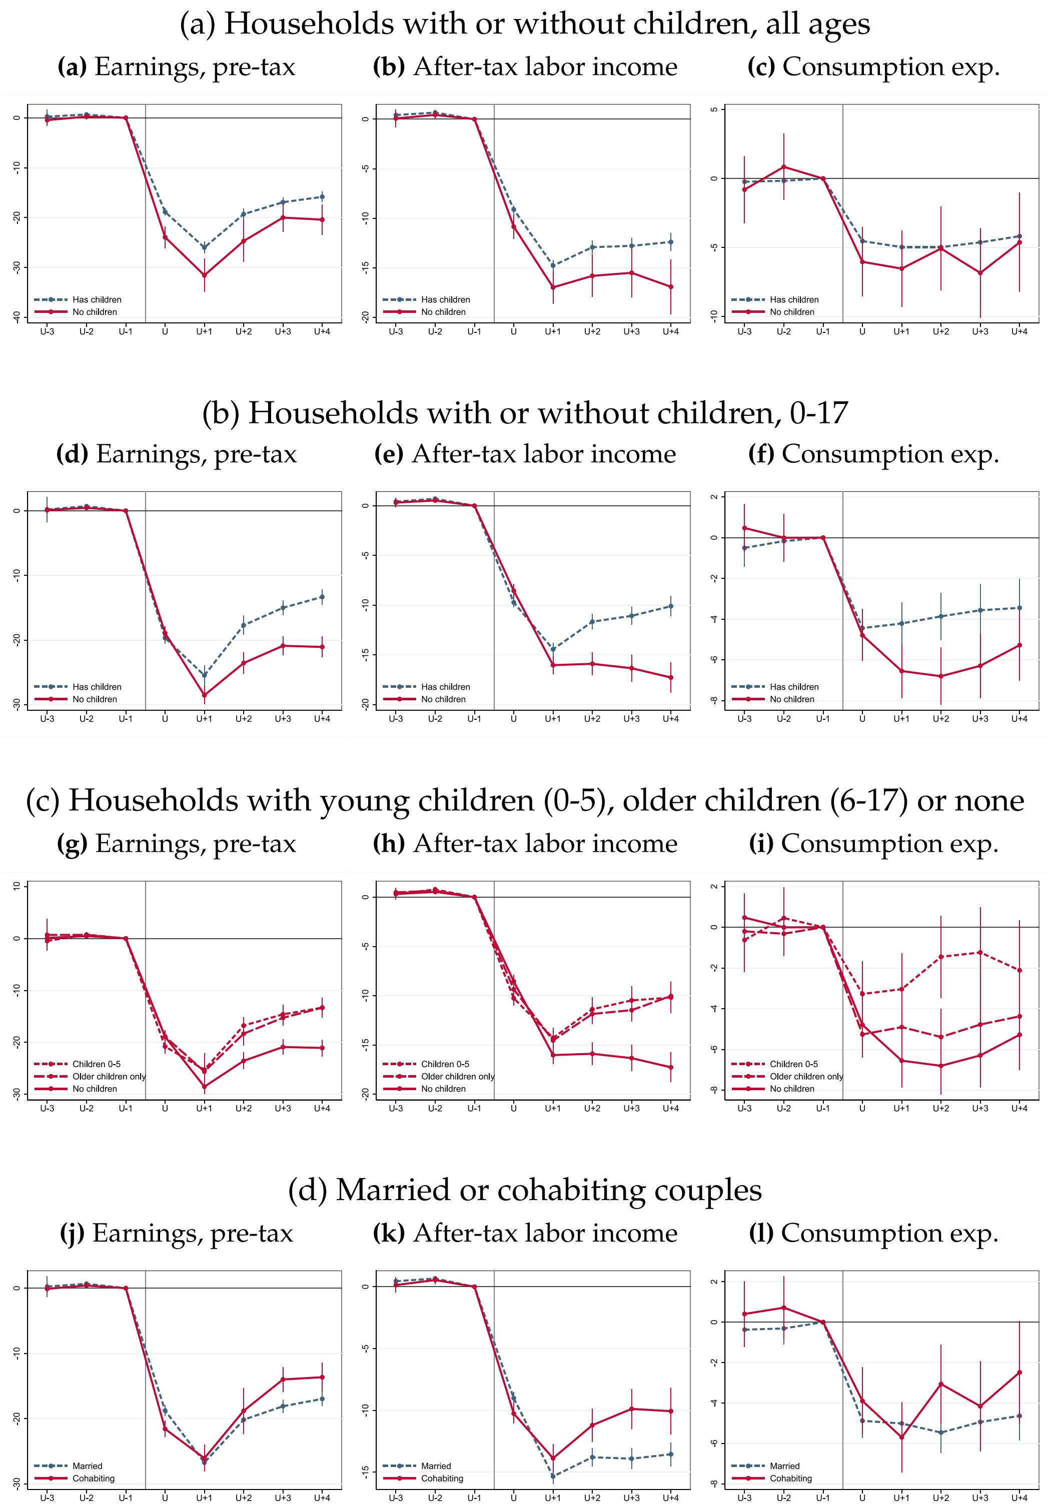

Fig. 4. Income loss and consumption expenditure responses by various margins.

Notes: All category variables are measured in the year of job loss, and groups are kept constant throughout the event window. All variables are measured as percentage change relative to the pre-job loss average of the treatment group (year $U-1$). Top and bottom 1% of observations are censored. 95% confidence intervals.

Table 3
The marginal propensity to consume within the year of job loss.

Panel (a)

	I	II	III	IV	V	VI
Income Shock$_t$	0.4420	0.4343	0.4332	0.4005	0.3993	0.4078
	(0.0230)	(0.0254)	(0.0247)	(0.0283)	(0.0276)	(0.0269)
Year fixed effects	No	Yes	Yes	Yes	Yes	Yes
Controls	No	No	Yes	Yes	Yes	Yes
HH income a.t.$_{t-1}$	No	No	No	Yes	Yes	Yes
Net wealth$_{t-1}$	No	No	No	No	Yes	No
Balance sheet$_{t-1}$	No	No	No	No	No	Yes
N	11,497	11,497	11,497	11,497	11,497	11,497

Panel (b)

	Low LTI	Medium LTI	High LTI
Income Shock$_t$	0.674***	0.310***	0.245**
	(0.149)	(0.0653)	(0.0803)
Income Shock$_t$*DTI$_{t-1}$=medium	−0.178	−0.0839	−0.148
	(0.135)	(0.0644)	(0.0912)
Income Shock$_t$*DTI$_{t-1}$=high	0.0610	0.213	0.238
	(0.174)	(0.125)	(0.142)
N	3,848	3,827	3,822

Notes: In Panel (a), controls include a fourth-order polynomial in age, a second-order polynomial of no. of children below age 18, and a dummy for marital status. "Balance sheet" in specification VI refers to conditioning on quartiles of debt, real assets, safe financial assets, and risky financial assets. We exclude observations with income shock in the top and bottom 1% and/or consumption response in the top and bottom 2.5%.
In panel (b), low (high) LTI refers to tertile 1 (2,3) of the distribution of liquid-assets-to-income. A high (low) DTI refers to tertile three (1,2) of the distribution of debt-to-income. Tertiles are computed for each calendar-year cohort of job losers using the mean of LTI or DTI in the two years preceeding job loss. The estimation in Panel (b) includes all controls from specification VI of Panel (a). *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively. Standard errors (in parenthesis) are clustered at the industry level.

variables set. Controlling for pre-job loss income level, in columns IV–VI, the coefficient shifts minimally across specifications. Column VI, our preferred specification, controls for pre-job loss debt, real assets, deposits, and risky assets.[20] We observe that for every dollar lost in the unemployment onset year, households, on average, curtail their spending by 41 cents, broadly consistent with existing literature.[21]

5.2. Heterogeneity in the MPC

To scrutinize the heterogeneity in the MPC, we do two key explorations: (i) segmenting the sample into tertiles based on distinct balance sheet components and worker characteristics, then re-estimating Eq. (5), integrating an interaction between the income shock and tertile dummies, and (ii) re-enacting (i) within each LTI-tertile, introducing controls for the tertiles of the DTI distribution, interacting with the income shock.[22] The insights, presented in Fig. 5, exhibit constrained variation in the MPC across male labor income (Panel 5a) and household income distributions (Panel 5b), and age (Panel 5c). Nevertheless, palpable heterogeneity in MPC comes to light upon examining net wealth, liquid assets, and debt.

Figs. 5e to 5h uncover a compelling, nearly linear distinction in MPCs among households with varying liquid assets and a U-shaped relationship in the distribution of debt-to-income. The non-monotonicity, though not dissected in detail, implies variations in credit access and refinancing options across debt levels. The second exercise, encapsulated in the second panel of Table 3, reinforces a persistent U-shaped relationship between DTI and the MPC across LTI distributions, with more elevated MPC evident in low-DTI and low-LTI households.

Delving into the practical implications, our findings again underline the imperative for macro models that encapsulate heterogeneity and potential asymmetry in consumer behaviors across different financial strata. By avoiding linear consumption function approximations and favoring models that highlight intersections between consumption, debt, and wealth, a more nuanced and realistic depiction of economic scenarios could be achieved. In the sphere of policy formulation, particularly in regard to unemployment insurance (UI) benefits and fiscal stimuli, it is crucial to acknowledge these MPC subtleties. As discussed in section 4.1.3 it remains important to use the findings as inputs for a comprehensive appraisal of policy alternatives.

[20] We control for each variable by dividing the sample into quartiles of the distribution and including a dummy variable for each quartile.

[21] Andersen et al. (2023) find that over a 24-month period Danish households reduce spending by 30% of their income loss from unemployment, whereas Landais and Spinnewijn (2021) exploit variation in welfare transfers across household types to find an estimate of the MPC when unemployed to be 55% in the Swedish setting, while Ganong and Noel (2019) find an MPC of 27% using variations in state-level UI benefits across US states.

[22] The estimation equation is given by $\widetilde{\Delta C}_i = \sum_{j=1}^{3} \left(\beta_j IncomeShock_i \cdot \mathbf{I}(Z_i = j) + \alpha_j \mathbf{I}(Z_i = j)\right) + \mathbf{X}_i'\delta + \epsilon_i$, where $I(Z_i = j)$ is an indicator function taking the value 1 if household i belong to tertile j of the variable of interest.

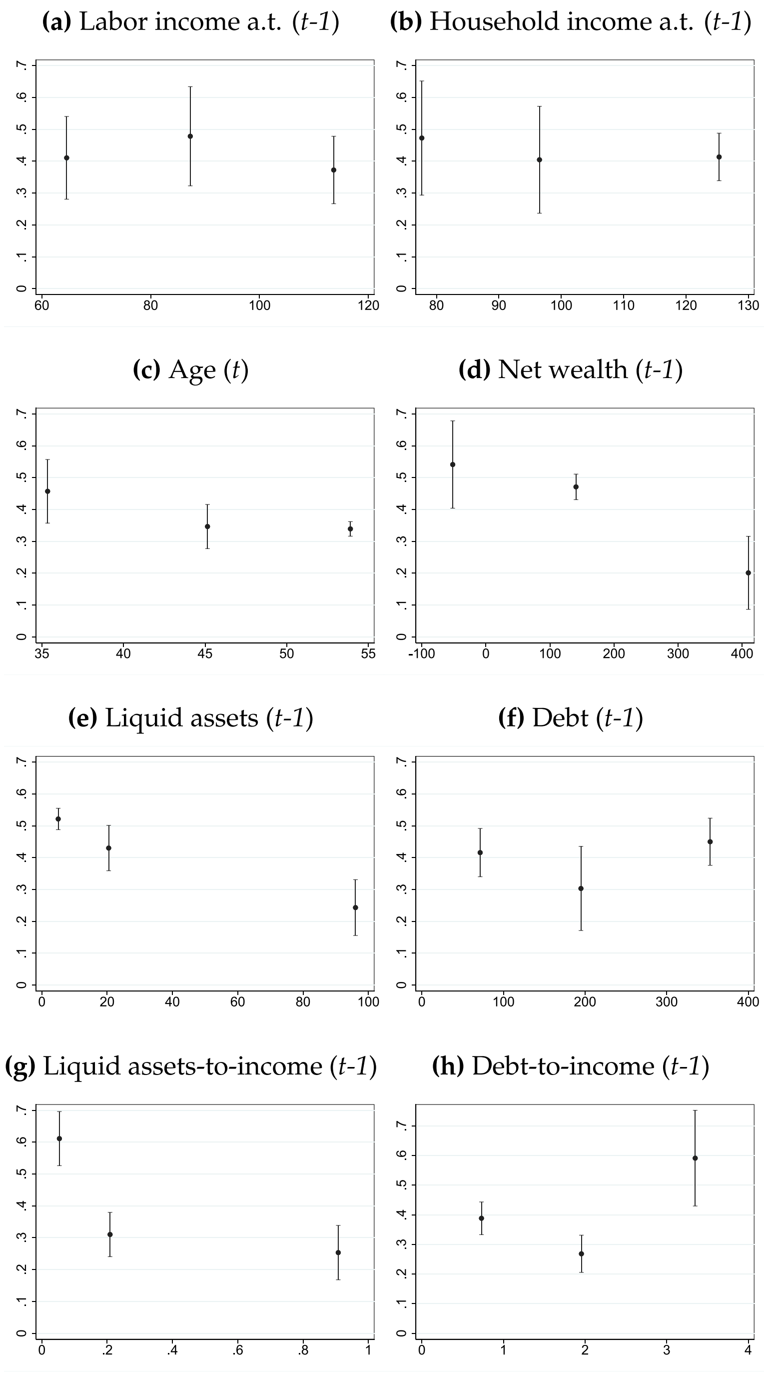

Fig. 5. The marginal propensity to consume across the distribution of income, age and wealth.
Notes The coefficients are obtained from regressions where the income shock is interacted with dummies for tertiles of the variable interest, including as controls a fourth-order polynomial in age, a second-order polynomial of no. of children, a dummy for marital status, a dummy for belonging to each tertile of the distribution, lagged household income, and lagged net wealth. The *y*-axis plots the mean of the variable within each tertile. All monetary values are CPI-adjusted with 2014 as base year, and measured in 1000 USD (NOK/USD=6.3019). Standard errors are clustered at the industry level, and vertical lines indicate 90% confidence intervals.

In Section 5.1, while our 'average MPC' findings resonate with established literature, a compelling narrative emerges upon contrasting our heterogeneity findings with those of others. Fig. 5 reveals a modest variance in the MPC across male labor income (Panel 5a) and household income distributions (Panel 5b), unexpectedly, given the escalation of income loss with income. Notably, Fagereng et al. (2021) observed the MPC varies with the size of an income shock, albeit in a different context: they explored

Table 4
The MPC and the importance of debt over the business cycle.

	I	II	III
Income Shock_t	0.421***	0.393***	0.394***
	(0.0451)	(0.0193)	(0.0206)
Income Shock_t*recession_t	0.0152		
	(0.125)		
Income Shock_t*$\text{DTI}_{t\text{-}1}$=medium		−0.126**	
		(0.0467)	
Income Shock_t*$\text{DTI}_{t\text{-}1}$=high		0.217**	
		(0.0704)	
Income Shock_t*Medium $\text{DTI}_{t\text{-}1}$*normal_t			−0.123**
			(0.0370)
Income Shock_t*High $\text{DTI}_{t\text{-}1}$*normal_t			0.205*
			(0.104)
Income Shock_t*Low $\text{DTI}_{t\text{-}1}$*recession_t			−0.00357
			(0.105)
Income Shock_t*Medium $\text{DTI}_{t\text{-}1}$*recession_t			−0.149
			(0.0802)
Income Shock_t*High $\text{DTI}_{t\text{-}1}$*recession_t			0.275
			(0.150)
N	11,497	11,497	11,497

Notes: We follow Aastveit et al. (2016), and define the following periods as recessions: 2001Q2–2001Q3, 2002Q3–2003Q1, 2008Q3–2010Q2. Controls include the interaction variable of interest, a fourth-order polynomial in age, a second-order polynomial of no. of children below age 18, year fixed effects, a dummy for marital status, and lagged household income, liquid assets, and net wealth. *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively. Standard errors are clustered at the industry level.

a positive, temporary income shock, contrasting with our focus on sustained, negative income loss due to unemployment. Thus, our unemployment-derived MPC notably diverges from the setting of a lottery win, prompting a question of alignment with previous MPC studies involving lottery winnings and tax rebates. For instance, our context links households with lower liquid assets to a higher MPC, mirroring findings from the previous consumption response studies using lotteries and tax rebates. Unlike the lottery study, however, we found a correlation between household leverage and consumption responses.

It is worthwhile to ponder the intrinsic differences between our study's context of persistent negative shocks—which potentially edge households closer to financial constraints—and those studies considering positive shocks, which typically ease financial burdens on households. There may be merit in dissecting the MPC's foundational factors and applications. For instance, it may be valuable to understand how or if insights derived regarding heterogeneity in responses from MPC studies focusing on positive (often temporary) shocks, such as lotteries and tax rebates, might be more pertinent to a setting of e.g. determining optimal stimulus packages.[23] Conversely, it would be intriguing to understand if MPCs (and their heterogeneous characteristics) stemming from unemployment may be more aptly applied to different model types and scenarios. We note this interesting contrast for future research consideration.

5.3. Variation over the business cycle

In our final exercise, we segregate the sample by employment status during a recession and investigate whether outcomes vary for job displacements in a recessionary context.[24] We investigate whether outcomes vary for job displacements in a recessionary context. Fig. 6 indicates that the reductions in earnings, after-tax labor income, and consumption expenditure are milder during recessions. This could imply that recessions affect a broad worker spectrum, while unemployment outside of recessions may more often result from, e.g., performance-based terminations (Gibbons and Katz, 1991). Job displacement during a recession might also send a less negative signal about worker quality to potential employers.

Next, we examine whether the MPC out of unemployment varies in a recession. Expanding baseline Eq. (5), we include a dummy variable for layoffs occurring during a recession, interacted with the income shock and DTI-tertile dummies. Recessionary household behavior could be influenced by various forces: mortgaged homeowners may benefit from falling policy rates due to floating interest rates, access to credit might be constrained, and plunging house prices may eradicate home equity for highly leveraged households. In this paper, we do not delve into the mechanisms behind potential heterogeneity in the MPC across business cycles, focusing instead on identifying observable differences in the MPC.

[23] The evidence to date stems from questionnaires of households MPC out of hypothetical transitory income shocks, see e.g. Bunn et al. (2018), Christelis et al. (2019) or Fuster et al. (2021).

[24] We use recession dates from Aastveit et al. (2016), as they use a "classical cycle" approach, using fluctuations in economic activity to classify recessions. This preferred approach is similar (though not identical) to the approach used by NBER, and "stricter" than the methods used by e.g. OECD. Thus, we define the following periods as recessions: 2001Q2–2001Q3, 2002Q3–2003Q1, 2008Q3–2010Q2. The results are robust to also defining the entire 2003 as a recession. Our results are partially robust to using OECD-defined recessions, but some conclusions differ since the recession dates only partially overlap.

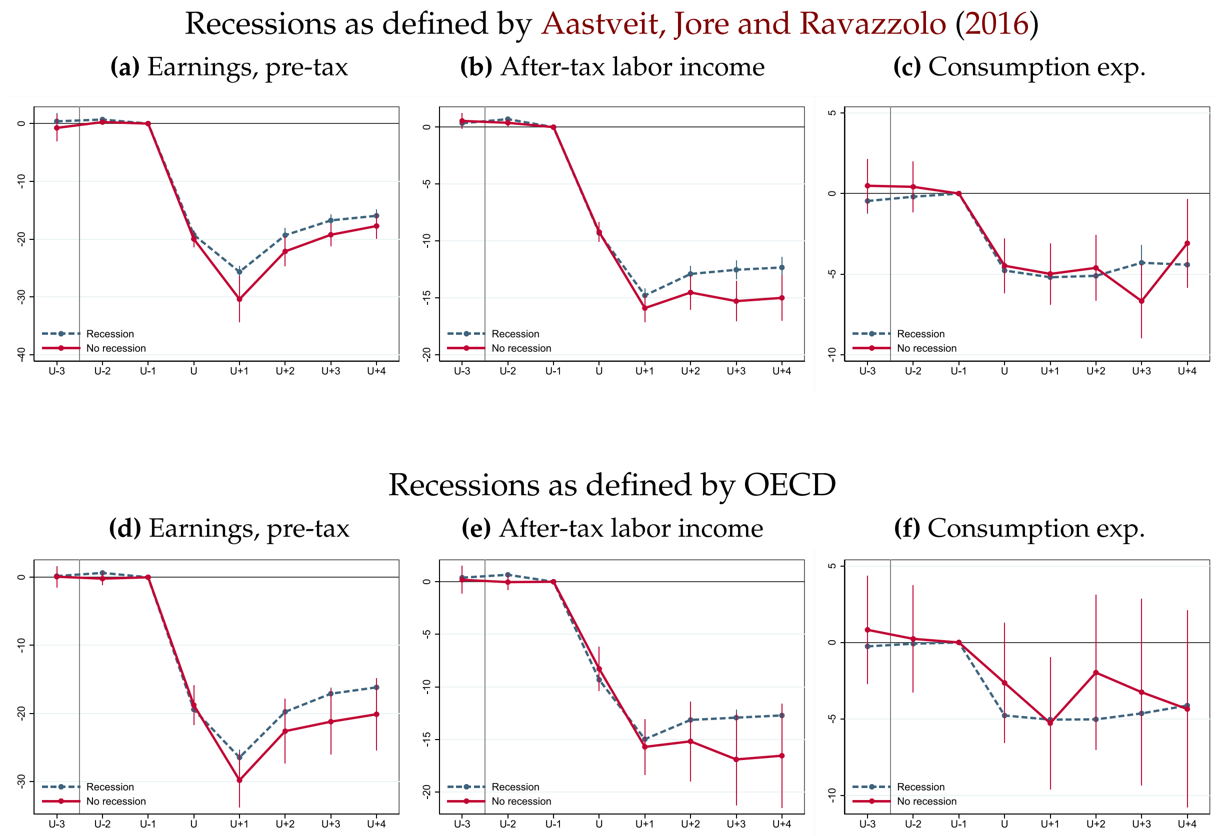

Fig. 6. Income loss and consumption expenditure responses over the business cycle.
Notes: Aastveit et al. (2016) define the following periods as recessions: 2001Q2–2001Q3, 2002Q3–2003Q1, 2008Q3–2010Q2. All category variables are measured in the year of job loss, and groups are kept constant throughout the event window. Top and bottom 1% of observations are censored. 95% confidence intervals.

Table 4 reports the estimation results. First, we find that the average MPC is somewhat higher for job losers in a recession, although the difference is small and the standard errors are wide.[25] This result is reported in the first column of Table 4, where the specification includes the recession dummy and an interaction between the income shock and the recession dummy. In the third column, we interact the income shock with dummies for DTI-tertile and the recession dummy. For households with a low DTI ratio, the MPC is virtually identical across the business cycle. For households in the mid-DTI group, the results indicate that those laid off in a recession have a somewhat lower MPC. High-DTI households, on the other hand, have a considerably higher MPC in recessions compared to normal times. However, the samples are too small for the differences to be statistically significant. Still, these results indicate that the U-shaped relationship between DTI and the MPC may be even stronger during recessions. Further, the results also warn us to be cautious about business cycle variations in household behavior when calibrating macroeconomic models.

6. Conclusion

In this research, we delved into household consumption patterns following job loss, harnessing detailed administrative records from Norway. These records, initially collated for tax purposes, enabled a comprehensive trace of household incomes, consumption expenditure patterns, and household balance sheet positions before, during, and after an unemployment spell.

Our analysis reveals that unemployment invariably leads to a persistent decline in income. Interestingly, households with minimal initial liquid assets or considerable debt generally witness gentler long-term downturns in both earnings and net income, even though the initial decline in earnings is consistent across all debt and asset ranges. This dip in income is paralleled by a marked decrease in consumption, equivalent to roughly one-third to one-half of the post-tax income drop. The reduction is milder for liquid households, but sharper for their indebted counterparts.

Many households with high liquidity also shoulder significant debt. Notably, while both debt and liquidity are associated with consumption responses, debt is found to wield a dominant influence among households with considerable amounts of both. Despite

[25] This is similar to Gross et al. (2020) who study the MPC out of liquidity and its variation over the business cycle.

their sizable holdings of liquid assets, these high-debt households noticeably curtail their consumption upon facing unemployment. Additionally, the data shows these households consistently allocate a significant portion of their disposable income to their mortgage commitments.

These revelations have ramifications for framing various economic policies, from unemployment benefits to fiscal stimuli (and implications for aggregate demand) and macroprudential regulations. While concerns such as moral hazard might make eligibility criteria problematic, our findings spotlight specific household segments with potentially heightened responsiveness to such policies. Hence, rather than being applied directly off the shelf, our findings should serve as input in a wider policy approach.

Within the realm of emerging macroeconomic models emphasizing the importance of micro heterogeneity, our results provide valuable insights. As underscored by Kaplan and Violante (2018) and Nakamura and Steinsson (2018), identifying key moments aids in distinguishing among competing models. Our exploration of consumption responses to unemployment and the role of debt and liquidity in household behavior post-job displacement enriches this discourse.

It is worthwhile to ponder the intrinsic differences between our study's context of persistent negative shocks—which potentially edge households closer to financial constraints—and those studies considering positive shocks (such as lotteries and tax rebates), which typically ease financial burdens on households. Our findings resonate with such studies underscoring the role of liquidity but differ somewhat regarding leverage. Our specific setting, including a persistent negative income shock, a mortgage market dominated by ARM mortgages, and mortgages with full recourse may influence our findings. Further understanding of the interplay of these conditions warrants exploration in upcoming studies.

Data availability

The authors do not have permission to share data.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jmoneco.2024.103578.

References

Aastveit, K.A., Jore, A.S., Ravazzolo, F., 2016. Identification and real-time forecasting of norwegian business cycles. Int. J. Forecast. 32 (2).

Addison, J.T., Teixeira, P., 2003. The economics of employment protection. J. Labor Res. 24 (1), 85–128.

Andersen, A.L., Duus, C., Jensen, T.L., 2016. Household debt and consumption during the financial crisis: Evidence from Danish micro data. Eur. Econ. Rev. 89 (October2016).

Andersen, A.L., Jensen, A.S., Johannesen, N., Kreiner, C.T., Leth-Petersen, S., Sheridan, A., 2023. How do households respond to job loss? Lessons from multiple high-frequency datasets. Am. Econ. J. Appl. Econ. 15 (4).

Baily, M.N., 1978. Some aspects of optimal unemployment insurance. J. Public Econ. 10 (3), 379–402.

Baker, S.R., 2018. Debt and the response to household income shocks: Validation and application of linked financial account data. J. Polit. Econ. 126 (4), 1504–1557.

Bertheau, A., Acabbi, E.M., Barceló, C., Gulyas, A., Lombardi, S., Saggio, R., 2023. The unequal consequences of job loss across countries. Am. Econ. Rev. Insights 5 (3), 393–408.

Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. Amer. Econ. Rev. 98 (5), 1887–1921.

Borusyak, K., Jaravel, X., Spiess, J., 2017. Revisiting Event-Study Designs: Robust and Efficient Estimation. The Review of Economic Studies rdae007.

Broda, C., Parker, J.A., 2014. The economic stimulus payments of 2008 and the aggregate demand for consumption. J. Monetary Econ. 68, S20–S36.

Browning, M., Crossley, T.F., 2001. Unemployment insurance benefit levels and consumption changes. J. Public Econ. 80 (1), 1–23.

Browning, M., Leth-Petersen, S., 2003. Imputing consumption from income and wealth information. Econom. J. 113 (488), 282–301.

Browning, M., Lusardi, A., 1996. Household saving: Micro theories and micro facts. J. Econ. Lit. 34 (4), 1797–1855.

Bunn, P., Roux, J. Le, Reinold, K., Surico, P., 2018. The consumption response to positive and negative income shocks. J. Monetary Econ. 96, 1–15.

Carneiro, P., García, I.L., Salvanes, K.G., Tominey, E., 2021. Intergenerational mobility and the timing of parental income. J. Polit. Econ. 129 (3), 757–788.

Carneiro, P., Ginja, R., 2016. Partial insurance and investments in children. Econ. J. 126 (596), F66–F95.

Carroll, C.D., 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. Q. J. Econ. 112 (1), 1–55.

Carroll, C.D., Slacalek, J., Tokuoka, K., 2017. The distribution of wealth and the marginal propensity to consume. Quant. Econ. 8 (3), 977–1020.

Chetty, R., 2006. A general formula for the optimal level of social insurance. J. Public Econ. 90 (10–11), 1879–1901.

Chetty, R., 2008. Moral hazard versus liquidity and optimal unemployment insurance. J. Political Econ. 116 (2), 173–234.

Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., VanRooij, M., 2019. Asymmetric consumption effects of transitory income shocks. Econ. J. 129 (622), 2322–2341.

Couch, K.A., Placzek, D.W., 2010. Earnings losses of displaced workers revisited. Amer. Econ. Rev. 100 (1), 572–589.

Davis, S.J., Von Wachter, T., et al., 2011. Recessions and the costs of job loss. Brook. Pap. Econ. Act. 42 (2 (Fall)), 1–72.

Deaton, A., 1991. Saving and liquidity constraints. Ecomometrica 59 (5), 221–248.

Dynan, K., Mian, A., Pence, K.M., 2012. Is a household debt overhang holding back consumption?[with comments and discussion]. Brook. Pap. Econ. Act..

Dynarski, M., Sheffrin, S.M., 1987. Consumption and unemployment. Q. J. Econ. 102 (2), 411–428.

Eika, L., Mogstad, M., Vestad, O.L., 2020. What can we learn about household consumption expenditure from data on income and assets? J. Public Econ. 104163.

Engen, E.M., Gruber, J., 2001. Unemployment insurance and precautionary saving. J. Monet. Econ. 47 (3), 545–579.

Fagereng, A., Halvorsen, E., 2017. Imputing consumption from norwegian income and wealth registry data. J. Econ. Soc. Meas. 42 (1), 67–100.

Fagereng, A., Holm, M.B., Natvik, G.J., 2021. MPC heterogeneity and household balance sheets. Am. Econ. J.: Macroecon. 13 (4), 1–56.

Farhi, E., Werning, I., 2016. A theory of macroprudential policies in the presence of nominal rigidities. Econometrica 84 (5), 1645–1704.

Flaaen, A., Shapiro, M.D., Sorkin, I., 2019. Reconsidering the consequences of worker displacements: Firm versus worker perspective. Am. Econ. J.: Macroecon. 11 (2), 193–227.

Fuster, A., Kaplan, G., Zafar, B., 2021. What would you do with $ 500? Spending responses to gains, losses, news, and loans. Rev. Econom. Stud. 88 (4).

Ganong, P., Noel, P., 2019. Consumer spending during unemployment: Positive and normative implications. Amer. Econ. Rev. 109 (7), 2383–2424.

Gerard, F., Naritomi, J., 2021. Job displacement insurance and (the lack of) consumption-smoothing. Amer. Econ. Rev. 111 (3), 899–942.

Gibbons, R., Katz, L.F., 1991. Layoffs and lemons. J. Labor Econ. 9 (4), 351–380.
Gross, T., Notowidigdo, M.J., Wang, J., 2020. The marginal propensity to consume over the business cycle. Am. Econ. J.: Macroecon. 12 (2), 351–384.
Gruber, J., 1997. The consumption smoothing benefits of unemployment insurance. Am. Econ. Rev. 87 (1), 192–205.
Hall, R.E., Mishkin, F.S., 1982. The sensitivity of consumption to transitory income: Estimates from panel data on households. Econometrica 50 (2), 461–481.
Halla, M., Schmieder, J., Weber, A., 2020. Job displacement, family dynamics, and spousal labor supply. Am. Econ. J. Appl. Econ. 12 (4), 253–287.
Hardoy, I., Schøne, P., 2014. Displacement and household adaptation: insured by the spouse or the state? J. Popul. Econ. 27 (3), 683–703.
Hendren, N., 2017. Knowledge of future job loss and implications for unemployment insurance. Amer. Econ. Rev. 107 (7), 1778–1823.
Huttunen, K., Møen, J., Salvanes, K.G., 2011. How destructive is creative destruction? Effects of job loss on job mobility, withdrawal and income. J. Eur. Econom. Assoc. 9 (5), 840–870.
Iacus, S.M., King, G., Porro, G., 2012. Causal inference without balance checking: Coarsened exact matching. Political Anal. 20 (1), 1–24.
Jacobson, L.S., LaLonde, R.J., Sullivan, D.G., 1993. Earnings losses of displaced workers. Am. Econ. Rev. 83 (4), 685–709.
Jappelli, T., Pistaferri, L., 2010. The consumption response to income changes. Annu. Rev. Econ. 2 (1), 479–506.
Jappelli, T., Pistaferri, L., 2014. Fiscal policy and MPC heterogeneity. Am. Econ. J.: Macroecon. 6 (4), 107–136.
Jarosch, G., 2023. Searching for job security and the consequences of job loss. Econometrica 91 (3), 903–942.
Johnson, D.S., Parker, J.A., Souleles, N.S., 2006. Household expenditure and the income tax rebates of 2001. Amer. Econ. Rev. 96 (5), 1589–1610.
Kaplan, G., Violante, G.L., 2014. A model of the consumption response to fiscal stimulus payments. Econometrica 82 (4), 1199–1239.
Kaplan, G., Violante, G.L., 2018. Microeconomic heterogeneity and macroeconomic shocks. J. Econ. Perspect. 32 (3), 167–194.
Kaplan, G., Violante, G.L., Weidner, J., 2014. The wealthy hand-to-mouth. Brook. Pap. Econ. Act. 7, 7–138.
Kawano, L., LaLumia, S., 2017. How income changes during unemployment evidence from tax return data. J. Hum. Resour. 52 (2), 418–456.
Kekre, R., 2022. Unemployment insurance in macroeconomic stabilization. Rev. Econom. Stud. rdac080.
Koijen, R., Van Nieuwerburgh, S., Vestman, R., 2015. 11. Judging the quality of survey data by comparison with truth as measured by administrative records. In: Improving the Measurement of Consumer Expenditures. University of Chicago Press, pp. 308–346.
Kolsrud, J., Landais, C., Nilsson, P., Spinnewijn, J., 2018. The optimal timing of unemployment benefits: Theory and evidence from Sweden. Amer. Econ. Rev. 108 (4–5), 985–1033.
Kolsrud, J., Landais, C., Spinnewijn, J., 2020. The value of registry data for consumption analysis: An application to health shocks. J. Public Econ. 189, 104088.
Korinek, A., Simsek, A., 2016. Liquidity trap and excessive leverage. Amer. Econ. Rev. 106 (3), 699–738.
Kreiner, C.T., Lassen, D.D., Leth-Petersen, S., 2014. Measuring the accuracy of survey responses using administrative register data: Evidence from Denmark. In: Carroll, C.D., Crossley, T.F., Sabelhaus, J. (Eds.), Improving the Measurement of Consumer Expenditures. University of Chicago Press, pp. 289–307.
Krolikowski, P., 2017. Job ladders and earnings of displaced workers. Am. Econ. J.: Macroecon. 9 (2), 1–31.
Krolikowski, P., 2018. Choosing a control group for displaced workers. Ind. Labor Relat. Rev. 71 (5), 1232–1254.
Krueger, D., Mitman, K., Perri, F., 2016. Macroeconomics and household heterogeneity. In: Handbook of macroeconomics, vol. 2, Elsevier, pp. 843–921.
Kueng, L., 2018. Excess sensitivity of high-income consumers. Q. J. Econ. 133 (4), 1693–1751.
Kyyrä, T., Wilke, R.A., 2007. Reduction in the long-term unemployment of the elderly: A success story from Finland. J. Eur. Econom. Assoc. 5 (1), 154–182.
Lachowska, M., Mas, A., Woodbury, S.A., 2020. Sources of displaced workers' long-term earnings losses. Amer. Econ. Rev. 110 (10), 3231–3266.
Landais, C., Spinnewijn, J., 2021. The value of unemployment insurance. Rev. Econom. Stud. 88 (6), 3041–3085.
McKay, A., Reis, R., 2021. Optimal automatic stabilizers. Rev. Econom. Stud. 88 (5), 2375–2406.
Meyer, B.D., Mok, W.K., Sullivan, J.X., 2015. Household surveys in crisis. J. Econ. Perspect. 29 (4), 199–226.
Mian, A., Rao, K., Sufi, A., 2013. Household balance sheets, consumption, and the economic slump. Q. J. Econ. 128 (4), 1687–1726.
Misra, K., Surico, P., 2014. Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. Am. Econ. J.: Macroecon. 6 (4), 84–106.
Nakamura, E., Steinsson, J., 2018. Identification in macroeconomics. J. Econ. Perspect. 32 (3), 59–86.
Olafsson, A., Pagel, M., 2018. The liquid hand-to-mouth: Evidence from personal finance management software. Rev. Financ. Stud. 31 (11), 4398–4446.
Schmieder, J.F., Von Wachter, T., Heining, J., 2023. The costs of job displacement over the business cycle and its sources: Evidence from Germany. Amer. Econ. Rev. 113 (5), 1208–1254.
Schwerdt, G., 2011. Labor turnover before plant closure: 'Leaving the sinking ship' vs. 'Captain throwing ballast overboard'. Labour Econ. 18 (1), 93–101.