



What can we learn about household consumption expenditure from data on income and assets?

TALL

SOM FORTELLER

DISCUSSION PAPERS

923

Lasse Eika, Magne Mogstad, and Ola L. Vestad

Lasse Eika, Magne Mogstad, and Ola L. Vestad

What can we learn about household consumption expenditure from data on income and assets?

Abstract:

A major difficulty faced by researchers who want to study the consumption and savings behavior of households is the lack of reliable panel data on household expenditures. One possibility is to use surveys that follow the same households over time, but such data are rare and they typically have small sample sizes and face significant measurement issues. An alternative approach is to use the accounting identity that total household spending is equal to income plus capital gains minus the change in wealth over the period. The goal of this paper is to examine the advantages and difficulties of using this accounting identity to construct a population panel data with information on household expenditure. To derive such measures of consumption expenditure, we combine several data sources from Norway over the period 1994–2014. This allows us to link tax records on income and wealth to other administrative data with information on financial and real estate transactions. Using this data, we derive household expenditure from the accounting identity, before assessing the sensitivity of this measure of consumption expenditure to the assumptions made and the data used. We then compare our measures of household expenditure to those reported in expenditure surveys and to the aggregates from national accounts. We also illustrate the research opportunities arising from the derived measures of consumption expenditure through two applications: the first is an examination of how relative wage movements among birth cohorts and education groups affected the distribution of household expenditure, while the second is a study of the transmission of income shocks to household consumption.

Keywords: administrative data, consumption measurement, income, wealth

JEL classification: C81, D12, D14, D31, D91, E21, G11

Acknowledgements: The project received financial support from the Norwegian Research Council (grant number 250516 and 227020). We are grateful for useful comments from three anonymous referees and the editor, from both the discussant and the participants at the Conference “New perspective on Consumption” at LSE, and from seminar participants at the Frisch Centre, the NHH, and the ISF. We are also grateful to Michael Graber for his help with the estimations in Section 4.2.

Address: Lasse Eika, Statistics Norway, Research Department. E-mail: lae@ssb.no

Magne Mogstad, University of Chicago Department of Economics; Statistics Norway, Research Department; NBER; IFS. E-mail: magne.mogstad@gmail.com

Ola L. Vestad, Statistics Norway, Research Department. E-mail: olv@ssb.no

Discussion Papers

comprise research papers intended for international journals or books. A preprint of a Discussion Paper may be longer and more elaborate than a standard journal article, as it may include intermediate calculations and background material etc.

© Statistics Norway
Abstracts with downloadable Discussion Papers
in PDF are available on the Internet:
<http://www.ssb.no/en/forskning/discussion-papers>
<http://ideas.repec.org/s/ssb/disppap.html>

ISSN 1892-753X (electronic)

Sammendrag

Mangel på paneldata med pålitelig informasjon om husholdningenes forbruk, sparing og formue er en utfordring for forskere som ønsker å studere husholdningenes forbruks- og spreadferd. Én mulighet er å bruke spørreundersøkelser som følger de samme husholdningene over tid, men slike undersøkelser er sjeldne og har små utvalg og store måleproblemer. En alternativ tilnærming er å beregne husholdningenes forbruk med utgangspunkt i registerdata, ved å bruke den regnskapsmessige identiteten som knytter forbruksutgifter sammen med inntekt, kapitalgevinster og endring i formue. Denne artikkelen undersøker muligheter og utfordringer ved en slik tilnærming. I tillegg til en rekke tradisjonelle registerdatakilder fra SSB brukes informasjon om eiendoms- og verdipapirtransaksjoner for å beregne årlige forbruksutgifter på husholdningsnivå for perioden 1994–2014. Artikkelen begynner med å beskrive hvordan vi går fram for å beregne et mål på husholdningenes forbruksutgifter, før vi analyserer følsomheten til dette målet for de ulike forutsetningene og datakildene som ligger til grunn for målet. Vi sammenligner deretter våre mål på forbruksutgifter med mål fra Forbruksundersøkelsen og med aggregerte mål fra nasjonalregnskapet. Til slutt illustrerer vi forskningsmulighetene som følger av de avledede målene for forbruksutgifter gjennom to anvendelser: den første er en analyse av hvordan endringer i relative lønninger mellom alders- og utdanningsgrupper samvarierer med endringer i relativt forbruk, mens den andre er en studie av overføringen av inntektssjokk til husholdningenes forbruk.

1 Introduction

A major difficulty faced by researchers who want to study the consumption and saving behavior of households is the lack of reliable panel data on household expenditures. The ideal data set is a long and representative panel of households with comprehensive information on consumption in conjunction with income, wealth, and other covariates. Unfortunately, such data are not available. For example, the Consumer Expenditure Survey (CEX) collects comprehensive information on the spending of US households, but it follows households for one year at most. Unlike the CEX, the Panel Study of Income Dynamics (PSID) follows households over longer periods of time, but it collects data only for a subset of consumption items. On top of this, expenditure surveys tend to suffer from problems such as small sample sizes, under-reporting of spending, and non-random response rates (Sabelhaus et al., 2015; Browning et al., 2014). Indeed, the correspondence between expenditure surveys and national accounts is deteriorating in many countries (Browning et al., 2014), and the time trend in consumption inequality is highly sensitive to the choice of data source and to how one attempts to correct for measurement problems (Attanasio and Pistaferri, 2016).

To overcome the lack of reliable panel data on household expenditures, Browning and Leth-Petersen (2003) suggest using administrative records on income and wealth to impute measures of consumption expenditure from the accounting identity that total household spending is equal to income plus capital gains minus the change in wealth over the period. Construction of such measures has a number of possible advantages. Tax records contain comprehensive information on income, and in some countries wealth information is recorded as well. Income and wealth components in tax records are often based on reports from third-parties such as employers, financial intermediaries, and government agencies. Administrative data often cover the entire population and follow households over time. But despite these advantages, it can still be difficult to derive a reliable measure of consumption expenditure.

A key challenge is that tax records rarely contain information about unrealized capital gains or the purchase and sales of assets. This means that additional data or strong assumptions are necessary to derive reliable measures of consumption expenditure from the accounting identity. One possibility is to try to distinguish between changes in net wealth due to unrealized capital gains (which do not change current consumption) and changes in net wealth due to a household saving some of its income (which reduces current consumption). To do so, existing research has used several alternative approaches that rest on different sets of assumptions about capital gains, asset prices, or investment behavior: zero capital gains on financial assets (Browning and Leth-Petersen, 2003; Kreiner et al., 2015); homogeneous returns and no within-year transactions (Fagereng and Halvorsen, 2017); or heterogeneous returns but no within-year transactions (e.g. Koijen et al., 2015). The first of these approaches is clearly at odds with the data, creating severe measurement errors in the measures of consumption expenditure (Browning and Leth-Petersen, 2003). The latter two approaches rely on weaker assumptions, but one may still be concerned with the reliability of the constructed measures of consumption expenditure. For example, Koijen et al. (2015) and Fagereng and Halvorsen (2017) drop nearly half of the population in Norway and Sweden to reduce the amount of errors in the measure of household consumption (e.g. households with negative values of the measure of household expenditure).

The primary goal of this paper is to examine the advantages and difficulties of using the accounting identity to construct a population panel data with information on household expenditure. To construct such measures of consumption expenditure, we combine several data sources from Norway over the period 1994–2014. This allows us to link tax records of the households’ income and wealth to data on their sales and purchases of assets. The key advantage of individual-level data on asset transactions is that they let the researcher directly observe net savings, i.e. capital gains minus the change in wealth. Thus, it is not necessary

to make assumptions about capital gains, asset prices, or investment behavior. We apply this approach to the Norwegian data, which provides detailed information about net savings in the two key assets of most households, real estate and stocks.¹ For other assets (bonds and equity funds), we follow Kojien et al. (2015) in assuming no within-year transactions while allowing for heterogeneous returns.

In Section 2, we set up the accounting identity framework, describe the procedure and data used to derive measures of household consumption expenditure, and illustrate the importance of distinguishing between expenditure on durable goods and the consumption flow that these goods provide. Section 3 explores the measurement errors that arise if one only uses tax records to derive measures of consumption expenditure, and compares our preferred measure of consumption expenditure to the measures based on expenditure surveys the national accounts. The findings from this section may be summarized in two broad conclusions. First, the combination of tax records on income and wealth and individual level information on financial and real estate transactions offers a unique opportunity to construct fairly reliable measures of household consumption expenditure. It is reassuring to find that the derived measure of household consumption expenditure compares well to aggregate measures from the national accounts. Second, deriving measures of household consumption expenditure based exclusively on tax records of income and wealth results in severe measurement errors. Without additional information on financial and real estate transactions, the derived measures of household consumption expenditure suffer from a large number of observations with non-positive consumption values as well as excessive dispersion and volatility in consumption. Short-cut approaches, such as excluding groups of households for which measuring consumption is likely to be particularly problematic, do not adequately resolve the problem of measurement errors.

In Section 4, we illustrate the research opportunities arising from the constructed population panel data on household expenditure through two applications. The first is a re-examination of the study of Attanasio and Davis (1996) of how relative wage movements among birth cohorts and education groups affected the distribution of household consumption. Their study uses data from expenditure surveys to construct a synthetic panel of consumption and wages. We show that the point estimates produced by a synthetic panel based on expenditure surveys are broadly similar to those one obtains using the population panel. However, the large sample sizes in the population data are needed to test and reject the hypothesis of full consumption insurance. The second application investigates the transmission of income shocks to household consumption using the partial insurance framework of Blundell et al. (2008). In this application we compare the results based on cross-sectional data, as used in Blundell and Preston (1998) and Blundell et al. (2013), to those one obtains using the population panel data. The results suggest the restrictions needed to use the cross-sectional data lead to significant biases both in the estimated variances of income shocks and in the transmission of these shocks to consumption.

2 Measuring household consumption expenditure

Our analysis considers the period 1994–2014. In every year, we observe the population of individuals registered as Norwegian residents. Using unique household identifiers, we construct household measures by aggregating information across all household members in a given year. Unless otherwise stated, the tables and figures

¹The derived measure of consumption expenditure does not depend on whether or not one observes or includes pension entitlements from (public or private) pension schemes. Pension entitlements are paid out as income after retirement, and included in the measure of income after tax and transfers. Social security contributions are included as taxes and contributions to employer-provided pension schemes are recorded as negative transfers.

are based on all residents.² To adjust for differences in size and composition, the household measures are equivalized.³ Each household member is then assigned an equal share of the equivalized household measure, abstracting from issues of unequal intrahousehold allocation of consumer goods.

2.1 Constructing a measure of household consumption expenditure

The household receives annual income E_{it} (labor income and cash transfers), and pays taxes τ_{it} . Let \mathbf{A}_{it-1} denote a portfolio (vector) of assets held by the household of individual i , where each asset component A_{ikt-1} is the level of asset k at the end of period $t-1$. To simplify notation, suppose that each A_{ikt-1} is held throughout the year, earning capital income $r_{kt}A_{ikt-1}$. Assume further that the household, at the end of the year, sells the assets \mathbf{A}_{it-1} at prices \mathbf{p}_t and buys a new portfolio \mathbf{A}_{it} at the same prices.

Given this notation, the following equation links current consumption and asset holdings to prior asset holdings and contemporaneous income net of taxes:

$$C_{it} + \sum_k p_{kt} A_{ikt} = \left(E_{it} - \tau_{it} + \sum_k r_{kt} A_{ikt-1} \right) + \sum_k p_{kt} A_{ikt-1}, \quad (1)$$

where C_{it} denotes the household's consumption expenditure throughout the year. We can rearrange (1) so that expenditure equals disposable income (labor income plus capital income minus taxes) minus the change in wealth plus capital gains:

$$C_{it} = \left(E_{it} - \tau_{it} + \sum_k r_{kt} A_{ikt-1} \right) - \sum_k (W_{ikt} - W_{ikt-1}) + \sum_k (p_{kt} - p_{kt-1}) A_{ikt-1}, \quad (2)$$

where $W_{ikt} = p_{kt} A_{ikt}$ and the final term on the right hand side is the capital gains on the portfolio held at the beginning of the year.

While our paper is centered on how to construct a measure of household consumption expenditure, it will also be useful to construct a corresponding measure of household savings. We define net savings S_{it} as the part of disposable income that is not used for consumption:

$$\begin{aligned} S_{it} &= (E_{it} - \tau_{it} + \sum_k r_{kt} A_{ikt-1}) - C_{it} \\ &= \sum_k (W_{ikt} - W_{ikt-1}) - \sum_k (p_{kt} - p_{kt-1}) A_{ikt-1} \\ &= \sum_k p_{kt} (A_{ikt} - A_{ikt-1}), \end{aligned} \quad (3)$$

where the second equality makes clear that net savings S_{it} equals the change in net wealth minus capital gains, and the third equality shows how S_{it} can be measured directly from data on sales and purchases of assets without having to measure or make assumptions about capital gains or changes in net wealth.

If all the components on the right hand side of equation (2) were observed, we could compute C_{it} directly. In reality, most but not all of these components are available in the Norwegian data. In the remainder of this section, we describe the empirical counterpart to equation (2), as summarized in Table 1, making precise what we can and cannot observe and how we proceed to measure C_{it} . Two central features of our measurement approach are particularly emphasized: the use of financial and real estate transactions to obtain direct measures of net savings, and the measures of consumption flows from housing and durable goods. The

²For each year, our full sample includes all individuals living in a household in which all members of age 18 and older filed a tax return in the current year and in the previous year.

³This is done by employing the so called EU equivalence scale. This scale divides the household measures by the sum of 1 for the first adult; 0.5 for each other adult; and 0.3 for each child under the age of 14.

latter allows us to construct measures that at least partly distinguish between the household’s spending and the consumption value of the flows of services from housing and durable goods.

2.2 Income components

Our data on income have several advantages over the income data available in many other countries. First, nearly all income components are third-party reported.⁴ Second, there is no attrition from the original sample because of the need to ask permission from individuals to access their tax records. In Norway, these records are in the public domain. Third, our income data pertain to all individuals, and not only to jobs covered by social security. Fourth, Norwegian administrative registers contain comprehensive and accurate information on most of the (taxable and non-taxable) components of income in Table 1. A notable exception is the value of owner-occupied housing services. Our procedure for estimating the value of owner-occupied housing services is described in Section 2.5.

Table 1 summarizes the key components of household income. In this table, we present each component as a share of gross household income (defined as the sum of labor income, cash transfers and capital income). On average, labor income makes up 66 percent of gross income, while government transfers and gross capital income account for around 20 and 11 percent, respectively. Figure 1 shows income shares at different percentiles of the disposable income distribution. As expected, government cash transfers are relatively important for the poor, labor income shares are increasing in disposable income until about the 90th percentile, and capital income becomes an increasingly important component at the top of the income distribution.

2.3 Measuring net savings

In order to derive reliable measures of household consumption expenditure, we not only need data on income but also information on either capital gains and changes in net wealth or on net savings. A key challenge is that the tax records contain information about the values of each asset at the beginning and end of the year, but not the unrealized capital gains or the net savings. In other words, these records offer information on W_{ikt} but not on A_{ikt} , p_{kt} , or S_{it} . To derive C_{it} from (2), however, it is necessary to distinguish between changes in net wealth due to unrealized capital gains (which do not change current consumption) and changes in net wealth due to a household saving some of its income (which reduce current consumption).

One possible way to address this missing data problem is to collect information on price changes ($p_{kt} - p_{kt-1}$) for each type of asset A_{ikt} . With data on changes in asset prices, it is possible to measure $\sum_k (p_{kt} - p_{kt-1})A_{ikt-1}$, provided that (i) every household earns the same $(p_{kt} - p_{kt-1})$ on a given type of asset A_{ikt-1} , and (ii) each A_{ikt-1} is held throughout the year. Together with tax data on $\sum_k (W_{ikt} - W_{ikt-1})$, it is then possible to distinguish between changes in net wealth due to unrealized capital gains and changes in net wealth due to a household saving some of its income. As a consequence, one may derive C_{it} from (2).

Another possibility is to measure net savings S_{it} , as defined in (3), from household level data on sales and purchases of assets. Since S_{it} equals the change in net wealth minus capital gains, it is then possible to derive C_{it} from (2), even in the absence of information on (changes in) asset prices and without invoking assumptions (i) and (ii). We now proceed to describe how we measure net savings in real estate and financial assets using a combination of these two approaches.

⁴Some of the components in “Other transfers” in Table 1 and Figure 1, such as lottery winnings and gifts and inheritances, are self-reported in the tax returns. Appendix Section D describes how we supplement the tax returns data with information from the inheritance tax registry to improve our measurement of gifts and inheritances.

Table 1: Income and Wealth Components Included in Equation (2)

Disposable Income $(E_{it} - \tau_{it} + \sum_k r_{kt} A_{ikt-1})$		Per capita: \$ 32,421
		Share of gross income (%)
=	Market Income Net of Taxes and Transfers $(E_{it} - \tau_{it})$	
=	Labor income	66.3
+	Government cash transfers	20.2
+	Other positive transfers	2.5
-	Other negative transfers	-1.3
-	Taxes	-22.7
+	Capital Income $(\sum_k r_{kt} A_{ikt-1})$	
=	Owner-occupied housing services	5.8
+	Dividends from securities	2.6
+	Other capital income	1.0
+	Interest on deposits	1.7
-	Interest on liabilities	-6.1
Capital Gains $\sum_k (p_{kt} - p_{kt-1}) A_{ikt-1}$		Per capita: \$ 10,560
		Share of gross capital gains (%)
=	Positive capital gains on real estate	71.0
-	Negative capital gains on real estate	-6.1
+	Positive capital gains on financial assets	29.0
-	Negative capital gains on financial assets	-20.3
Wealth (W_{ikt})		Per capita: \$ 122,638
		Share of gross wealth (%)
=	Value of real estate	70.4
+	Value of other real capital	3.7
+	Value of securities	12.4
+	Value of other financial assets	1.6
+	Value of liquid assets	11.9
-	Value of liabilities	-31.3

Notes: *Labor income* includes wage income, income from self-employment, sickness allowance, and parental leave benefits. *Government cash transfers* include pensions, unemployment benefits, housing allowance, educational grants, child benefits, and social assistance. *Other transfers* include inheritance, gifts, lottery winnings, alimony, and contributions to pension schemes. *Other capital income* includes returns on life insurance, taxable rental income, and capital income from abroad. *Other real capital* includes cars, caravans, boats, fixtures, chattels, and capital for business activities. *Other financial assets* include personal pension plans and life insurance. *Liquid assets* include cash and bank deposits. Each income (wealth) component is summed over all households and divided by the corresponding measure of gross income (wealth). The table reports averages over 1994–2014 of these ratios, along with per capita measures of disposable income, capital gains, and net wealth measured in 2014 USD.

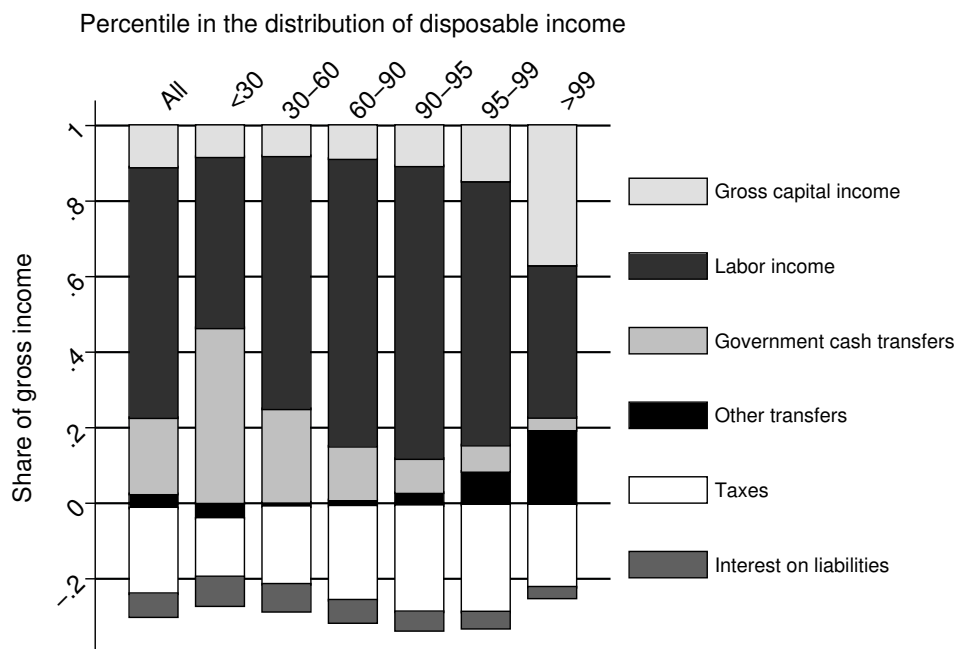


Figure 1: Income Components by Rank in the Income Distribution

Notes: This figure shows the relative size of different income components by rank in the distribution of disposable income. Each income component is summed over all households in each group, and then divided by total gross income (the sum of all positive income components) within the group. The figure reports averages over 1994-2014 of these ratios.

Our measure of net savings in real estate is based on transactions data from the Norwegian Land Register. For nearly all properties in Norway, this data set contains information on the last transaction prior to 1993. In addition, it records nearly all real estate transactions during the period 1993–2014. The data set provides detailed information about the transactions, including unique identifiers for both the seller, the buyer, and the property; the selling price; and characteristics of the property. Using this information, we construct measures of net savings (and capital gains) in real estate.

In addition to the data on real estate transactions, we have access to data covering all transactions of listed and unlisted stocks over the period 2003–2014. The data set includes unique identifiers for both the seller, the buyer, and the stock, as well as the transaction price. Using this information, we construct measures of net savings (and capital gains) in stocks.

The data on transactions in financial assets have two key limitations: They do not cover all types of financial assets, and stock transactions are not observed before 2003. To address these limitations, we use the tax data in combination with information on asset specific price changes as reported in the Norwegian Financial Accounts (used by Statistics Norway in the national account systems).

In Norway, the tax authorities collect information on the values of the vast majority of assets at the beginning and end of the year. Nearly all components of financial wealth are third-party reported (e.g., from banks and financial intermediaries). We are therefore able to accurately measure the values of most components of financial wealth, such as bank deposits, liabilities, and most securities. As shown in Figure 2, the tax data on financial wealth mirrors closely the aggregates from the Norwegian Financial Accounts.

To obtain measures of asset specific price changes, we consider four types of assets: bonds, equity funds, other listed securities, and unlisted securities. Assuming the same increase in prices within each type of

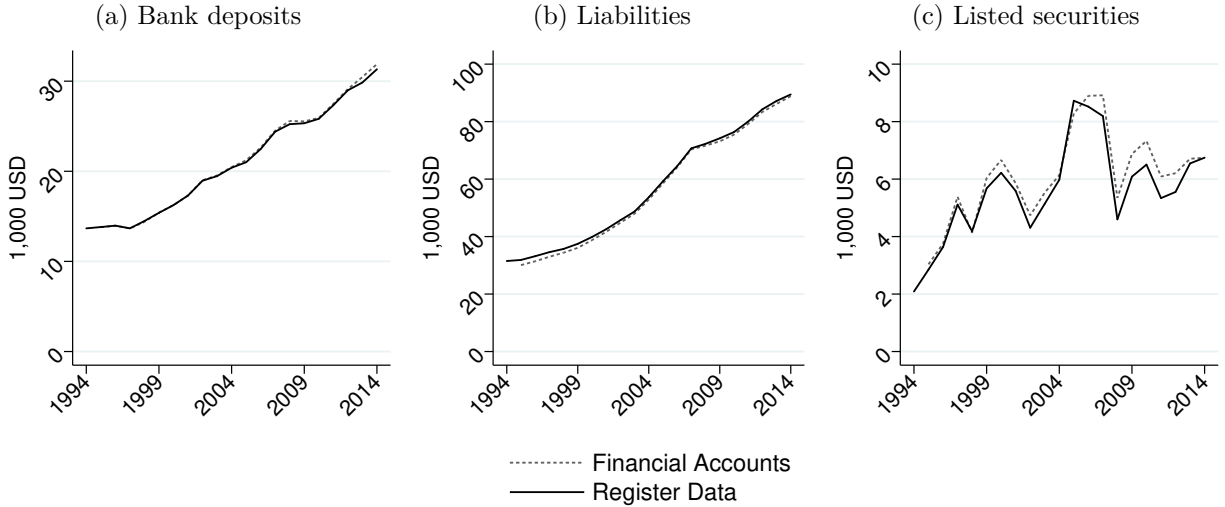


Figure 2: Per Capita Holdings in Bank Deposits, Liabilities, and Listed Securities

Notes: All measures are expressed in 2014 USD. Data from the Financial Accounts are not available for 1994.

assets, we can measure $(p_{kt} - p_{kt-1})$ from the Financial Accounts. For every household, we can then impute net savings (and capital gains) for each type of assets from information in the tax data on the values of assets held at the beginning and end of the year.⁵ Figure 3 compares our measures of capital gains (left panel) and net savings (right panel) in financial assets to aggregate measures reported in the Financial Accounts. It is reassuring to find a strong similarity both in levels and in the time trends, even prior to 2003 (the first year of transactions data). The discrepancies in 2002 and 2008 are almost entirely due to measurement error in capital gains for unlisted securities.⁶

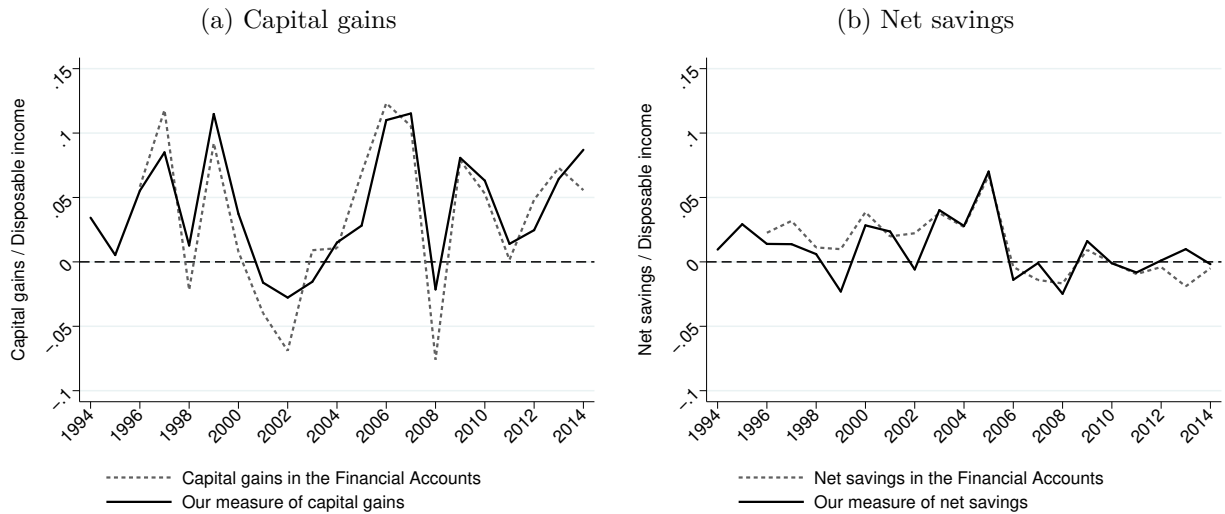


Figure 3: Capital Gains and Net Savings in Financial Assets against the Financial Accounts

Notes: This figure shows total capital gains on financial assets and total net savings in financial assets, both as fractions of total disposable income.

A natural question is how the measures of net savings and capital gains based on tax records differ

⁵We refer to Appendix Section B for further details on how we measure net investments in financial assets.

⁶Share ownership in a private company is usually difficult to value due to the absence of a public market for the shares. In Norway, the tax assessment value of an unlisted stock depends on the book value per share, reflecting the company's assets minus its liabilities (but not necessarily future earnings and risks).

from those we obtain by combining the tax records with transactions data. Panel (a) of Figure 4 shows the distribution of net savings, as defined in equation (3), if we measure savings in real estate based on tax records only and when we take advantage of the data on real estate transactions. This figure illustrates how the tax data overstate the variability in net savings across households. The primary reason is that the tax data record the full mortgage amount but not necessarily the actual market value of the property.

Panel (b) of Figure 4 shows distributions of net savings among the top 5% of the financial wealth distribution for three different approaches to handling capital gains on financial assets: using our preferred approach based on tax records, transactions, and prices; approximating capital gains with a measure of taxable capital gains from tax records;⁷ and assuming no capital gains. Going from one of the alternative approaches to our preferred measure of capital gains significantly reduces the dispersion in net savings, and shifts the distribution of net savings towards the left.

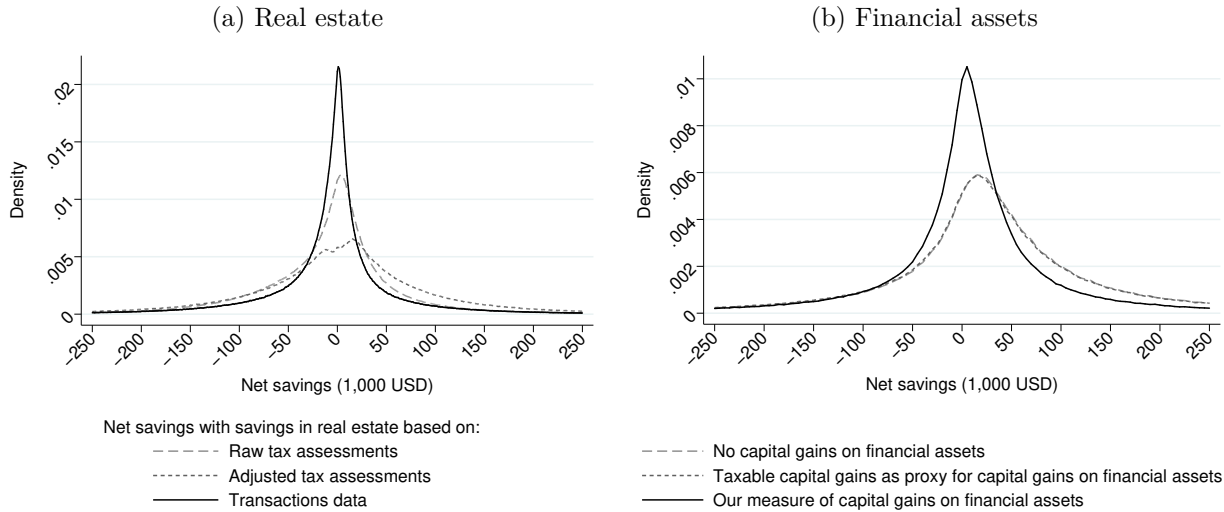


Figure 4: Distributions of Net Savings Using Different Measures of (a) Savings in Real Estate and (b) Capital Gains on Financial Assets

Notes: Net savings is defined as disposable income minus consumption. Measures are expressed in 2014 USD and adjusted for household size using the EU scale. Panel (a) displays the distribution of net savings when savings in real estate are derived from transactions data, and when savings in real estate are derived from tax assessments. “Adjusted tax assessments” are raw tax assessment values adjusted according to the aggregate ratio of selling prices to tax assessments. The sample is restricted to households trading real estate, and the figure is based on pooled cross-sections over the period 1994–2009. Panel (b) displays the distribution of net savings when capital gains on financial assets are measured based on tax records, transactions, and price changes; when taxable capital gains are used as a proxy for capital gains on financial assets; and when capital gains on financial assets are assumed to be zero. The sample is restricted to the top 5% of the financial wealth distribution, and the figure is based on pooled cross-sections over the period 2003–2014.

2.4 Measuring wealth

To derive household expenditure from equation (3), it is necessary to observe income and net savings, but not wealth. However, for many other purposes (e.g. to study consumption and savings behavior), it can be useful to know the joint distribution of income, consumption, and wealth. As explained above, the Norwegian tax records are accurately measuring the values of most components of financial wealth, such as bank deposits,

⁷The tax returns data contains a measure of realized capital gains and losses defined as the selling price minus the original purchase price, net of a tax deduction reflecting the returns that would have accrued had the same amount instead been invested in a safe asset. Hence, in a given year, taxable realized capital gains include asset appreciation that may have occurred years or decades earlier. For the purpose of deriving a measure of annual consumption expenditure, however, one needs to measure capital gains at accrual. That is, capital gains should be measured as the increase or decrease in the value of capital assets in each year regardless of whether assets were sold for a (taxable) realized capital gain.

liabilities, and most securities. The key exception is real estate. We now describe how we supplement the tax records with other data sources to infer the market values of real estate.

In principle, the Norwegian tax authorities are supposed to assess a property at a certain percentage of its fair market value.⁸ Prior to 2010, however, the tax assessment values differ significantly from the actual market values, and these differences vary considerably across properties depending on a wide range of factors such as area, year of construction, and housing type. As part of a tax reform in 2010, the Norwegian Tax Administration reassessed all residence values based on a price per square meter calculated by Statistics Norway (using hedonic pricing models with information on property type, size, geographic regions, last sales date, age of building). While this improved the quality of the tax assessment values on residences, differences between tax assessments and market values for individual dwellings remain a serious concern. Moreover, tax assessments of most real estate other than residences, such as recreational properties, farms, and land, have not been revised.

The starting point for measuring the market value of real estate is our transactions data. These data give information on the market prices of most properties at one or several points in time. At other points in time, however, it is necessary to make assumptions about how prices evolve. To do so, we combine our data on the characteristics of the properties with house price indices for specific regions and types of properties. These price indices allow us to interpolate and extrapolate the market value to years in which a property is not traded. There are, however, properties that are either (a) never sold during the period 1986–2015 or (b) not included in the transactions data. Prior to 2007, for example, we do not observe the transactions of units in housing cooperatives. About 13% of households own a unit in a housing cooperative. For primary residences subject to (a) or (b), we use the post 2010 tax assessment values to get one or more measures of the market values for the years 2010–2014. Once we have these measures, we apply the price indices to estimate market values for the previous years. We refer to Appendix Section A.1 for further details on our procedure to measure the market value of real estate.

Figure 5 compares our estimates of the market values of households' primary residences to those reported in the Survey on Living Conditions. In 2004 the survey asked a representative sample of households about the expected market value of their primary residence. It is reassuring to find that our estimates mirror closely the self-reported values, both across the distribution of the value of primary residences (Panel (a)) and across the disposable income distribution (Panel (b)). By comparison, tax assessment values differ significantly, even if we adjust the tax assessment values according to the aggregate differences between selling prices and tax assessments in 2004. The tax assessment values are especially inaccurate in the middle and upper parts of the income distribution.

Table 1 and Figure 6 show that real estate is the key component of gross wealth for most households. As expected, financial assets become increasingly important in the upper part of the wealth distribution. For households in the top percentile, financial assets account for more than 50 percent of gross wealth.

⁸In 2014, for example, the tax assessment value was supposed to be 25 percent of the property's value for a primary residence, and 60 percent of the property's value for secondary residences.

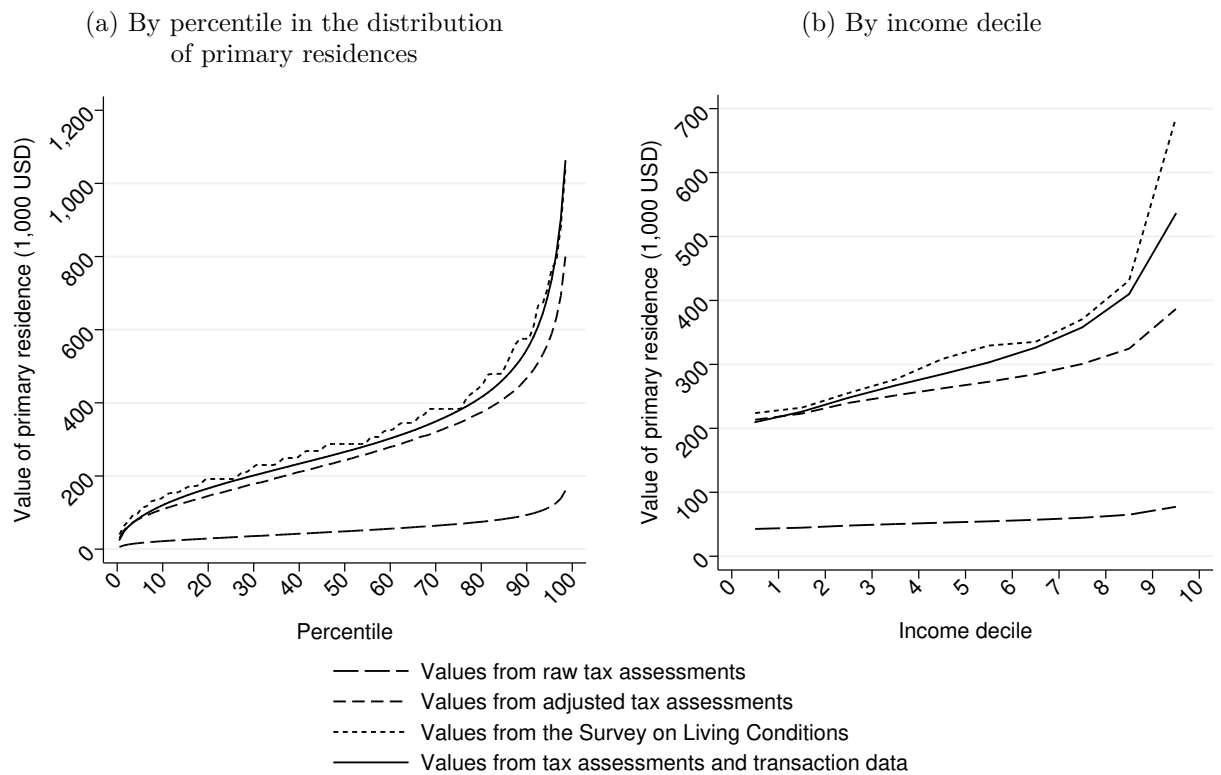


Figure 5: The Value of Primary Residences Based on Different Sources of Data

Notes: This figure displays the average value of primary residences; by percentile in the distribution of primary residences in Panel (a) and by income decile in Panel (b). The value of primary residences is measured based on (i) tax assessments only (raw and adjusted); (ii) the 2004 Survey on Living Conditions; and (iii) tax assessments and transactions data. “Adjusted tax assessments” are raw tax assessment values adjusted according to the aggregate ratio of selling prices to tax assessments. The percentage of households owning a residence is 75.7, 82.0, and 77.5, according to (i), (ii) and (iii), respectively. The sample includes all households owning a residence in 2004. In Panel (a), the top percentile is dropped.

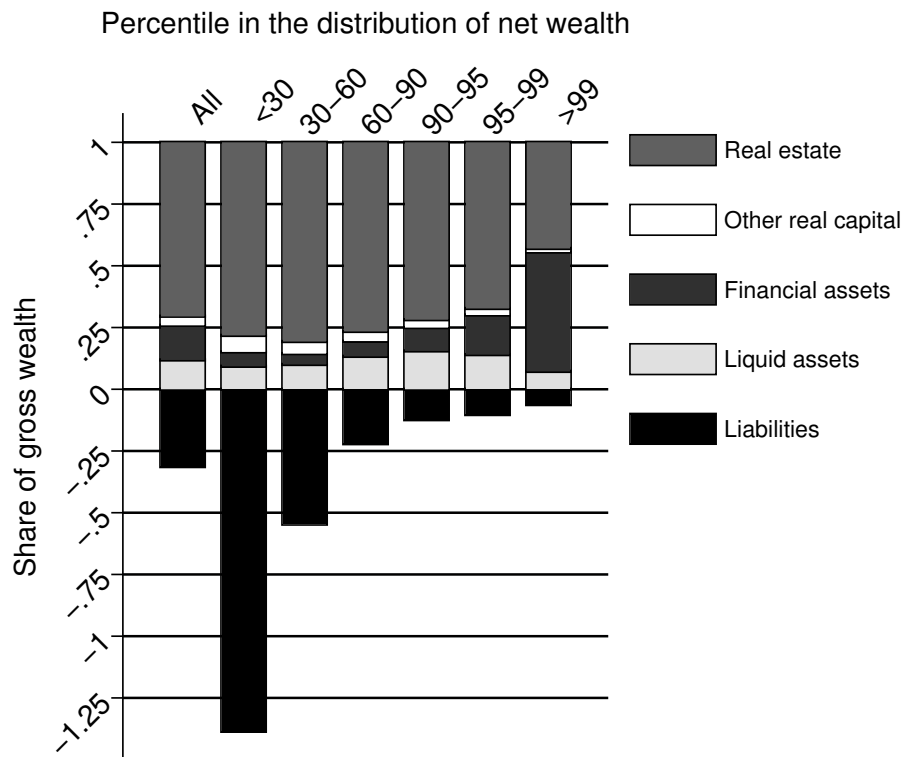


Figure 6: **Wealth Components by Rank in the Wealth Distribution**

Notes: This figure shows the relative size of different wealth components by rank in the distribution of net wealth. Each wealth component is summed over all households in each group, and then divided by total gross wealth (the sum of all positive wealth components) within the group. The figure reports averages over 1994–2014 of these ratios.

2.5 Consumption versus Expenditure

So far, we have used the terms “consumption expenditure” and “consumption” interchangeably. However, these terms are distinct economic concepts: Consumption expenditure refers to the money spent on consumer goods and services, whereas a broader concept of consumption would also include goods and services produced at home, in-kind transfers, and the flow of services generated by housing and durable goods. While our data do not allow us to capture all these components, we are able to construct measures of consumption that reflect the flow of services from certain durable goods and the housing services consumed by homeowners.

2.5.1 Value of owner-occupied housing services

If a household buys a housing unit that they previously had occupied as renters, the household’s consumption expenditure will go down by the amount of rent they used to pay to the former landlord, while the consumption of housing services remains unchanged. By including owner-occupied housing services, measured as the amount a homeowner household would pay to rent their house in the market, we obtain a measure of consumption that is comparable across renter and homeowner households.⁹ Specifically, we use the rental equivalence approach, which attempts to estimate the flow of services for an owner dwelling based on information about market rents for rented dwellings. Statistics Norway estimates the value of owner-occupied housing services in the national accounts as the rental price (net of expenses for maintenance and insurance) for all owner-occupied residences, based on a representative sample of renter-occupied housing units. We distribute this aggregate value across households according to each household’s share of the total value of primary residences. As a result, all variation in the value of owner-occupied housing services across homeowner households within a given year is due to variation in the value of primary residences, while the rent-to-value ratio is a year-specific constant common to all homeowner households.¹⁰ The average rent-to-value ratio for 1994–2014 is 2.88 percent.

Figure 7 shows the impact of including the value of owner-occupied housing services in our measure of capital income, with homeowners in panel (a) and renters in panel (b). This figure shows that including the value of owner-occupied housing services significantly reduces the proportion of homeowners with negative capital income (which includes interest payments on debt). Including the value of owner-occupied housing services also eliminates the bunching of homeowners at zero capital income (consisting of homeowners who have paid off their mortgage debt). Comparing panels (a) and (b), it is evident that renters have relatively low levels of capital income.

The treatment of owner-occupied housing services also has visible impacts on the cross-sectional distribution of consumption. This can be seen in Figure 8, where our preferred measure is denoted “Consumption”. Excluding the value of owner-occupied housing services reduces measured income and consumption for all homeowners, and results in a shift towards the left in the distribution of consumption.

Panel (a) of Figure 9 shows time series of consumption, with and without the value of owner-occupied housing services, for households that are initially renters but purchase a house in year 0. When including the value of owner-occupied housing services, median consumption increases steadily throughout the period. For consumption measured without the value of owner-occupied housing services, in contrast, a steady increase in median consumption is interrupted by a decline in the year of purchase, reflecting a reduction in rent

⁹An alternative to including owner-occupied housing services in imputed consumption is to perform separate analyses by homeowner status, as in Browning and Leth-Petersen (2003), Kreiner et al. (2015), and Koijen et al. (2015).

¹⁰In Appendix Section A.2, we show that switching to heterogeneous rent-to-value ratios based on actual rents and residence characteristics reported in the 2015 Rental Market Survey has very little impact on the cross-sectional distribution of consumption.

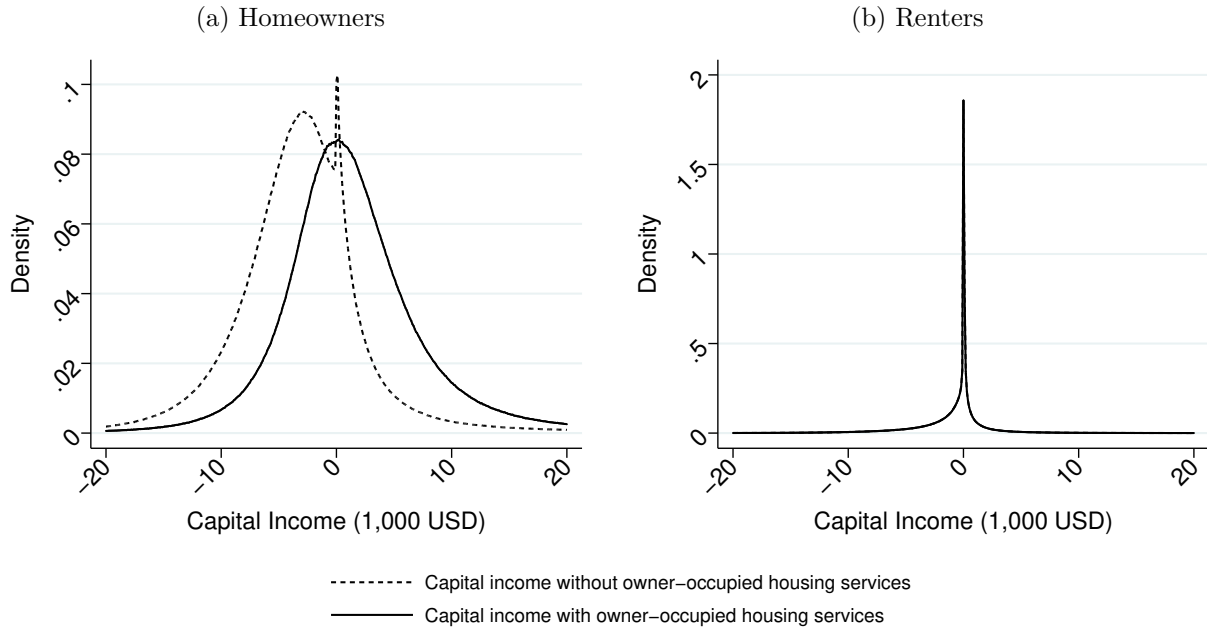


Figure 7: Distributions of Capital Income with and without Owner-Occupied Housing Services
 Notes: This figure displays the distribution of capital income for homeowners and renters with and without owner-occupied housing services. Measures are expressed in 2014 USD and adjusted for household size using the EU scale, and households are weighted by the number of household members. The figure is based on pooled cross-sections over the period 1994-2014.

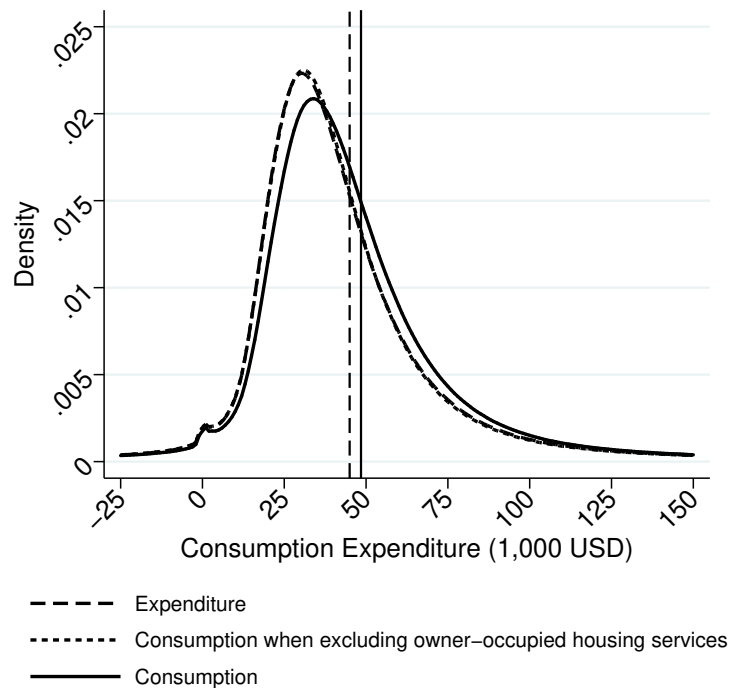


Figure 8: Distributions of Consumption and Expenditure
 Notes: “Consumption” corresponds to our preferred measure of consumption expenditure, while “Expenditure” refers to our preferred measure after excluding owner-occupied housing services and treating durable goods as fully consumed in the year of purchase. Averages are represented by the vertical lines. Consumption is measured in 2014 USD and adjusted for household size using the EU scale, and households are weighted by the number of household members. The figure is based on pooled cross-sections over the period 1994-2014.

payments and an increase in interest expenses.

2.5.2 Durables

Durable goods, such as cars and household furnishings, generate an expenditure mainly in the period when they are purchased, but they generate flows of consumption services until they are replaced or scrapped. Since data from tax records are available at an annual frequency, we define durable goods as consumption goods that might be used several times over a period of more than one year. While we do not have detailed data on all such goods, the types of durables recorded in the tax returns (reported as “Other real capital” in Table 1) include cars, caravans, boats, fixtures, and chattels, with cars being the most important component.¹¹ We assume that the consumption flow from a durable good equals the depreciation in value.¹²

Panel (b) of Figure 9 considers households that purchase their first car in year 0, comparing our preferred measure of consumption to one treating durable goods as fully consumed in the year of purchase. While our preferred measure of consumption increases steadily throughout, the alternative measure shows a pronounced transitory jump in consumption in the year of purchase.¹³ However, the treatment of expenditure on durables has little impact on the cross-sectional distribution of consumption. This can be seen in Figure 8, by comparing the measure of consumption without the value of owner-occupied housing services with one that also treats durable goods as fully consumed in the year of purchase (denoted “Expenditure”).

¹¹In the national accounts for 2012, net purchases of vehicles amount to 6 percent of total household consumption; 27 percent of total purchases of durables; and 47 percent of total purchases of durables excluding semi-durables (such as clothing, footwear, household utensils, equipment for sports, and books).

¹²For example, the tax appraisal of a new car is 75 percent of the store price. This means that we assume the purchase of a new car increases current consumption by 25 percent of the store price. The annual depreciation rates for used cars, and hence the measured consumption flows from used cars, vary between 5 and 15 percent of the store price depending on the age of the car.

¹³If a household finances the purchase of a new car by withdrawing funds from bank accounts, increasing debt, and/or selling other assets, the purchase price will be fully reflected in a reduction in net wealth and a corresponding increase in current consumption. Similarly, a purchase financed out of current income will also be fully reflected in an increase in current consumption. By including the purchase price net of annual depreciation in the “Other real capital” component of net wealth, however, the purchase of a new car will instead lead to an increase in current consumption corresponding to the annual depreciation. More generally, for all durable goods recorded in the tax returns and for any given year, we assume that the consumption flow from a durable good equals the depreciation in value.

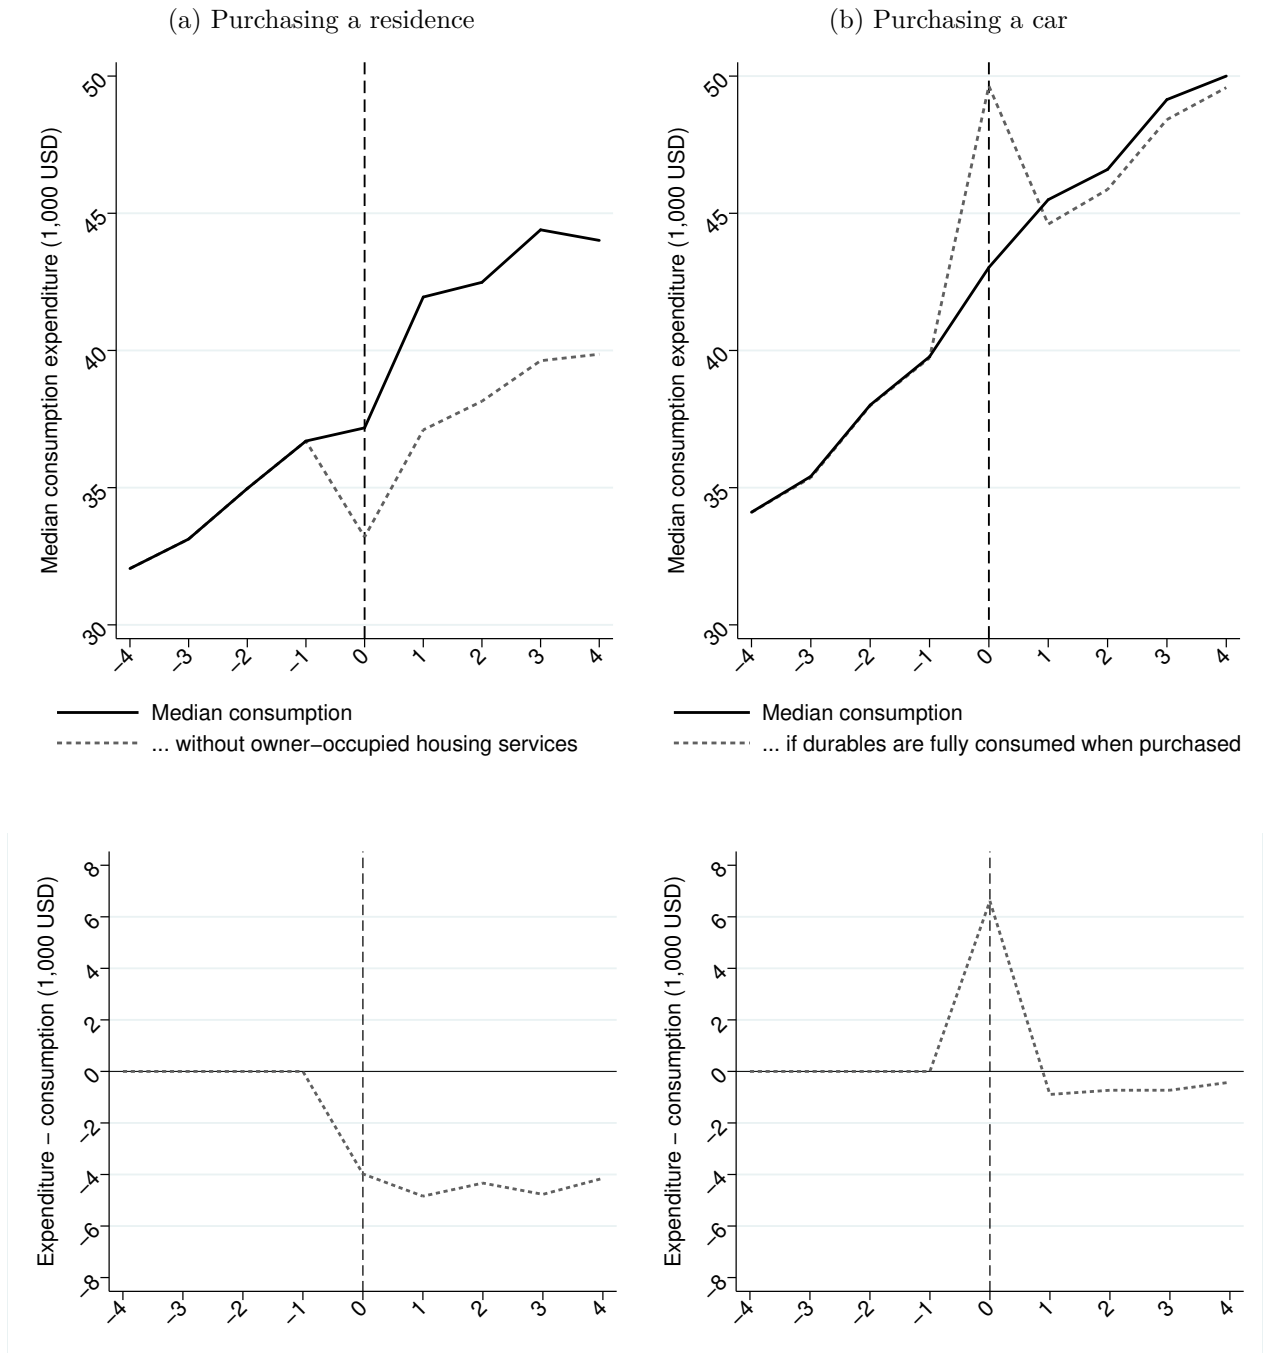


Figure 9: Median Consumption around the Purchase of a Residence or a Car

Notes: This figure compares median consumption among households buying a residence or a car in year 0 (2004) with median consumption if we (a) do not include owner-occupied housing services in consumption, or (b) are counting durables as fully consumed in the year of purchase. The sample for Panel (a) is restricted to households that do not own a residence before buying one in 2004, while the sample for Panel (b) is restricted to households that do not own a car before buying one in 2004. Consumption is measured in 2014 USD and adjusted for household size using the EU scale, and households are weighted by the number of household members. The upper panels show the levels of the two measures of consumption over time, while the lower panels show the differences between the two measures of consumption for each year.

3 Comparing alternative measures of consumption expenditure

3.1 Measures of consumption expenditure based on administrative data

When using administrative records on income and wealth to impute measures of consumption expenditure, a key challenge is that tax records typically do not contain information about the stock of each asset, but (at best) the values of each asset at the beginning and end of the year. This means that additional data or strong assumptions are necessary to impute reliable measures of consumption expenditure.¹⁴ In particular, it is necessary to distinguish changes in net wealth due to unrealized capital gains from changes in net wealth due to a household saving some of its income. As described in Section 2.3, we address this missing data problem by linking tax records on income and wealth to other administrative data on financial and real estate transactions and by using information on asset specific prices.

In this section, we assess the measurement errors in and limitations of consumption measures relying on tax records only. To this end, we construct a benchmark measure of consumption subject to the following restrictions:

- (a) Savings in real estate are calculated based on tax assessments only¹⁵
- (b) Capital gains on financial assets are set to zero
- (c) Data from the inheritance tax registry is not used
- (d) The value of owner-occupied housing services is set to zero
- (e) Durables are counted as fully consumed in the year of purchase

Our “preferred” measure and the benchmark measure would coincide if restrictions (a)–(e) were dropped. Note that restrictions (a)–(c) represent sources of measurement error for measures of consumption expenditure based exclusively on tax records, without detailed information about the stock of each asset. Restrictions (d) and (e) are not about measurement error per se, but represent refinements that can be made if the purpose is to measure household consumption as opposed to expenditure. Thus, whether to invoke (d) and (e) depends on the question of interest.

Table 2 shows the share of all households with non-positive consumption and selected distributional measures of consumption when restrictions (a)–(e) are relaxed in a sequential manner. We first relax assumption (a), i.e. we improve the measurement of savings in real estate by calculating net savings based on transactions data. Improving the measurement of savings in real estate has sizable impacts both on the fraction of households with non-positive consumption and on the mean and median levels of consumption. The dispersion in consumption is also greatly reduced, especially in the lower parts of the distribution, as measured by the 50/10 percentile ratio. Accounting for capital gains on financial assets (restriction (b)) also contributes to reducing the fraction of households with non-positive consumption values, to reducing the dispersion in consumption, and to raising the mean and median levels of consumption. Including the value of owner-occupied

¹⁴Up until 2007, Swedish tax records did contain information about the stock of each financial asset at the end of each year, for most types of financial assets; see Kojien et al. (2015). Assuming that transactions take place only at the end of each year, capital gains and net savings in listed securities (i.e. securities for which end-year prices are observed) can be measured using information on the stock of each asset in combination with end-year prices. In Appendix Section B.2, we investigate the consequences of assuming no within-year transactions for measures of net savings by contrasting our preferred transactions-based approach with a measure of net savings based on end-year prices and the stock of each asset.

¹⁵We adjust the raw tax assessment values according to the aggregate ratio of selling prices to tax assessments. For instance, when tax assessments of primary residences are set at 25 percent of estimated market values (from 2010 onward), we multiply tax assessments of primary residences by four.

housing services (restriction (d)) reduces the fraction of households with non-positive consumption values and increases the mean and median consumption levels, while restrictions (c) and (e) have only minor impacts on the cross-sectional distribution of consumption.

Table 2: Distributional Measures of Consumption by Imputation Procedure

	Benchmark measure	Relax restriction (a)	Relax restrictions (a) and (b)	Relax restrictions (a) - (c)	Relax restrictions (a) - (d)	Preferred measure
Non-positive consumption	0.085	0.047	0.039	0.039	0.035	0.036
50th / 10th Percentile	7.69	2.33	2.17	2.16	2.12	2.13
90th / 50th Percentile	2.37	2.07	2.04	2.06	2.01	1.97
Mean	40,185	42,517	44,675	44,979	48,839	48,454
50th Percentile	36,777	37,548	38,026	38,202	41,889	41,913

Notes: This table shows distributional measures of consumption when consumption is measured under different subsets of the following set of restrictions: (a) Savings in real estate are calculated based on tax assessments only; (b) Capital gains on financial assets are set to zero; (c) Data from the inheritance tax registry is not used; (d) The value of owner-occupied housing services is set to zero; (e) Durables are counted as fully consumed in the year of purchase. The benchmark measure is subject to all five restrictions, while relaxing all five restrictions results in our preferred measure. Consumption is measured in 2014 USD and adjusted for household size using the EU scale, and households are weighted by the number of household members. The table is based on pooled cross-sections over the period 1994–2014.

Panel (a) of Figure 10 shows the distributions of our preferred measure and the benchmark measure of consumption. The preferred measure distribution is centered to the right of the benchmark measure distribution and has a smaller left tail. Panels (b)-(d) of Figure 10 perform the same comparison after excluding groups of households for which measuring consumption is likely to be particularly problematic without having good measures of savings in real estate and financial assets. Large differences remain even after excluding households in the top of the financial assets distribution (Panel (b)) or households buying or selling real estate (Panel (d)). Excluding all homeowner households (70 percent of the households), on the other hand, appears to eliminate much of the differences across the two measures.

Using data other than tax records has important consequences not only in terms of the cross-sectional distribution of consumption, but also in terms of reducing the year-to-year volatility in measured consumption. This is visualized by Figure 11, in terms of consumption and savings per capita. The differences in income levels across the two measures are due to our preferred measure including the value of owner-occupied housing services as one of the income components.

3.2 Comparison with the National Accounts and the Survey of Consumer Expenditures

This section investigates the consistency between our data and two other sources of information on household consumption, income, and savings in Norway. The Survey of Consumer Expenditures (SCE) measures annual household consumption expenditure for a representative sample of households,¹⁶ with a definition of

¹⁶Households with members above age 80 are underrepresented, and individuals living under special arrangements (for instance at hospitals) are excluded. Respondents keep a diary for two weeks in which all the household’s purchases are noted in detail, while purchases of goods and services bought infrequently are registered in an interview covering the last twelve months. Annual household consumption expenditure is measured as the sum of annualized expenditure from the diaries and expenses covered by the interview.

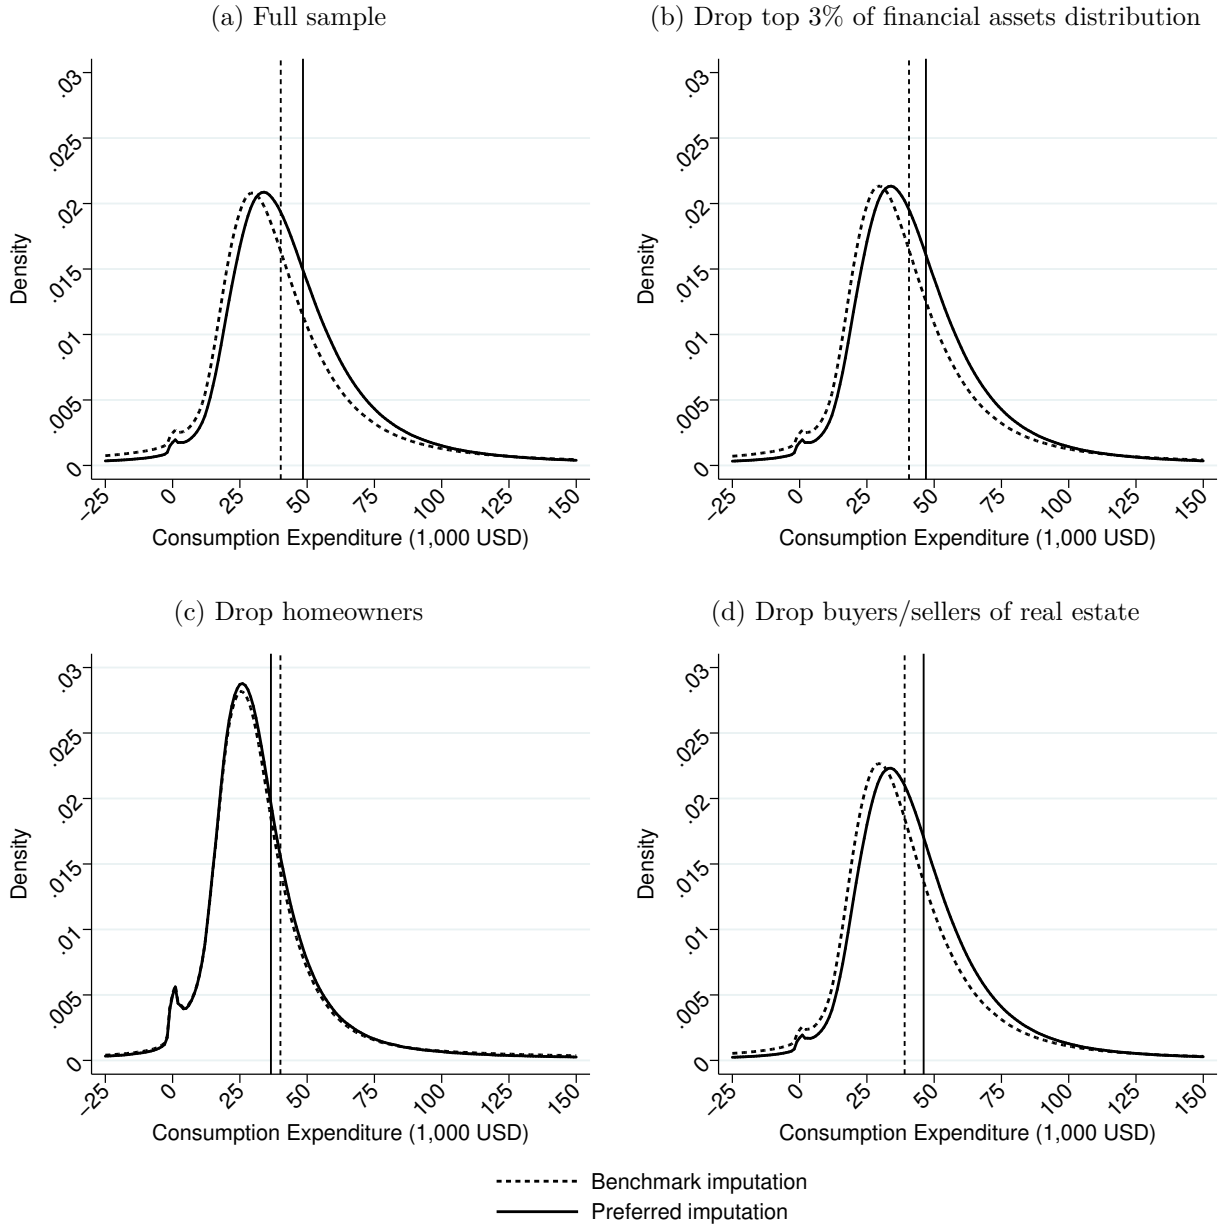


Figure 10: Distributions of Consumption by Imputation Procedure

Notes: This figure compares the distribution of consumption resulting from our preferred imputation procedure to the benchmark imputation distribution. Averages are represented by the vertical lines. Consumption is measured in 2014 USD and adjusted for household size using the EU scale, and households are weighted by the number of household members. The figure is based on pooled cross-sections over the period 1994–2014.

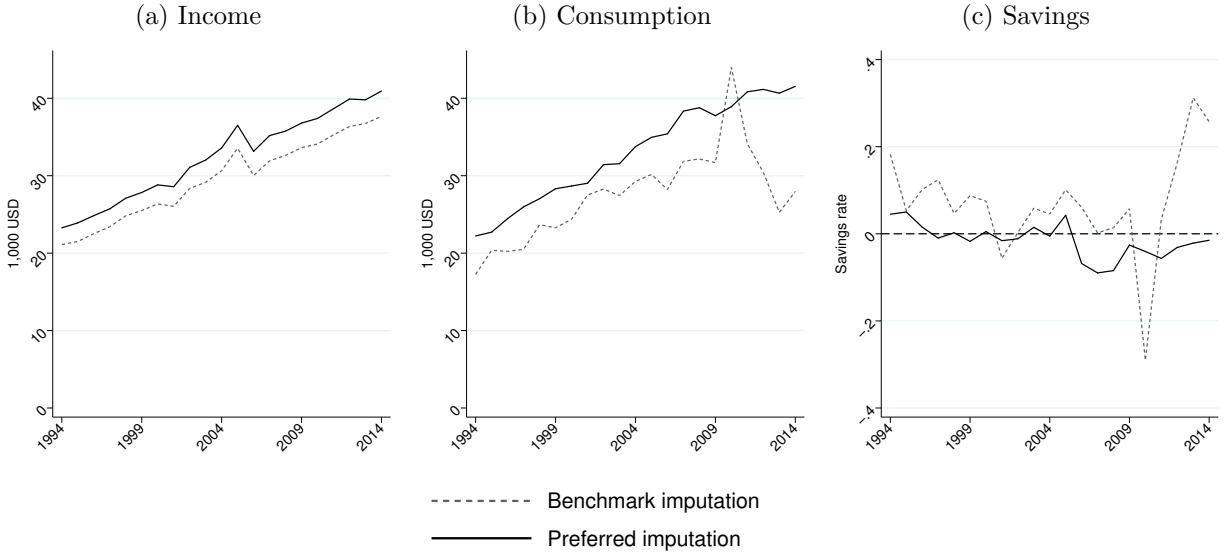


Figure 11: Income, Consumption, and Savings per Capita by Imputation Procedure

Notes: The savings rate is defined as the difference between disposable income and consumption per capita, divided by disposable income per capita. Income and consumption are measured in 2014 USD.

consumption expenditure that is broadly consistent with our register based measure starting in 1996.¹⁷ The national accounts (NA) measure aggregate consumption, income, and savings for all individuals residing in Norway. Measures are based on a number of register and survey based data sources, for instance on sales and production of goods and services. The definitions of household consumption and income are broadly similar to the definitions we use in our imputation, but we nevertheless modify the NA time series slightly, so as to make them more comparable to our register based measures and the SCE. Appendix Section C describes the main differences between the register based measures and those in the NA as well as our modifications of the NA time series.

Figure 12 shows the time series of disposable income, consumption, and savings per capita. There is a close correspondence between all three sources in terms of disposable income, and between the NA and our register based measure in terms of consumption and savings. The SCE measure of consumption diverges from the other two consumption measures from the early 2000s. Just like in many other countries (Browning et al., 2014), consumption expenditure is believed to be under-estimated in the SCE due to under-reporting and dropout bias (Mørk and Willand-Evensen, 2004). The cross-sectional distributions in Figure 13 show a close resemblance between the SCE and our register based measure of consumption in 1996, and growing discrepancies in more recent years.

3.3 Under-reporting and dropout bias in the SCE

Koijen et al. (2015) compare their register-based consumption measure for Sweden with a measure of consumption from the Swedish expenditure survey by linking survey participants with the register data. They find that the survey understates consumption and that the understatement is larger for the rich, as measured by either wealth or income. We are not able to link the SCE with the register data, but instead compare our register based measure of consumption expenditure with consumption expenditure in the SCE by matching

¹⁷One difference is that renovation expenses are counted as consumption in our register based measure, but not in the SCE measure of consumption. The SCE also counts durables as fully consumed in the year of purchase.

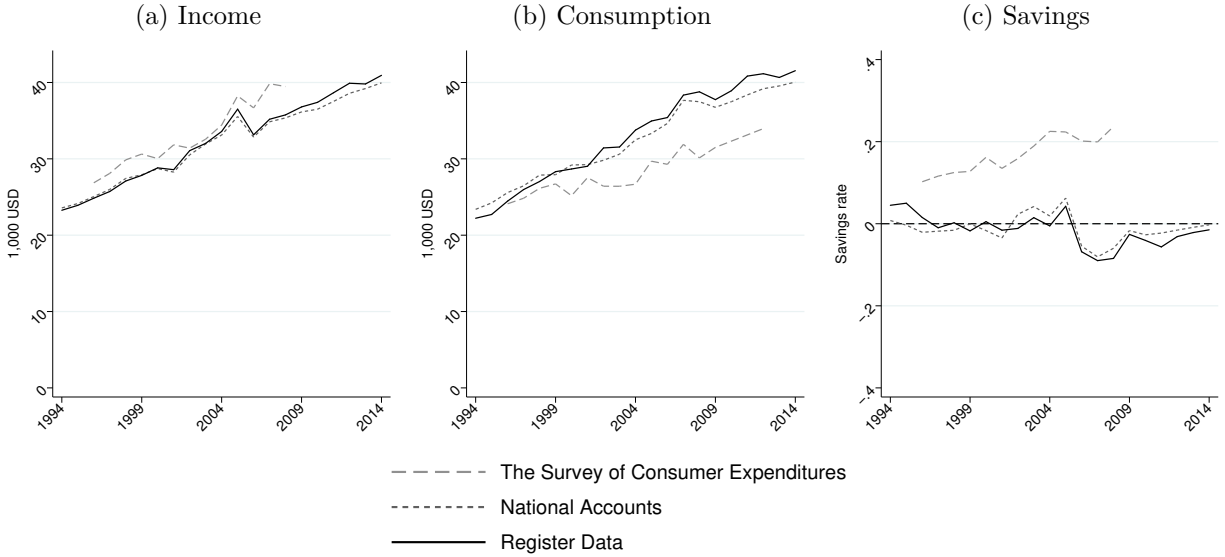


Figure 12: Income, Consumption, and Savings per Capita - Register Based Measures against the Survey of Consumer Expenditures and the National Accounts

Notes: Consistent data from the Survey of Consumer Expenditures are available between 1996 and 2009 and in 2012 for consumption, and between 1996 and 2008 for income and savings. We have modified the national accounts time series so as to make them more comparable to our register based measures, as described in Appendix Section C. The savings rate is defined as the difference between disposable income and consumption per capita, divided by disposable income per capita. Income and consumption are measured in 2014 USD.

on income. Figure 14 compares the register-based consumption measure with the SCE for the years 1996 and 2009, in terms of mean consumption expenditure by percentiles of the respective income distributions. Except for the first quartile, mean consumption is lower in the SCE than in the register data, and the relative differences are increasing over time. In both years, the percentage difference between register- and survey-based consumption is larger at the top of the income distribution than in the middle, but the increase in differences over time is most pronounced for the second and third quartiles. We now investigate whether the increasing discrepancies between register- and survey-based consumption are due to (i) changes in response rates by income or (ii) changes in reporting behavior by income.

The response rate in the SCE decreased from 62 percent in 1996 to 54 percent in 2009 (Holmøy and Lillegård, 2014). Figure 15 shows the response rate by income percentile relative to the average response rate, for 1996 and 2009. For both years, households in the first income quartile are underrepresented in the SCE by about 20 percent. The response rate at the top of the income distribution, however, has decreased significantly: the top five percent income earners went from being overrepresented by about five percent to being underrepresented by about 30 percent.

To adjust for differences in response rates by income, we construct a counterfactual distribution of register-based consumption by reweighting the observations in the register data to match the distribution of income in the SCE.¹⁸ Means from this counterfactual distribution of register-based consumption are reported as “Register data reweighted” in Figure 14. The reweighting has little impact on the differences between register- and survey-based consumption, and the differences are increasing over time (for all quartiles but the first) also after adjusting for differences in response rates. This suggests that changes in response rates by income are not a main driver of the increased discrepancy between the SCE and the register data, while

¹⁸Ideally, we would rather have reweighted the SCE, but the SCE does not have support for the entire distribution of income in the register data.

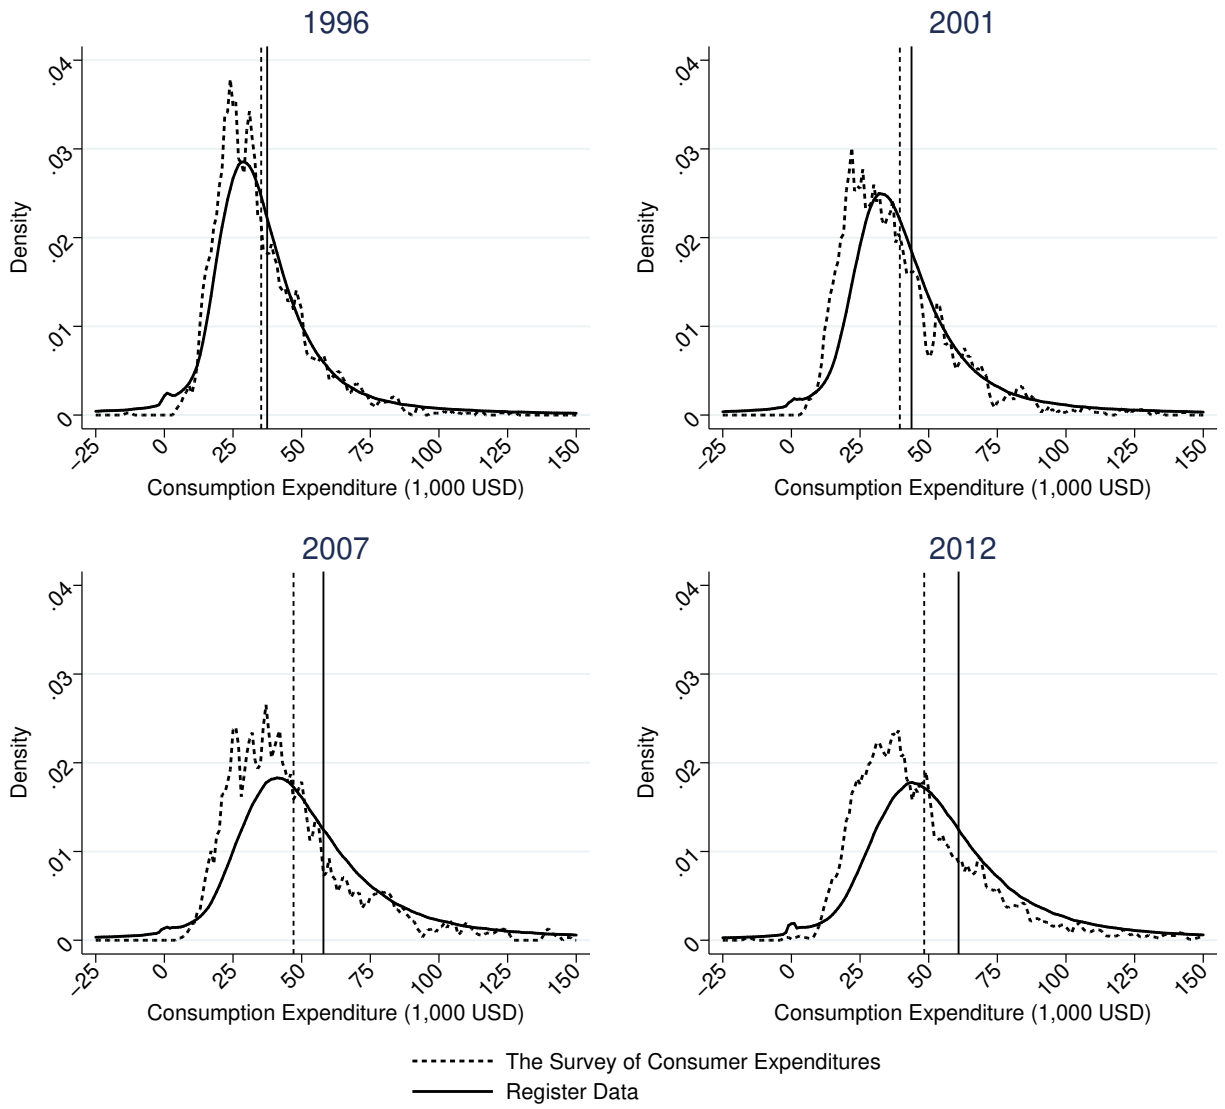


Figure 13: Distributions of Consumption Measures - Register Based Measure against the Survey of Consumer Expenditures

Notes: Measures are expressed in 2014 USD and adjusted for household size using the EU scale. Vertical lines denote sample averages. Households are weighted by the number of household members below the age of 80 (and SCE sample weights).

the problem of under-reporting in the SCE appears to have worsened at all parts of the income distribution except at the bottom.

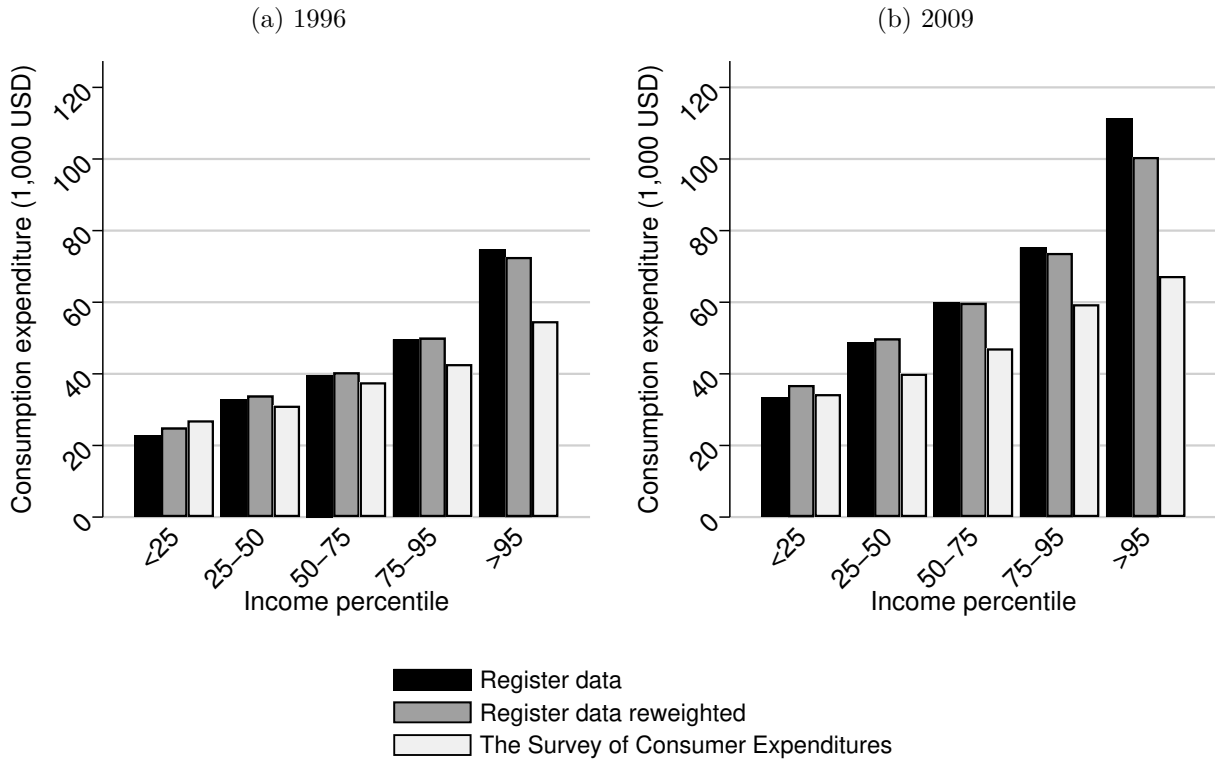


Figure 14: **Consumption by Income Percentile - Register-Based Measure against the Survey of Consumer Expenditures**

Notes: The figure displays mean consumption by income percentile in our register-based data and the SCE. When reweighting the register data, each household is assigned a weight proportional to the fraction of households with a given income level in the SCE relative to the fraction of households with the same income level in the register. The income measure used is disposable income (as defined in Section 2) less owner-occupied housing services, interest on liabilities, inheritance, gifts, and lottery winnings. Measures are expressed in 2014 USD and adjusted for household size using the EU scale, and households are weighted by the number of household members below age 80 (and SCE sample weights).

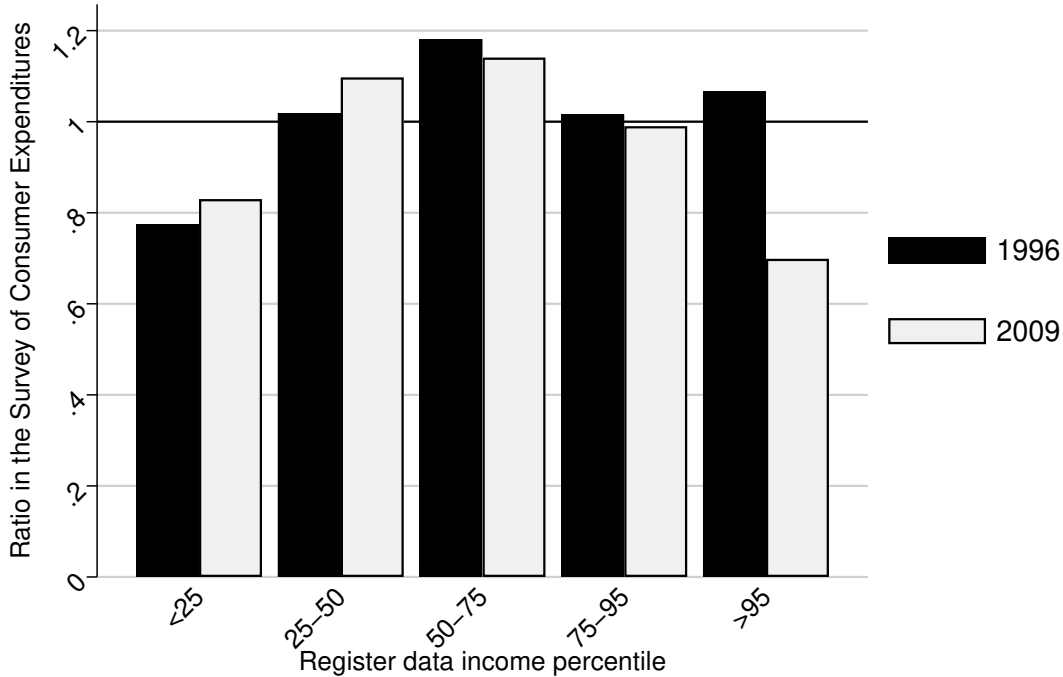


Figure 15: **Income Density in the Survey of Consumer Expenditures by Income Percentile in the Register Data**

Notes: The figure displays the (weighted) fraction of observations in the SCE in different income groups, relative to population shares in the same groups. Groups are defined by income percentiles in the register data. The income measure used is disposable income (as defined in Section 2) less owner-occupied housing services, interest on liabilities, inheritance, gifts, and lottery winnings. Income is adjusted for household size using the EU scale, and households are weighted by the number of household members below the age of 80 (and SCE sample weights).

4 Evidence on consumption and income inequality

We now illustrate the research opportunities arising from the derived measures of household consumption expenditure through two applications. The first application is a re-examination of the study of Attanasio and Davis (1996) of how relative wage movements among birth cohorts and education groups affected the distribution of household consumption. Their study uses data from expenditure surveys to construct a synthetic panel of consumption and wages. We show that the findings produced by a synthetic panel based on expenditure surveys are broadly similar to those one obtains using the population panel. The second application investigates the transmission of income shocks to consumption using the partial insurance framework of Blundell, Pistaferri, and Preston (2008). In this application we show the biases that arise if one uses cross-sectional data as compared to population panel data with information on household expenditure and income.

4.1 Relative wage and consumption movements

Using synthetic panel data on US consumption and wages during the 1980s, Attanasio and Davis (1996) study the relationship between relative wages and relative consumption across different groups in the US population, where groups are defined on the basis of the year of birth and educational attainment of the household head. They do not find a significant relationship between one-year changes in consumption and one-year wage changes. However, when considering longer (eight- and ten-year) horizons, for which persistent

wage factors are more likely to be important, the relationship between changes in consumption and changes in wages becomes strongly significant. Attanasio and Davis (1996) interpret the close relationship between low-frequency relative wage and consumption movements as a failure of the consumption insurance hypothesis.

We replicate the analyses in Attanasio and Davis (1996) using the derived measure of household consumption expenditure and a measure of hourly wages from the Norwegian Wage Statistics (covering the period 1997–2014). We consider households in which the oldest household member is a non-immigrant male born between 1945 and 1984. For each year t , we exclude households with significant income from self-employment, and households with a male head older than 60 or younger than 25.¹⁹ We align our measure of consumption with the one used in Attanasio and Davis (1996) by excluding owner-occupied housing services, treating durable goods as fully consumed in the year of purchase, and by measuring consumption without adjusting for household size and composition. In line with their study, we define groups by crossing four education groups (Less than High School; High School; Post secondary education; and College) with 5-year birth cohort groups. Since the analysis is based on synthetic panels, we are able to contrast some of the results based on our derived measure of consumption expenditure with corresponding results based on the SCE.²⁰

Table 3 summarizes real hourly wage movements between 1997 and 2012 by birth cohorts and education groups. The table has one panel for each education group, with each panel showing log deviations from the 2002 value for the 1965–69 cohorts of men with the same level of education. Each row outlines the cross-sectional age profile of real wages for a given year and education group, while each column delineates the evolution of real wages for a particular cohort-education group. Comparing Table 3 with the corresponding table in Attanasio and Davis (1996), an interesting feature common to the US in the 1980s and Norway in the 2000s is the sharp increase in wages for highly educated men. Within-cohort group comparisons of wage growth between 2002 and 2012 reveal that real wages of college educated men rose by about 20 log points relative to the wages of men without a high school degree.

Movements in household consumption expenditure are summarized in Figure 16, with our register based measure of consumption in the left panel and the SCE measure of consumption in the right panel. The figure is organized by birth cohorts, education groups, and years in the same way as Table 3: each line shows the cross-sectional age profile of consumption for a given year and education group, while the evolution of consumption for a particular cohort-education group is given by the vertical differences within each sub-figure. For the register based measure in the left panel, the overall patterns of consumption movements are similar to the patterns of wage growth apparent from Table 3: While all groups experienced quite substantial growth in real consumption expenditure over the period, the consumption growth was most pronounced for higher educated men. The age profiles based on the SCE are rather noisy, with irregular consumption movements both across years within cohorts and across cohorts within years.

Figure 17 examines high- and low-frequency comovements between relative wages and consumption. In the upper panels we plot annual differences of mean log consumption against mean log wages for the cohort-education groups in our sample, controlling for age (a cubic polynomial) and year fixed effects. Although the slope coefficient is much less precisely estimated for the SCE measure than for the register based measure of consumption, neither of the two measures show evidence of a systematic relationship between year-to-year relative wage and consumption movements among cohort-education groups. This conforms well to the consumption insurance hypothesis and to other theories predicting consumption smoothing.

¹⁹We follow Attanasio and Davis (1996) in using only the observations for which all members of the group satisfy the age restrictions in the indicated year.

²⁰The SCE was only conducted until 2012, with a two-year gap between 2009 and 2012, and it contains educational attainment only from 2002 onward. Hence, with the SCE data we can construct one ten-year difference and up to seven one-year differences per group.

Table 3: Real Hourly Wage Movements for Men by Birth Cohort and Education Group

Year	5-Year Birth Cohort							
	1980-84	1975-79	1970-74	1965-69	1960-64	1955-59	1950-54	1945-49
Less than High School Education								
1997	.	.	.	-0.11	-0.06	-0.04	-0.03	-0.05
2002	.	.	-0.05	0.00	0.05	0.06	0.05	0.04
2007	.	0.12	0.19	0.21	0.24	0.24	0.21	.
2012	0.20	0.28	0.32	0.33	0.35	0.35	.	.
High School Education								
1997	.	.	.	-0.15	-0.08	-0.05	-0.04	-0.02
2002	.	.	-0.06	0.00	0.05	0.06	0.06	0.07
2007	.	0.13	0.19	0.22	0.25	0.24	0.23	.
2012	0.25	0.30	0.34	0.36	0.38	0.36	.	.
Some Postsecondary Education								
1997	.	.	.	-0.11	0.02	0.05	0.03	0.05
2002	.	.	-0.12	0.00	0.11	0.12	0.13	0.16
2007	.	0.13	0.25	0.32	0.39	0.38	0.34	.
2012	0.23	0.35	0.43	0.47	0.52	0.50	.	.
College Education								
1997	.	.	.	-0.13	-0.01	0.06	0.08	0.08
2002	.	.	-0.13	0.00	0.06	0.07	0.06	0.05
2007	.	0.12	0.26	0.36	0.40	0.42	0.39	.
2012	0.21	0.36	0.46	0.53	0.54	0.54	.	.

Notes: Each panel shows real hourly wage movements by birth cohort for a given education group, in terms of deviations from the mean of log hourly wages in 2002 for the 1965-69 cohorts of men with the same level of education. Hourly wages are measured in 2014 USD.

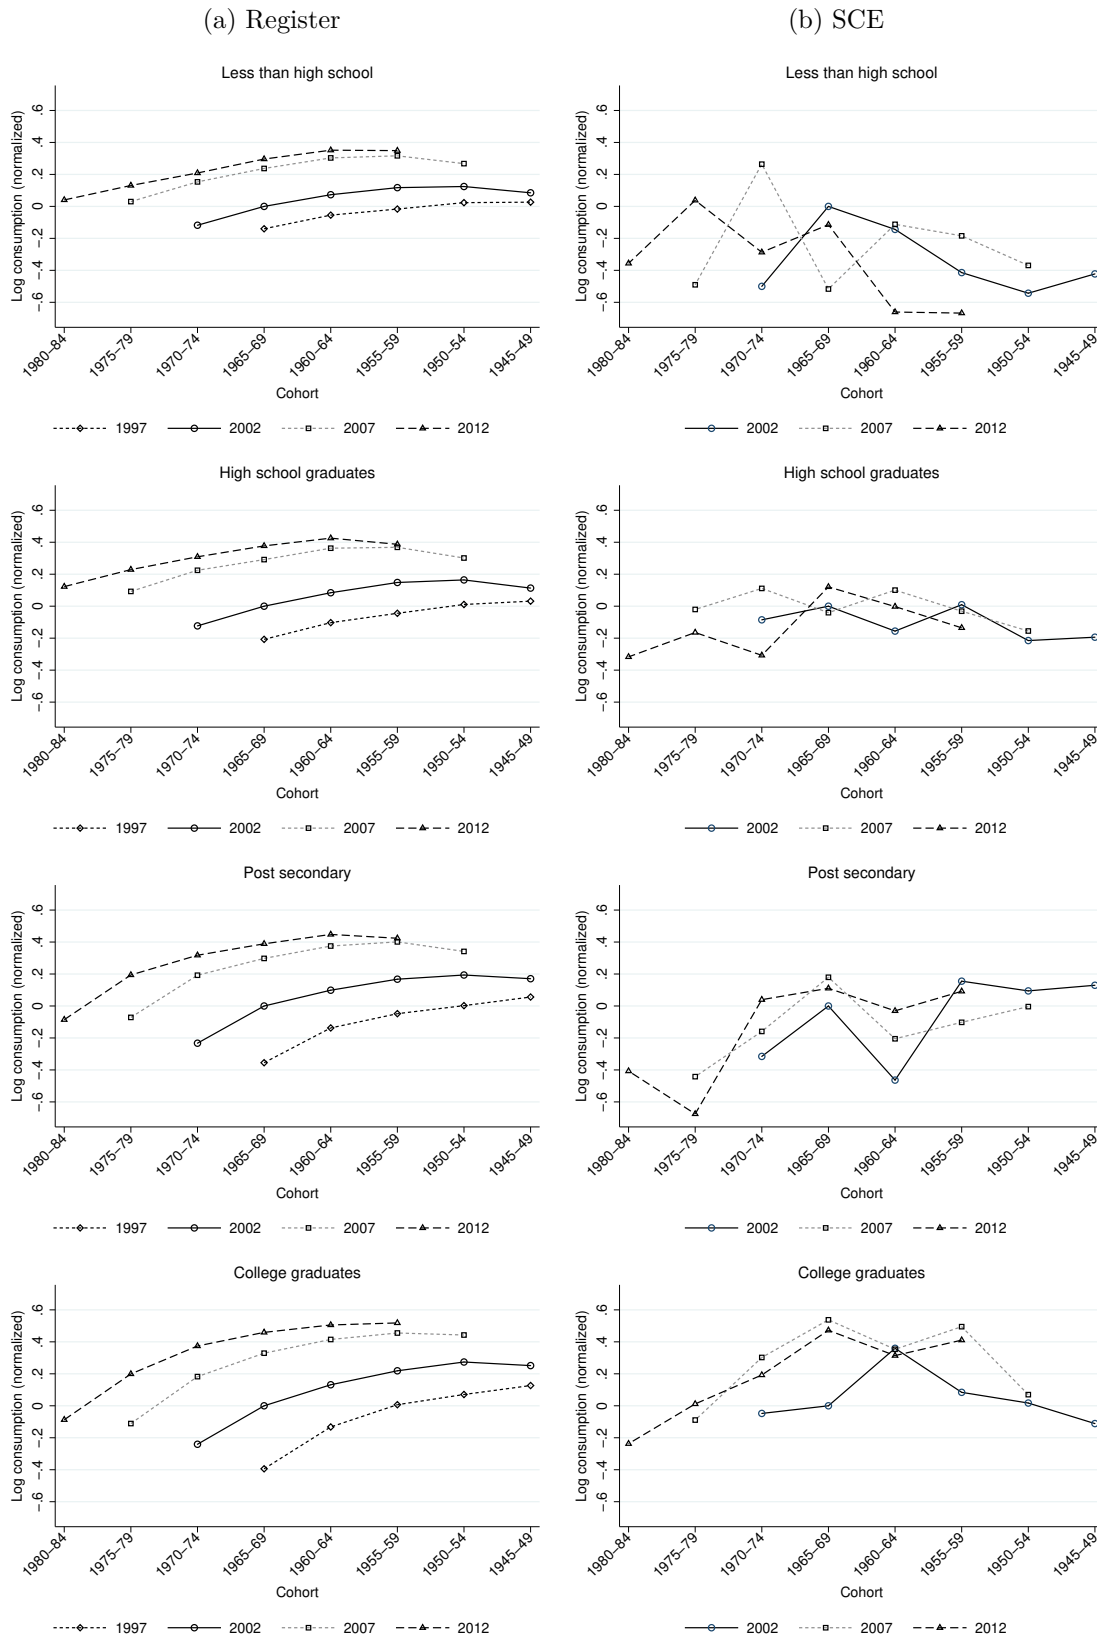


Figure 16: Cross-Sectional Age Profiles in Consumption - Register vs the Survey of Consumer Expenditures

Notes: Each panel shows cross-sectional age profiles in real consumption for a given education group, in terms of deviations from the mean of log consumption in 2002 for the 1965-69 cohorts of men with the same level of education. The value of owner-occupied housing services is excluded from the register based consumption measure, and durable goods are treated as fully consumed in the year of purchase. Consumption is measured in 2014 USD.

The lower panels of Figure 17 show analogous scatterplots for 10-year first differences. For both the register based measure and for the SCE measure of consumption, the figure reveals a close relationship between low-frequency relative wage and consumption movements. The clustering of education groups along the regression line suggests that between-education components are important drivers of the close relationship between relative wage and consumption movements. And, despite considerably more noise in the SCE than in the register based measure of consumption, the pooled regression results suggest that moving from the register based measure of consumption to the SCE measure does not necessarily give rise to entirely different conclusions.

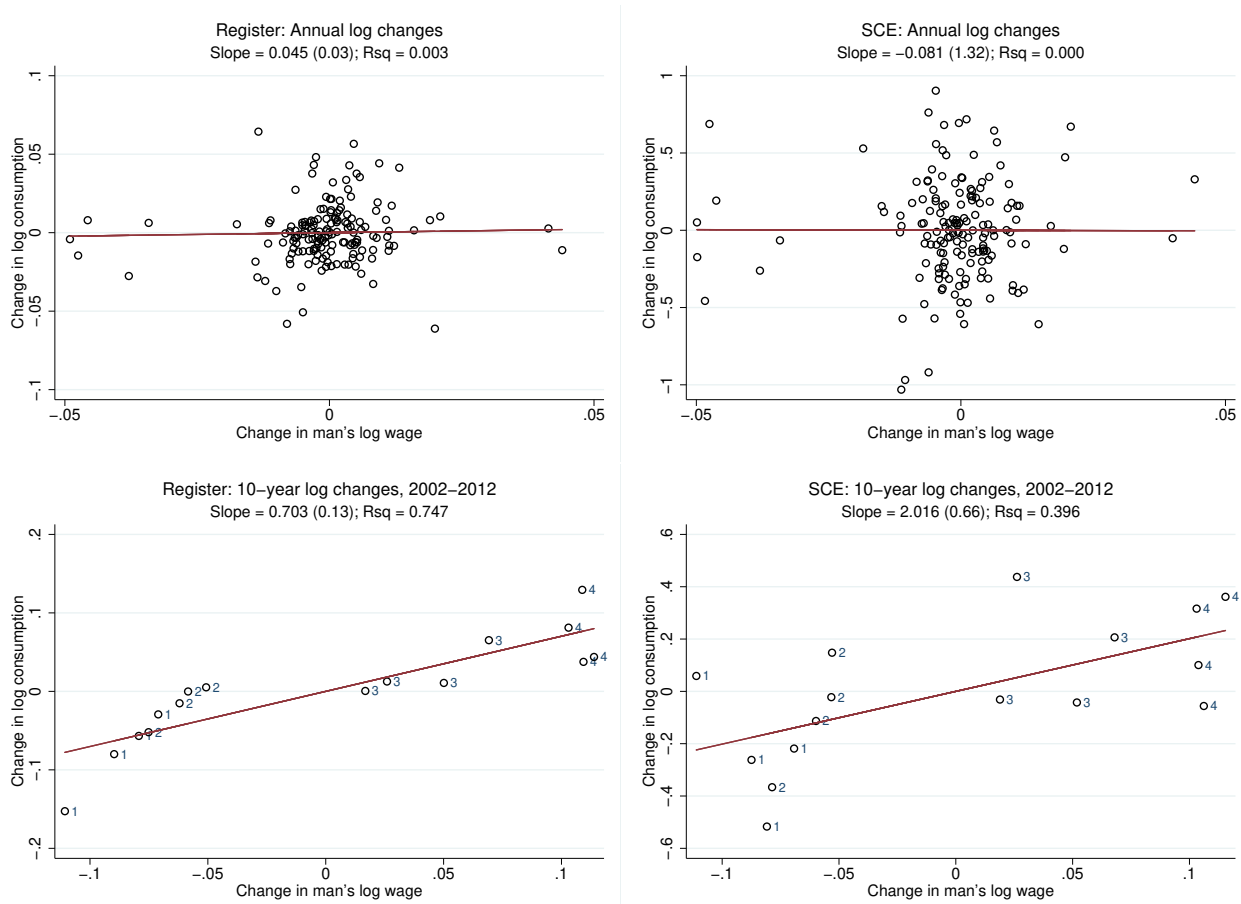


Figure 17: Household Consumption vs. Male Wages, Log-Change Residuals

Notes: For annual log changes, the plotted values are residuals from regressions on year effects and a cubic in age, while for 10-year log changes, the plotted values are residuals from a regression on a cubic in age. Annual changes are measured over the period 2002–2009. In the lower panels, 1 denotes Less than High School education, 2 denotes High School education, 3 denotes Post secondary education, and 4 denotes College education. The value of owner-occupied housing services is excluded from the register based consumption measure, and durable goods are treated as fully consumed in the year of purchase.

We now examine how the relationship between high- and low-frequency relative wage and consumption movements changes when controlling for potential determinants of group-level differences in consumption growth that are also correlated with group level differences in real wage growth. We consider one-, five-, and ten-year differences for three different samples: one based on the SCE; one using the register based measure but only the years for which the SCE is also available (“restricted sample”); and one using the register based measure of consumption and including the full set of one- and ten-year differences for the years 1997–2014 (“full sample”). In the first column of Table 4 we report the specification corresponding to that of Figure 17, which contains the mean log pre-tax hourly wage among men, year effects, and a cubic polynomial in age.

We then investigate the sensitivity of the coefficient on male wages to adding controls for family size and composition (column 2), log pre-tax wages among women (column 3), the mean of log gross capital income (column 4), and the mean of log gross wealth (column 5). While the high-frequency regression coefficient on male wages is small and statistically indistinguishable from zero in all specifications and for all samples, the low-frequency coefficients appear somewhat more sensitive to variations in sample size and to the inclusion of additional controls. Controls for family size and wealth have the largest negative impact on the coefficients on male wages. Using the register based measure of consumption with the full set of five- and ten-year differences, the null hypothesis of full consumption insurance can be rejected based on all specifications, including the one containing the full set of controls (column 5). The results from the restricted samples are less robust - especially those based on the SCE measure of consumption.

Table 4: Synthetic Panel Regressions

	(1)	(2)	(3)	(4)	(5)
Dependent variable: $\Delta_\tau \log c_t^g$					
Register data, full sample					
$\Delta_1 \log w_t^g$ $N = 408$	0.05 (0.04)	0.02 (0.04)	0.03 (0.04)	0.01 (0.04)	-0.04 (0.03)
$\Delta_5 \log w_t^g$ $N = 280$	0.94*** (0.11)	0.67*** (0.12)	0.72*** (0.15)	0.59*** (0.14)	0.15** (0.07)
$\Delta_{10} \log w_t^g$ $N = 140$	1.39*** (0.11)	0.81*** (0.19)	1.00*** (0.23)	0.99*** (0.22)	0.36** (0.14)
Register data, reduced sample					
$\Delta_1 \log w_t^g$ $N = 172$	0.04 (0.04)	0.00 (0.04)	-0.01 (0.04)	-0.02 (0.05)	-0.06 (0.04)
$\Delta_5 \log w_t^g$ $N = 84$	0.72*** (0.16)	0.47*** (0.18)	0.50** (0.22)	0.50** (0.22)	0.08 (0.11)
$\Delta_{10} \log w_t^g$ $N = 16$	0.70*** (0.15)	0.23 (0.22)	0.28 (0.23)	0.28 (0.19)	-0.27* (0.14)
Survey of Consumer Expenditures					
$\Delta_1 \log w_t^g$ $N = 172$	-0.08 (1.35)	-0.42 (0.99)	-0.09 (1.11)		
$\Delta_5 \log w_t^g$ $N = 84$	0.54 (0.82)	-0.12 (0.80)	0.26 (0.88)		
$\Delta_{10} \log w_t^g$ $N = 16$	2.02** (0.74)	0.53 (1.22)	0.66 (1.43)		
Controls:					
Age	x	x	x	x	x
Year	x	x	x	x	x
Family size		x	x	x	x
Female wages			x	x	x
Capital income				x	x
Wealth					x

The value of owner-occupied housing services is excluded from the register based consumption measure, and durable goods are treated as fully consumed in the year of purchase. The family size controls are the number of adults and the number of children under 14. Heteroskedasticity robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.2 The partial insurance framework

Blundell and Preston (1998) and Blundell et al. (2013) use repeated cross-sections to estimate the transmission of income shocks to consumption. We now examine the biases that arise from the use of such cross-sectional data. We begin by presenting the partial insurance framework of Blundell, Pistaferri, and Preston (2008) for the transmission of income shocks to consumption. We then use our panel data on consumption and income to identify and estimate the key parameters. Next, we show the additional restrictions needed if one only has access to repeated cross-sections, and examine empirically the biases that arise due to these restrictions.

4.2.1 Estimation with panel data

Log income net of observable household characteristics (y_{it}) is assumed to be the sum of a permanent component (y_{it}^P) which follows a random walk process and a transitory component (y_{it}^T) which follows a MA(1) process. The income process for household i is specified as follows;

$$\begin{aligned} y_{it} &= y_{it}^P + y_{it}^T \\ y_{it}^P &= y_{it-1}^P + \eta_{it} \\ y_{it}^T &= \varepsilon_{it} + \theta\varepsilon_{it-1}, \end{aligned} \tag{4}$$

where η_{it} and ε_{it} have mean zero and variances $\sigma_{\eta_t}^2$ and $\sigma_{\varepsilon_t}^2$, and are serially uncorrelated and independent of each other. From expression (4) it follows that (unexplained) income growth is determined by the current permanent shock and the current and previous period's change in the transitory shock:

$$\Delta y_{it} = \eta_{it} + \Delta\varepsilon_{it} + \theta\Delta\varepsilon_{it-1}. \tag{5}$$

BPP propose a reduced-form equation that allows to study the degree of transmission of income shocks into consumption, where the permanent income shock has an impact on consumption with a transmission parameter ϕ_t , and the impact of the transitory shock is measured by the transmission parameter ψ_t :

$$\Delta c_{it} = \phi_t\eta_{it} + \psi_t\varepsilon_{it} + \xi_{it}. \tag{6}$$

The random term ξ_{it} permits for shifts in consumption that are independent of those to income. Note that the BPP framework nests the two extreme cases of full insurance of income shocks (complete markets hypothesis) and no insurance (as in autarky, with hand-to-mouth consumers), as well as intermediate cases with both transmission parameters strictly between zero and one (partial insurance). Transmission parameters closer to zero indicate higher degrees of insurance.

From the bivariate process (4) and (6), one can derive the following covariance restrictions:

$$\begin{aligned}
\text{cov}(\Delta y_t, \Delta y_{t+\tau}) &= \begin{cases} \sigma_{\eta t}^2 + \sigma_{\varepsilon t}^2 + [\theta - 1]^2 \sigma_{\varepsilon t-1}^2 + \theta^2 \sigma_{\varepsilon t-2}^2 & \text{for } \tau = 0 \\ [\theta - 1] \sigma_{\varepsilon t}^2 + \theta [1 - \theta] \sigma_{\varepsilon t-1}^2 & \text{for } \tau = 1 \\ -\theta \sigma_{\varepsilon t}^2 & \text{for } \tau = 2 \\ 0 & \text{for } |\tau| > 2 \end{cases} \\
\text{cov}(\Delta c_t, \Delta c_{t+\tau}) &= \begin{cases} \phi_t^2 \sigma_{\eta t}^2 + \psi_t^2 \sigma_{\varepsilon t}^2 + \sigma_{\xi t}^2 & \text{for } \tau = 0 \\ 0 & \text{for } \tau \neq 0 \end{cases} \\
\text{cov}(\Delta c_t, \Delta y_{t+\tau}) &= \begin{cases} \phi_t \sigma_{\eta t}^2 + \psi_t \sigma_{\varepsilon t}^2 & \text{for } \tau = 0 \\ -\psi_t^2 \sigma_{\varepsilon t}^2 & \text{for } \tau = 1 \\ 0 & \text{otherwise,} \end{cases}
\end{aligned} \tag{7}$$

where $\text{cov}(\cdot, \cdot)$ denotes a cross-sectional covariance.

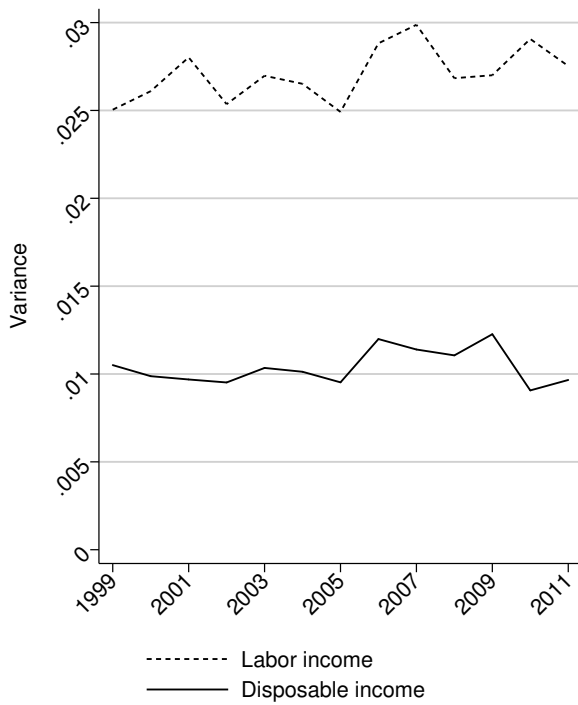
Suppose we have data for households $i = 1, \dots, N$ of a given birth cohort at time $t = 0, 1, \dots, T$. The parameters to identify are $\{\sigma_{\eta t}^2, \sigma_{\varepsilon t}^2, \sigma_{\xi t}^2, \phi_t, \psi_t\}_{t=1}^T$ and θ . For any given period $T - 2 \geq t \geq 1$, the above system of nonlinear equations has 5 unknowns in addition to θ and 6 non-redundant equations and is therefore (over)determined.

To estimate the model we employ a sample of individuals that are between the ages of 25 and 60 in both 1994 and 2014 (i.e. cohorts 1954 to 1969). We restrict the sample to individuals that live in a household which for at least six consecutive years have positive labor income, disposable income, and imputed consumption, and in which no adult enters or exits and there is no more than two labor income recipients. We exclude the household-year observations corresponding to the first and the last year in which the household satisfies all sample criteria. We consider two income measures: *Labor income* includes wage income, income from self-employment, sickness allowance, and parental leave benefits. *Disposable income* is as defined in Section 2, less dividends, gifts, and inheritance.

Estimated variances of permanent income shocks and permanent transmission parameters, for the years 1999–2011, are reported in Figure 18, while Table 5 displays the means of the variances of the permanent and transitory shocks and the transmission parameters, and the MA(1) parameter. As expected, the variance of permanent shocks to labor income is much larger than the variance of permanent shocks to disposable income. The estimated transmission parameters suggest that about 46% of a permanent shock to labor income is carried over to consumption, while the transmission parameter for permanent shocks to disposable income is closer to two thirds. Together with the estimated variances, the transmission parameters suggest an important role for insurance provided by the government through taxes and transfers.

In terms of the original BPP estimates, the estimated variance of permanent shocks to disposable income in the US (for the period 1979-1992) is 0.022 and hence closer to labor income than to disposable income for Norway, which is consistent with a much higher level of income inequality in the US than in Norway. Note also that our estimated transmission parameter for permanent shocks to disposable income is remarkably close to the corresponding parameter for the US: 0.658 and 0.642, respectively. These estimates suggest that a 10 percent permanent income shock induces a 6.5 percent permanent change in consumption.

(a) Variances of permanent income shocks ($\sigma_{\eta t}^2$)



(b) Permanent transmission parameters (ϕ_t)

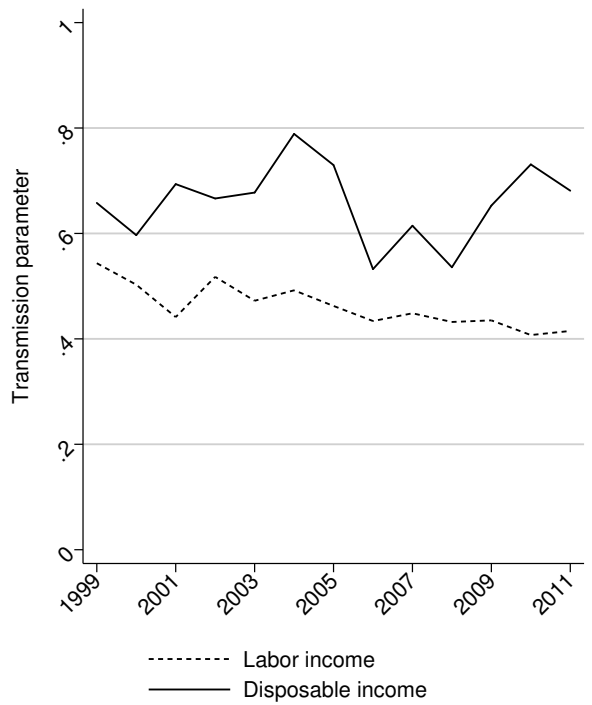


Figure 18: Variances of Permanent Income Shocks and Transmission Parameters

Notes: The sample consists of individuals born between 1954 and 1969 that live in a household which for at least six consecutive years have positive labor income, disposable income, and imputed consumption, and in which no adult enters or exits and there is no more than two labor income recipients. *Labor income* includes wage income, income from self-employment, sickness allowance, and parental leave benefits. *Disposable income* is as defined in Section 2, less dividends, gifts, and inheritance. As in BPP, we estimate the parameters by diagonally weighted minimum distance and allow consumption to be measured with measurement error.

Table 5: **Mean Variance of Income Shocks and Transmission Parameters**

	Panel data		Repeated cross-section			
	Labor income (1)	Disposable income (2)	BP (1998)		BLP (2013)	
			Labor income (3)	Disposable income (4)	Labor income (5)	Disposable income (6)
Permanent component						
Variance of shocks (σ_η^2)	0.0271 (0.0003)	0.0104 (0.0002)	0.0008 (0.0001)	0.0005 (0.0001)	0.0037 (0.1279)	0.0011 (0.0879)
Transmission parameter (ϕ)	0.4618 (0.0062)	0.6577 (0.0117)			0.2748 (0.1658)	0.5037 (0.3173)
Transitory component						
Variance of shocks (σ_ε^2)	0.0500 (0.0003)	0.0136 (0.0001)				
Transmission parameter (ψ)	0.1233 (0.0027)	0.4783 (0.0068)				
MA(1) parameter (θ)	0.1356 (0.0023)	0.1464 (0.0024)				

Notes: Columns (1) and (2) display means of the estimates for the years from 1999 to 2011. As in BPP, we estimate the parameters by diagonally weighted minimum distance and allow consumption to be measured with measurement error. Columns (3)–(6) display results from specifications assuming that the variances of permanent shocks and transmission parameters are constant over time. Standard errors (in parentheses) are based on bootstrapping with 250 replications. The sample consists of individuals born between 1954 and 1969 that live in a household which for at least six consecutive years have positive labor income, disposable income, and imputed consumption, and in which no adult enters or exits and there is no more than two labor income recipients. *Labor income* includes wage income, income from self-employment, sickness allowance, and parental leave benefits. *Disposable income* is as defined in Section 2, less dividends, gifts, and inheritance.

4.2.2 Estimation with repeated cross-sectional data

Suppose we only have access to repeated cross-sectional data on income and consumption. In this case, we do not observe the growth in income and consumption for individual households but instead rely on year-to-year changes in the variances of income and consumption and cross-sectional covariances between income and consumption to estimate the process defined by (4) and (6). Assuming that $\theta = 0$, the key empirical moments are now

$$\begin{aligned}
 \Delta \text{var}(y_t) &= \sigma_{\eta t}^2 + \Delta \sigma_{\varepsilon t}^2 \\
 \Delta \text{var}(c_t) &= \phi_t^2 \sigma_{\eta t}^2 + \psi_t^2 \sigma_{\varepsilon t}^2 + \sigma_{\xi t}^2 \\
 \Delta \text{cov}(c_t, y_t) &= \phi_t \sigma_{\eta t}^2 + \Delta[\psi_t \sigma_{\varepsilon t}^2].
 \end{aligned} \tag{8}$$

The parameters to identify are $\{\sigma_{\eta t}^2, \sigma_{\varepsilon t}^2, \sigma_{\xi t}^2, \phi_t, \psi_t\}_{t=1}^T$. In any given period t , the above system of nonlinear equations has 5 unknowns and 3 non-redundant equations and is therefore underdetermined. Additional restrictions must therefore be imposed to achieve identification. Blundell and Preston (1998) assume that there is no insurance against permanent shocks and that transitory shocks are fully insured, i.e. $\phi_t = 1$ and $\psi_t = 0$. Furthermore they assume that $\text{var}(\xi_t) = 0$, i.e. there are no innovations to consumption that are independent of innovations to income (due to preference shocks, for example). With these additional

assumptions, the above system of equations (8) simplifies to

$$\begin{aligned}\Delta\text{var}(y_t) &= \sigma_{\eta t}^2 + \Delta\sigma_{\varepsilon t}^2 \\ \Delta\text{var}(c_t) &= \sigma_{\eta t}^2 \\ \Delta\text{cov}(c_t, y_t) &= \sigma_{\eta t}^2.\end{aligned}\tag{9}$$

The system now provides three equations to identify the variance of permanent shocks and the year-to-year change in the variance of transitory shocks. Table 5 provides GMM estimates of the variances of the permanent shocks, where the variances are assumed to be constant over time, as they otherwise are very noisy. The estimates of the variances of permanent shocks to labor income and disposable income are very low (below 0.001) compared to the panel data estimates, reflecting advantages of the more flexible panel data estimation.

Blundell et al. (2013) use a similar procedure as Blundell and Preston (1998), but do not restrict the permanent shock transmission parameter to unity. The key empirical moments are now

$$\begin{aligned}\Delta\text{var}(y_t) &= \sigma_{\eta t}^2 + \Delta\sigma_{\varepsilon t}^2 \\ \Delta\text{var}(c_t) &= \phi_t^2 \sigma_{\eta t}^2 \\ \Delta\text{cov}(c_t, y_t) &= \phi_t \sigma_{\eta t}^2.\end{aligned}\tag{10}$$

With three equations and three unknowns, this is an exactly determined system at a given period $t \geq 1$. Table 5 provides GMM estimates of the variances of permanent shocks and transmission parameters, once again assuming that the variances and transmission parameters are constant over time. The variances of the permanent shocks to labor income and disposable income are 0.004 and 0.001, and hence larger than with ϕ_t set to unity but still significantly lower than the baseline estimates based on the BPP specification and estimated with panel data. Moreover, the transmission parameters are significantly lower than those from the BPP specification, and all parameters are imprecisely estimated.

5 Conclusion

A major difficulty faced by researchers who want to study the consumption and savings behavior of households is the lack of reliable panel data on household expenditures. One possibility is to use surveys that follow the same households over time, but such data are rare and they typically have small sample sizes and face significant measurement issues. An alternative approach is to use the accounting identity that total consumption expenditure is equal to income plus capital gains minus the change in wealth over the period. The primary goal of this paper was to examine the advantages and difficulties of deriving household expenditure measures from this accounting identity. To construct such measures of consumption expenditure, we combined several data sources from Norway over the period 1994–2014, linking tax records on income and wealth to other administrative data with household level information on financial and real estate transactions. We then compared measures of consumption expenditure derived from the accounting identity to those reported in expenditure surveys and to aggregates from the national accounts. We also illustrated the research opportunities arising from the derived measures of household consumption expenditure through two applications: an examination of how relative wage movements among birth cohorts and education groups affected the distribution of household consumption, and an investigation of the pass through of income shocks

to consumption.

Our findings may be summarized in three broad conclusions. First, the combination of tax records on income and wealth and household level information on financial and real estate transactions offers a unique opportunity to construct reliable measures of household consumption expenditure. It is reassuring to find that the derived measure of household consumption expenditure compares well to aggregate measures from the national accounts. Second, deriving measures of household consumption expenditure based exclusively on tax records on income and wealth is likely to result in severe measurement errors. Without additional information on financial and real estate transactions, the derived measures of consumption suffer from a large number of observations with non-positive consumption values as well as excessive dispersion and volatility in consumption. Short-cut approaches, such as excluding groups of households for which measuring consumption is likely to be particularly problematic, seem insufficient to resolve the measurement errors. Third, population panel data on income and consumption may be important for the conclusions drawn about income uncertainty, consumption and savings behavior. This is especially true in the application where we study the transmission of income shocks to consumption using the partial insurance framework of Blundell et al. (2008). To make this clear, we compared the results based on cross-sectional data, as used in Blundell and Preston (1998) and Blundell et al. (2013), to those one obtain using the population panel data. The results suggest that the restrictions needed to use the cross-sectional data lead to significant biases both in the estimated variances of income shocks and in the transmission of these shocks to consumption.

Acknowledgments

The project received financial support from the Norwegian Research Council (grant number 250516 and 227020). We are grateful for useful comments from three anonymous referees and the editor, from both the discussant and the participants at the Conference “New perspective on Consumption” at LSE, and from seminar participants at the Frisch Centre, the NHH, and the ISF. We are also grateful to Michael Graber for his help with the estimations in Section 4.2.

References

- Attanasio, O. and S. J. Davis (1996). Relative wage movements and the distribution of consumption. *J. Polit. Econ.* 104(6), 1227–1262.
- Attanasio, O. P. and L. Pistaferri (2016). Consumption inequality. *J. Econ. Perspect.* 30(2), 3–28.
- Blundell, R., H. Low, and I. Preston (2013). Decomposing changes in income risk using consumption data. *Quant. Econ.* 4(1), 1–37.
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *Am. Econ. Rev.* 98(5), 1887–1921.
- Blundell, R. and I. Preston (1998). Consumption inequality and income uncertainty. *Q. J. Econ.* 113(2), 603–640.
- Browning, M., T. F. Crossley, and J. Winter (2014). The measurement of household consumption expenditures. *Annu. Rev. Econ.* 6(1), 475–501.
- Browning, M. and S. Leth-Petersen (2003). Imputing consumption from income and wealth information. *Econ. J.* 113(4), F282–F301.
- Eitrheim, Ø. and S. K. Erlandsen (2004). House price indices for Norway 1819–2003. In Ø. Eitrheim, J. T. Klovland, and J. F. Qvigstad (Eds.), *Historical Monetary Statistics for Norway 1819-2003*, pp. 349–375. Norges Bank Occasional Papers No 35, Norges Bank, Oslo.
- Fagereng, A. and E. Halvorsen (2017). Imputing consumption from Norwegian income and wealth registry data. *J. Econ. Soc. Meas.* 42(1), 67–100.
- Holmøy, A. and M. Lillegård (2014). Forbruksundersøkelsen 2012, dokumentasjonsrapport (in Norwegian). Report: Documents 2014/17, Statistics Norway, Oslo.
- Koijen, R., S. Van Nieuwerburgh, and R. Vestman (2015). Judging the quality of survey data by comparison with "truth" as measured by administrative records: Evidence from Sweden. In C. D. Carroll, T. F. Crossley, and J. Sabelhaus (Eds.), *Improving the Measurement of Consumer Expenditures*, pp. 308–346. University of Chicago Press, Chicago.
- Kreiner, C. T., D. D. Lassen, and S. Leth-Petersen (2015). Measuring the accuracy of survey responses using administrative register data: Evidence from Denmark. In C. D. Carroll, T. F. Crossley, and J. Sabelhaus (Eds.), *Improving the Measurement of Consumer Expenditures*, pp. 289–307. University of Chicago Press, Chicago.
- Mørk, E. and E. Willand-Evensen (2004). Husholdningers forbruk: En sammenlikning av forbruksundersøkelsen og nasjonalregnskapet (in Norwegian). Report: Documents 2004/73, Statistics Norway, Oslo.
- Sabelhaus, J., D. Johnson, S. Ash, D. Swanson, T. I. Garner, J. Greenlees, and S. Henderson (2015). Is the consumer expenditure survey representative by income? In C. D. Carroll, T. F. Crossley, and J. Sabelhaus (Eds.), *Improving the Measurement of Consumer Expenditures*, pp. 241–262. University of Chicago Press, Chicago.

Appendix

A Measuring the value of real estate and the value of owner-occupied housing services

In this section we provide further details regarding two components that are included in our imputation and either (i) poorly measured in the tax returns (the value of real estate) or (ii) not reported in the tax returns (the value of owner-occupied housing services).

A.1 Value of real estate

The tax returns registry contains tax assessment values of real estate: the sum of tax assessments within categories of real estate, and, starting in 2010, separate tax assessment values for all residences. Prior to 2010, the tax assessment values differ significantly from actual market values, and these differences vary considerably across residences depending on a wide range of factors such as geographic region, year of construction, and property type. In 2010, the Norwegian Tax Administration reassessed all residence values based on a price per square meter calculated by Statistics Norway (using hedonic pricing models with information on size, geographic region, year of construction, and property type). Although the post-2010 tax assessment values on average are well aligned with actual market values, differences between tax assessments and market values for individual residences remain a serious concern. Moreover, tax assessments of most real estate other than residences, such as recreational properties, farms, and land, have not been revised.

To estimate the market value of real estate (owned directly or through a housing cooperative) for each household, we therefore supplement the tax records with information on transactions in real estate from the Land Registry. For each property, the register covers the last transaction between 1986 and 1993, and all transactions from 1993 onward. The register provides detailed information about the transactions, including unique identifiers for both the seller, the buyer, and the property; the selling price; an indicator for whether the property was traded on the open market; and characteristics of the property.

The transactions data is our main source of information about market values of real estate. For properties traded on the open market at least once between 1986 and 2015, we set the market value at the time of the transaction equal to the selling price. To obtain estimates of market values at other points in time, it is necessary to make assumptions about the evolution of prices. To this end, we combine our data on the characteristics of the properties with house price indices for specific regions and types of properties.²¹ These price indices allow us to interpolate and extrapolate the market values to years in which a property is not traded.²²

There are, however, properties that are either (a) never sold during the period 1986–2015 or (b) not included in the transactions data. Norway has a high stamp duty on real estate transactions which is (legally) avoided if the transaction is not registered in the Land Registry. This can be the case for contractors buying old houses to renovate or remodel before selling. Another issue is that units in housing cooperatives are

²¹For 2005–2014 we use 33 indexes covering all combinations of 3 different housing types and 11 regions. For 1992–2004 we use 20 indexes covering all combinations of 5 different housing types and 4 regions. Prior to 1992 we use a house price index constructed by Eitrheim and Erlandsen (2004).

²²For properties observed with selling prices in several years, we estimate the price for years between two sales as a weighted average of the two successive selling prices, making sure that the estimated prices are matching the selling price for years in which a selling price is observed.

not covered by the transactions data prior to 2007.²³ For primary residences for which no market value is observed in the transactions data, we use the post 2010 tax assessment values to get one or more measures of market values for the years 2010–2014.²⁴ Once we have these measures, we apply the price indices to estimate market values for earlier years. If we cannot assign a market value to a primary residence, neither from the transactions data nor from the tax returns registry in 2010 or later, the residence is assigned a value proportional to its contemporaneous tax assessment value.^{25,26}

Each year, between 12 and 13 percent of the households are owning properties other than primary residences for which the value is unknown (and set to zero). This problem is to a large extent driven by inherited recreational properties and commercial properties (including parts of a farm other than the farmhouse) that are not registered as residences in 2010 or later.

A.2 Owner-occupied housing services

In theory, the value of owner-occupied housing services should be equal to the rental income the homeowner could have received if the house had been rented to a tenant.²⁷ Statistics Norway estimates the value of owner-occupied housing services in the national accounts as the rental price (net of expenses for maintenance and insurance) for all owner-occupied residences, based on a representative sample of renter-occupied housing units. We distribute this aggregate value across households according to each household’s share of the total value of primary residences. As a result, all variation in the value of owner-occupied housing services across homeowner households within a given year is due to variation in the value of primary residences, while the rent-to-value ratio is a constant common to all homeowner households. The average rent-to-value ratio for 1994–2014 is 2.88 percent. Fagereng and Halvorsen (2017) use a capital market approach and find a long run real rate of return to housing in Norway of 3 percent.

To investigate the sensitivity of our measures with respect to assuming a common rent-to-value ratio within each year, we follow the procedure used for the national accounts to calculate an alternative measure of the value of owner-occupied housing services. In particular, we calculate the rental value for each owner-occupied residence based on characteristics of the residence and actual rents observed for residences with the same characteristics.²⁸ The upper panels of Figure A1 show how the resulting rent-to-value ratio is

²³Note that we do observe tax assessment values of each household’s units in housing cooperatives. And until 2010, all tax assessments of real estate were adjusted annually by the same factor. We can therefore infer from the tax records whether a household traded units in housing cooperatives in a given year, also in years for which units in housing cooperatives are not covered by the transactions data.

²⁴The post-2010 tax assessments are directly linked to addresses. We can therefore use the household’s address of residence from the National Population Register to link households to post-2010 tax assessments. However, an issue for early years (especially prior to 2001), is that the information on residential addresses from the National Population Register is not always sufficiently detailed to separate between different apartments within an apartment building. This is not a problem when the values of all apartments in an apartment building are similar. When there are large differences in values across apartments within a building, we make use of information from 2007 or later to link households to unique units. Starting in 2007, we observe the complete transactions history for all units, including units in housing cooperatives, and most addresses are sufficiently detailed to separate between different apartments within an apartment building.

²⁵A year t tax assessment is multiplied with the median of ratios between year t tax assessments and year t imputed values.

²⁶For 1994 (2014), the values of primary residences not owned through a housing cooperative are measured as follows: 54 (66) percent are based on selling prices associated with transactions on the open market; 15 (12) percent are based on other selling prices; 26 (22) percent are based on post 2010 tax assessment values; and 5 (0) percent are set proportional to the tax assessment value.

²⁷Our use of a rental equivalence method is in line with current practice by U.S. government agencies: A similar approach is used by the Bureau of Labor Statistics for the Consumer Price Index, and by the Bureau of Economic Analysis in constructing the National Income and Product Accounts.

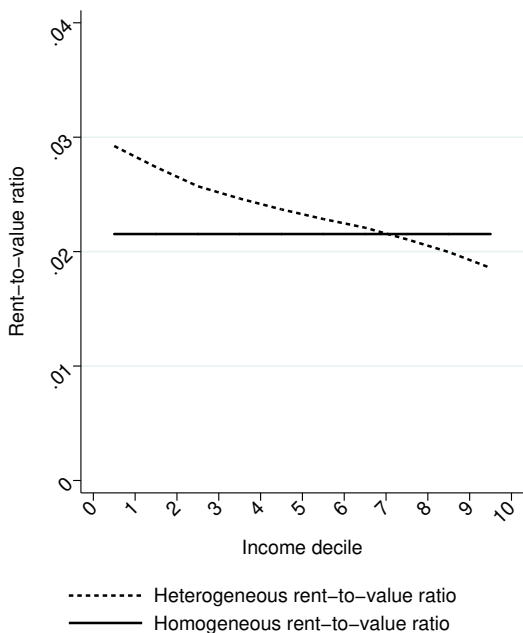
²⁸Our alternative approach to estimate the value of owner-occupied housing services uses residence characteristics reported in the 2011 Population and housing census (which are linked to our household identifiers), and regression coefficients based on actual rents and residence characteristics reported in the 2015 Rental Market Survey. Residences are characterized by size, number of rooms, two residence types, and five geographic regions. To ease comparison, we normalize the 2015 Rental Market

decreasing both with disposable income and with the value of primary residences. Except at the tails of the distributions of income and the value of primary residences, respectively, differences between the homogeneous and the heterogeneous rent-to-value ratios are moderate. In the lower panels of Figure A1 we plot the two measures of the value of owner-occupied housing services against disposable income and the value of primary residences. Differences across the two measures are moderate, except at very high levels of income and for the most valuable residences, where the national accounts procedure produces lower values due to a lack of heterogeneity across residences with the same characteristics. Figure A2 shows the cross-sectional distributions of consumption associated with the two approaches to estimating the value of owner-occupied housing services. It is reassuring to find that the two distributions are very much alike.

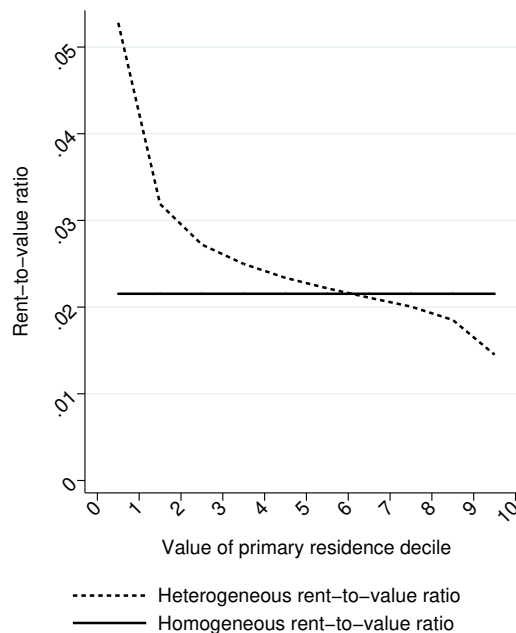
Survey rents to the level of the derived rents for 2011.

Median Rent-to-Value Ratio

(a) By income

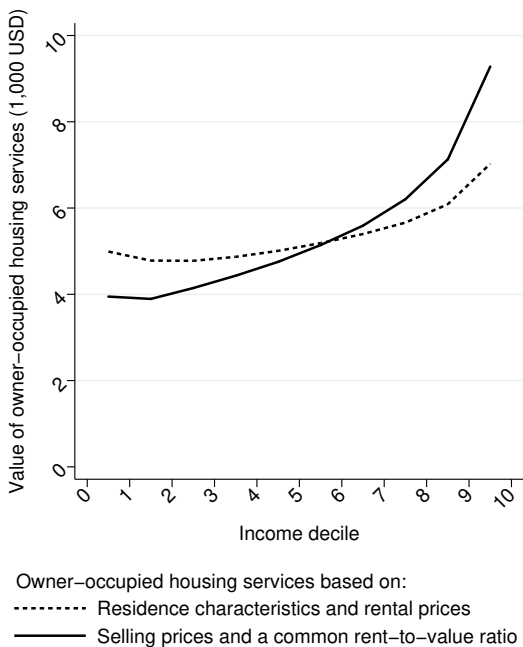


(b) By value of primary residence



Mean Value of Owner-Occupied Housing Services

(c) By income



(d) By value of primary residence

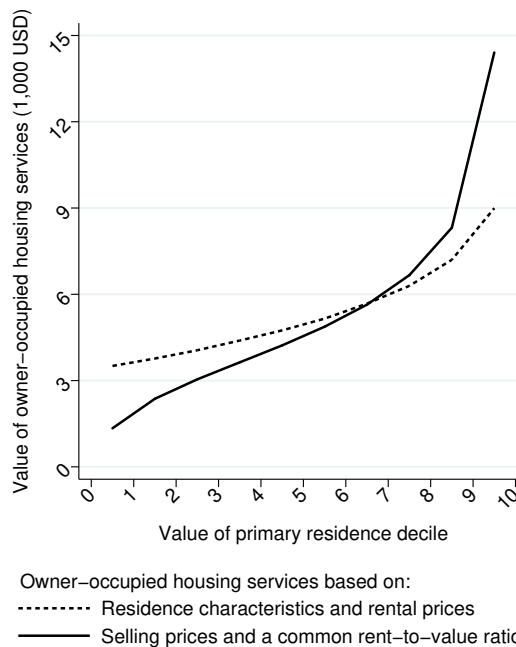


Figure A1: Homogeneous vs Heterogeneous Rent-to-Value Ratios

Notes: This figure displays median rent-to-value ratios (upper panels) and the average value of owner-occupied housing services (lower panels) among homeowner households in 2011. Values are expressed in 2014 USD and adjusted for household size using the EU scale, and households are weighted by the number of household members.

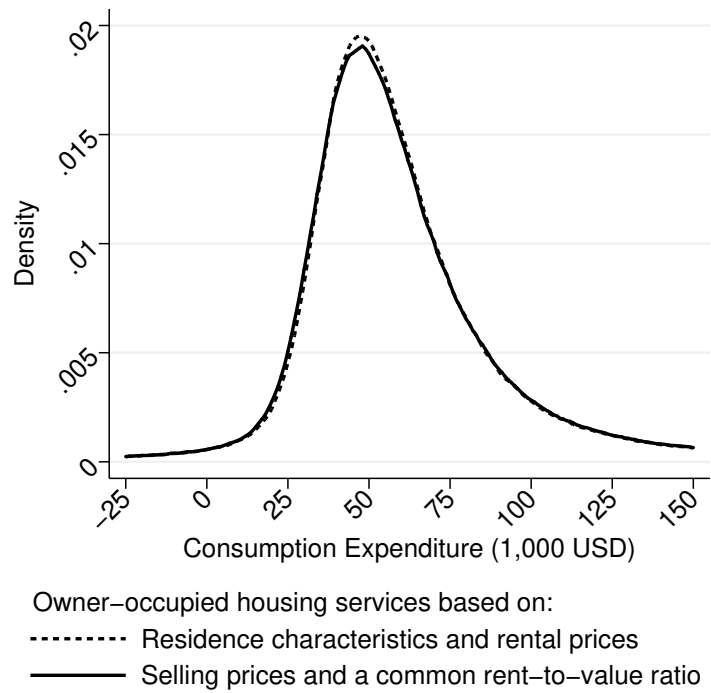


Figure A2: Distributions of Consumption Using Different Rent-to-Value Ratios

Notes: This figure displays distributions of consumption among homeowner households in 2011. Consumption is measured in 2014 USD and adjusted for household size using the EU scale, and households are weighted by the number of household members.

B Capital gains on financial assets

B.1 Different measures of capital gains and net savings

In order to derive reliable measures of household consumption expenditure, we need to separate between shifts in asset values caused by savings and shifts caused by changes in asset prices (capital gains), as only the former is relevant when inferring how much the household consumes out of its income. The Norwegian tax returns data contain end-year asset values for different types of assets (bonds, equity funds, other listed securities, and unlisted securities), but not the unrealized capital gains or the net savings. This section describes how we supplement the tax returns data with information on price changes for each type of asset and on net savings in listed and unlisted stocks to derive reliable measures of net savings in financial assets.

We have access to transactions data covering transactions in listed and unlisted stocks over the period 2003–2014, that can be linked with the tax returns data using unique identifiers for individual buyers and sellers. For each individual we observe the number of shares purchased or sold, $A_{ikt} - A_{ikt-1}$, and the price p_{kt} . For the asset types other listed securities and unlisted securities, therefore, we can compute precise measures of annual net savings directly, and define capital gains residually as the change in end-year asset values minus net savings.²⁹

For the years and types of assets for which we do not observe transactions, we rely on a combination of end-year asset values from tax returns and changes in asset prices from the Financial Accounts. Assuming the same increase in prices within each type of assets, we assign annual price changes from the Financial Accounts to the household’s asset positions at the beginning and end of the year. The mean of the two resulting asset returns measures is used as our measure of capital gains.³⁰ Net savings are then defined residually as the change in end-year asset values minus capital gains. For unlisted securities, we assume that capital gains equal the change in end-year asset values; that is, we assume that stocks in private companies are not traded. This is partly justified by unlisted securities being traded less frequently than other assets,³¹ and partly by price shifts being likely to deviate quite substantially from the common price trend suggested by the Financial Accounts.³²

Figure B1 compares our preferred measure of net savings in financial assets to measures resulting from three different alternative approaches to handling capital gains on financial assets: (i) measuring capital gains on all types of assets (except unlisted securities) using common price trends from the Financial Accounts, and assuming no trade in unlisted securities; (ii) approximating capital gains with a measure of taxable capital gains from tax records; and (iii) assuming no capital gains. Alternative (i) is our preferred approach for years

²⁹We calculate net savings from the transactions data, and derive a measure of capital gains as the change in end-year asset values less net savings. Capital gains on listed securities are accurately measured with this procedure, since end-year market values can be inferred from end-year tax assessments. For unlisted securities, however, our measure of capital gains contains not only capital gains, but also differences between selling prices and tax assessments. However, measurement error in capital gains on unlisted securities has no impact on our derived measure of consumption expenditure, since the transactions data allow us to measure net savings directly.

³⁰We also estimated capital gains based on asset positions at the beginning of the year only. A transition to this approach has only minor impacts on the derived measures of savings and consumption expenditures, relative to our preferred approach that partly accounts for within-year transactions.

³¹15 percent of households with unlisted securities traded unlisted securities in 2009, while 39 percent of those with listed securities traded listed securities.

³²Share ownership in a private company is usually difficult to value due to the absence of a public market for the shares. In Norway, the tax assessment value of an unlisted stock depends on the book value per share, reflecting the company’s assets minus its liabilities (but not necessarily future earnings and risks). Moreover, it is the book value for year $t - 1$ that enters the year t tax returns of a personal shareholder. Private companies that own stocks in public companies are therefore valued according to the end-year value of the stocks one year earlier, and a private company that owns a private company that owns a public company is valued according to the end-year value two years earlier, etc. The tax assessment values of private companies are therefore lagging behind the actual values.

prior to 2003, for which we do not have data on transactions in stocks. The figure shows that alternative (i) is reasonably close to our preferred measure, while alternatives (ii) and (iii) are rather far off for most years, both in levels and in terms of the time trends.

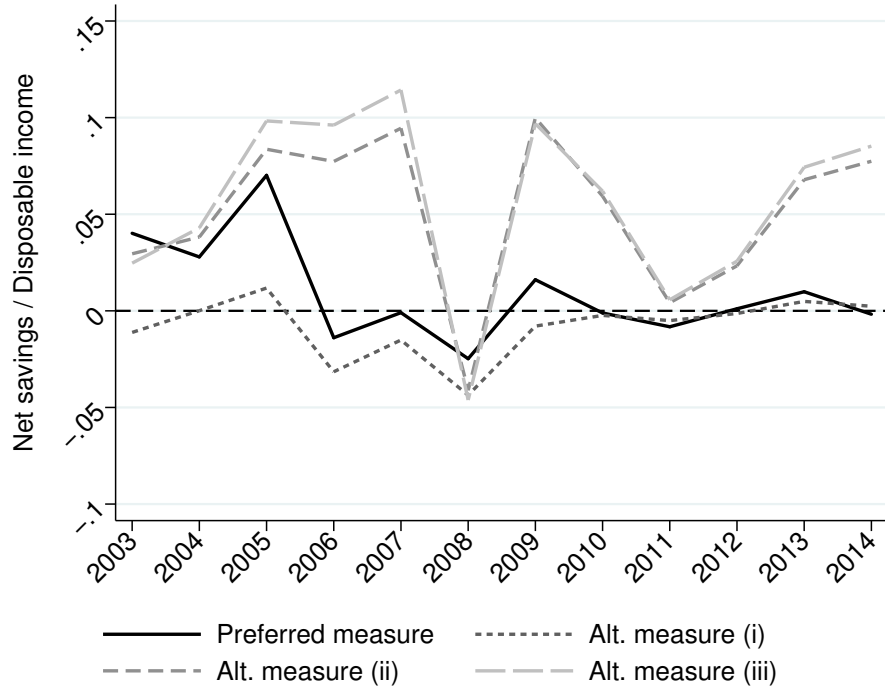


Figure B1: Different Measures of Net Savings in Financial Assets

Notes: This figure shows total net savings in financial assets as fractions of total disposable income. It compares our preferred measure of net savings in financial assets to measures based on three different approaches to handling capital gains on financial assets: (i) measuring capital gains on all types of assets except unlisted securities using common price trends from the Financial Accounts, and assuming no trade in unlisted securities; (ii) approximating capital gains with a measure of taxable capital gains from tax records; and (iii) assuming no capital gains.

B.2 Net savings based on end-year prices

In this section, we investigate the consequences of assuming no within-year transactions for measures of net savings by contrasting our preferred transactions-based approach with a measure of net savings based on end-year prices and the stock of each asset. Unlike the Swedish tax records, the Norwegian tax records do not contain information about the stock of each asset at the end of each year. However, we collect information on the number of listed and unlisted stocks held at the end of the year from the Shareholder register, and asset specific prices for listed stocks from Datastream. The Shareholder register is available from 2004 onward. To calculate an alternative measure of net savings in listed stocks, we proceed as follows: We calculate capital gains by assuming that stocks held at the end of year $t - 1$ are held until the end of year t ; $(p_{kt} - p_{kt-1})A_{ikt-1}$. Equipped with a measure of capital gains, we calculate net savings residually as the change in wealth minus capital gains: $S_{it} = p_{kt}A_{ikt} - p_{kt-1}A_{ikt-1} - (p_{kt} - p_{kt-1})A_{ikt-1}$.

The measurement error in measures of net savings based on end-year prices and holdings arises from two sources: within-year transactions, and changes in prices between the time of transaction and the end of the year. To highlight each of these two sources of measurement error, we construct a simple example where we split the year in two: $t = 0$ refers to the beginning of the year (January 1), $t = 2$ refers to the end of the year

(December 31), and $t = 1$ is somewhere in-between (July). Consider an individual who owns one asset worth p_0A_0 in January, trades for $p_1(A_1 - A_0)$ in July, and ends the year with p_2A_2 worth of the asset. (Since we consider one individual owning a single asset, we omit the i and the k subscripts.) Net savings for the entire year are given by

$$\begin{aligned} S &= p_2A_2 - p_1A_1 - (p_2 - p_1)A_1 + p_1A_1 - p_0A_0 - (p_1 - p_0)A_0 \\ &= p_2(A_2 - A_0) - (p_2 - p_1)(A_1 - A_0), \end{aligned} \tag{B1}$$

while a measure of savings based on end-year prices and holdings would be calculated as follows:

$$\begin{aligned} S^m &= p_2A_2 - p_0A_0 - (p_2 - p_0)A_0 \\ &= p_2(A_2 - A_0). \end{aligned} \tag{B2}$$

Equations (B1) and (B2) make clear that the two measures will coincide only if (i) there are no within-year transactions ($A_1 = A_0$), or (ii) prices do not change between the time of transaction and the end of the year ($p_2 = p_1$). We now proceed by assessing the empirical relevance of assumptions (i) and (ii), before we compare our preferred measure of net savings (based on transactions) to one based on end-year prices and holdings.

Panel (a) of Figure B2 shows how transactions in listed stocks by individuals in our sample are spread across the months of the year, in terms of the average fraction of yearly transaction volumes occurring in each calendar month. Although transaction volumes tend to be higher at the beginning and end of the year than in the middle of the year, a substantial fraction of yearly transactions is occurring within the year rather than at the end of the year. Panel (b) of Figure B2 shows the distribution of differences between observed prices and end-year prices for all sales and purchases of listed stocks by individuals in our sample. Considering only the price differences between -100 and 100%, the distribution looks fairly symmetric around zero, but with substantial mass in each of the two tails. For about 5% of the transactions, the price more than doubled between the time of transaction and the end of the year.

Figure B3 compares two measures of net savings in listed stocks: our preferred measure calculated using transactions, and an alternative measure based on end-year prices and holdings. The two measures are not too far apart for the full sample of households (Panel (a)). This is mainly because only 13% of the households in our sample own listed stocks, and only about half of those owning listed stocks have non-zero net savings in listed stocks in a given year. The differences between the two measures are more pronounced when we restrict the sample to “big traders” in Panel (b), with “big traders” being defined as households in the top 1% of the distribution of the absolute value of net savings in a given year.

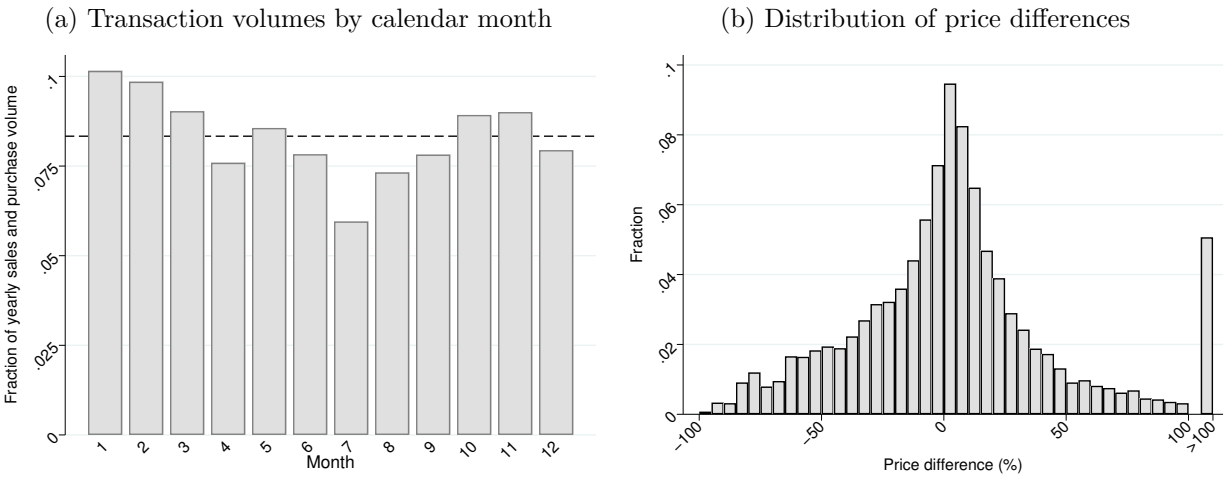


Figure B2: Transactions by Calendar Month and Differences Between Observed Prices and End-Year Prices

Notes: Panel (a) shows average monthly transaction volumes as fractions of yearly transactions, for transactions in listed stocks by individuals in our sample. The horizontal line indicates a hypothetical uniform distribution. Panel (b) shows the distribution of differences between observed prices and end-year prices for all sales and purchases of listed stocks by individuals in our sample. Price differences are reported as $100 * (\text{End-year price} - \text{Transaction price}) / (\text{Transaction price})$. Both figures are based on data for the years 2004–2014.

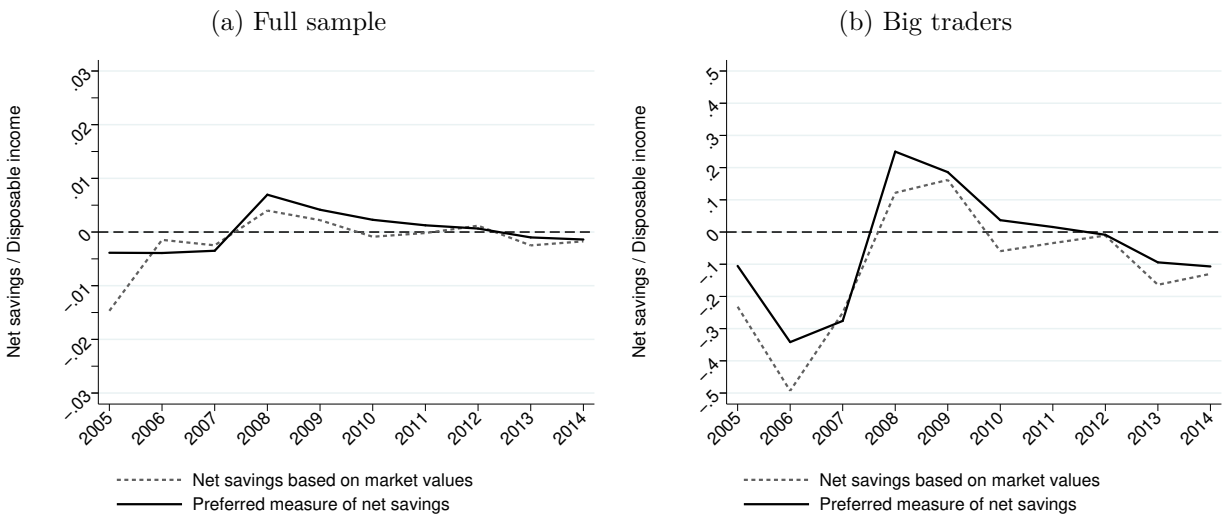


Figure B3: Different Measures of Net Savings in Listed Stocks

Notes: This figure shows total net savings in listed stocks as a fraction of total disposable income, for the full sample of households in Panel (a) and for households in the top 1% of the distribution of the absolute value of net savings in a given year in Panel (b).

C The National Accounts

The national accounts (NA) measure aggregate output, consumption, and income for various institutional sectors in a country, with individuals residing in the country constituting the household sector. Income, consumption, and savings in fixed capital formation net of depreciation are estimated explicitly, while net financial investments are set as the residual. Measures are based on a number of register- and survey-based data sources. In particular, the measure of consumption is based on register data on sales and production of goods and services, and on surveys such as the Rental Market Survey. The definitions of household consumption, income, and savings in the NA are broadly similar to our definitions, with only a small number of notable differences:

- (1) Home renovation expenses are for the most part considered to be investments in the NA, and will be counted as consumption in our data.
- (2) Estimated depreciation of real estate counts as (negative) income and savings in the NA but is not included in our measures.
- (3) The households' net savings in pension funds are counted as income and savings in the NA but are not included in our measures.³³
- (4) The part of FISIM that is not related to home mortgages is added to income and consumption in the NA but is not included in our measures.³⁴
- (5) In the NA, durables other than real estate and capital for business activities are counted as fully consumed in the year of purchase.³⁵

Table C1 shows how components (1)–(5) give rise to differences between our measures of household consumption, income, and savings and the corresponding measures in the NA. The first and last rows show that adjusting for differences due to (1)–(4) reduces the gaps between the NA and our measures of consumption and savings, while the gap in disposable income per capita is somewhat larger after adjusting the NA measure of income accordingly. Figure C1 shows the time series of disposable income, consumption, and savings per capita for both the unadjusted and the adjusted NA measures, while Figure 12 in Section 3.2 shows the adjusted measures only.

³³Income invested in pension funds is not included in the measure of disposable income in the NA but rather as a separate adjustment term. We initially define income in the NA as the sum of disposable income and investments in pension funds (such that $income = consumption + investments$).

³⁴FISIM stands for Financial Intermediation Services Indirectly Measured, and is an estimate of the value of services provided by financial intermediaries without explicit charges. In the NA, FISIM related to home mortgages counts as inputs in the production of housing services and is therefore crossed out of the account.

³⁵While potentially important on the disaggregate level, smoothing durable consumption intertemporally has little impact on the aggregate level as long as the stock of durables does not change significantly.

Table C1: Differences between the National Accounts and Our Data in 2014 (in USD/Capita)

	Income	Consumption	Savings
NA unadjusted - register data	821	-3,351	4,172
Adjustments:			
(1) Home renovation		2,902	-2,902
(2) Depreciation	1,942		1,942
(3) Pension funds	-2,714		-2,714
(4) FISIM	-1,029	-1,029	
(5) Durables		-5	5
NA adjusted - register data	-980	-1,484	503

Notes: The first and last rows show differences between measures of income, consumption, and savings per capita in the 2014 NA and the corresponding register based measures, before and after adjusting the NA measures to account for differences due to (1)–(5). The middle rows show how each of the adjustments matter for the NA measures of income, consumption, and savings. The adjustment for (1) Home renovation is obtained as the difference between investments in real estate as measured in the NA and investments in real estate as measured in our data. The adjustment for (5) Durables is calculated using our data, and defined as the difference in consumption and savings that occurs when we change from treating durables (other than real estate) as fully consumed in the year of purchase to our preferred treatment of durables. Adjustments (2)–(4) are taken from the NA. All measures are reported in 2014 USD.

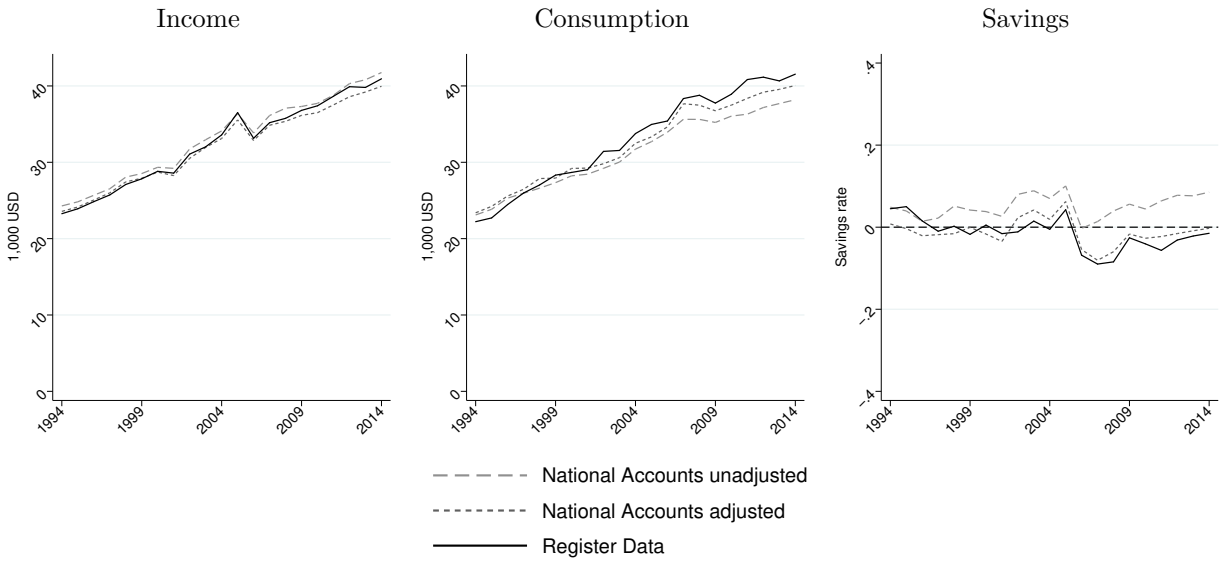


Figure C1: Income, Consumption, and Savings per Capita - Register Based Measures against Unadjusted and Adjusted National Accounts Measures

Notes: The savings rate is defined as the difference between disposable income and consumption per capita, divided by disposable income per capita. Income and consumption are measured in 2014 USD.

D Gifts and inheritances

Gifts and (advances of) inheritances are self-reported in the tax returns, but reporting is required only for the receiver. To improve our measurement of gifts and inheritances we therefore supplement the tax returns data with information from the inheritance tax registry, which contains personal identifiers for both donors and receivers. The inheritance tax registry is available from 1995 to 2013 and includes gifts and inheritances that are relevant for tax purposes. Gifts and inheritances are taxable only when the values received from the same person exceed a certain tax-free amount (USD 80,000 in 2013). The inheritance tax registry nevertheless also includes transfers below the tax-free amount, as multiple transfers may exceed the threshold collectively. Taxes on gifts and inheritance, and therefore also the inheritance tax registry, were abolished in 2014.