DISCUSSION PAPER SERIES

DP12402

STUDYING CONSUMPTION PATTERNS USING REGISTRY DATA: LESSONS FROM SWEDISH ADMINISTRATIVE DATA

Jonas Kolsrud, Camille Landais and Johannes Spinnewijn

LABOUR ECONOMICS and PUBLIC ECONOMICS

ISSN 0265-8003

STUDYING CONSUMPTION PATTERNS USING REGISTRY DATA: LESSONS FROM SWEDISH ADMINISTRATIVE DATA

Jonas Kolsrud, Camille Landais and Johannes Spinnewijn

Discussion Paper DP12402
Published 29 October 2017
Submitted 14 November 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

STUDYING CONSUMPTION PATTERNS USING REGISTRY DATA: LESSONS FROM SWEDISH ADMINISTRATIVE DATA

Abstract

This paper measures consumption expenditures using registry data on income and asset holdings in Sweden and illustrates how a registry-based measure can alleviate some critical limitations of traditional survey measures in capturing changes in consumption inequality and consumption responses to shocks. In the construction of our measure, we build on previous work exploiting the identity coming from the household budget constraint between consumption expenditures and income net of savings. We try to improve this measure using more registry information to account for the contribution of both financial and real assets to consumption flows. We demonstrate the power of the registry-based measure to study the relationship between income and consumption inequality, especially at the top of the income distribution. We also exploit the longitudinal dimension to study consumption responses to important life-time events and the different means used to smooth consumption.

JEL Classification: N/A

Keywords: N/A

Jonas Kolsrud - jonas@kolsrud.se
Uppsala University

Camille Landais - camillelandais81@gmail.com
LSE and CEPR

Johannes Spinnewijn - j.spinnewijn@lse.ac.uk
LSE and CEPR

Studying Consumption Patterns using Registry Data: Lessons From Swedish Administrative Data

Jonas Kolsrud* Uppsala University

Camille Landais† London School of Economics

Johannes Spinnewijn‡§ London School of Economics

November 8, 2017

Abstract

This paper measures consumption expenditures using registry data on income and asset holdings in Sweden and illustrates how a registry-based measure can alleviate some critical limitations of traditional survey measures in capturing changes in consumption inequality and consumption responses to shocks. In the construction of our measure, we build on previous work exploiting the identity coming from the household budget constraint between consumption expenditures and income net of savings. We try to improve this measure using more registry information to account for the contribution of both financial and real assets to consumption flows. We demonstrate the power of the registry-based measure to study the relationship between income and consumption inequality, especially at the top of the income distribution. We also exploit the longitudinal dimension to study consumption responses to important life-time events and the different means used to smooth consumption.

Keywords: Consumption Measurement, Registry Data, Inequality and Smoothing

1 Introduction

The access to administrative data from tax, social security and other government registries has been the driver of a real data revolution in economic research and in the field of public economics in particular. Researchers around the world have been doing enormous data collection efforts using these registers to document trends in inequality in income, wealth, health, social mobility, etc which are having an important impact on the public debate. This data revolution has lead to new empirical breakthroughs that should help evaluating policy: the analysis of income responses to tax incentives, unemployment responses to unemployment benefits, health expenditure changes in response to health coverage, etc. While the analysis of household consumption - understanding

*Institute for Housing and Urban Research, P O Box 514, SE-751 20 Uppsala, jonas.kolsrud@nek.uu.se

†Dept of Economics, Houghton Street London, WC2A 2AE +44(0)20-7955-7864, c.landais@lse.ac.uk

‡Dept of Economics, Houghton Street London, WC2A 2AE +44(0)20-7955-7022, j.spinnewijn@lse.ac.uk

§We thank Arnaud Dyèvre and Mathilde Muñoz for excellent research assistance. We also acknowledge financial support from the ERC grants 679704 and 716485.

how it responds to income shocks and how these responses affect consumption inequality - has remained central in economic modelling, with a few recent exceptions, the ever-growing literature on consumption has been mostly confined to the use of survey data on consumption.

Traditional measures of consumption based on surveys have some clear advantages and can provide information that may be impossible to attain otherwise, but they suffer from important limitations as well. Most evidently, the relative small samples of traditional budget surveys make any aggregate or distributional analysis challenging and affect the precision with which consumption responses to adverse events or income shocks can be estimated. This paper builds on recent work to construct an alternative, registry-based measure of consumption using comprehensive, detailed information on income and wealth from Swedish tax registers for the *universe* of households in Sweden. The basic idea underlying the construction is that household consumption equals household income minus changes in the household's asset positions. Importantly, the universal coverage of the registry data clearly helps overcoming the small sample limitation of survey data, the advantages of which we demonstrate in a series of applications.[1]

The registry-based measure we construct has additional advantages that can contribute positively to the quality of empirical research on consumption. First, beyond the universal coverage, administrative records allow for the construction of long panels, which we exploit for analyzing consumption dynamics in response to particular events. Second, the measure allows to analyze precisely how different means are used to smooth consumption in response to shocks and reduce the variance in consumption. We can shed light on the use of debt and different types of assets, but also on the role played by other members of the household or by the government through taxes and transfers, etc. Third, the registers allow to identify income shocks or specific events underlying income shocks like unemployment or retirement, by linking tax registers to social security registers. Finally, the registry-based measure benefits from the accuracy and reliability of the third-party reported administrative data.

Still, the quality of the registry-based measure critically depends on the comprehensiveness of the available data on income and changes in the asset portfolio. Indeed, to identify consumption expenditures from a household's budget constraint one requires data on total income and all changes in asset holdings. For example, the consumption out of bank accounts from one period to the next equals the interests earned minus the change in the account holdings, both of which are precisely recorded in the Swedish administrative registers. However, measuring the contribution of changes in other asset types to consumption can be challenging. In an ideal setting all asset transactions are observed and one can simply subtract purchases from sales to calculate the corresponding flow of consumption expenditures. In practice, the way asset variables are recorded in tax registers may not be directly relevant for determining the consumption flow, like when levying a tax on realized capital gains or on wealth. Sweden had a wealth tax in place until 2007 and the tax registry actually recorded information on the end-of-year balances (i.e., quantities) of all listed financial

[1]See Pistaferri [2015] for a broader discussion of the measurement issues with household consumption and the use of other data sources (e.g., scanner data, data from financial aggregators)

assets, which allows us to calculate the re-balancing from one year to the other. Even when this data is missing, once can still try to exploit information on prices to impute the contribution from changes in balance values to the consumption flow.

We discuss different approaches to use registry information in constructing a measure of consumption and explain in detail how these can be applied in Sweden. In addition to the challenges with the financial assets, we also discuss issues around durable goods and in particular the role of real estate and the imputation of rents for the measurement of consumption. We also exploit the availability of the annual Household Budget Survey (HUT, Hushllens Utgifter), which can be matched to the registry data and thus allows us to compare our registry-based measure with survey-based measures of consumption.

As a first check, we confirm that aggregate consumption using the registry-based measure closely follows the national accounts. Moreover, the distribution of the registry-based measure is similar to the distribution of the survey measure and the correlation between the two is very strong. Still, the measures tend to differ more systematically in the tails. Importantly, the survey measure tends to underestimate consumption levels relative to the registry-based measure at the high end of the distribution. This reflects the well-known challenge for surveys to incentivize high-income households to accurately report their expenditures on the one hand and to capture the tail of Pareto distributions on the other hand. As a consequence the registry-based measure seems to dominate the survey measure in capturing how income inequality translates into consumption inequality, especially at the top of the income distribution. In particular, we show that five-year differences in income are strongly correlated with five-year differences in consumption and the correlation is highest across the high income percentiles. Not only the income increases, but also the consumption increases are the largest at the top of the income distribution. These patterns are completely missed using the consumption surveys. The registry-based measure thus offers quite a different perspective on recent arguments in the US that the increase in consumption inequality has been much less pronounced than the increase in income inequality, especially at the top (see Krueger and Perri [2006]).

We also demonstrate the value of the registry-based measure and its longitudinal dimension to analyze how consumption responds to income and wealth shocks. A large body of economic research has analysed how idiosyncratic income changes affect household consumption (e.g., Blundell et al. [2008], Heathcote et al. [2014], Blundell et al. [2016]). Stemming back to the Life Cycle-Permanent Income Hypothesis of Modigliani-Brumberg and Friedman, this literature relates movements in consumption, mostly relying on survey data, to anticipated and unanticipated income changes as well as persistent and non-persistent shocks to resources. Instead we focus on three well-specified events during a lifetime that have a substantial impact on individuals' income and expenditures: parenthood, health shocks and retirement. Our focus on these specific events circumvents the challenging identification of the specific nature of income shocks and is inspired by the direct relevance for policies that condition on these events (e.g., Baily [1978], Chetty [2006]). We use event-study designs which complement our related work studying how household consumption responds

to job loss (see Landais and Spinnewijn [in progress]) and evolves during the unemployment spell (see Kolsrud et al. [2015]).

For each of the three events that we study, we find a substantial drop in earnings that persists over several years and is exacerbated by a drop in earnings for other household members. The drop in earnings corresponds to substantially smaller drops in consumption, but the means of consumption smoothing are very different depending on the event. For the health shocks, we find that most cushion is provided by sickness and disability benefits, while financial assets play a very limited role in smoothing consumption, leading to (or consistent with) a permanent decline in consumption expenditures of around 5%. We find a more substantial role for retirement savings in addition to the pension benefits in smoothing consumption around retirement, but consumption expenditures still drop by around 7 to 10%. For the arrival of the first child, social insurance again provides substantial smoothing, but we find a quite different and smoother pattern for consumption expenditures driven by consumption out of assets. Leading up to the birth of the child, households increasingly become home owners and 'consume' more housing services in the form of higher imputed rents. Interestingly, we find similar patterns of consumption out of other durable goods. In particular, we find similar anticipatory behaviour in the purchase of new cars.

Overall, we are not only exploiting the panel nature and the possibility of demonstrating the different means of consumption smoothing in these applications, but we are also very much relying on the power of the registry-based measure. In comparison, the estimates obtained using consumption surveys are much more noisy and an order of magnitude less precise.

All our programs and detailed documentation of the data are available online and hopefully help other researchers in their analysis of household consumption. While the Swedish context is quite ideal for the construction of a registry-based measure, we hope that our efforts can also spur further research to improve our measurement, extend it beyond the period 1999-2007 in Sweden, but also extend it to other countries with potentially less precise information on income or wealth.

The remainder of our paper is organized as follows. In the final part of this introduction, we explain how this paper relates to the existing literature dedicated to building consumption meaures out of registry data. In section 2, we lay out the methodology used to build residual measures of expenditures out of household budget constraints and information on income and wealth. In section 3, we explain the way we implemented this methodology in the Swedish context given the data available. In section 4 we show various descriptive statistics for our registry measure of consumption expenditures. We show how it compares to other measures of consumption, from national accounts or survey data, and how our measure proves better at capturing consumption inequality and dynamics at the top end of the distribution. In section 5, we further demonstrate the value of our measure in a series of applications devoted to the analysis of consumption responses to various life cycle shocks. Section 6 concludes.

Related Literature In our attempt to construct an alternative, registry-based measure of household consumption, our work relates directly to a literature initiated by Browning and Leth-Petersen

[2003]. The original goal of this stream of papers was to investigate issues with existing consumption survey measures and how to improve upon them using registry data (cf. for instance Kreiner et al. [2014]). The comparison of consumption survey measures and registry-based measures of course is facilitated when the survey data and registers can be linked, like in Denmark and Sweden. In this context, Browning and Leth-Petersen [2003] were the first to measure consumption expenditures as the residual from what one earns minus what one saves, exploiting information on both income and wealth in Denmark. Being unable to properly impute unrealized capital gains, Browning and Leth-Petersen [2003] assumed away capital gains in the construction of their registry-based measure and that the benefits of this measure of consumption were limited.

To overcome this limitation, subsequent work has tried to obtain additional data in order to properly impute capital gains. In the context of Sweden, Koijen et al. [2014] use disaggregated information on financial assets (stocks, bonds, mutual funds,etc) to determine the proper contribution of changes in these financial assets to consumption. They also use some information on real estate transactions in order to disentangle unrealized capital gains from active rebalancing of households' real assets. Our paper directly builds on Koijen et al. [2014], using similar data in Sweden, but extending the analysis for a longer time span and for the universe of the Swedish population.[2] We also improve their measure of consumption expenditures by properly imputing rents to homeowners and by refining the computation of real estate transaction flows. Since Koijen et al. [2014] do not consider imputed rents in households' capital income, they have to operate under sample restrictions to address the non-comparability of homeowners' and renters' expenditures.[3]

Our paper also relates to two recent papers trying to build similar consumption measures from administrative data in Norway, the only Scandinavian country that still has a wealth tax in place. Fagereng and Halvorsen [2017] use account balances for individual stocks similar to Koijen et al. [2014].[4] Compared to Koijen et al. [2014], they include imputed rents from owner-occupied housing, but they also resort to sample restrictions to mitigate measurement error concerns regarding the used market value of real estate wealth.[5] Interestingly, Fagereng and Halvorsen [2017] also use data on inheritances and gifts, which is available in Sweden as well. Most related to our paper is Eika et al. [2017], who further improve the comprehensiveness of the registry-based consumption measure. Compared to Fagereng and Halvorsen [2017], they use data on real estate transactions and also impute consumption flows for cars using registry data. Our paper is similar in spirit, both in the construction of the registry-based measure and the use of applications to demonstrate the value-added of registry-based measures of consumption expenditrues. Moreover, since in Sweden the Budget Surveys are linkeable to the registry data, we can revisit our empirical applications using

[2]Koijen et al. [2014] focus on a subsample of roughly 11,000 households that are both present in the LINDA panel ($\approx 300,000$ individuals) and in one cross-section of the consumption expenditure survey over the period 2003 to 2007.

[3]Similarly to Kreiner et al. [2014], Koijen et al. [2014] use registry information to inform the magnitude of potential biases in survey data. Using car registers, they investigate the importance of recall errors in survey data regarding purchases of large durables.

[4]The data on financial asset holdings is limited to stocks publicly traded in Norway, and is only available for years 2004-2011.

[5]They use the so-called "capital market approach" and assume a single long-term rate of return to the value of real estate investment.

the survey measure and thus make the comparison with this traditional approach more explicit.

2 Residual Measure of Expenditure: Methodology

Like previous work measuring consumption with registry data (e.g. Browning and Leth-Petersen [2003], Koijen et al. [2014] and Eika et al. [2017]), we start from the observation that any krona of income is either spent or invested. A household's budget constraint indeed leads to an accounting identity between expenditures on the one hand and income net of the change in assets on the other hand. Consumption expenditures C_{it} by household i in period t can thus be written as

$$C_{it} = Z_{it} - \sum_k p_{kt} \left[A_{ikt} - A_{ikt-1} \right], \tag{1}$$

where Z_{it} captures all sources of income and transfers, $\mathbf{A}_{it}$ denotes the portfolio of assets and $\mathbf{p}_t$ the corresponding vector of prices at which they are traded. The practical challenge in using this identity is twofold:

First, we require comprehensive information on all sources of income of a household and on all assets its members own. This challenge can be mostly overcome by the quality and scope of registry data. Indeed, data from administrative registries can provide detailed information on labor income, capital income, taxes paid and transfers received, including even private transfers like bequests. This information is mostly third-party reported by employers, financial institutions, government administrations, etc. The registries miss out on the income and wealth that is unreported, but the resulting measure error is small when tax enforcement is strong and the role of cash holdings is minor.

Second, to construct how a household's asset portfolio contributes to the consumption by its members, one needs to know the asset transactions, $A_{ikt-1} - A_{ikt}$, and the corresponding transaction price p_{kt} for all assets k. Transaction data is hard to come by. Exceptions exist when transactions are subject to a specific tax (e.g., housing) or recorded for other reasons (e.g., cars). Depending on the tax treatment of capital, different types of data may be available at different levels of aggregation, including on capital income $Y_{ikt} = r_{kt} A_{ikt}$ (e.g., interests, dividends), on realized capital gains $G_{ikt} = \left[p_{kt} - p^0_{kt} \right] \left[A_{ikt-1} - A_{ikt} \right]$ (e.g., on sales of stocks), where p^0_{kt} is the original purchase price of the asset sold at t, and/or on the balance value $W_{ikt} = p_{kt} A_{ikt}$ of an asset. We include capital income in our measure of income Z_{it}, but not the realized capital gains. While a krona of capital income translates one-to-one in a krona of expenditures, a krona of realized capital gains does not, neither does a krona change in capital wealth. To see this, we can re-write the asset component of the identity in (1) as

$$p_{kt} \Delta A_{ikt} = \Delta W_{ikt} - \Delta p_{kt} A_{ikt-1},$$

where $\Delta X_{ikt} = X_{ikt} - X_{ikt-1}$. Since only the active re-balancing of assets contributes to the flow of consumption, this suggests two approaches to measure the contribution of assets to the consumption

flow: a price approach and a quantity approach.

Price Approach The first approach we call the price approach as we use the change in asset prices to impute the consumption flow from the change in asset wealth. When the price of an asset changes, one cannot simply use changes in the balance value to construct the consumption flow out of assets. Indeed, any unrealized capital gain will lead to an identical increase in wealth W without contributing to the consumption flow. However, when we have information on prices p_{kt}, we can simply impute the increase in wealth due to the price change and subtract this from the overall increase in wealth,

$$p_{kt}\Delta A_{ikt} = \Delta W_{ikt} - \frac{\Delta p_{kt}}{p_{kt-1}} W_{ikt-1}.$$

Note that this is exactly how financial assets like bank accounts are treated in measuring the corresponding consumption flow: we subtract interest payments $r_{kt}W_{ikt-1}$ (considered as capital income) from the change in the asset balance ΔW_{ikt}. Hence, a krona of savings interests accrued on the savings account does not contribute to the consumption flow.

The accuracy of the price approach will depend on the level at which information is aggregated. When observing only broad asset classes, we are forced to impute the unrealized gains using a price index for that asset class, not accounting for the exact holdings and price variation within that class. Heterogeneity in asset prices and portfolio holdings is substantial and missing out on this variation by only observing broader asset classes could be problematic, since

$$\frac{\Delta \bar{p}_{Kt}}{\bar{p}_{Kt-1}} W_{iKt-1} \neq \Sigma_{k\in K}[\frac{\Delta p_{kt}}{p_{kt-1}} W_{ikt-1}],$$

for asset types k in asset class K.

Quantity Approach The second approach we call the quantity approach as we start from observed differences in quantities of assets held to measure the corresponding consumption flow. Even when for tax purposes only aggregate wealth is relevant, banks or other instances may still provide more detailed information on asset quantities. Linking a change in quantities with the asset price allows approximating the value of the underlying transaction $p_{kt}\left[A_{ikt} - A_{ikt-1}\right]$. The shortcoming of this approach is that in practice such information is available at low frequency and the change in observed quantities may correspond to several transactions made in between the times of observation. The accounting identity, however, presumes that an individual holds on to her assets throughout the period and trades at the end of the period at price p_{kt}. Using end-of-the-year balances thus introduces measurement error if transactions are timed in between or when multiple transactions happen within a year (potentially to exploit price variation at higher frequency). That is,

$$p_{kT}\left[A_{ikT} - A_{ikT-F}\right] \neq \Sigma_{t=T-365}^{T} p_{kt}\left[A_{ikt} - A_{ikt-1}\right].$$

Comparing the price and quantity approach, a key advantage of the latter is that active re-

balancing is required for asset holdings to affect the imputed consumption flow. For the price approach, however, the imputed consumption flow could be fully driven by price variation. Note that additional information can be used to identify active re-balancing. For example, realized capital gains in a given year do comprise the sum of all transactions. However, the orginal purchase value is subtracted from the sales. That is, the realized capital gains at time T over the past year equals

$$G_{ikT,365} = \Sigma_{t=T-365}^{T} \left[p_{kt} - p_{kt}^{0}\right] \left[A_{ikt-1} - A_{ikt}\right].$$

Hence, while realized capital gains may seem useful, without information on the original purchasing price, we can not back out the resulting consumption flow.

Durable Consumption An important issue in the measurement of consumption is the treatment of investments in durable goods. While such investments generate a consumption flow for several periods, our approach assigns the entire consumption value to the period of purchase. The importance of making the distinction between durable and non-durable consumption will depend on the application one has in mind. For example, the ability to delay investments in durable goods in response to a temporary shock will improve the smoothing of the actual consumption flow and thus reduce the welfare impact of a shock. However, the distinction seems less relevant when studying consumption inequality. While consumption expenditures surveys often ask directly about durable investments, registry data may also provide information on particular durable assets owned (e.g., cars) and thus allow to impute the corresponding consumption flow. This seems to be particularly important for the case of real estate, where we account for the corresponding consumption flow by calculating the 'imputed rents'. This also improves the comparability of consumption expenditures for homeowners and renters.

3 Implementation in Sweden

We now describe the different components we use for the construction of a registry-based measure of consumption in Sweden. The main source of data is the longitudinal dataset LISA, which merges several administrative and tax registers for the universe of Swedish individuals aged 16 and above. In addition to socio-demographic information (such as age, family situation, education, county of residence), LISA contains exhaustive information on income and transfers. We link these data to other registers, including data on asset holdings. This data is collected by the tax administration as Sweden levied a tax on wealth (until 2007). Our data thus provides comprehensive information on income, transfers and assets, which are reported to the tax administration at the individual level, but can be aggregated up at the household level using household identifiers. The measure we construct based on this information captures annual expenditures between December of each year and is computable for the universe of households between 1999 and 2007. Table 5 provides an overview of the different income and wealth components and their corresponding magnitudes.

3.1 Income and Transfers

The dataset LISA contains exhaustive information on income from labour and capital, transfers and taxes on an annual basis. For labour earnings, we use a measure of disposable income with contains all employment-related earnings including wages, vacation pay and severance payments, but also wage income for self-employed and business owners. Business losses for small business owners and independent contractors are included in LISA from 2004 onward.[6] For capital income, we include interests earned or paid and dividends received on all assets, including banks accounts, bonds, listed securities and closely-held businesses. Capital income also includes rental income, but we explicitly take out the realized capital gains (or losses) from selling financial or real assets. To the income from labour and capital, we also add all government transfers received, which include unemployment benefits, sickness and disability benefits, pension benefits, housing benefits, child benefits and other smaller programs. In our application we single out the receipt of different benefits depending on the event we study. Any decomposition analysis is considered before taxes as we only have information on aggregate taxes paid. The aggregate taxes are, however, subtracted from the total income and transfers received.

3.2 Financial Assets

Data on wealth comes from the wealth register (Förmögenhetsregistret), which comprehensively covers the asset portfolios for the universe of Swedish individuals from 1999 to 2007. The register contains information on the balance value of all financial assets and listed securities without any top-coding of wealth. This includes information on aggregate holdings by asset class (bank accounts, bonds, stocks, mutual funds, etc.), but also on total outstanding debt, which includes mortgage debt, consumer credit, student debt, etc. All financial institutions were compelled to report this information directly to the tax administration for the purpose of the wealth tax.

The amount of bank account holdings reported in the wealth registry change over our sample period. Before 2006 only bank holdings which carry returns exceeding 100 kronor are reported.[7] After 2006, all bank holdings exceeding 10,000 kronor, irrespective of the amount of interest they carry, are reported to the tax administration.[8] We may also miss some financial wealth held in foreign banks to the extent that these banks do not comply with the requirement to transmit information on the financial wealth of their Swedish customers to the Swedish tax authority.[9]

[6]This results in slightly higher labour earnings before 2004. For instance, in 2004 labour earnings excluding business losses are 0.5% lower compared to the labour earnings measure which includes business losses. Throughout our sample period business losses can be carried forward for tax purposes. Hence, losses in one year affect tax payments and therefore disposable income in subsequent years.

[7]The exchange rate is 1 SEK≈0.11 USD

[8]Before 2006 about 70% of all bank holdings reported in national accounts are reported in the wealth statistics. The new censoring rule of 2006 results in 97% of all bank holdings found in national accounts to be found in the wealth register. How the change in censor rule is handled when estimating consumption expenditures is described in detail in Appendix B.

[9]Note that the fraction of foreign-held assets is low (about 3% of household assets according to the Savings Barometer of Statistics Sweden), and likely to be held by the wealthiest households. Foreign assets held via Swedish financial institutions are required to be reported while assets held by foreign financial institutions do not meet the

The balance values in the wealth registers allow us to compute the contribution to consumption from asset re-balancing for the asset classes with fixed prices. For example, for bank holdings, we calculate the change in bank holdings between year $t-1$ and t, which together with earned interests (included in our income measure) determines the consumption out of bank holdings. Similarly, for debt, we calculate the change in outstanding debt between year t and $t-1$, which net of paid interests (again included in our income measure) determines the consumption out of debt. The data is complemented with disaggregated data (KURU) on the balance positions for all financial securities by their International Securities Identification Number (ISIN). Swedish financial institutions are required to report the quantity and ISIN code of their clients' security holdings.[10] However, instead of using the reported prices, the wealth register held by Statistics Sweden only aggregates the listed securities for which it can find a trading price. The register thus excludes non-listed securities, like the ownership of small business or stocks in companies that are not publicly traded. Lundberg and Waldenström [2017] estimate that the unlisted securities amount to approximately 5% of all wealth held in securities. For the listed securities, we use closing prices from the last day of trade in December 2000 to 2007 based on the SIX data, which covers security prices from about 1,500 different securities exchanges and contributors. Just like Statistics Sweden, we currently ignore the financial securities from the KURU data for which we can not link the security to a price.[11] For options, warrants and the tax-favored capital insurance accounts, we do not observe the quantities owned, but only the balance values and use the price approach with an aggregate price-index to impute the consumption flows. Note that for the tax-deferred private pension accounts, we do not observe the balance value, but only the contributions and withdrawals. This is, however, exactly what we need to compute the corresponding consumption flows.

3.3 Real Assets

There are two main forms of real estate in Sweden that can be held by individuals: real estate that individuals own and real estate for which individuals own the occupation right. The first category includes all stand alone housing, semi-detached, bungalows, terrace housing and second homes, but also commercial real estate such as apartment buildings, industrial plants and farming property. The second category includes all apartments. In Sweden it is uncommon that individuals own their apartments. Rather individuals who live in apartments buy a share in a housing cooperation (co-op henceforth). Apartments can be bought and sold like any other real estate but from a legal point

same requirement, since the third-party reporting obligation only applies to Swedish financial institutions.

[10] This information is pre-printed on every tax payer's income statement, but can be amended by the tax payers.

[11] The SIX data has wider coverage than the Swedish Stock Exchanges used by Statistics Sweden to assign prices to the individual securities. Hence, we manage to retrieve more financial wealth than in the wealth register. Note that for about 1% of stock holders in the wealth register, we are not able to match their records to the KURU data. For these individuals we use the balance values in the wealth registers and a price-index to impute the consumption flows. Since our coverage of financial assets is not complete, we are worried that recorded changes in asset positions may be offset by corresponding changes in unrecorded assets positions. For 0.03% of individuals, we have trimmed too large transactions without offsetting changes in calculating the consumption flow. In principle, we could could further extend our coverage by using the self-reported prices for the securities for which we cannot assign a price. Appendix B provides more detail on the KURU data, its relation with the wealth register and the trimming procedure.

of view co-ops are different from regular housing as the occupier does not legally own the property, but only the right to occupy it.

The wealth tax register (Förmögenhetsregistret) contains detailed information on the stock of real estate wealth, estimated at market value as of December 31 of each year.[12] In principle, this allows for an imputation of the consumption flows following the price approach using housing price indices. For owned housing, however, we have data on all transactions from the housing price registry (Fastighetsprisstatistiken) from 1999 to 2007. Since we observe all real estate purchased and sold, we can directly calculate the contribution from re-balancing real estate holdings to the annual consumption flow for each household.[13]

A separate register records the sales of co-ops, but not the purchases. However, the purchasing price of the sold apartment is registered.[14] For individuals buying and selling coop apartments we combine the information from the wealth tax registers and the purchasing price of sold apartments to impute the corresponding consumption flows.[15] For individuals buying a coop without having owned one the previous year we use the coop wealth variable from the wealth registry.

Like for financial assets, investments in real estate generate capital gains (or losses), but can also generate other income. Rental income is part of the disposable income measure we use. For real estate that is not rented out, we impute rents to account for the consumption services a house provides. As discussed before, imputing such rents may be desirable for all durables which generate consumption flows beyond the moment of purchase, but clearly this value will be highest for housing services. Different approaches have been used for imputing rents, ranging from estimating the alternative income that would be generated when investing in a low-risk asset to estimating the potential rental income when subletting the house (see Poterba [1992]). The latter approach is followed by Statistics Sweden and used for the national accounts.[16] The imputed rent is calculated as the square-meter living space multiplied by a square-meter rental value which depends on the geographical region, construction period and dwelling category. While we do not have individual data on the imputed rents, we follow Eika et al. [2017] and assign the aggregate imputed rents to

[12]Statistics Sweden use data on the real estate tax value and regional sales statistics to compute market values for housing wealth. The real estate tax value, as used for determining the wealth tax, is a function of characteristics of the property and updated frequently. See Lundberg and Waldenström [2017], but also Appendix B for further details.

[13]From 2004 onward we observe all parties involved in a transaction. Before 2004 only one seller and one buyer is recorded, even if more individuals are involved. Before 2004 we calculate housing transaction flows on the household level under the assumption that only two households are involved in most transactions. After 2004 we calculate these flows on an individual basis. More detail is provided in Appendix B.

[14]The purchasing price of the sold apartment is recorded so that the tax administration can calculate the tax on sales profits.

[15]In particular, we observe individuals who buy a new coop apartment during the years for which we compute the residual consumption measure. Then we estimate the probability of buying a coop apartment based on variables such as household characteristics and changes in debt. We also use the same covariates to predict purchasing prices. For individuals where the probability of buying a coop apartment is sufficiently high we impute the purchasing value.

[16]We note that rents are controlled in Sweden according to a user cost principle. That is, rents are negotiated each year between the tenants association and the landlords and rents are only supposed to vary based on apartment characteristics. Thus, two apartments with equivalent standard within a geographical area cost the same regardless of the location in that area.

the households based on the value of their primary real estate in the wealth tax register.[17]

Finally, we have data on all car transactions from the vehicle register (Fordonsregistret). This data allows us to connect cars to their owners and includes information on the car's general status (registration date, when the car became the current owner's, in use or not, etc.) and various specifics (make, model, odometer reading, fuel consumption per kilometer, etc.). However, the data does not contain any value assessment of the cars. While we currently do not use this information in the construction of our consumption measure, we illustrate the use of this data to highlight patterns in the type of spending responses.

3.4 Extending beyond the 1999 - 2007 period

All our data ranges from at least 1999 until 2007. In principle, the measure could be extended for a longer time period. However, the wealth tax was abolished in Sweden in 2007, after which the government collected only limited information on assets. Hence, any imputation of the consumption flows would need to rely on information from realized capital gains and income (see Saez and Zucman [2016]). The data can be extended more easily to earlier years. The LISA data dates back to 1990. The wealth register (Förmögenhetsregistret) started only in 1999, but the comprehensive information on income and wealth can still be retained from the Income and Taxation Register (Inkomst- och taxeringsregistret) from 1988 onwards.

4 Registry-Based Measure: Descriptive Statistics

4.1 Comparison to National Accounts and to HUT Surveys

We start by comparing the evolution and the distribution of our registry-based measure of consumption expenditures to consumption measured in the national accounts (NA) and in HUT surveys of consumption expenditures. There are significant conceptual differences between what is considered consumption expenditures in the national accounts, in the consumption surveys, and in our registry-based measure.

The HUT surveys provide direct measures of bi-weekly consumption expenditures at the moment the household is surveyed. Information on consumption comes from shopping books where surveyed households are asked to report all their expenditures during the 2 weeks preceding their interview date. For a few durable goods, extra information is added from answers to interview questions, where individuals are asked to recollect some expenditures for up to 12 months prior to the interview date. In the national accounts, consumption is computed from various sources. Information from the HUT surveys is used, but complemented with various information from the registry data and with VAT receipt data. One of the main difference between our registry measure of expenditures and the NA measure pertains to the treatment of some investments. Large renovations, home refurbishments, and replacement of home appliances is treated as investment

[17] This imputation assumes that the primary real estate is not (partially) subletted and ignores potential consumption flows from secondary real estates.

(rather than consumption) in the national accounts, while they show up as expenditures in our registry measure. But compared with the survey data, both the NA and the registry measure of consumption include imputed rents. Both measures also include expenditures in the black market or on illegal goods and services.[18] In Appendix C we describe in more details the sources and concepts used for measuring consumption in the NA and in the HUT.

Figure 1 compares the evolution of total consumption expenditures per adult in Sweden in the HUT, the NA and our registry measure over the period 2000 to 2007.[19] For NA, we display two series: the official series, and one in which investments made by households (refurbishing of real estate, etc.) are added back to NA consumption. The Figure shows that the NA and the registry measures compare very well over the period, especially when household investments are added to the NA series of consumption expenditures. The HUT measure, which does not include imputed rents, is systematically lower.

While the NA measure of consumption relies on many different sources and possibly offers the most comprehensive and conceptually sound measure of aggregate consumption, it does not offer any guidance in terms of the distribution of consumption and its evolution. Economists therefore have to mostly rely on consumption survey measures to investigate consumption distribution and consumption dynamics. Figure 2 compares the density of the distribution of individual consumption expenditures in the HUT and with our registry-measure for all years 2003-2007 for all adults between 25 and 55. For both the HUT and the registry, individual consumption is computed by dividing total household consumption by the number of adults in the household. The vertical lines denote average consumption expenditures for each measure. A well-known issue with the registry-based measure (e.g., Eika et al. [2017]) is that for a substantial share of households the imputed consumption expenditures are negative.[20] This can be due misattributing wealth increases to savings out of income when they are due to price increases or due to inheritances or gifts.[21] The density of the registry measure, which includes imputed rents, is also slightly skewed to the right compared to the HUT distribution, and has a fatter right tail. We also report the density of a registry measure excluding imputed rents, which lines up very well with the HUT distribution. The averages of the two distributions are very close.

Furthermore, Sweden offers the possibility to match HUT survey to the registry data for years 2003 to 2009. This gives the opportunity to directly compare consumption measures in the HUT and in the registry measure for the same households. In Figure 3, we plot log household consumption in HUT vs log expenditures from the registry measure. We restrict the sample to all household from the HUT surveys from 2003 to 2007 that are surveyed in the month of December to make sure that the annual expenditure measure from the HUT (which includes recall items such as durables

[18] This type of consumption is recorded in the residual measure as long as the resources that are used to pay for it are accounted for.

[19] Adults are defined as individuals aged 16 or more

[20] The consumption measure equals zero for about 30.000 households per year. These are virtually all single-person households of which and most of them are immigrants (about 85%) who recently moved to Sweden.

[21] Note that the registry data can be merged with a register containing all inheritances and gifts, which we will pursue in future work.

from the last 12 months) corresponds to the annual measure in the registry data which spans all expenditures from December of year $n-1$ to December of year n. When the household composition in the HUT and registry data differs, we reconstruct households for the registry measure using the exact same individuals that are listed as members of the household in the HUT. To make the expenditure measure comparable, we exclude imputed rents from the registry measure. Results show that, within households, both measures of consumption expenditures compare quite well, with most observations scattered closely along the 45 degree line.

4.2 Consumption At the Top End of the Income Distribution

While the distribution of the HUT and the registry measures seem to compare reasonably well at first glance, we show that the granularity of the registry measure can overcome some critical limitations of the HUT measures when trying to identify certain fundamental aspects of consumption inequality and its dynamics.

A particularly important limitation of consumption surveys is that they have limited ability to identify consumption inequality at the top end of the distribution. To exemplify this issue, Figure 4 compares the income and consumption shares by deciles of the earnings distribution. We focus on the population of individuals between 25 and 55 years old and rank individuals by their gross income (including capital income) in the HUT and in the registry data respectively. We then plot earnings shares (of total earnings) and individual consumption shares (of total consumption) of each gross income decile. Not surprisingly, consumption is more equally distributed than gross earnings, in large part due to the tax and transfer system. The striking fact is that, although earnings and consumption shares from the HUT and the registry compare quite well for much of the distribution, they differ significantly at the top end of the distribution. The HUT underestimates the consumption share of the top decile by about a fourth.

There are two reasons for this. First, as is well known from research on top income inequality, the Pareto structure of the income (and consumption) distribution at the top, means that survey sampling will tend to miss out much of what is happening in the fat tails of the distribution, irrespective of the quality of the income or consumption measure. This fact can be seen by comparing the earnings share for both the HUT sample and the registry sample, where we used the same registry-measure of earnings in both samples. The large discrepancy between the top decile earnings share in the HUT and in the registry simply stems from the fact that too few very high earnings individuals are sampled in the HUT to detect the amount of inequality at the top end of the income distribution.

The second reason is that, in consumption surveys, mismeasurement of consumption (and income) seems to be more severe at top end of income distribution. As noted by Bee et al. [2013], in the U.S. Consumption Expenditure Survey (CEX), "there is strong evidence of under-representation at the top of the income distribution and under-reporting of income and expenditures at the top." Among the many reasons for such mismeasurement at the top, it is often argued that consumption surveys have more difficulty getting top income individuals to participate in the survey, not drop

out, fill properly their diaries and track their expenditures, especially irregular or infrequent ones.[22] In our context, the issue of mismeasurement of consumption at the top is quite visible on Figure 3. While the HUT and registry measures of consumption compare very well within household for most observations, the graph shows that at the top of the distribution of household consumption, most observations are below the 45 degree line. In other words, for these households at the top of the distribution, consumption expenditures are systematically underestimated in the survey compared to the registry data.

A critical consequence of the poorer measurement of income and expenditures in consumption surveys is that survey data will typically miss the dynamics of consumption inequality at the top, and in particular the relationship between income inequality dynamics and consumption inequality dynamics. The question of knowing how much the rise in income inequality at the top of the income distribution in the past 40 years has been accompanied by a similar surge in consumption inequality has been at the centre of a lot of recent debates. Krueger and Perri [2006], using consumption survey data, argued that inequality in consumption has not increased much. Aguiar and Bils [2015] to the contrary, argue that the lack of increase in consumption inequality in survey data is due to measurement issues. Using various ways to try to correct for these biases, they suggest that consumption inequality has increased almost as much as income inequality.

The main object of interest in this heated debate is the elasticity of consumption with respect to income. Krueger and Perri [2006] argue that this elasticity is low because inequality is driven mostly by rise in the dispersion of the transitory component of the income process. Comparing our registry measure to the HUT survey measure enables to shed new light on this debate. In Figure 5, we compare the elasticity of log consumption to log income, by income fractiles, measured in the registry data versus the HUT. For the registry data in panel A, we rank individuals by percentile of distribution of disposable income in 2003. Then, for each percentile, we compute the log change real disposable income between 2003 and 2007 (as we only have the HUT since 2003) and plot it against the log change in real consumption expenditures over the same period.[23] Each dot on the graph is labelled with the percentile of the income distribution to which it corresponds. We do the same exercise for the HUT data in panel B, using now the survey measure of consumption and the survey measure of disposble income. Panel A of Figure 5 shows clearly a strong and positive relationship between log changes in consumption expenditures and log changes in income across income percentiles. This relationship is particularly significant for higher fractiles in the income distribution, while the relationship between log changes in income and consumption is much flatter at lower level of the income distribution. That is, not only are the highest log income increases observed at the top of the income distribution, but these also correspond to relatively higher increases in log consumption, translating into higher consumption inequality at the top.[24] In panel

[22] On the many issues of measuring consumption inequality in surveys see for instance Attanasio and Pistaferri [2016]

[23] To compute log changes in real earnings and consumption, we express everything in terms of 2003 SEK using CPI

[24] Note that if the consumption elasticity were constant (and no other confounding factors would affect income and consumption), then all log consumption changes would be on a linear line through the origin with slope equal to the

B, however, the graph is very noisy, and does not seem to exhibit any clear positive elasticity of consumption with respect to income. This evidence suggests that the low elasticity of consumption with respect to income found in survey data (e.g. Krueger and Perri [2006]) is not structural, but a consequence of the poor measurement of income and consumption in survey data, especially at the top end of the income distribution. As a consequence, survey data will tend to underestimate the rise in consumption inequality following the surge in income inequality of the last 40 years. More generally, this suggests that there are critical limitations to the estimation of consumption responses to income and wealth variations in the survey data, and that these limitations may be overcome by the use of registry measures of consumption.

5 Consumption Responses to Income Shocks

We demonstrate the research value of using registry-based measures, by applying our measure to the study of consumption responses to shocks. An extensive literature has studied, both theoretically and empirically, the responses of household consumption expenditures to anticipated or unanticipated changes in their available resources (see for instance Jappelli and Pistaferri [2010] for a recent review). The empirical literature has relied so far mostly on two types of data: survey data and proprietary data (Nielsen scanner data, credit card expenditures data, bank data, etc.). Registry-based measures offer four main advantages over existing data for the purpose of analysing consumption responses to shocks: comprehensiveness, universal coverage, panel dimension, and efficiency.

First, registry-based measures are comprehensive in the sense that they capture all household consumption expenditures. To the contrary, survey measures sometimes capture only part of total household consumption expenditures (e.g. food expenditures in the PSID) while proprietary data are often even more limited (credit card expenditures, checkout scanner data, etc.). The comprehensiveness of the registry-based measure also offers the possibility to analyze the anatomy of consumption responses by decomposing expenditure changes into various components (income, transfers, debt, consumption out of assets, durables, etc.). Second, registry-based measures have a universal coverage. This obviously offers statistical power, compared to the small samples from surveys. More importantly, when focusing on subsamples of individuals experiencing specific shocks, there is no reason for survey samples or proprietary data samples to be representative of the full population of individuals experiencing these shocks. This is particularly true for bank data or credit card users data. Third, the panel dimension of registry-based measures offers a critical advantage over cross-sectional survey data for the study of consumption responses. With the latter, identification relies on strong pseudo-panel assumptions regarding cross-sectional heterogeneity, while panel data enables to control for individual fixed-effects. Finally, registry-based measures offers efficiency over survey data for it is relies on third-party reported, registry-information, which is not subject to various biases (recall errors, attrition and non-response, etc.) found in survey data,

consumption elasticity.

and that might be correlated with time. It is important to stress that registry-based measures are obviously also subject to measurement error. Yet, in a panel specification with individual fixed effect, it is enough that within individual variation in measurement error is uncorrelated with the timing of the shock to be able to identify the dynamic effect of the shock on consumption. This is the assumption that we maintain throughout this section.

We focus on three particular types of events: health shocks, retirement, and parenthood (i.e. the arrival of the first child). These events offer well-identified sources of temporary or permanent variations in earnings. Interestingly, some of these events are (at least partially) anticipated (e.g. retirement, parenthood) while others may be totally unanticipated (health shocks). For each of these events, a well-defined sub-literature exists, looking at their impact on household labor supply and consumption using survey data, often reaching ambiguous conclusions. This motivates our revisiting the impact of these shocks on household expenditures using our registry measure.

5.1 Methodology

For each event type (health shock, retirement, parenthood) we are interested in identifying the dynamics of some outcome Y_{it} around the time of the event. Households are indexed by i and $t = 1, .., T$ denote the calendar year of observation. For each household, the event happens at some time E_i. We restrict our attention to the following models:

$$Y_{it} = \alpha_i + \nu_t + \sum_{j=-N_0}^{N_1} \beta_j \cdot \mathbb{1}[J_{it} = j] + \varepsilon_{it} \tag{2}$$

where $[-N_0; N_1]$ is the window of dynamics effects around the event, α_i is a household fixed-effect and ν_t is a time effect. $J_{it} = t - E_{it}$ denotes event time, that is the time in year relative to the occurrence of the event. We are interested in retrieving the full path of dynamic effects $\{\beta_j\}_{j=-N_0}^{N_1}$, which means we are interested in pre-event as much as in post-event effects.[25]

If we restrict our attention to a sample of households who all receive treatment at some point in time E_i, an OLS specification with two-way fixed effects (household and time fixed effects) can only identify the true dynamic effects $\{\beta_j\}_{j=-N_0}^{N_1}$ up to a linear trend (see Borusyak and Jaravel [2016] for a recent discussion). More fundamentally, when introducing household fixed effects, and focusing only on treated households, it is impossible to jointly and non-parametrically identify the full sequence of dynamic effects $\{\beta_j\}_{j=-N_0}^{N_1}$ and of any other time-related variable, such as calendar time, or age. The reason is that, within household, the passage of event time and of any other time related variable is collinear![26]

We deal with this identification issue in two different ways: first by using control groups, second

[25]The model in equation (2) abstracts from heterogeneity in the dynamic treatment effects but such effects are certainly important, and the framework can be extended to accommodate such heterogeneity.

[26]In equation (2) we have focused on calendar time fixed effects, but in some instances, one might be more interested in controlling for other time related variables such as age, tenure, or experience fixed effects. This is easily done by denoting age (or tenure or experience) by $t = 1, .., T$, and age (or tenure or experience) fixed effects by ν_t. The same collinearity problem obviously arises in such settings.

by getting rid of household fixed effects.

Control groups For certain events, such as health shocks, it is possible to construct a reasonable control group of individuals who never experience the event. The introduction of a control group that never experiences treatment allows to identify the time effects ν_t independently of the dynamic treatment effect of the event $\{\beta_j\}_{j=-N_0}^{N_1}$, because control individuals never experience treatment. Identification relies on the assumption that the control individuals offer a good counterfactual of the time trends of the treated individuals absent dynamic treatment effects. In settings where the timing of the event is random and anticipatory effects can reasonably assumed to be zero (i.e. all coefficients $\{\beta_j\}_{j=-N_0}^{0}$ are zero), this identification assumption can partially be tested by looking at parallel pre-trends between the control and treatment groups prior to the event. But it should be noted that in fully flexible settings where dynamic treatment effects can be non zero prior to the event, the identification assumption of parallel trends is fundamentally untestable.

We create control groups using nearest-neighbor matching based on pre-event characteristics. One of the main challenges is to select observable characteristics that can reasonably be assumed to be unaffected by the event prior to its occurrence. This is more easily done in settings where the timing of the event is random and fully unanticipated. We adopt the following matching strategy: for each calendar year t, we take all individuals who receive the event in that particular year ($E_{it} = t$), and find a nearest neighbor from the sample of all individuals who never receive treatment. Individuals are matched exactly on age, gender, region of residence in $t-1$ (21 cells), level of education in $t-1$ (10 cells) and family structure in $t-1$ (12 cells), and by propensity score on their number of dependent children in $t-1$, 12 industry dummies in $t-1$ and their earnings in $t-1$, $t-2$ and $t-3$. Note that by matching exactly on age, we make sure that age effects, which could not be identified in specification (2), are not confounding our estimated dynamic effects.

For some events though, almost everybody gets treated (e.g. retirement) and/or the never treated do not constitute a credible counterfactual for the relevant time trends absent treatment (e.g. parenthood, where individuals who never have children are somewhat different from individuals who do end up becoming parents). For these types of events, one solution is to focus the control group on individuals who do not receive treatment within an event time window of observation $[-R_0; R_1]$ rather than individuals who do not receive treatment at all. Ideally, the time window $[-R_0; R_1]$ should be much larger than the time window of interest for the dynamic effects $[-N_0; N_1]$.[27] In most cases, this is not doable, and this solution comes at the cost of stringent assumptions on the dynamic effects $\{\beta_j\}_{j=-N_0}^{N_1}$.

Dropping individual fixed effects For the events where no natural control group is available, an alternative is to abandon individual fixed effects and control directly for time related variables

[27] If we assume all dynamic effects to be zero outside the range $[-N_0; N_1]$, then we need for identification that $R_0 > 2N_0 + N_1$ and $R_1 > 2N_1 + N_0$ so that control units are not contaminated by their own potential treatment effects.

instead.[28] We therefore also report estimates of pure event-study specification without individual fixed effects, but with fully flexible age effects, using the following specification similar to Kleven et al. [2015]:

$$Y_{it} = \sum_a \gamma_a \cdot \mathbb{1}[A_{it} = a] + \nu_t + \sum_{j=-N_0}^{N_1} \beta_j \cdot \mathbb{1}[J_{it} = j] + \varepsilon_{it} \tag{3}$$

where $\{\gamma_a\}$ capture a full set of age fixed effects. In the absence of individual fixed effects, identification is not brought by within individual variation but by variation in event time across individuals conditional on calendar time and age. Identification therefore now relies on the assumption that the timing of the event is random across individuals and therefore uncorrelated with individual fixed effects. While the assumption that there is no selection on receiving the event at age a or at age a' might be valid for a close to a', it might be more problematic when $a' - a$ is large. For instance, there might be some randomness in having a first child at age 25 versus age 27, yet, it is likely that women who have their first kid at 36 are quite different from women who have their first kid at 20. As a consequence, one might also need to apply sample restrictions and run specification (3) on a restricted time window with a limited set of dynamic effects $\{\beta_j\}$.

Household composition For household level outcomes around the health and retirement shocks, we fix the composition of the household as of the last year prior to event, meaning that household growth and re-composition is not accounted for, neither before nor after the event year. Regarding the parenthood event, we hold household composition constant as of the event year: Note also that we do not adjust household consumption for household equivalence scales. Because household structure is fixed, household size differences are subsumed in the fixed effects in specification (2).

Event studies using consumption surveys We also provide comparisons of registry-based estimates with survey data estimates. As is often the case with consumption expenditure surveys, the HUT is a cross-section and the lack of a panel dimension creates issues when trying to identify consumption dynamics around specific events. Yet, HUT surveys can be matched to registry data based on personal identifiers, so that we know for each individual surveyed in the HUT the timing of the event E_i (if it occurs) and can therefore compute their event time J_{it} at the moment their consumption is observed. With the HUT data, we follow a control group approach. Using nearest-neighbor matching based on pre-event characteristics, we create for each event a control group of individuals who are not subject to the event within an event time window of observation $[-R_0; R_1]$.[29] For individuals in the control group, we allocate a placebo event time E_i corresponding to the event time experienced by their nearest neighbor in the treated group. We then estimate the following

[28] An interesting third alternative proposed by Borusyak and Jaravel [2016] is to use individual random effects instead of individual fixed effects.

[29] For health shocks, we take individuals who are never subject to health shock in the entire period of our data. For retirement, we take individuals who do not experience the event within 3 years of the event experienced by the nearest neighbor in the treated group. For parenthood, control individuals do not experience the shock before 4 years.

specifications to retrieve the dynamic effects:

$$Y_{it} = \sum_{j=-N_0}^{N_1} \beta_j \cdot \mathbb{1}[J_{it} = j] \cdot \mathbb{1}[T_i = 1] + \gamma \cdot \mathbb{1}[T_i = 1] + \sum_{j=-N_0}^{N_1} \kappa_j \cdot \mathbb{1}[J_{it} = j] + \nu_t + X'_{it}\alpha + \varepsilon_{it} \quad (4)$$

where T is an indicator for being in the group of treated individuals experiencing the event.

5.2 Health Shocks

A long literature has investigated the effect of health shocks on labour supply and consumption. Identification of health shocks relies on self-reported measures of health (e.g. Meyer and Mok [2009]), objective measures of health, or both (e.g. Blundell et al. [2017]). A related literature has also investigated how social insurance programs, such as workers' compensation or disability insurance (DI), affect the earnings and consumption responses to health shocks, using rejected applicants to insurance programs for identification (Bound [1989]). Almost all papers use survey measures to capture consumption responses, at the exception of Kostøl and Mogstad [2015] who use a registry-based measure of consumption to analyse the effect of DI receipt on household consumption. We revisit this literature using the specificity of the sickness benefit system in Sweden to identify health shocks.

Identifying Health Shocks In Sweden, employees who become sick are eligible to a paid sick leave. To this effect, they need to obtain a certificate from a physician providing evidence of their incapacitation. For the first 14 days of incapacitation, individuals get a "sickness wage" of 80% of their actual wage, paid directly by the employer. When incapacitation lasts longer than 14 days, individuals start receiving benefits from the Social Insurance Agency (SIA). These sickness benefits amount to 80% of actual wage, up to a cap.[30] Sickness benefit receipts therefore identify the occurrence of a health shock large enough to involve a work incapacitation larger than 14 days. We define an health shock event as the first occurrence of sickness benefits receipt by an individual over the period 1997-2011.[31] We restrict the sample to individuals aged 25 to 55 at the time of the event. Appendix Table 2 displays descriptive statistics for our sample of analysis, breaking it down between treated individuals and the control group, using nearest-neighbor matching, of individuals who are never treated. Individuals experiencing a health shock are around 38 years old on average, and have little net wealth. The median individual has only 77k SEK in total net wealth ($\approx$ 25% of

[30] Individuals with partial incapacitation after 14 days can work part time and will receive reduced sickness benefits.

[31] Before 2008, there was no clear rule regarding the time individuals could receive sickness benefits, or on the issue of knowing when they should be transferred to the disability insurance system instead. In 2008 the sickness insurance (SI) system and its transition of individuals from SI to DI was made more rigorous and strict with the chain of rehabitalization. After 90 days of SI benefit take-up the individual must be deemed unfit to carry out any work at the current employer in order to keep SI benefits. After 180 days the individual must be deemed unfit to carry out any work on the labor market as a whole to keep SI benefits. Individuals who are assessed to be able to work are transferred to the public employment service. Individuals with a permanent inability to work are given DI. The assessment of individuals' ability to work is done by a specialized officer of the SIA based on a physician's statement.

total annual household earnings), and zero liquid assets (bank holdings) prior to the onset of the health shock. A long tail of individuals have already negative net wealth when hit by the shock. This obviously limits consumption smoothing opportunities.[32]

Results Figure 6 reports event study estimates of the earnings, benefits, and consumption responses to a health shock. We use specification (2) on the sample of treated individuals and a control group of individuals obtained from nearest-neighbor matching on pre-event characteristics.[33] Panel A shows that health shocks are on average associated with a permanent decline in earnings equivalent to $\approx 10\%$ of household consumption levels in year -1. Interestingly, there seems to be very little anticipatory effects on earnings prior to the health shock. Panel B shows that earnings from other members of the households does not provide much opportunities for consumption smoothing. This could be due to correlated shocks to labor supply: other members of the household may need more time to take care of the individual hit by the health shock. Alternatively, this could be due to limited added worker effects in a context where health shocks are covered by relatively high sickness benefits. Panel C shows that sickness benefits cover indeed a large fraction ($\approx 80\%$) of the drop in earnings in the immediate aftermath of the health shock, but slowly decline over time after the initial shock. A small fraction of individuals ends up transitioning into disability insurance. This is visible in panel D, where we see a small increase in the amount of disability benefits received two to five years after the initial health shock. Finally, panel E provides evidence of a significant and permanent decline of consumption expenditures, of around 5% following a health shock. Panel A suggests that health shocks are associated with a quite permanent decrease in earnings. In that sense, the evidence from panel E of a permanent decline in consumption following a health shock is not so surprising in light of a standard dynamic model of consumption.

In Figure 11 panel A, we provide a full decomposition of the consumption response into six main components: own earnings, earnings of other members of the household, sickness and disability benefits, consumption out of debt, consumption out of assets, and other transfer and taxes. The graph suggests that most of the consumption smoothing is taken care of by the social safety net, through an increase in benefits and transfers and the decline in taxes, while assets, debt, and household labor supply contribute very little to consumption smoothing on average.

Finally, we compare in Figure 12 panel A the estimates of consumption responses to health shocks obtained from our registry-based measure of consumption with the estimates from the consumption surveys. To make the consumption measures comparable, we take out imputed rents from our registry-based measure of consumption. The Figure clearly highlights the superiority of the registry measure over the survey measure. Using the latter measure, we find a drop of around 5% in consumption following the event, but estimates are very imprecise and noisy. The registry measures estimates deliver a similar drop, but with very tight confidence intervals and a much

[32] Treatment and control groups are well balanced in terms of demographics although treated individuals have larger debt levels and lower level of net wealth on average.

[33] Appendix Figure A.1 displays the evolution of the average level of earnings, benefits and consumption for both the treated and control group as a function of event time.

less noisy dynamic pattern. We report on the graph of test of equality of the estimated average drop in consumption over the four years following the health shock, estimated with the registry measure (β^R) and with the HUT survey (β^S). Interestingly, the t-stat suggests that the estimates are consistent across both sources. But the registry-based measure delivers a much higher degree of precision.

5.3 Retirement

Consumption patterns around retirement have been investigated in a vast stream of papers starting from Hamermesh [1984], sparking intense debates over the alleged existence of a "retirement consumption puzzle". For example, Banks et al. [1998], using cross sectional data from the British consumer expenditure survey, documented the presence of a drop in total non-durable consumption expenditures around retirement. Several papers using the panel structure of the PSID in the US, have documented similar drops in *food* expenditures around retirement. Aguiar and Hurst [2005], however, show using cross-sectional data that while food expenditures decline at retirement, food intake does not, consistent with a substitution away from food expenditures and towards home-production at retirement. Due to data availability constraints, there is still remarkably little evidence on the evolution of broad measures of total expenditures around retirement using high-quality panel data.[34] We contribute to this literature by providing compelling evidence of the evolution of total household expenditures around retirement using our registry measure on a very large panel of Swedish retirees. Thanks to the universal coverage of our registry measure, it delivers much greater precision in identifying expenditure drops than existing survey measures. We then decompose the response of total expenditures to the substantial drop in earnings around retirement and document the respective role of pension, transfers and asset consumption in smoothing consumption at retirement.

Identifying Retirement Our definition of retirement in the Swedish context follows closely previous work (see Palme and Svensson [2004]). In particular, restricting the sample to Swedes aged 50 or more, we identify retirement as (i) the first year in which the individuals' labour income is inferior to the base amount (BA), (ii) immediately followed by another year of income less than BA. The base amount is a yearly income threshold, annually set by the Swedish parliament, and used to determine if an individual is employed in a given year. If the sum of an individual's income derived from employment and self-employment exceeds the BA, the person is considered employed.[35] Using this definition of retirement year, we find that approximately 1.5% of the Swedish adult population enters retirement every year. Appendix Table 3 displays descriptive statistics for our sample of analysis. Individuals around retirement are obviously much older than in the health shocks sample. While their annual consumption expenditures prior to retirement are quite similar to those in the

[34] Aguila et al. [2011], using the CEX in the US, provide some evidence of a drop in total expenditures at retirement, but the panel dimension is very short (1 year) and the estimates quite imprecise.

[35] The BA closely follows inflation and pensions are indexed on it. In 2007, the last year in our consumption panel, the BA was 40,300 SEK (6,000 USD).

health shocks sample, they have a much larger level of net wealth ($\approx$ 4 times larger), much more liquid assets, and much less debt. They are also much more likely to be homeowners.

Results The evolution of earnings, pension and consumption around retirement is reported in Figure 7. Because retirement is an event that affects everyone, and because the distribution of the timing of retirement is very compressed around ages 60 to 65, creating a control group is not straightforward. We therefore adopt as a baseline the event study specification (3). Figure 7 displays these event study estimates and all coefficients are expressed as a fraction of total household consumption level in the year prior to retirement.[36] Panel A shows, unsurprisingly, that retirement is associated with a massive drop in individual earnings, equivalent to $\approx$ 50% of pre-retirement household consumption levels. There seem to be anticipatory effects, with earnings dropping by about 10% of pre-retirement household consumption levels in the five years leading to retirement. Such anticipatory effects might be driven by gradual labor supply reduction prior to full retirement or by productivity and health shocks that affect the timing of retirement (cf. Banks et al. [1998], or Smith [2006]). Panel B shows that earnings by other members of the household also decline around retirement. In Appendix Figure A.2, we provide evidence that other household members are also more likely to retire and to start drawing pensions, pointing to the existence of some coordination in the timing of retirement between spouses. Panel C shows that individuals start drawing pension benefits at the time of retirement, which provide significant means to smooth consumption expenditures. Yet, panel D provides compelling evidence that total expenditures experience a significant decline of around 7 to 10% at retirement.

How do these results compare to existing studies using panel data? Aguila et al. [2011] estimate food spending to decline by about 6% at retirement and do not find any evidence of a consumption drop in non-durable spending. But their panel dimension is very short so that the effects are only estimated for the first couple of quarters after retirement, and their estimates are quite imprecise. Hurd and Rohwedder [2008] find that spending declines at a small rate, with an average pre-post retirement difference of 6% for total spending. Yet, given large standard errors, they cannot reject zero effect. Our results of a 7% decline are in line with these estimates, but are much more precisely estimated.

In Figure 11, we provide a full decomposition of the consumption response into six main components: own earnings, earnings of other members of the household, pension benefits, consumption out of debt, consumption out of assets, and other transfer and taxes. The social insurance and transfer system does contribute very significantly to consumption smoothing around retirement. But compared to the health shock event, consumption out of assets does now offer a significant

[36] We also provide for comparison results from a strategy where we create a control group of individuals not experiencing retirement in a 4 year window around the time of the event of the treated individual. Results are shown in Appendix Figure A.2 and are very consistent with the event study results using flexible age controls. The short term effects of retirement are relatively well identified, and consistent across both strategies. Nevertheless, it should be noted that identifying longer term dynamic effects of retirement is complicated by the fact that it is hard to observe individuals that are old, not yet retired and who would provide a proper counterfactual for individuals who have retired much earlier.

additional source of consumption smoothing. Individuals do draw down their assets when they retire (and in particular their private retirement pension accounts) so that consumption out of assets offers a 7 to 8% increase in consumption over the first five years after retirement.

To further exemplify the limitations of survey data for estimating consumption dynamics around retirement, we compare our results using registry data to results from the HUT survey of consumption expenditures in Figure 12 panel B. The estimated dynamics of consumption expenditures is quite similar in the HUT survey and the registry measure. We report on the graph of test of equality of the estimated average drop in consumption over the four years following retirement, estimated with the registry measure (β^R) and with the HUT survey (β^S), which suggests that the estimates are consistent across both sources. Strikingly though, the registry-based measure delivers a much higher degree of precision compared to the HUT where the standard errors on the estimated coefficients are close to 10 times larger.

5.4 Parenthood

A recent literature has documented the stark impact of the arrival of the first child on earnings, labor supply and wage rates of women within the household (e.g. Kleven et al. [2015], Angelov et al. [2016]). Like retirement, parenthood tends not to be a fully unanticipated shock. Households may therefore find more easily means to smooth the earnings shock represented by parenthood. Interestingly, the arrival of children is also complementary with a certain number of child care expenditures and child care related leisure, which may impact the volume and structure of household expenditures around the time of parenthood. We contribute by studying how consumption expenditures and the structure of expenditures dynamically evolve around the arrival of children in the household.

Definition of the Event We focus on the arrival of the first child. For all women above 16 years of age, we define the time of the arrival of the first child as the first year in which she is ever reported as having a dependent biological or adopted child under the age of three.

We study consumption patterns at the household level, fixing the household composition as observed at time E_i. That is, we aggregate for all event years the consumption, debt, income and transfers of all members who are in the treated woman's household in the year of the event.[37]

Results To identify the dynamic patterns of earnings and consumption around parenthood, we follow specification (2) using a control group to fully identify time effects. While women who never have children would offer a natural control group, they appear to be a quite selected sample based on observable characteristics. For each treated woman at time E_i, we instead chose a nearest-neighbor

[37] A household can be created in two ways in the LISA register: either when a couple is married or when it gets its first child. Bearing this statistical definition in mind, household structure needs to be held fixed at event time to allow for meaningful event studies at the household level. Household composition is held fixed at year $E_i - 1$ for other shocks, but doing so for the first child event would imply that households with unmarried mothers would be constituted of the mother only. Household-level consumption would not include the father's consumption as he would be considered a household member only in year E_i.

from a control group of women who have not yet experienced parenthood at time E_i and who will not be experiencing parenthood in the following four years. Conditional on age and time, we are therefore comparing treated women to similar women who happen to have children later in life.[38] The evolution of earnings, transfers and consumption around parenthood is reported in Figure 8. In line with evidence from Denmark in Kleven et al. [2015], Panel A shows that parenthood is associated with a large drop in earnings for women right at the arrival of the first child, equivalent to $\approx 25\%$ of pre-child household consumption levels. But interestingly, panel B shows that also men experience a (slight) drop in earnings after the arrival of children, which is not the case in Denmark. We hypothesize that this may be due to the larger fraction of fathers taking parental leave in Sweden compared to Denmark. Panel C shows that the system of taxes and transfers does provide significant cushion against the overall drop in earnings at the arrival of the first child, resulting in a more moderate drop in consumption after the arrival, as shown in Panel D.

Consistent with evidence from Denmark, parenthood is not associated with large pre-trends in earnings: in the years leading to the event, earnings of men and women increase at just a slightly faster rate than comparable men and women of similar age and characteristics, but then drop significantly after the child arrival and remain significantly lower for the mother in the years after. Panel D shows that consumption expenditures increase by about 10% in the five years prior to the first child birth. Consumption expenditures then drop in the year of the child arrival, but do not fall below the level of consumption five years before and quickly ick up again in the years after.

Durables, and housing in particular, contribute to the increase in consumption expenditures leading up to the arrival of the first child. In Figure 9 panel A, we show that prior to having a first child, individuals are much more likely to move houses or apartments. They are also much more likely to become homeowners around the time of parenthood, as shown in panel B of Figure 9. This translates into a large and significant increase in consumption out of debt (mostly mortgages) as shown in panel C. This is partly offset by the decrease in consumption flow due to the increase in real estate assets, which is shown net of the increasing imputed rents in panel D. In Appendix Figure A.4, we show that for individuals who remain renters (i.e., who do not have any owner-occupied real estate assets) throughout the time window around parenthood, consumption expenditures exhibit a similar but much milder dynamic pattern around the arrival of the first child, and that in particular, consumption out of debt is much smaller around the time of the event. Housing is not the only durable expenditure that households adjust around the arrival of the first child. In Figure 10, using data from the car registry, we show that the probability of buying a new car increases significantly in the last two years before the arrival of the first child.

Like for the previous events, we provide a full decomposition of the consumption response after the arrival of the first child (see Panel C of Figure 11), further illustrating the substantial role played by the portfolio changes. We separate out the consumption out of debt, the consumption of real estate (including imputed rents) and the consumption out of other (mainly financial) assets.

[38] We provide in Appendix Figure A.3 the evolution of the average level of all our outcomes of interest for the treated and control group.

Interestingly, also these other assets are used to significantly increase the consumption expenditures around child birth.

Finally, in Figure 12 panel C, we compare the estimated dynamic patterns of expenditures from the registry measure with evidence from the HUT survey measure. The survey estimates are again very imprecise, but the overall pattern matches quite well the dynamic effects estimated in the registry data.

6 Conclusion

The Scandinavian countries received a lot of attention in recent economic research, not just because of their interesting history, culture and policies, but also because of the exceptionally detailed and comprehensive administrative data available for research. Individuals can be linked to other members of the household and across all registries through the same personal identification number. In addition, Scandinavian countries levy or levied a tax on wealth, which makes that excellent data on asset holdings tends to be available as well. So this offers the best context to construct a residual measure of expenditures from information on income and wealth. The data requirements for costructing valuable registry-based measure are still quite formidable. So, even within Scandinavia, important differences exist across countries in terms of data coverage that have important implications for the computation of residual measure of consumption.

This paper has focused on the Swedish context, providing excellent data on financial and real assets to credibly measure their contribution to consumption flows, but only over a small time period, for example compared to Norway. Having laid out the potential challenges and remedies, we have attempted to further improve the registry-based measurement of consumption and illustrated how powerful this consumption measure can be to study consumption inequality and consumption smoothing. We hope that future research will build on our efforts to extend the measurement to other time periods and countries, when the challenges may be even more binding. The ultimate goal is not just to revisit questions previously approached using consumption surveys, but also to exploit complementarities between the two alternative measure and hopefully also answer questions that could not be addressed before due to the limitation of consumption surveys.

The role of expenditures and how they respond to changes in income and wealth is at the heart of policy design, ranging from the design of fiscal stimuli to the design of social insurance programs. Still, how expenditures translate into welfare remains an underexplored question for which the rich information from registry data can provide new complementary insights to standard surveys. An important next step for the evaluation of welfare is also to take the role of public goods seriously and account for the consumption of public goods in the measurement of consumption. Registry data on education, public tansportation, health expenditures, etc can again be particular useful here and further improve our understanding of consumption smoothing and inequality.

References

Aguiar, Mark and Erik Hurst, "Consumption versus Expenditure," *Journal of Political Economy*, 2005.

_ and Mark Bils, "Has consumption inequality mirrored income inequality?," *The American Economic Review*, 2015, *105* (9), 2725–2756.

Aguila, Emma, Orazio Attanasio, and Costas Meghir, "Changes in consumption at retirement: evidence from panel data," *Review of Economics and Statistics*, 2011, *93* (3), 1094–1099.

Angelov, Nikolay, Per Johansson, and Erica Lindahl, "Parenthood and the Gender Gap in Pay," *Journal of Labor Economics*, 2016, *34* (3), 545–579.

Attanasio, Orazio P. and Luigi Pistaferri, "Consumption inequality," *The Journal of Economic Perspectives*, 2016, *30* (2), 3–28.

Baily, Martin N., "Some Aspects of Optimal Unemployment Insurance," *Journal of Public Economics*, 1978, *10* (3), 379–402.

Banks, James, Richard Blundell, and Sarah Tanner, "Is there a retirement-savings puzzle?," *American Economic Review*, 1998, pp. 769–788.

Bee, Adam, Bruce D. Meyer, and James X. Sullivan, "The Validity of Consumption Data: Are the Consumer Expenditure Interview and Diary Surveys Informative?," in "Improving the Measurement of Consumer Expenditures" NBER Book Series Studies in Income and Wealth, University of Chicago Press, October 2013, pp. 204–240.

Blundell, Richard, Jack Britton, Monica Costa Dias, and Eric French, "The impact of health on labour supply near retirement," 2017. IFS Working Paper W17/18.

_ , Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *The American Economic Review*, 2008, *98* (5), 1887–1921.

_ , _ , and Itay Saporta-Eksten, "Consumption Inequality and Family Labor Supply," *American Economic Review*, February 2016, *106* (2), 387–435.

Borusyak, Kirill and Xavier Jaravel, "Revisiting Event Study Designs," 2016. Available at SSRN: https://ssrn.com/abstract=2826228.

Bound, John, "The health and earnings of rejected disability insurance applicants," *The American Economic Review*, 1989, *79*, 482–503.

Browning, Martin and Søren Leth-Petersen, "Imputing consumption from income and wealth information," *The Economic Journal*, 2003, *113* (488), F282–F301.

Chetty, Raj, "A General Formula for the Optimal Level of Social Insurance," *Journal of Public Economics*, 2006, *90* (10-11), 1879–1901.

Eika, Lasse, Magne Mogstad, and Ola Vestad, "What Can We Learn About Household Consumption From Information on Income and Wealth," 2017. Working Paper.

Fagereng, Andreas and Elin Halvorsen, "Imputing consumption from Norwegian income and wealth

registry data," *Journal of Economic and Social Measurement*, 2017, *42* (1), 67–100.

Hamermesh, Daniel S., "Life-cycle effects on consumption and retirement," *Journal of Labor Economics*, 1984, *2* (3), 353–370.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante, "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," *American Economic Review*, July 2014, *104* (7), 2075–2126.

Hurd, Michael D. and Susann Rohwedder, "The Retirement Consumption Puzzle: Actual Spending Change in Panel Data," Working Paper 13929, National Bureau of Economic Research April 2008.

Jappelli, Tullio and Luigi Pistaferri, "The Consumption Response to Income Changes," *Annual Review of Economics*, 2010, *2* (1), 479–506.

Kleven, Henrik J., Camille Landais, and Jacob E. Sogaard, "Children and gender inequality: Evidence from Denmark," 2015. Working Paper.

Koijen, Ralph, Stijn Van Nieuwerburgh, and Roine Vestman, "Judging the Quality of Survey Data by Comparison with Truth as Measured by Administrative Records: Evidence From Sweden," in "Improving the Measurement of Consumer Expenditures," University of Chicago Press, July 2014, pp. 308–346.

Kolsrud, Jonas, Camille Landais, Peter Nilsson, and Johannes Spinnewijn, "The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden," 2015. forthcoming, American Economic Review.

Kostøl, Andreas Ravndal and Magne Mogstad, "Earnings, Disposable Income, and Consumption of Allowed and Rejected Disability Insurance Applicants," *The American Economic Review*, 2015, *105* (5), 137–141.

Kreiner, Claus Thustrup, David Dreyer Lassen, and Søren Leth-Petersen, "Measuring the accuracy of survey responses using administrative register data: evidence from Denmark," in "Improving the measurement of consumer expenditures," University of Chicago Press, 2014, pp. 289–307.

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

Landais, Camille and Johannes Spinnewijn, "The Value of Unemployment Insurance," in progress.

Lundberg, Jacob and Daniel Waldenström, "Wealth Inequality in Sweden: What can we Learn from Capitalized Income Tax Data?," *Review of Income and Wealth*, 2017, pp. n/a–n/a.

Meyer, Bruce and Wallace K.C. Mok, "Disability, earnings, income and consumption," Working Paper 18869, National Bureau of Economic Research 2009.

Palme, Mårten and Ingemar Svensson, "Income security programs and retirement in Sweden," in "Social Security Programs and Retirement around the World: Micro-Estimation," University of Chicago Press, 2004, pp. 579–642.

Pistaferri, Luigi, "Household Consumption: Research Questions, Measurement Issues, and Data Collection Strategies," Mimeo, Stanford University 2015.

Poterba, James M., "Taxation and Housing: Old Questions, New Answers," *The American Economic Review*, 1992, *82* (2), 237–242.

Saez, Emmanuel and Gabriel Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *The Quarterly Journal of Economics*, 2016, *131* (2), 519–578.

Smith, Sarah, "The retirement-consumption puzzle and involuntary early retirement: Evidence from the British Household Panel Survey," *The Economic Journal*, 2006, *116* (510).

Figure 1: NA, HUT and Registry-Based Consumption Expenditures

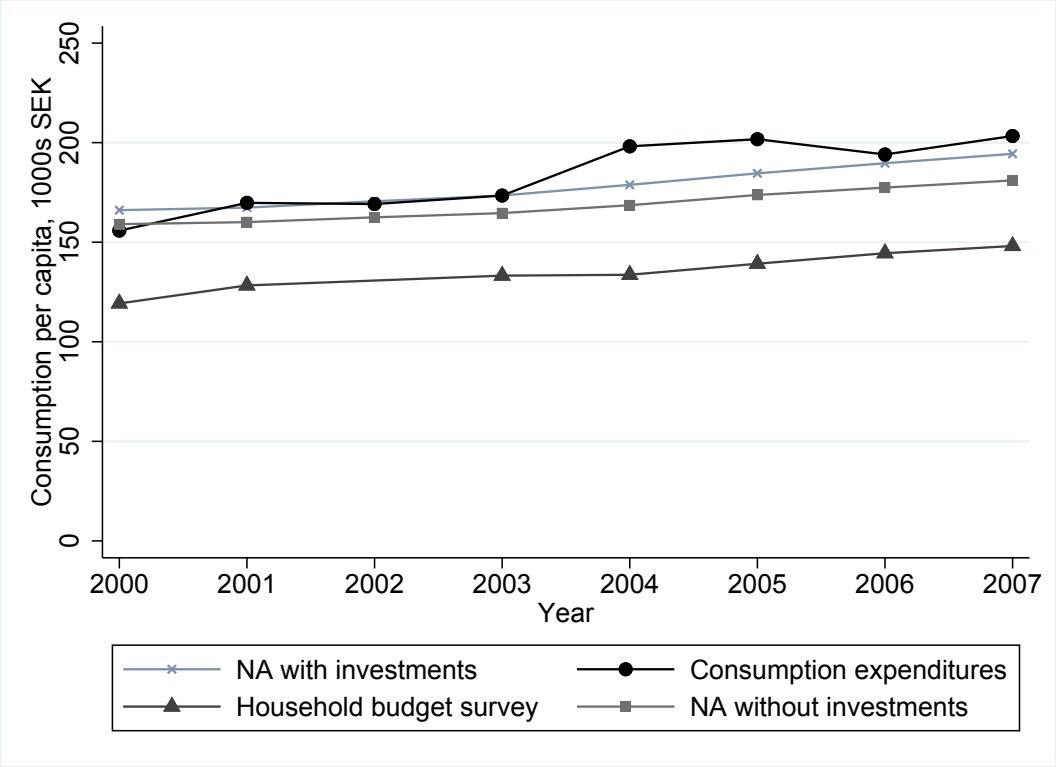

Notes: The Figure reports the evolution of consumption expenditures per adult (defined as being 16 or over) in the national accounts (NA), in the HUT consumption surveys, and in our registry-based measure. For NA, we display two series: the official series, and one in which investments made by households (refurbishing of real estate, etc.) are added back to NA consumption. The HUT survey measure of consumption does not include imputed rents while both the registry measure and the NA measure do. In Appendix Appendix C we describe in more details the sources and concepts used for measuring consumption in the NA and in the HUT.

Figure 2: Distribution of HUT and Registry-Based Consumption Expenditures: 2003-2007

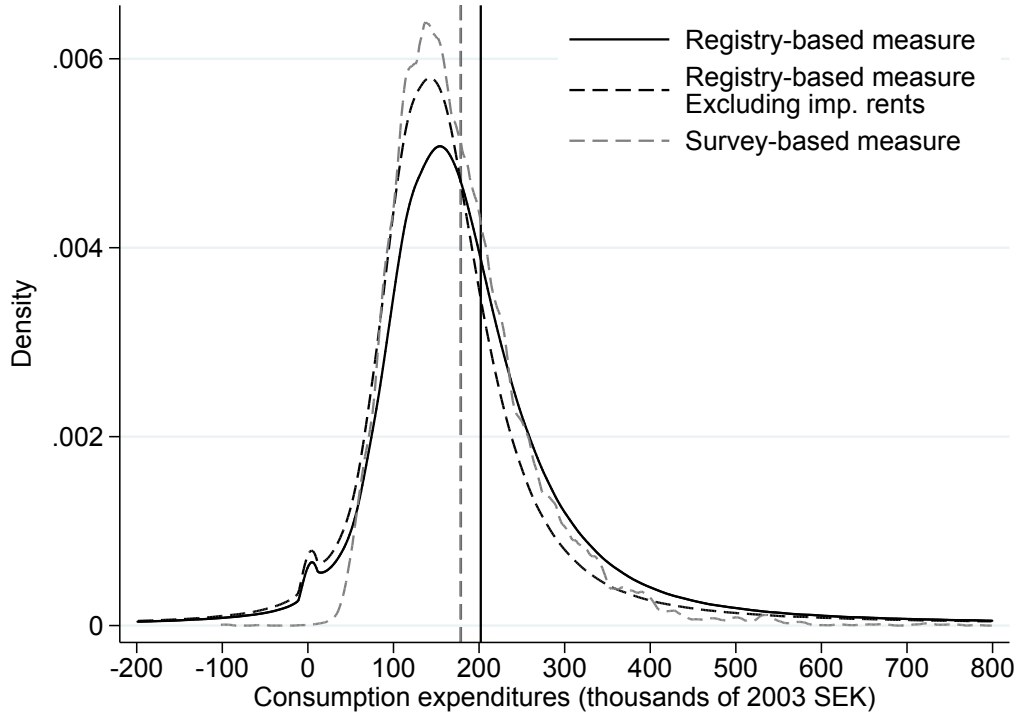

Notes: The Figure compares the distribution of consumption expenditures per adult aged 25 to 55 in the survey and in the registry data over the period 2003 to 2007. For all adults aged 25 to 55, consumption is computed as total household consumption divided by the number of adults in the household. To make expenditure measure comparable in both samples, we also show the distribution of the registry based measure when we exclude imputed rents from household expenditures. All figures are in thousands of constant 2003 SEK. The vertical lines denote the average for each distribution.

Figure 3: HUT and Registry-Based Consumption Expenditures

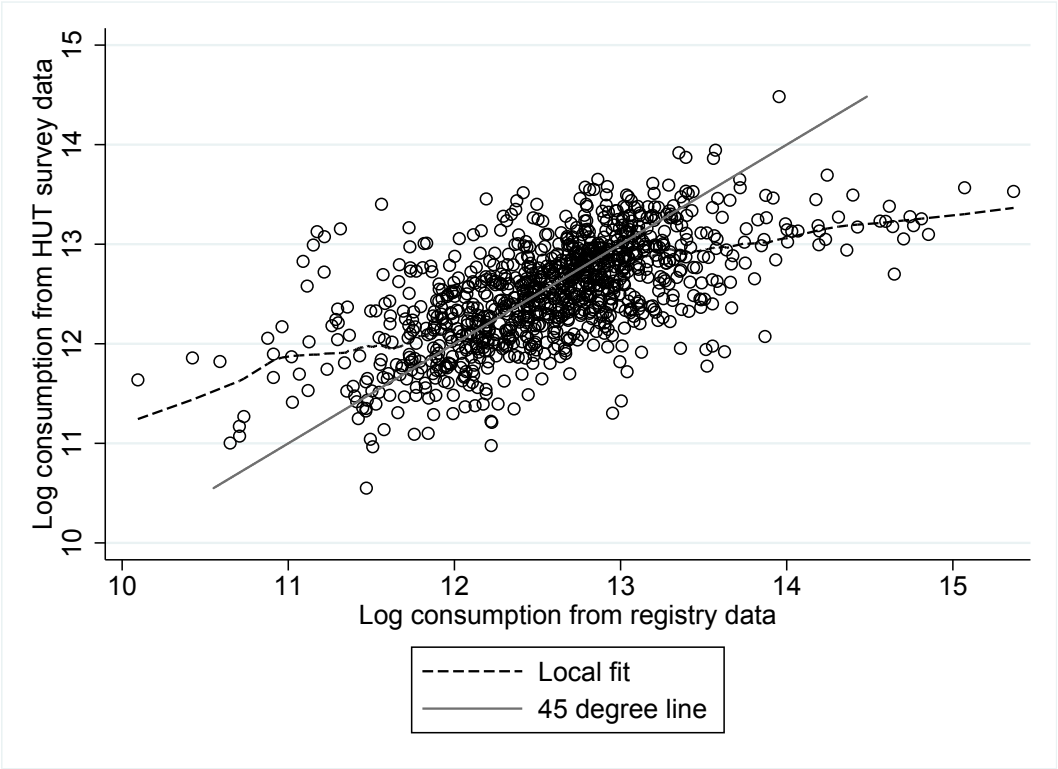

Notes: For years 2003 to 2009, the identifiers in the HUT survey can be used to match surveyed individuals to the registry data. The Figure directly compares consumption measures in the HUT and in the registry measure for the same households. It plots the log of total household consumption in HUT vs log of household expenditures in the registry. We restrict the sample to all household from the HUT surveys from 2003 to 2007 that are surveyed in the month of December to make sure that the annual expenditure measure from the HUT (which includes recall items such as durables from the last 12 months) corresponds to the annual measure in the registry data which spans all expenditures from December of year $n-1$ to December of year n. When the household composition in the HUT and registry data differs, we reconstruct households for the registry measure using the exact same individuals that are listed as members of the household in the HUT. To make expenditure measure comparable in both samples, we exclude imputed rents from the registry measure.

Figure 4: Income and Consumption Shares by Deciles of Gross Income - 2003

A. Registry Measure

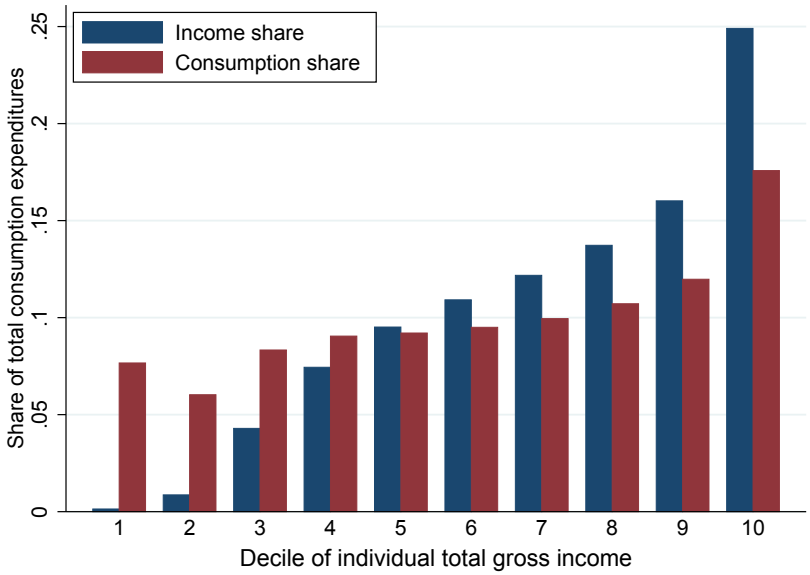

B. HUT Survey Measure

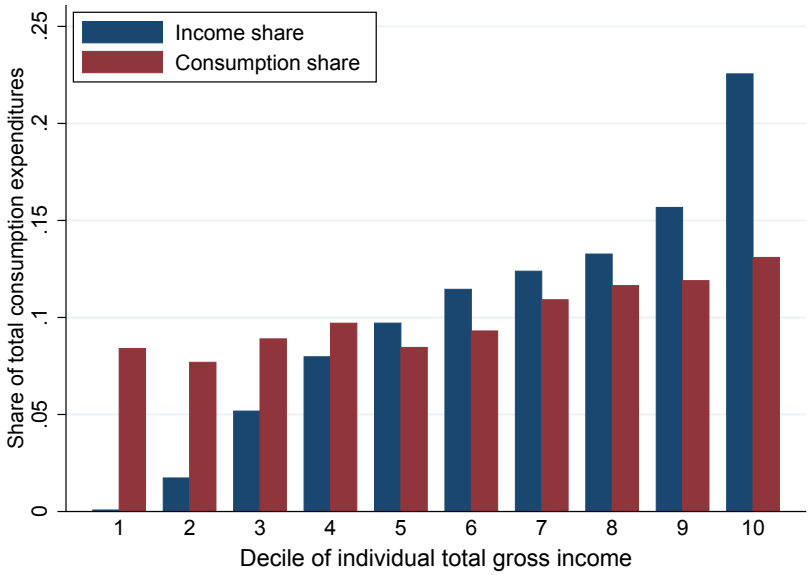

Notes: The Figure reports the income and consumption shares by deciles of total gross income. The sample is restricted to individuals between age 25 and 55 in 2003. Individuals are ranked in deciles of the distribution of gross individual income including capital gains. The income share is computed as the sum of gross earnings (not including capital gains) by decile, divided by the sum of all gross earnings in the sample. The consumption share is computed as the sum of individual consumption by decile, divided by the sum of all consumption in the sample. In panel A, we exclude imputed rents from the registry measure, to make it comparable to the HUT measure. For household consumption expenditures in the registry data, we split total expenditures among household members using the OECD equivalence scale. In panel B, we use the survey sampling weights and the registry measure of gross income to rank individuals and compute income shares. Discrepancies in income shares between panel A and panel B stem from sampling in the HUT. For household expenditures, we split total expenditures among household members using the OECD equivalence scale.

Figure 5: Log consumption changes vs log income changes by percentiles of the income distribution

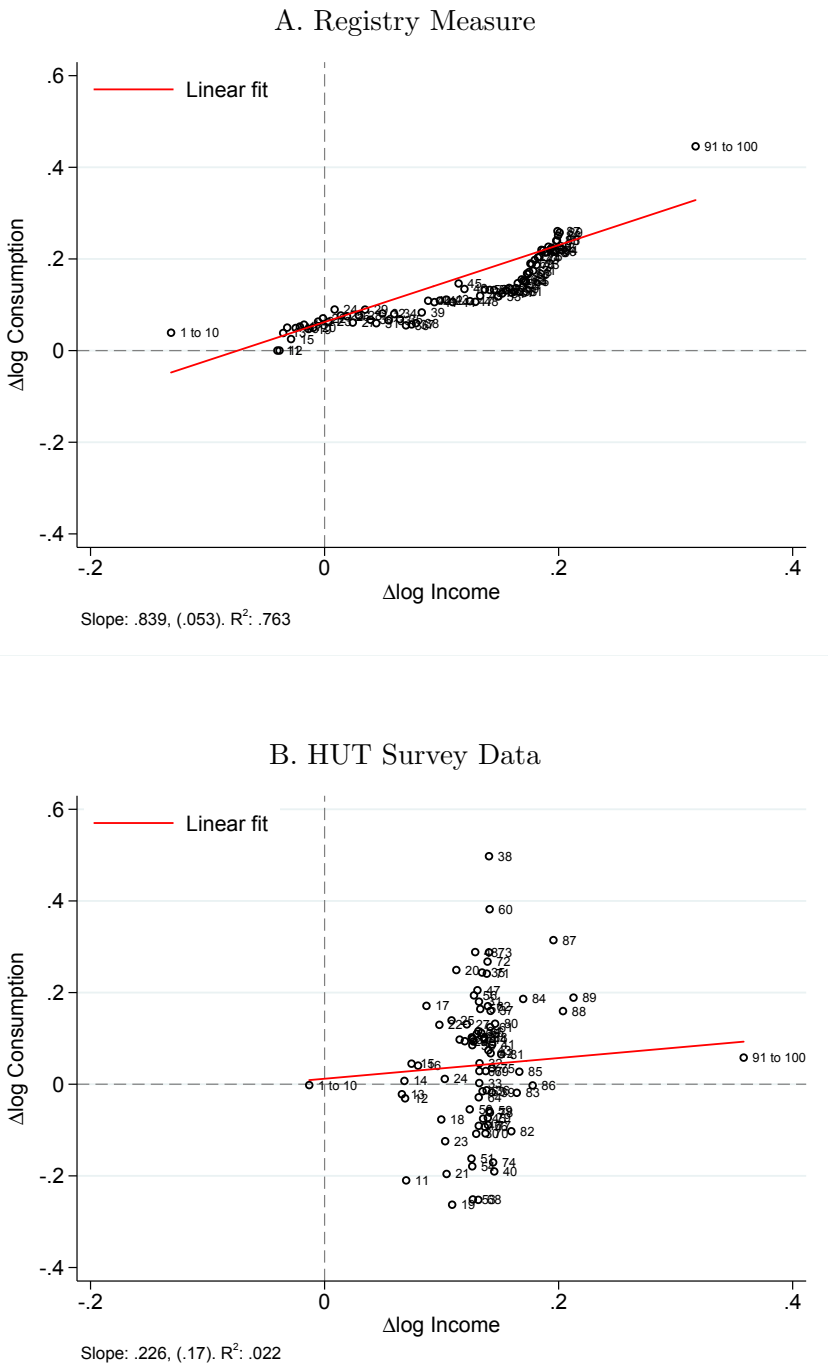

Notes: The Figure compares the elasticity of log consumption to log income, across income fractiles, measured in the registry data versus the HUT survey. For the registry data in panel A, we rank individuals by percentile of the disposable income distribution in 2003. Then, for each percentile, we compute the log change in disposable income between 2003 and 2007 and plot it against the log change in real consumption expenditures over the same period. Each dot on the graph is labelled with the percentile of the disposable income distribution to which it corresponds. For the HUT data in panel B, we do the same exercise using the survey measure of disposable income and the survey measure of consumption, except we do log changes over the 2003 to 2007 period, as we only have the HUT since 2003. Panel A shows clearly a strong and positive elasticity of consumption expenditures with respect to income. This relationship is particularly significant for higher fractiles in the income distribution, while the relationship between log changes in income and consumption is much flatter at lower level of the income distribution. Panel B does not show any sign of significant correlation between income and consumption changes across income fractiles.

Figure 6: Earnings, Benefits, and Consumption Responses to Health Shocks

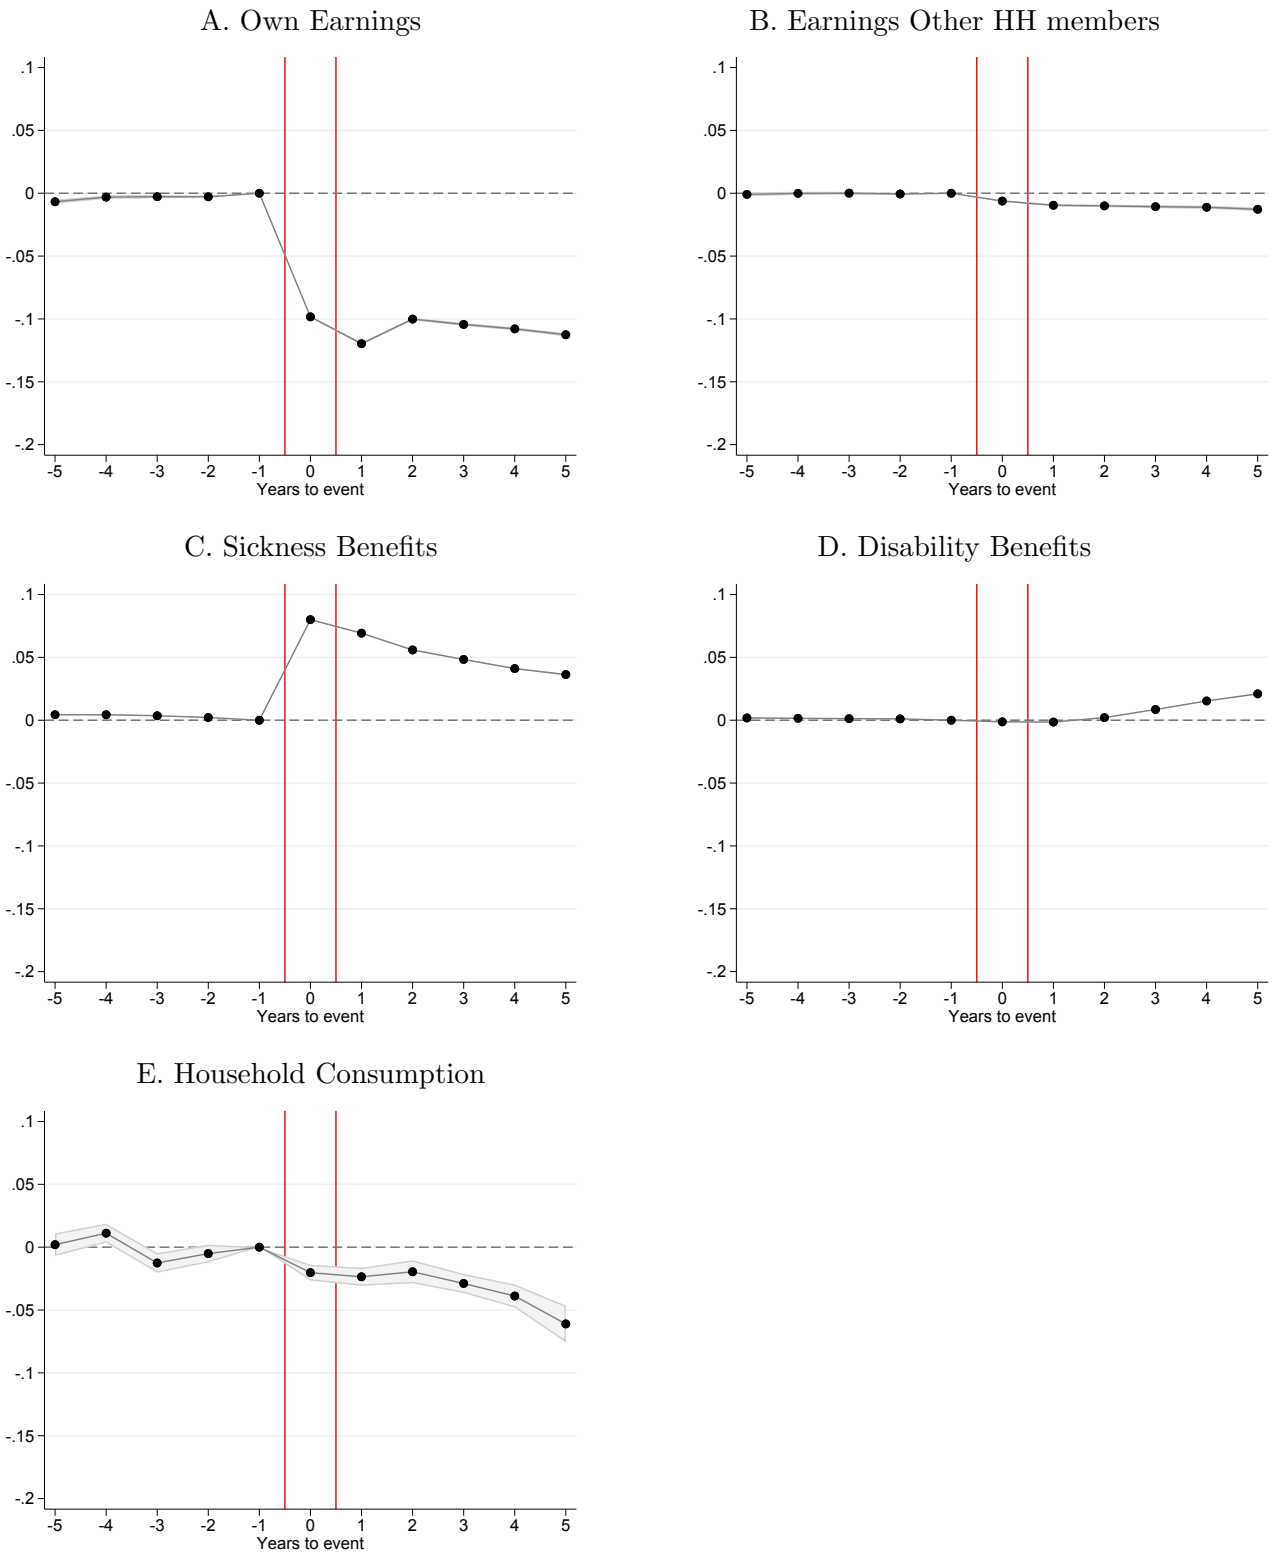

Notes: The Figure reports event study estimates of the earnings, benefits, and consumption responses to a health shock, using specification (2) on the sample of treated individuals and a control group of individuals obtained from nearest-neighbor matching on pre-event characteristics. All estimates are expressed as a fraction of total household expenditures as of event year -1. Health shocks are defined as the first occurrence of sickness benefits receipt by an individual (i.e. a work incapacitation larger than 14 days). We restrict the sample to individuals aged 25 to 55 at the time of the event. Individuals may transition to disability benefits after a year of sickness benefit receipt. In Panel A, earnings are the gross earnings of the individual experiencing the health shock. In Panel B, we sum the gross earnings of all the other members of the household. The composition of the household is fixed as of the last year prior to event. See text for details.

Figure 7: Evolution of Earnings, Pension, and Consumption Around Retirement

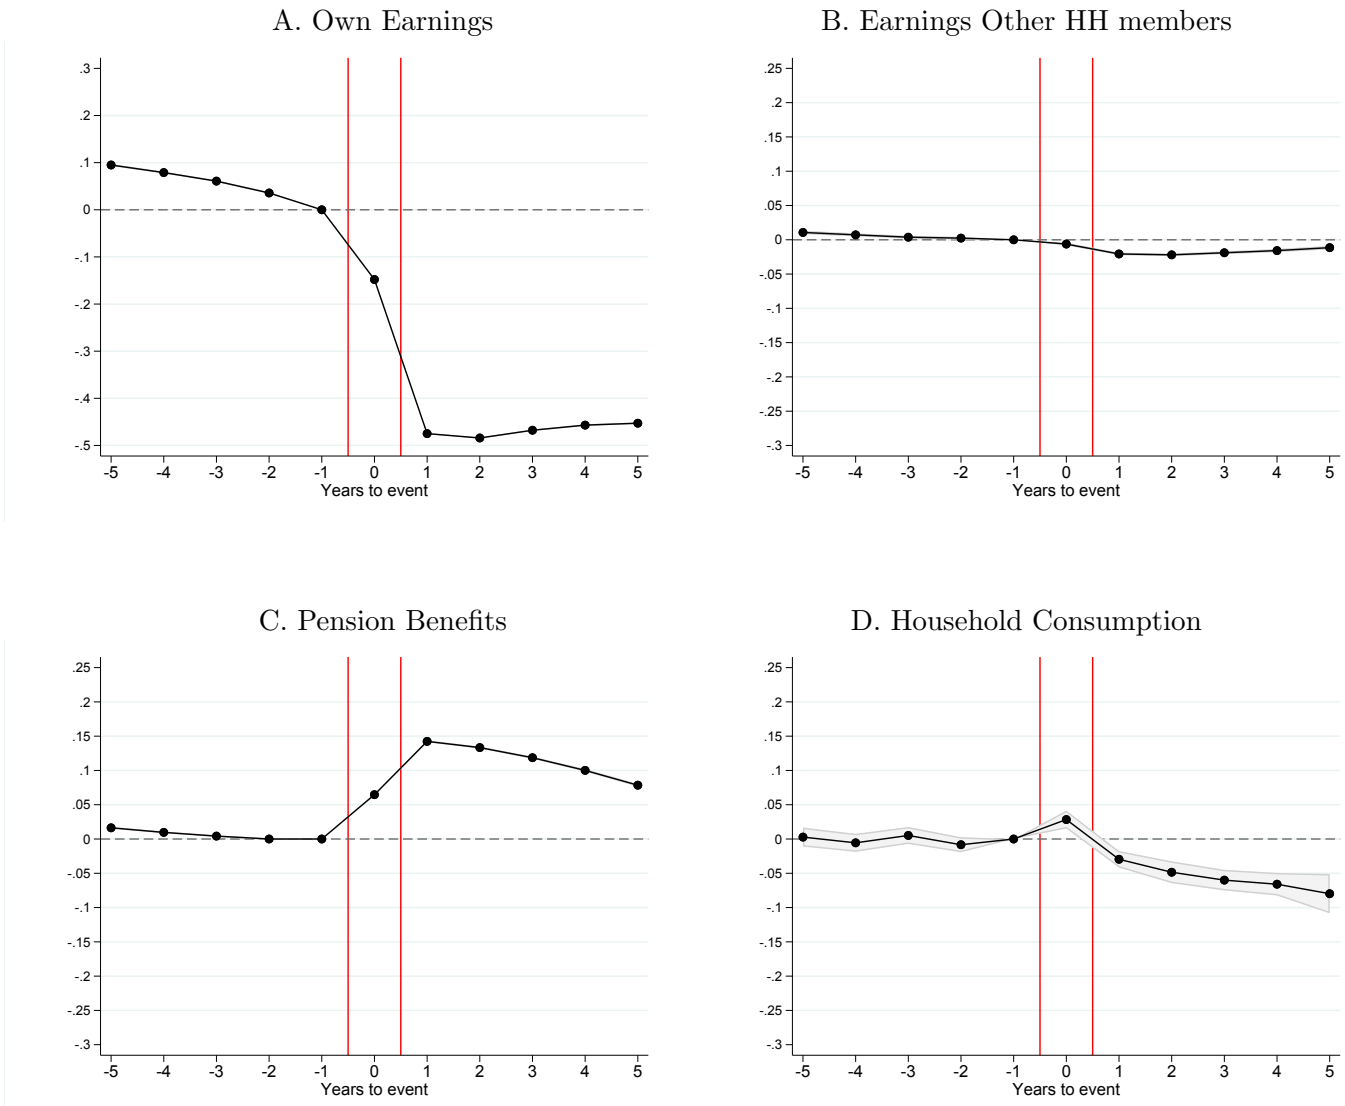

Notes: The Figure reports event study estimates of the earnings, pension, and consumption dynamics around retirement, using specification (3). All estimates are expressed as a fraction of total household expenditures as of event year -1. The sample is restricted to individuals aged 50 or more and retirement is defined as (i) the first year in which the individuals' labour income is inferior to the base amount (BA), (ii) immediately followed by another year of income less than BA. In Panel A, earnings are the gross earnings of the individual experiencing retirement. In Panel B, we sum the gross earnings of all the other members of the household. The composition of the household is fixed as of the last year prior to retirement. Pension benefits is the sum of all public pensions and occupational pensions, but does not include withdrawals from individual pension accounts, which are counted as consumption out of assets. See text for details.

Figure 8: Evolution of Earnings, Transfers, and Consumption Around Parenthood

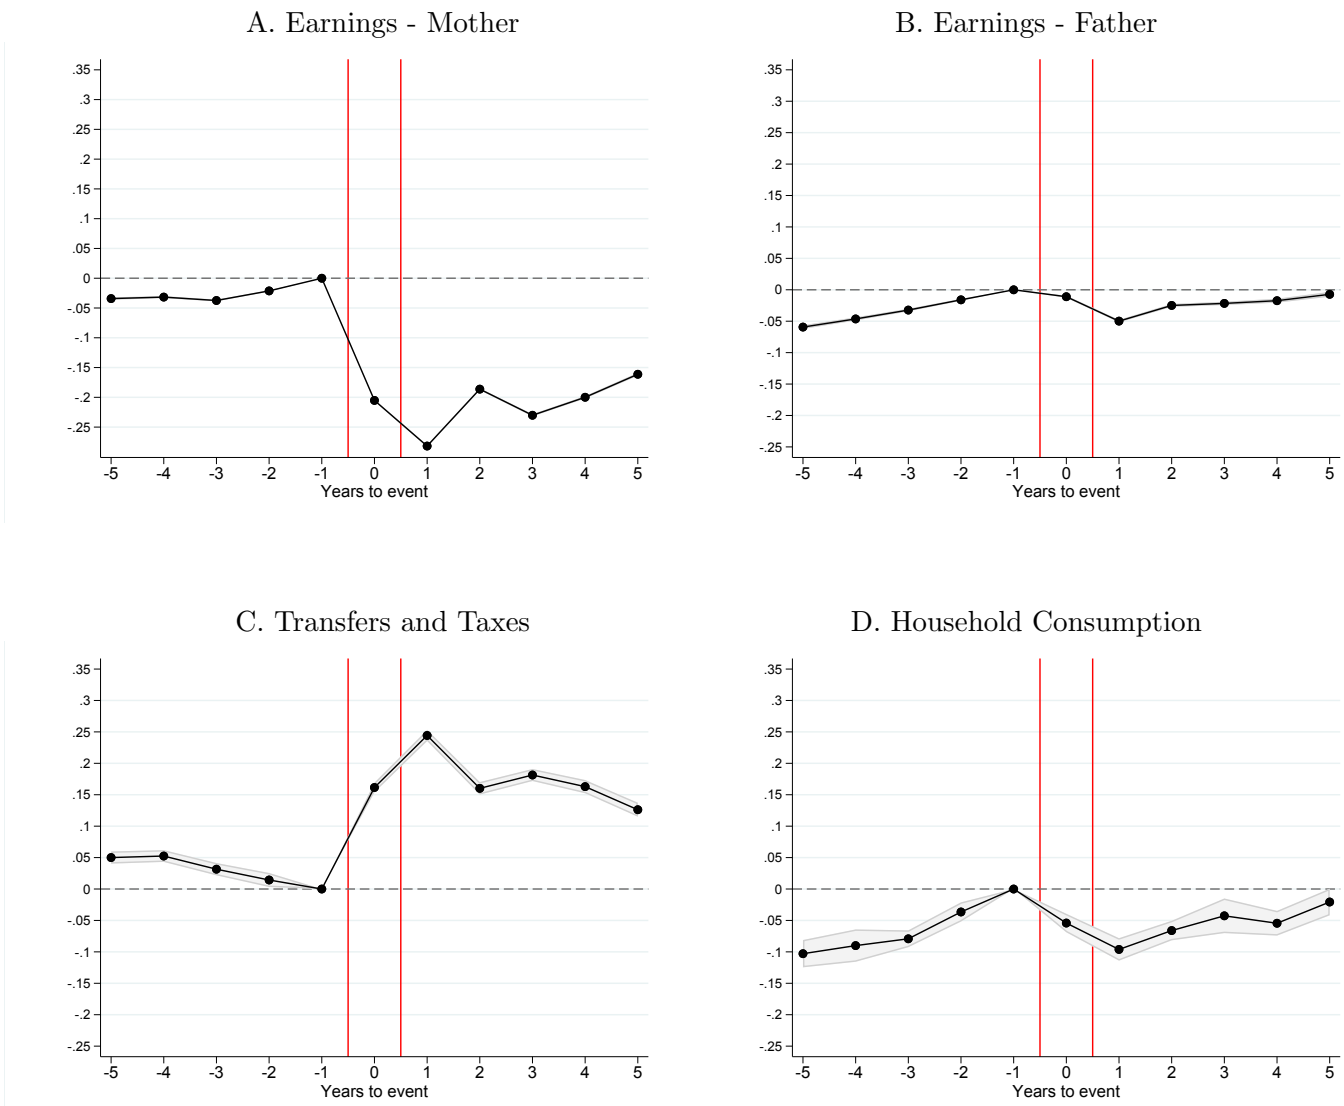

Notes: The Figure reports event study estimates of the earnings, net transfers, and consumption dynamics around the arrival of the first child, using specification (2) on the sample of treated individuals and a control group of individuals obtained from nearest-neighbor matching on pre-event characteristics. All estimates are expressed as a fraction of total household expenditures as of event year -1. The composition of the household is fixed as of the year of the arrival of the first child. See text for details.

Figure 9: Housing & Consumption Around Parenthood

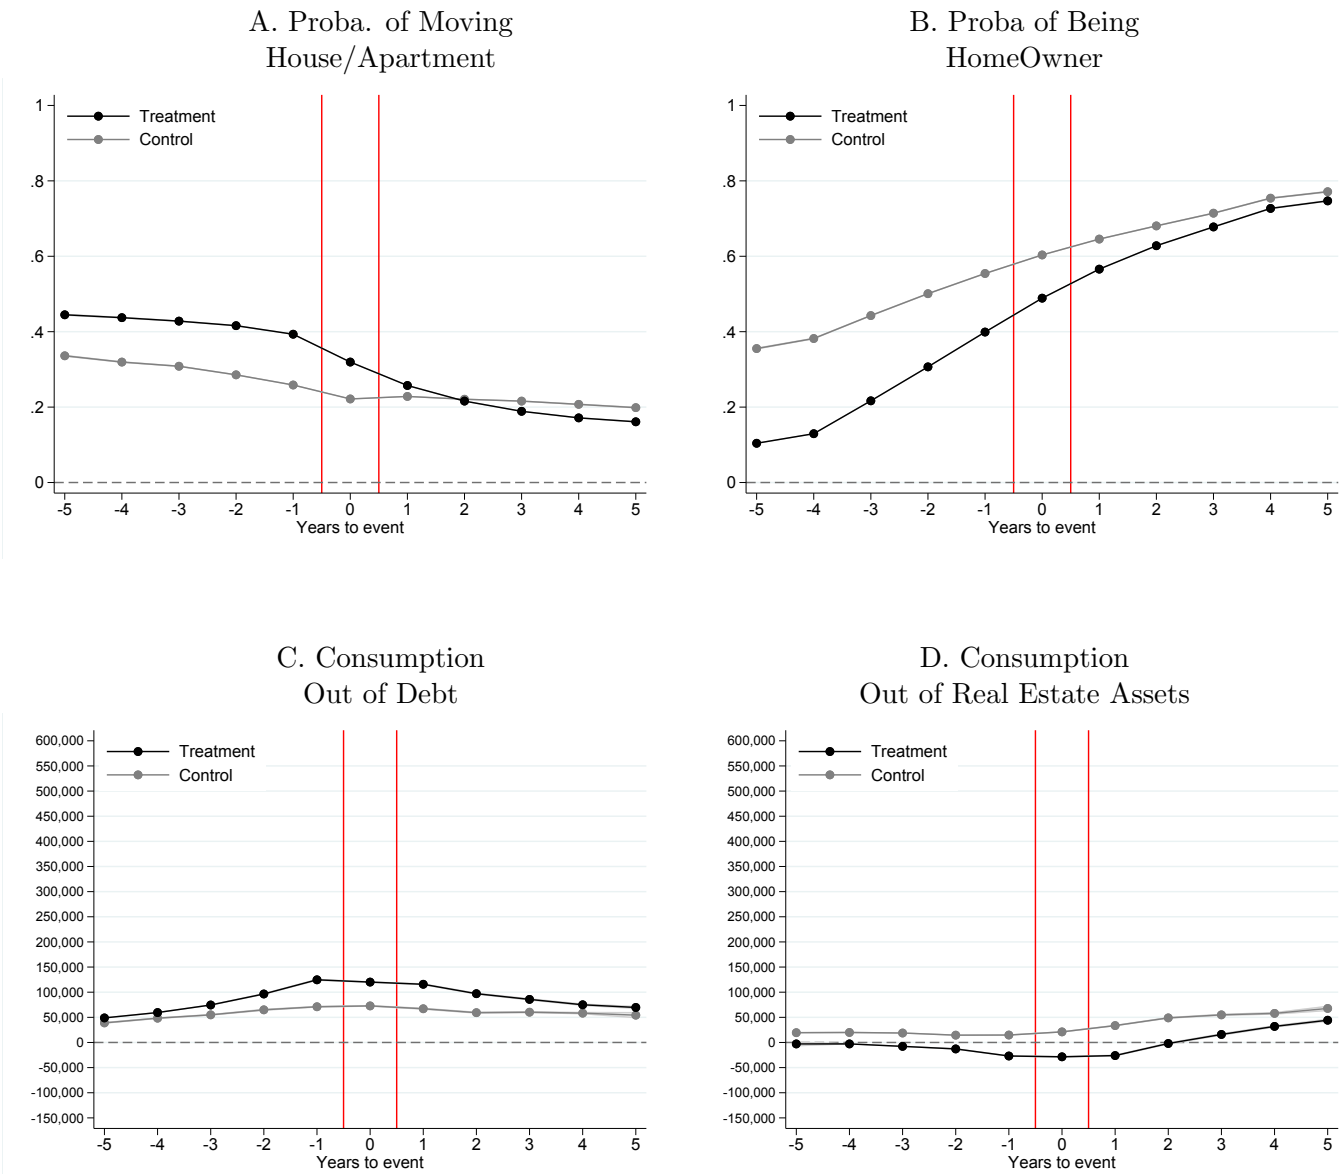

Notes: The Figure highlights the role of housing in determining consumption dynamics around the arrival of the first child. For each panel, we report the average level of the outcome as a function of event time in the treated group and in the control group of women who do not experience a first child birth in the three years following the year when their treated counterpart does, obtained from nearest-neighbor matching on pre-event characteristics. The control women are allocated the placebo event time of their nearest-neighbor in the treated group. See text for details. In panel A, we report the probability to move house or apartment, and in panel B the probability to become a homeowner measured by having positive "owner occupied real estate wealth" in the wealth tax data. In panel C, we report consumption out of debt, which is equal to the change in debt between year t and $t-1$ minus all interests payments on existing debt. In panel D, we report consumption out of real estate assets, which includes imputed rents. The composition of the household is fixed as of the year of the arrival of the first child.

Figure 10: Durable Consumption Around Parenthood: Probability of Buying a New Car

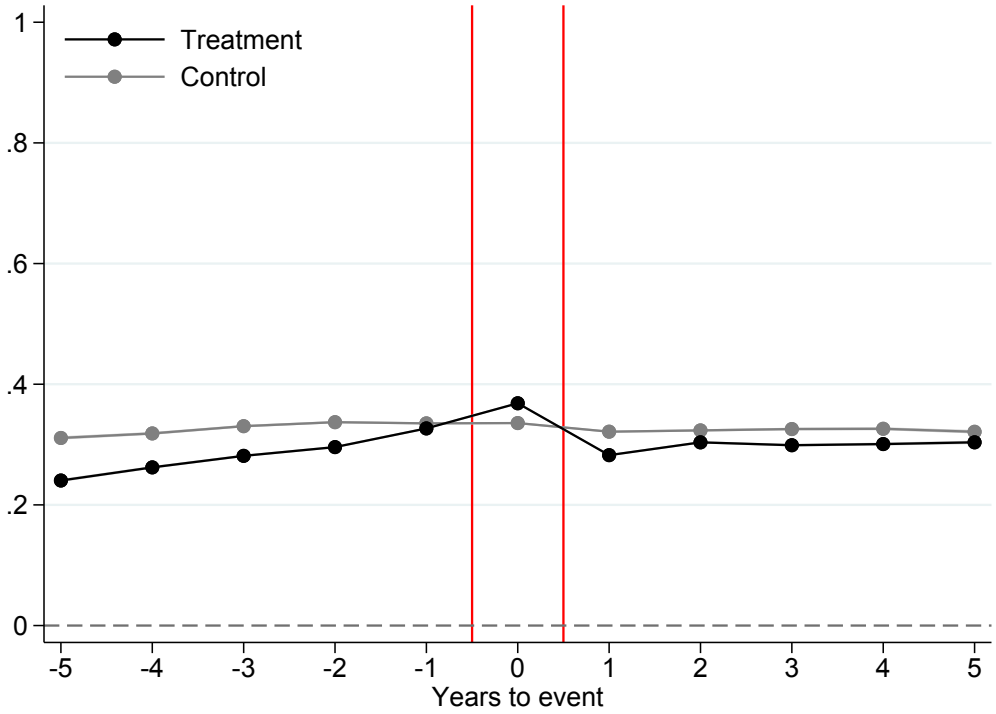

Notes: The Figure highlights the role of durable consumption in determining consumption dynamics around the arrival of the first child. We report the average probability for the household to buy a new car, from the car registry data, as a function of event time in the treated group and in the control group of women who do not experience a first child birth in the three years following the year when their treated counterpart does, obtained from nearest-neighbor matching on pre-event characteristics. The control women are allocated the placebo event time of their nearest-neighbor in the treated group. See text for details. The composition of the household is fixed as of the year of the arrival of the first child.

Figure 11: Decomposition of Consumption Expenditures Responses to Events

A. Health Shock

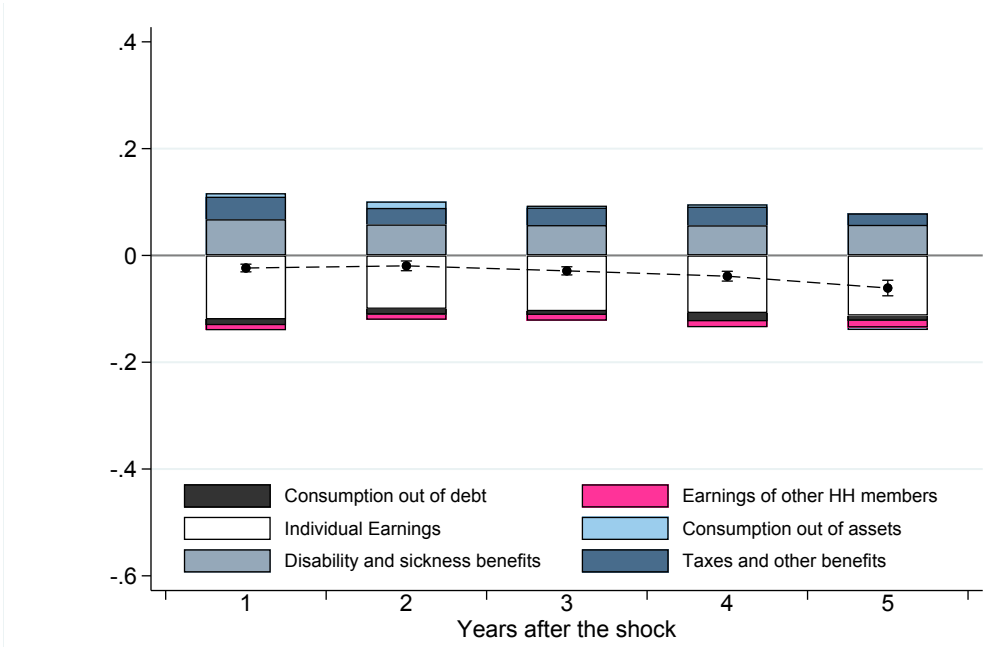

B. Retirement

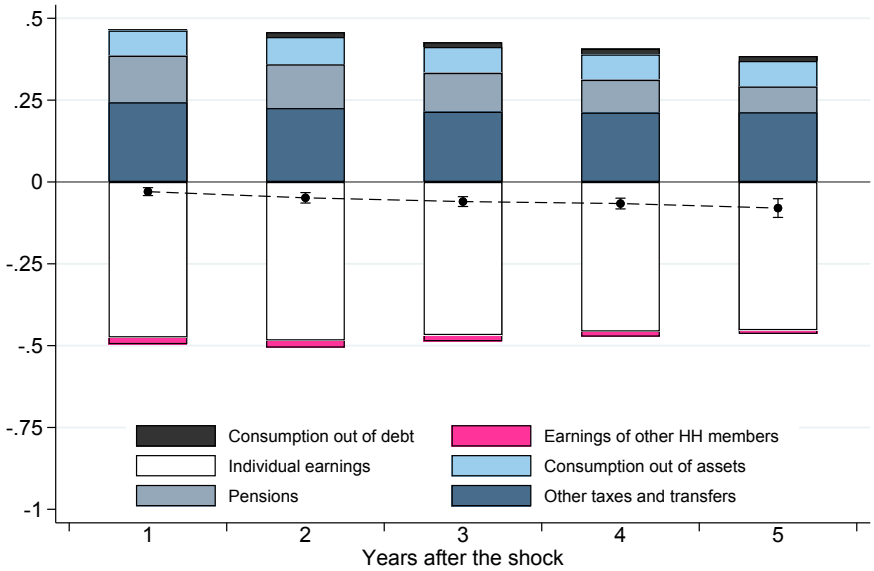

Figure 11: Decomposition of Consumption Expenditures Responses to Events (*continued)*

C. Parenthood

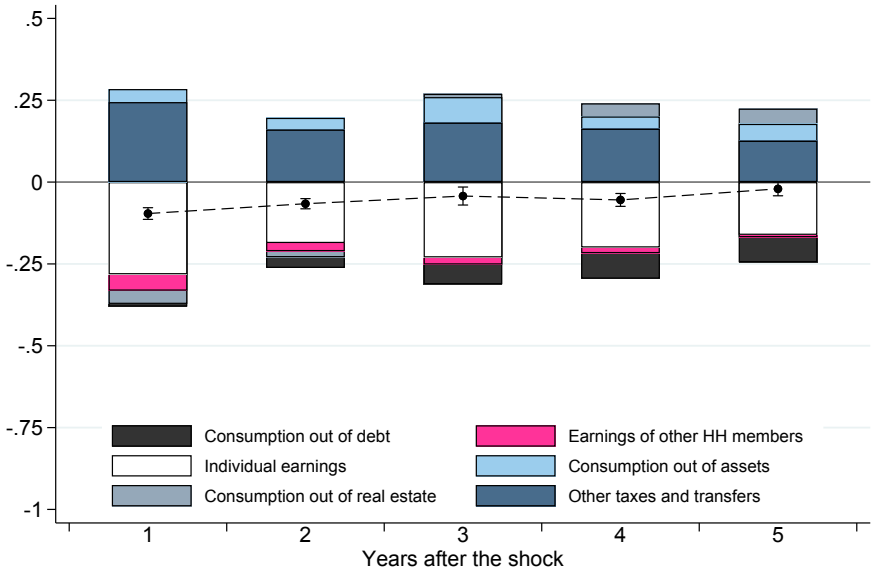

Notes: The Figure decomposes the consumption expenditure response to health shocks, parenthood and retirement for event time years $J_{it} = 1$ to $J_{it} = 5$. For panel A, we decompose consumption into six main components: own earnings, earnings of other members of the household, sickness and disability benefits, consumption out of debt, consumption out of assets, and other transfer and taxes. For panel B, we focus on the same components and pension benefits instead of sickness and disability benefits. For panel C, we focus on the same components and we extract from the consumption out of assets the consumption out of real estate, while it was not the case in Panel A and Panel B. The dot represents the total change in consumption expenditures relative to event year -1. Each component of the bar represents the contribution of each component to the total change in consumption expenditures. For panel A and C, these components are estimated using specification (2) on the sample of treated individuals and a control group of individuals obtained from nearest-neighbor matching on pre-event characteristics. For panel B, these components are estimated using specification (3) on the sample of treated individuals. All estimates are expressed as a fraction of total household expenditures as of event year -1.

Figure 12: Consumption Dynamics Around Event: Survey vs Registry measure estimates

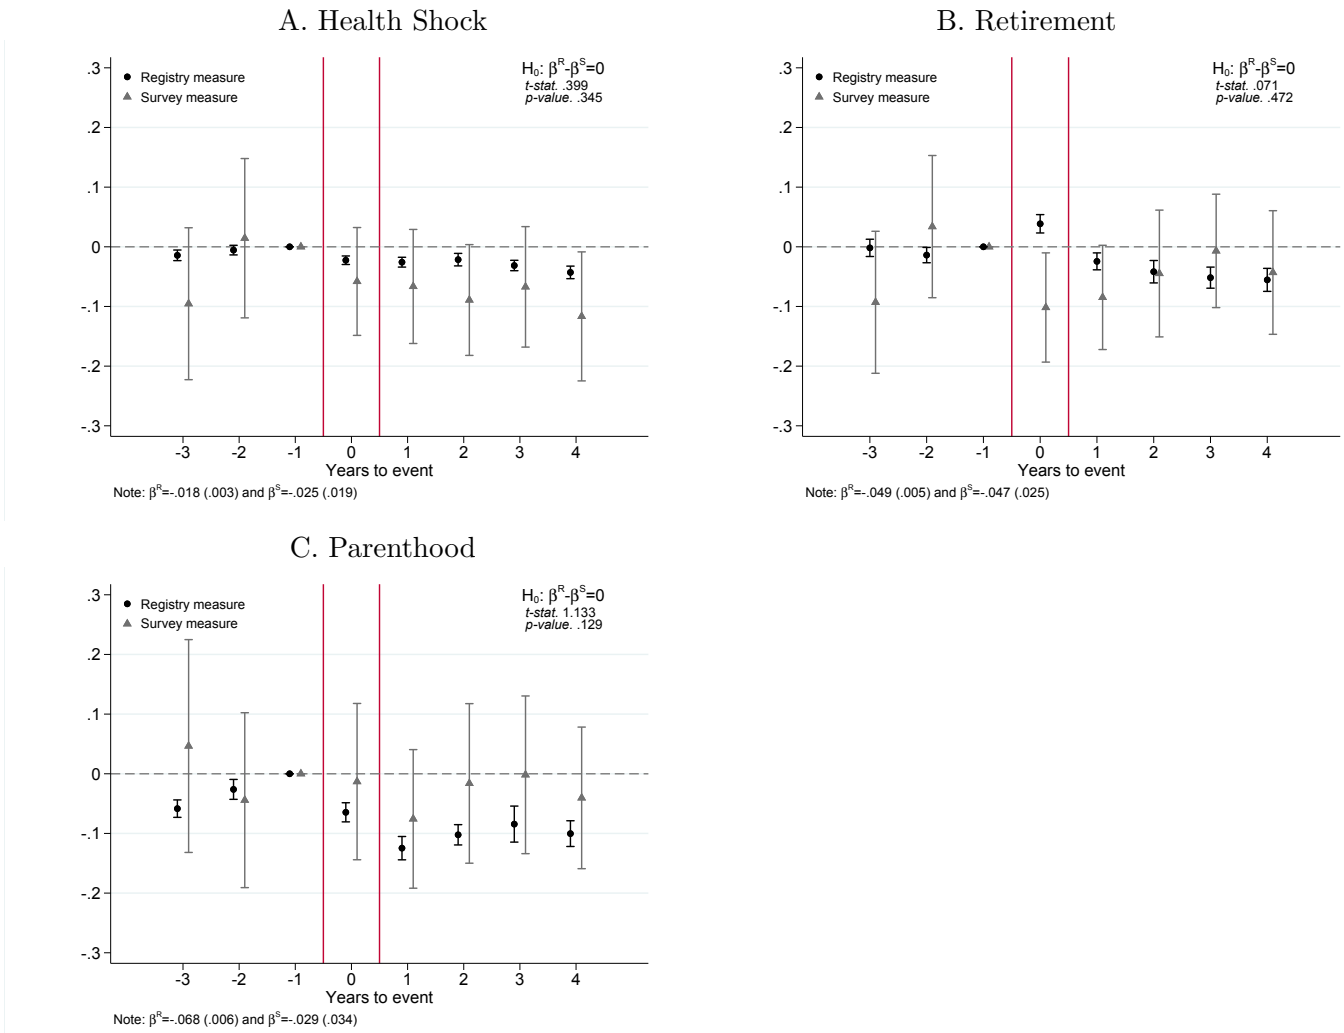

Notes: The Figure reports event study estimates of consumption dynamics around the three events: health shocks in panel A, retirement in panel B and parenthood in panel C. For each panel, we compare estimates from the HUT survey with estimates from the registry measure of consumption where we remove imputed rents to make consumption measure consistent with the HUT survey. For the HUT estimates, we use specification (4). For the registry measure the event study estimates are obtained using specification (2) for panel A and C, and specification (3) for panel B. In each panel, we report the t stat and p value of a test of equality of the survey estimate (β^s) and the registry estimate (β^r) of the average change in consumption in the 5 years post event. See text for details.

Table 1: Overview of income and wealth components

Disposable income (per capita)	160,930 kronor
Income from labor and transfers	Percentage of gross income
Labor income	55.66
Benefits replacing labor income	9.27
Pensions and means tested benefits	15.95
Tax deductions and other transfers	6.67
Taxes	-27.91
Income from holding capital	
Imputed rents	9.00
Dividends from securities	2.18
Other capital income	0.83
Interest on deposits	0.44
Interest on liabilities	-3.77
Net capital gains (per capita)	56,263 kronor
Capital gains components	Percentage of gross capital gains
Positive gains from real assets	66.09
Negative gains from real assets	-1.39
Positive gains from financial assets	33.91
Negative gains from financial assets	-10.27
Net wealth (per capita)	523,119 kronor
Wealth components	Percentage of gross wealth
Real wealth	69.32
Securities	19.84
Bank account holdings	9.07
Other wealth	1.77
Debt	-29.24
Observations	58,396,928

Notes: Income and wealth are measured years 2000-2007 in 2003 year's Swedish kronor (1 SEK $\approx$ 0.13 USD in June 2003). The sample is all individuals residing in Sweden aged 16 years and above. "Gross income" is disposable income from labor and capital plus paid taxes and imputed rents. "Labor income" is income from employment and self-employment. "Benefits replacing labor income" are benefits suchs as unemployment and sickness insurance and parental leave benefits. In "Pensions and mean tested benefits" social aid and housing aid are included. "Tax deductions and other transfers" includes benefits such as child support. "Other capital income" includes income from subletting homes, gains from selling real estate and tax on lottery wins (primarily wins from outside of the EEA). "Other wealth" consists of cars, boats, art etc. that individuals have reported to the tax authority.

Appendix A Additional Graphs And Tables

Figure A.1: Consumption, earnings and asset Responses to Health Shocks: Treated vs Control

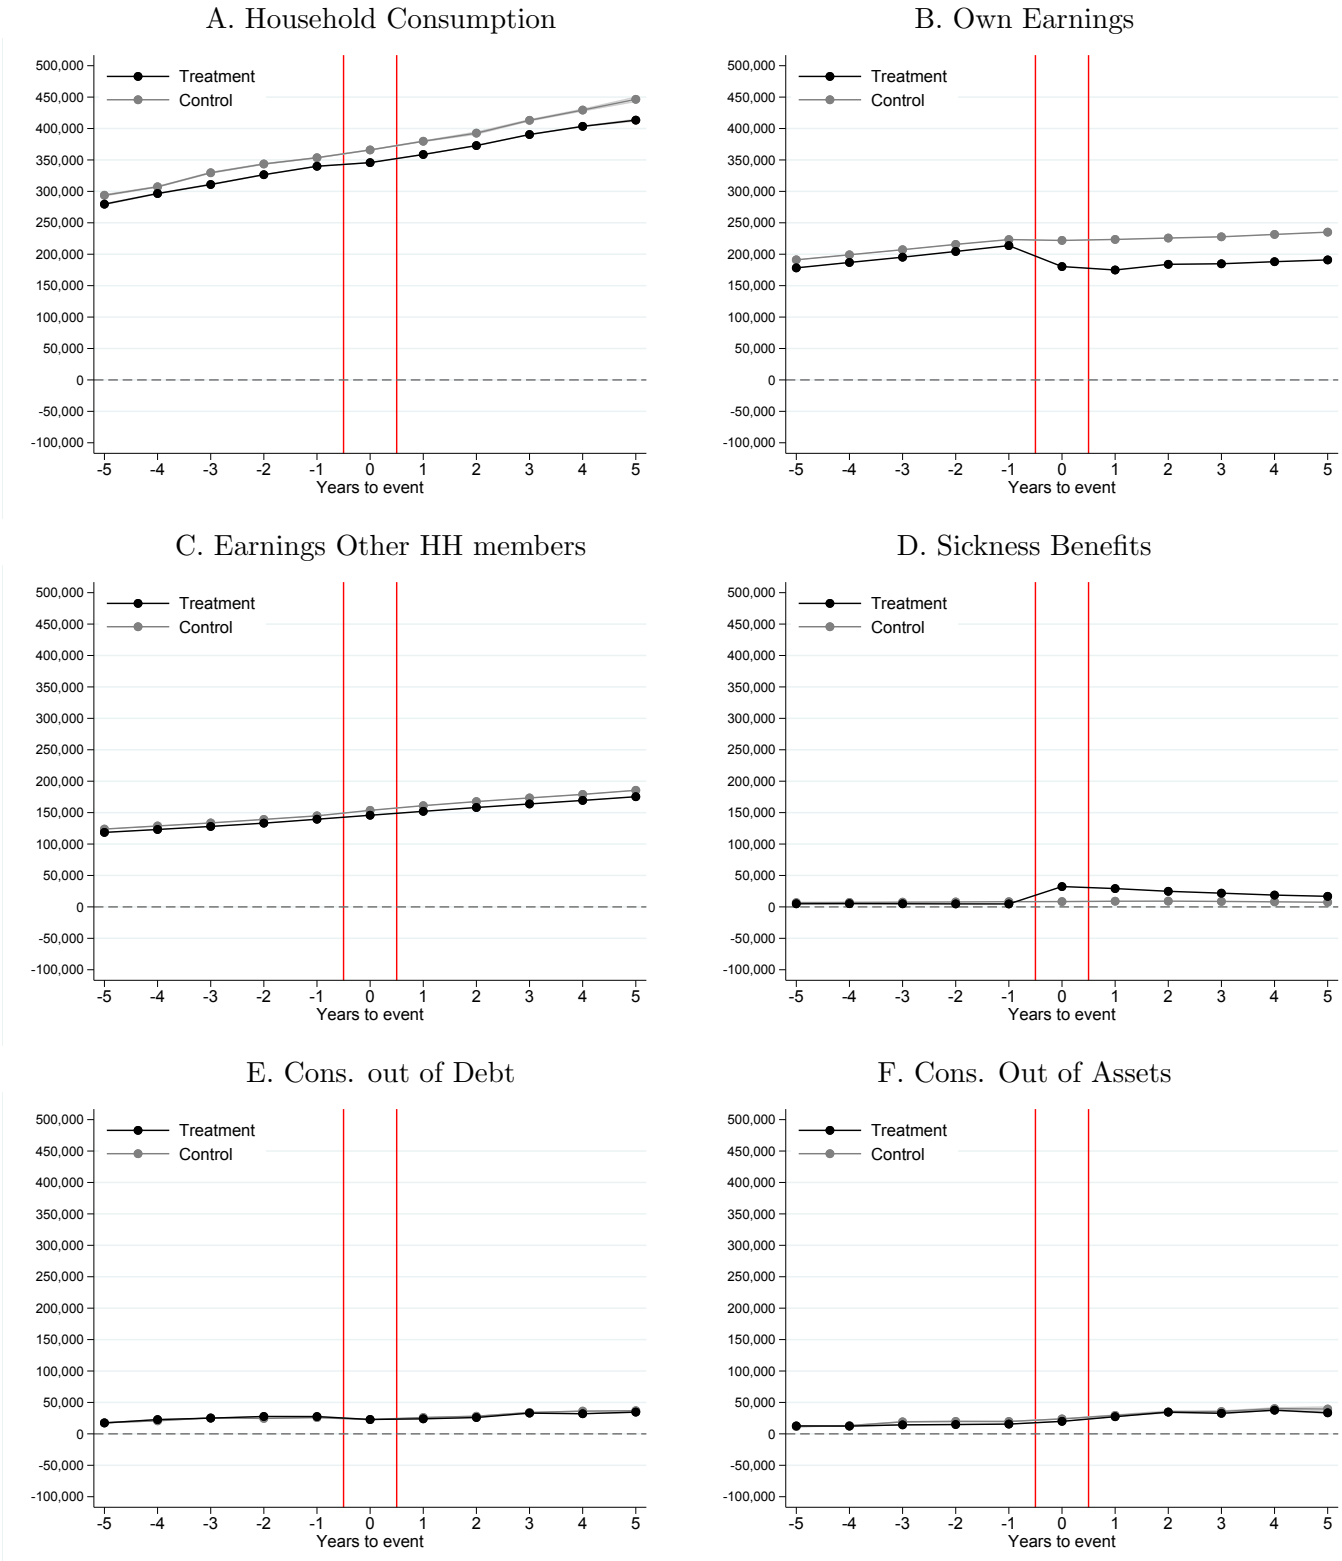

Notes: These Figures report the coefficients and confidence intervals of the time dummies in event study-type regressions. A period of five years before and after the first sickness spell is studied. Total average household consumption, earnings, sickness benefits, consumption out of debt and consumption out of assets are shown. The control group is created by propensity score matching on a wide array of covariates including education level, gender, age, region of residence, marital status, three lags of income, the number of dependent children and industrial sectors of activity. Household composition is held fixed as of the last year prior to event.

Figure A.2: Earnings and Consumption Expenditures Responses to Retirement: Treated vs Control

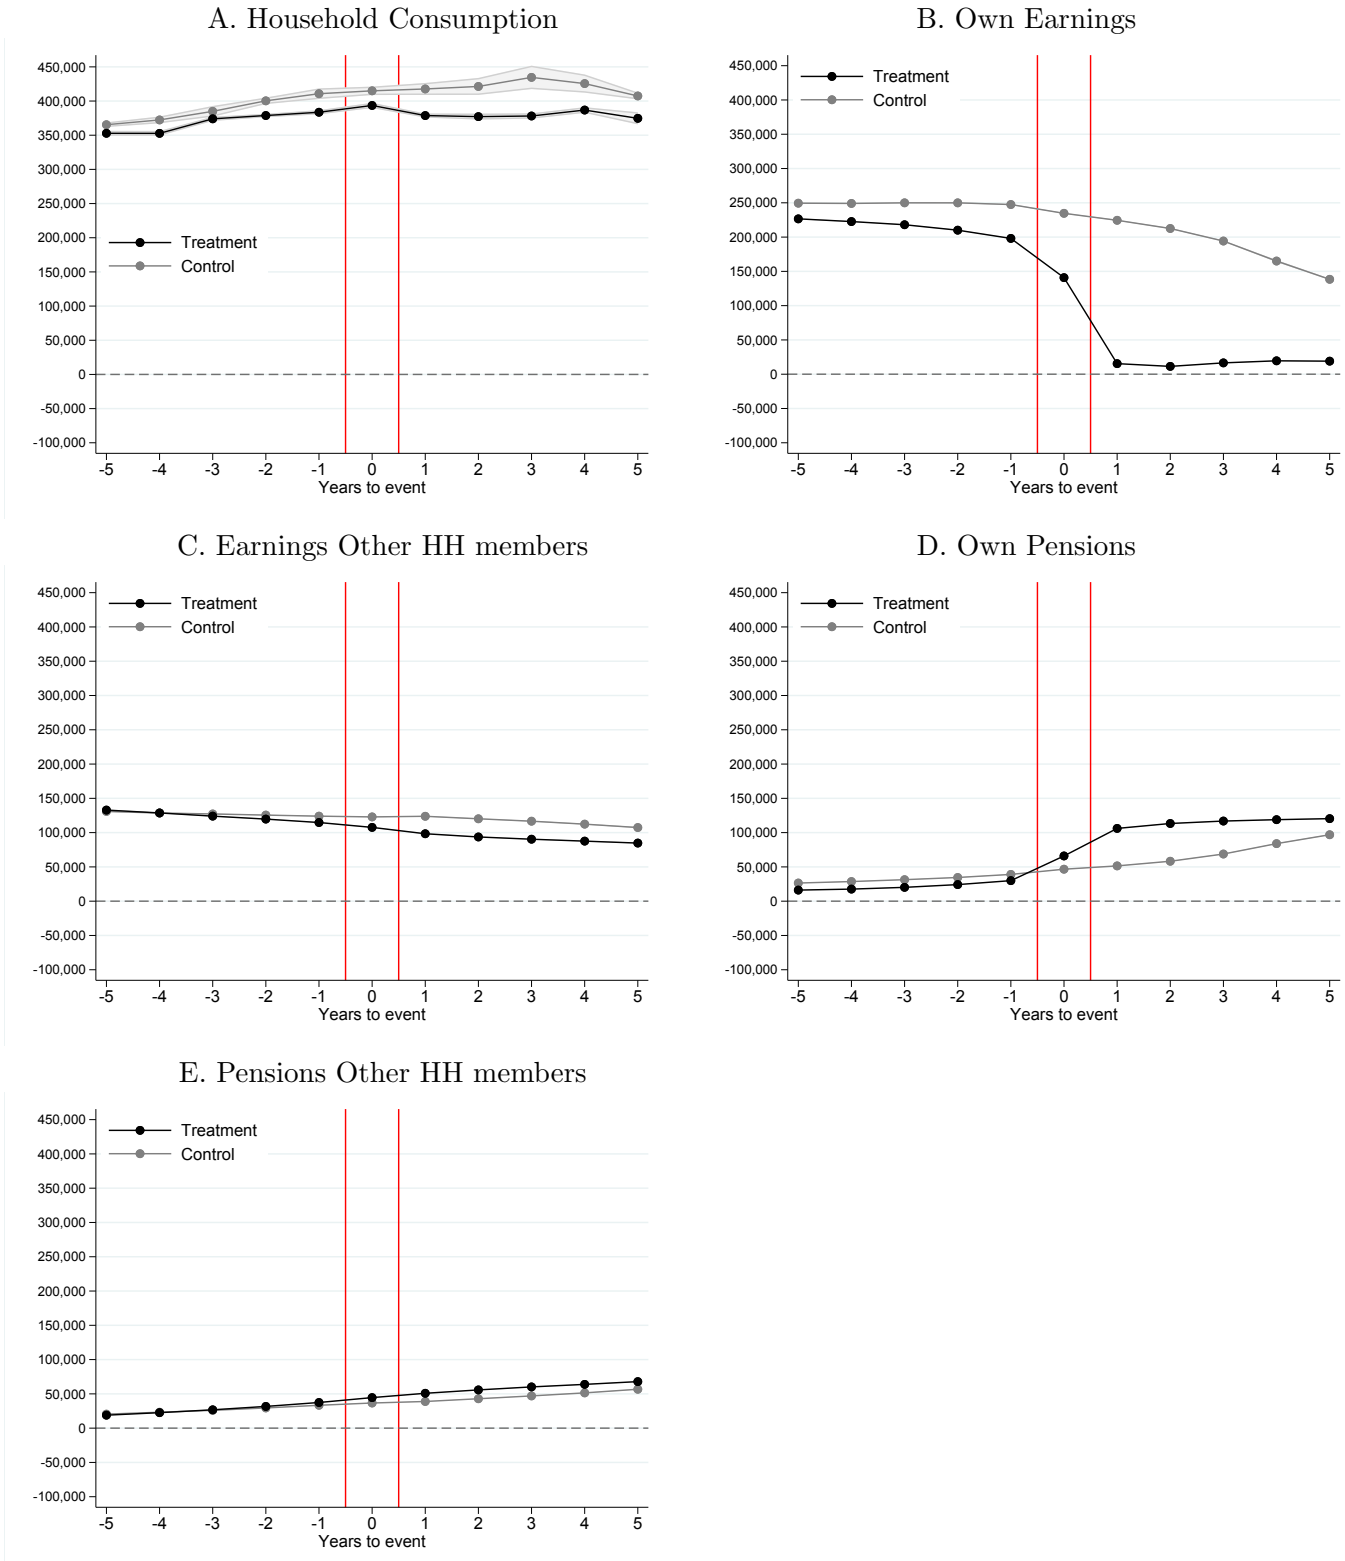

Notes: These Figures report the coefficients and confidence intervals of the time dummies in event study-type regressions. A period of five years before and after the year when individuals retire is studied. Total average household consumption, earnings, and pensions are shown. The control group is created by propensity score matching on a wide array of covariates including education level, gender, age, region of residence, marital status, three lags of income, the number of dependent children and industrial sectors of activity. Household composition is held fixed as of the last year prior to event.

Figure A.3: Earnings and Consumption Expenditures Responses to Parenthood: Treated vs Control

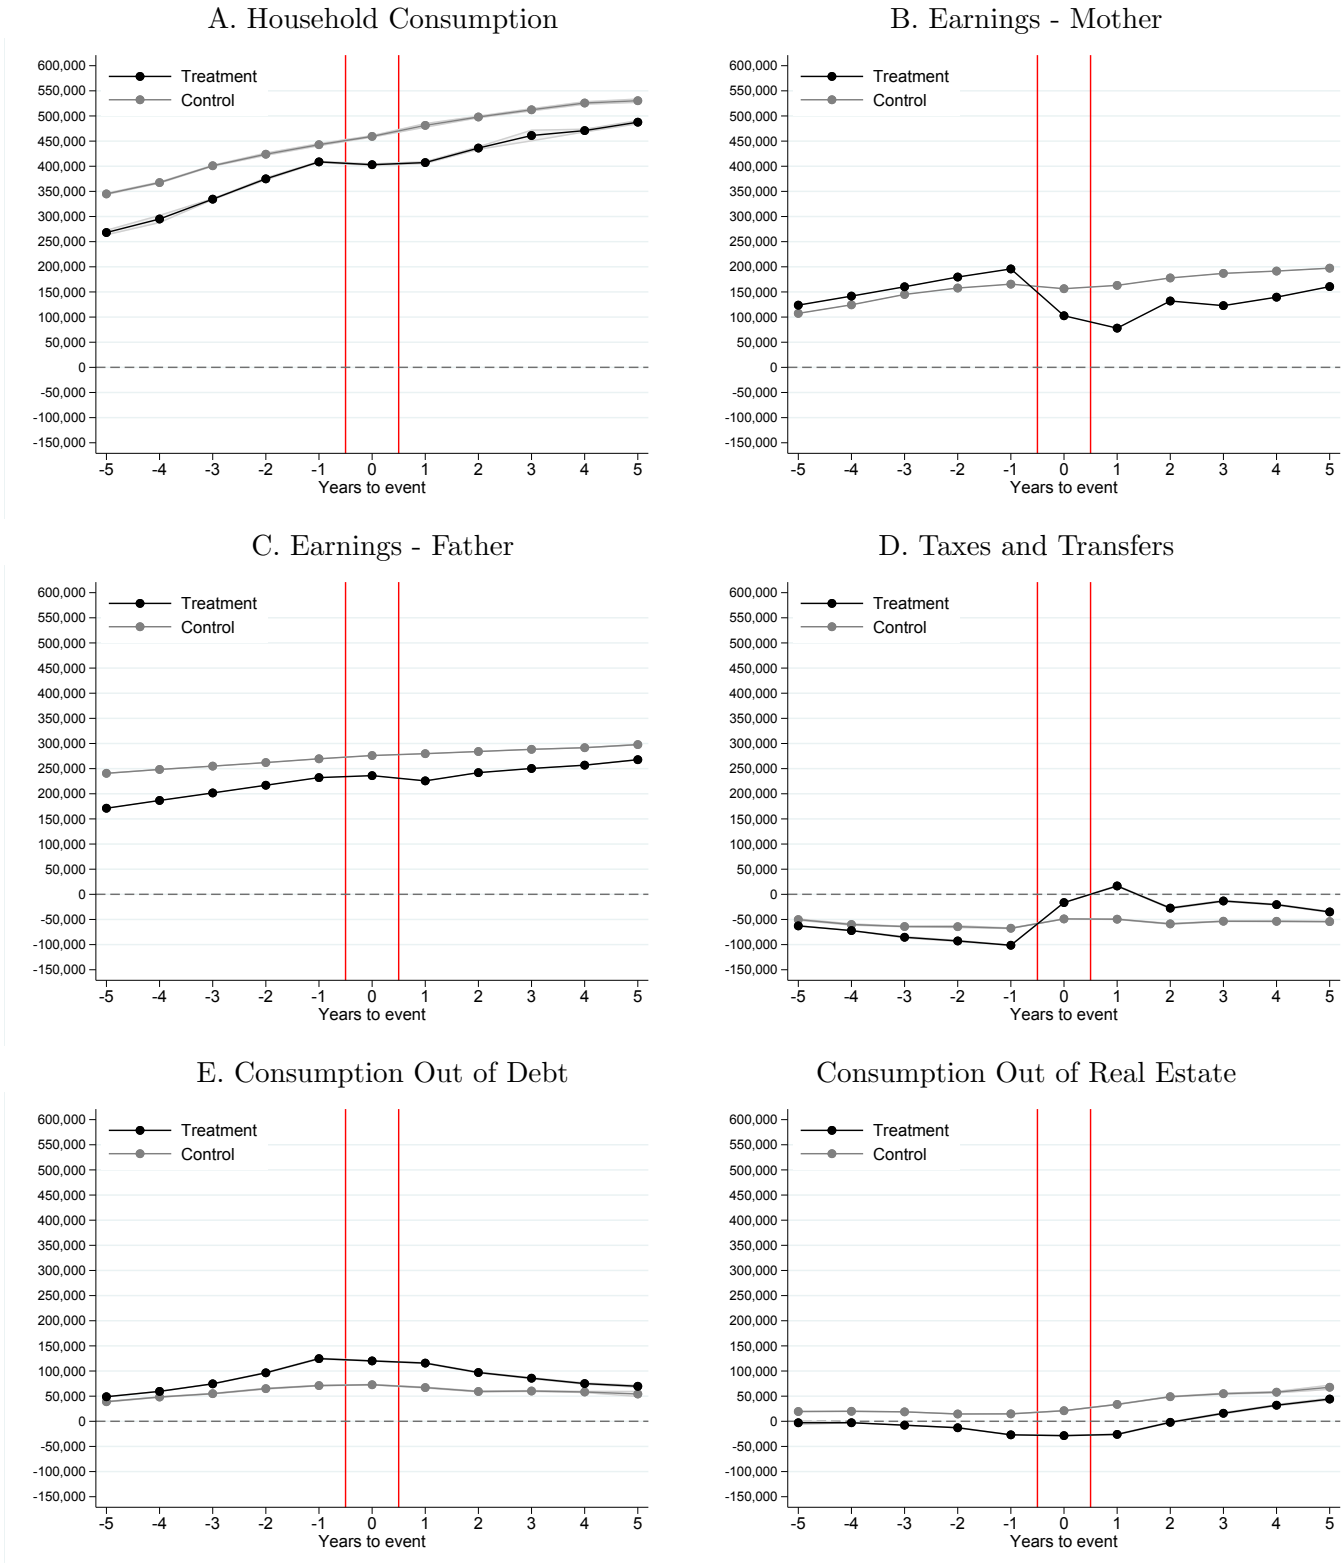

Notes: These Figures report the coefficients and confidence intervals of the time dummies in event study-type regressions. A period of five years before and after the year when a woman gets her first child is studied. Total average household consumption, earnings, taxes and transfers (including parental benefits), consumption out of debt, and consumption of real estate assets are shown. The control group is created by propensity score matching on a wide array of covariates including education level, age, region of residence, marital status, three lags of income, and industrial sectors of activity. We impose that women in the control group do not have their first child in the three

Figure A.4: Evolution of Consumption Around Parenthood: Always Renters

A. Household Consumption

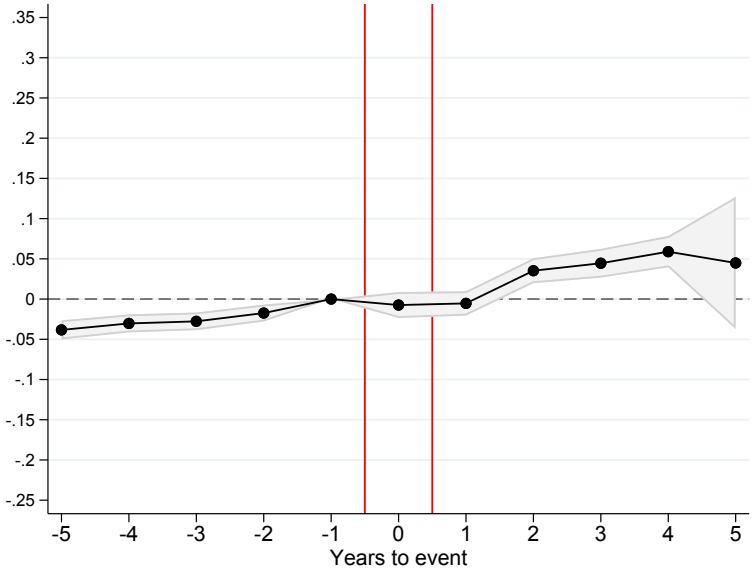

B. Consumption Out of Debt

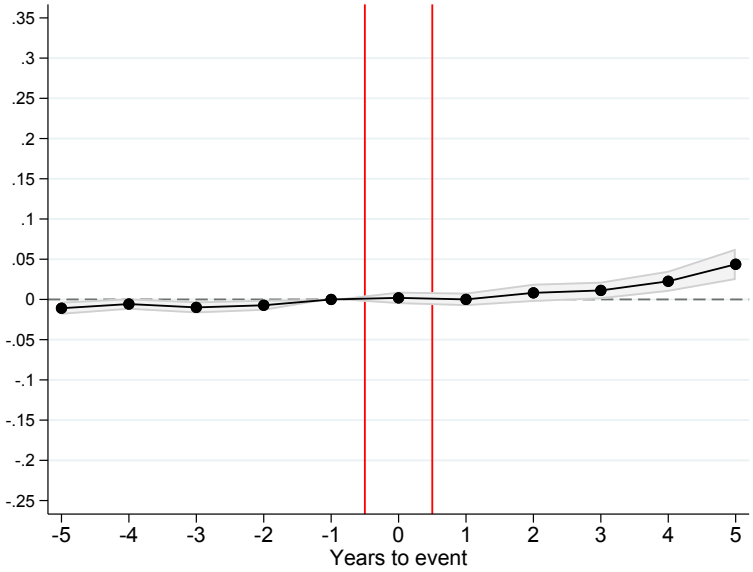

Notes: These Figures report the coefficients and confidence intervals of the time dummies in event study-type regressions. A period of five years before and after the year when a woman gets her first child is studied. Household consumption and consumption out of debt are shown as differences between the average value in the control group and the average value in the treatment group, expressed as a percentage of the average treated household consumption in year -1. This new value is centered at 0 in year -1. The control group is created by propensity score matching on a wide array of covariates including education level, age, region of residence, marital status, three lags of income, and industrial sectors of activity. We impose that women in the control group do not have their first child in the three years following the year when their treated counterpart does. Household composition is held fixed as the year of the event. Renter status is defined at the level of the household: a household needs to have zero housing wealth throughout the window of observation to be considered renter.

Table 2: Summary Statistics: Health Shock Sample

	Mean	P10	P50	P90
Panel A : Treated individuals				
	I. Demographics			
Age	38.9	27	38	52
Fraction with tertiary education	0.32	-	-	-
Fraction men	0.42	-	-	-
Fraction married	0.40	-	-	-
	II. Income and Wealth SEK 2003(K)			
Gross earnings	214	53	215	337
Total household gross earnings	353	104	305	630
Consumption expenditures	340	110	288	622
Net wealth	499	-231	87	1,582
Bank holdings	59	0	0	154
Debt	482	1	278	1,151
Sickness benefits received	63	3	23	192
Panel B : Control Group				
	I. Demographics			
Age	38.9	27	38	52
Fraction with tertiary education	0.32	-	-	-
Fraction men	0.42	-	-	-
Fraction married	0.43	-	-	-
	II. Income and Wealth SEK 2003(K)			
Gross earnings	223	42	219	368
Total household gross earnings	368	95	318	664
Consumption expenditures	354	107	295	647
Net wealth	687	-207	150	1,820
Bank holdings	73	0	2	179
Debt	505	0	289	1,201
Sickness benefits received	24	0	0	58

Notes: All values, except sickness benefits, are reported as averages in the year before the shock, expressed in thousands of kronor, in 2003 values. Consumption expenditures, net wealth, bank holdings and debt are expressed at the household level, gross earnings and sickness benefits are those of the treated individual (or its matched counterpart). For sickness benefits, we report the sum of the benefits received the year of the shock and the year after. Higher education refers to college/university less than 2 years, college/university of 2 years and more and PhD.

Table 3: Summary Statistics: Retirement Sample

	Mean	P10	P50	P90
	I. Demographics			
Age	60.5	53	61	65
Fraction with tertiary education	0.26	-	-	-
Fraction men	0.51	-	-	-
Fraction married	0.63	-	-	-
	II. Income and Wealth SEK 2003(K)			
Gross earnings	198	45	183	347
Total household gross earnings	313	69	257	605
Consumption expenditures	383	114	310	687
Net wealth	1,629	-37	873	3,648
Bank holdings	179	0	49	448
Debt	372	0	153	863
Pensions received (0 and 1)	114	0	59	302

Notes: All values, except pensions, are reported as averages in the year before the shock, expressed in thousands of kronor, in 2003 values. Consumption expenditures, net wealth, bank holdings and debt are expressed at the household level, gross earnings and pensions are those of the treated individual. For pensions, we report the sum of the benefits received the year of retirement and the year after. Higher education refers to college/university less than 2 years, college/university of 2 years and more and PhD.

Table 4: Summary Statistics: Parenthood sample

	Mean	P10	P50	P90
Panel A : Treated individuals				
	I. Demographics			
Age	28.5	22	28	35
Fraction with tertiary education	0.49	-	-	-
Fraction married	0.29	-	-	-
	II. Income and Wealth SEK 2003(K)			
Gross earnings	196	23	204	330
Total household gross earnings	428	112	429	708
Consumption expenditures	409	109	327	748
Net wealth	354	-374	6	1,310
Bank holdings	74	0	11	194
Debt	623	6	352	1,611
Share of households buying a car	0.33	-	-	-
Share of households changing house	0.39	-	-	-
Share of homeowners	0.40	-	-	-
Panel B : Control Group				
	I. Demographics			
Age	28.5	22	28	35
Fraction with tertiary education	0.49	-	-	-
Fraction married	0.31	-	-	-
	II. Income and Wealth SEK 2003(K)			
Gross earnings	165	10	169	301
Total household gross earnings	435	125	418	735
Consumption expenditures	442	136	374	759
Net wealth	583	-340	68	1,797
Bank holdings	80	0	4	202
Debt	694	19	456	1,596
Share of households buying a car	0.34	-	-	-
Share of households changing house	0.26	-	-	-
Share of homeowners	0.55	-	-	-

Notes: All values are reported as averages in the year before the shock, expressed in thousands of kronors, in 2003 values. Consumption expenditures, net wealth, bank holdings and debt are expressed at the household level, gross earnings are those of the treated individual. Higher education refers to college/university less than 2 years, college/university of 2 years and more and PhD.

Appendix B Further Details on the Construction of the Registry-Based Measure

Disposable Income Our consumption measure build on two main blocks. One is net-of-tax income from labor and transfers and the other is the flow from asset holdings. Net-of-tax income from labor and transfers are stored in the variable disposable income (DispInk) in LISA. The disposable income measure includes all labor income, both from employment ans self-employment, all public transfers and transfers from collective agreements such as occupational pensions and workers' compensations and annuities. The disposable income measure also includes capital income and student loans. From the sum of gross-of-tax labor income and transfers, Statistics Sweden subtracts the sum of taxes paid.

To avoid double counting we subtract changes in student loans and most of the capital income. Specifically we subtract interest payments from bank account holdings, dividends, gains from selling real estate and negative capital income from paid interest rates. However, the tax paid on these gains are kept in the disposable income measure. The only income capital from capital to still be included are incomes from subletting and lottery wins, primarily such from outside the EEA.

Bank Account Holdings Before 2006 bank account holdings are only reported if their interest payments exceed 100 kr. In 2006 and 2007 all bank account holdings surpassing 10,000 kr are reported, irrespective of the interest payments. Smaller amounts are also reported if their interest payments surpass 100 kr. The regime change in 2006 led to a large increase in bank account holdings: the variable increases on average by about 35,000 kr. In 2005 about 70% of the bank account holdings that can be found in national accounts data can be found in the wealth register. In 2006 the bank account holdings in the wealth data is 97% of the holdings reported in national accounts. To avoid a large drop in the residual consumption measure in 2006 bank account holdings are censored in the same way as in 2005. This leads to a smaller increase in bank account holdings, about 22,000 kr, and to consumption increasing by about 10,000 kr on average. This adjustment improves the correspondence between consumption in national accounts and the residual consumption measure significantly. Using the 2006 censoring rule on the 2005 data does not address the problem that a much larger share of bank account holdings are reported in the data. As interest rates in 2005 was below 1% means that almost all reported holdings surpass 10,000 kr. The 2006 censoring rule is applied for 2007s consumption measure. Using the 2005 bank account censoring rule also for 2006 improves the match between survey consumption in HUT and the residual consumption measure.

KURU data To assign a price on the individual securities found in the KURU data we use closing prices from the last day of trade in December 2000 to 2007. We use data provided by SIX which covers listed securities from about 1,500 different securities exchanges and contributors.[39]

[39]Securities which are not traded on Swedish stock exchanges, but are traded abroad, are reported in their respective currencies. To convert these prices into Swedish kronor we use exchange rates for the last day of trade each year.

This ensures that not just listed securities found on Swedish stock exchanges such as the Stockholm Stock Exchange (SSE) and First North (an exchange for smaller firms and start-ups) are used. In the data we find 40% of all the stocks and mutual funds from the KURU data and 43% of all bonds.[40] When Statistics Sweden constructed the wealth statistics they used price information only on securities listed on the SSE and first North. This means that listed securities which were not traded at either SSE or First North were not included in the wealth statistics.

Table (5) reports some descriptive statistics from the KURU data. The first column show the share of securities held by the individuals in our sample that we find in the asset price data. As mentioned previously we find about 13 percentage points or almost 50% more securities compared to Statistics Sweden. The main explanation for this is that we use data on securities that are not exclusively traded on Swedish stock exchanges. The second column show the share of individual-security cells we can match to our price data. For stocks and funds these are about the same as the numbers in column 1. For bonds, however, we are able to assign a price to almost 90% of all bonds held by the individuals in our sample. In column 3 we compare the total worth of price assigned securities to the value of securities found by Statistics Sweden. Notably we can attach a price to a significantly larger amount of stocks and bonds which translates into a much larger reported stock and bond wealth compared to the original wealth statistics. The sum of stock wealth is almost 35 times higher compared to what is found by Statistics Sweden. However, for funds we only find 8% of the wealth that Statistics Sweden found. Finally, column 4 show the share of individuals with imputed security flows for each type of security. 9.4% of the individuals have imputed consumption expenditure flows from stocks while the corresponding figures for bonds and funds are 2.7 and 47.4% respectively. Note that the imputations are only made for cases where the sum of held security wealth in the wealth statistics is larger compared to what we find.

The significant surge in stock and bond wealth is primarily due to large amounts of foreign traded stocks held by a small number of wealthy individuals. But when comparing our calculations with statistics Sweden's we find about twice as many stock owners and twice as many bond owners (for which the securities can be assigned with a market price). All else being equal, this would however only explain a small share of the large difference in stock and bond wealth found by us and Statistics Sweden.

Outliers Wealthy individuals sometimes make very large changes in their asset positions without them being offset by corresponding changes in other types of assets. We are worried that these are assets we miss out on because of the incomplete coverage of the KURU data. As a result, these non-offset position changes affects aggregate consumption. For instance, in 2006 a single individual sells stocks to a worth of 3 billion kr which is not offset by a comparable change in real estate or bonds. Just one such transaction adds about 500 kr to average consumption. To manage

The currency data is provided by Sveriges Riksbank (the Swedish central bank) and are available for download on its site: `www.riksbank.se/en`.

[40] We have also used data from ThomsonReuters' service DataStream. In this data we found 11% of all stocks and mutual funds and 8% of all bonds. The main explanation seems to be that ISIN codes which are not currently used are also not retained in the data.

Table 5: Coverage ratios and wealth assessments

	Share of ISINs found	Share of individual-security cells found	Security wealth found relative to SCB	Share with imputed security flows
Stocks	0.400	0.433	34.54	0.094
Bonds	0.429	0.886	7.73	0.027
Funds	0.400	0.350	0.08	0.474

Notes: Stocks and funds cannot be separated in the KURU data. Therefore the shares found in column 1 are the same. Data spans 2000-2007.

very large changes in wealth the following rule for outliers is applied: 1. Outliers are defined as individuals with changes in stock or bond wealth surpassing three standard deviations; i.e, extreme outliers. However, some years there are only a few very large position changes taking place which results in a low standard deviation of the change in stocks and bonds. Therefore an absolute rule is also used: the standard deviation must be larger than 500,000 kr for the observation to be considered as an outlier. The number is chosen as this is what the standard error is for those years which have a larger number of large trades that seem to affect the average residual consumption. 2. To make sure that there is no offsetting change in another asset class lagged consumption is used to see if a change in, for instance, stock wealth leads to a similar change in consumption. For very large amounts (above 1.5 million kr) it is likely that the individual has invested in some asset which is not picked up in the data. Therefore an outlier must also have a change in consumption which is at least 75% of the change in either stocks or bonds. When computing consumption, the outliers will have a transaction amount of zero instead. This means that we retain the observation but leave out the large non-offset change in asset holdings. For the whole period 2000-2007 the share of outliers is on average 0.03% or about 2,200 individuals annually. The years 2006 stands out as mentioned. This year 0.14% or 10,734 individuals made large transactions which has an impact on overall consumption. When the asset changes of all individuals in 2006 are considered the average change in bonds is -10,441 kr. When outliers are left out, the change is instead 2,345 kr. All else being equal, excluding these outliers removes more than 12,500 kr from average consumption, or a decrease by about 6.5%. Removing outliers leads to an improved correspondence between national accounts and the residual consumption measure.

Housing Values Housing in Sweden can be divided into two broad categories: owned housing and coop apartments. Below we describe how these two categories of housing are valuated in the wealth registry. We also describe in detail how housing wealth flows are calculated for owned housing and coop apartments.

Owned housing consists of several types of property, mainly primary residences (stand-alone houses, terrace buildings, bungalows etc.), vacation homes, apartment buildings, commercial property and farms. Their value is assessed for taxation purpose every sixth year. When tax values are assessed the tax administration collects characteristics of the property such as living space, size of the lot, construction year, when larger refurbishments were made etc. The taxation value is

determined by hedonic regressions run on actual sales prices using the collected characteristics as explanatory variables. Tax values also consider the market value of properties in the neighborhood. For primary housing the tax administration divide Sweden into about 9000 different "value areas". A simplified tax value assessment is also made three years after the assessment described above. This is also made every sixth year which means that the tax vale of a house is assessed every third year. The simplified tax value assessment only consider how prices in the value area has evolved the last three years, it does not collect any property specific characteristics.

The tax value is supposed to reflect 75% of the market value two years back. To determine market values statistics Sweden calculates municipal-level sales price ratios by dividing sales prices by taxation values for sold houses. For other categories of owned housing such as apartment buildings and commercial property this is done at the county level as there are often not a sufficient amount of transactions taking place in a single municipality during a year.

The other broad category of housing is coop apartments. Coop apartments are bought and sold just like owned housing but the buyer does not buy the apartment per se, he/she only buys the right to occupy the apartment. The apartment, as the whole apartment building (a coop can consist of several apartment buildings as well), is instead owned by a coop housing association. The coop housing association is run by a board, mostly consisting of individuals living in the building. Data on coop wealth is available 2004-2007.

Coop apartment wealth is determined on the apartment building level in the same way as the market value of primary housing determined. Statistics Sweden then divides the worth of the coop apartment building and divides it by the number of households residing in the apartment building. This overstates coop wealth for households living in small apartments and understates coop wealth for households living in large apartments (relative to the average apartment size of the building). In addition, individuals who rent apartments in coops wrongly get coop wealth assigned to them. These different measurement errors are caused by the lack of an apartment registry which assigns each individual to a specific apartment. Such a register is available in Sweden from 2013.

Consumption flows from housing Whenever possible we use data on actual housing transactions to determine the flow from housing wealth to consumption. For owned housing we use data on actual transactions from the housing price registry (Fastighetsprisregistret) to calculate these flows. From 2004 we observe all parties involved in a transaction and can calculate ownership shares depending on the number of buying or selling parties involved in each transaction to assign the appropriate transaction value to each individual. The ownership shares assume that each involved party owns or buys an equal share. According to the wealth register, which records ownership shares for each type of real estate (primary housing, second homes, farming properties and apartment buildings), 96.1% of all jointly owned primary housing is owned at 50%. This suggests that most ownership is split equally among the owners.

Before 2004, only one buyer and seller per transaction is recorded. Who is recorded as buyer and seller is random. For individuals of the same households this is not an issue. But for transactions

involving two or more households the censoring of buying and selling parties results in housing transactions on the individual level being unrecorded. For instance, consider two siblings selling a jointly owned house. Only sibling A is recorded as a seller while sibling B is not, which means that sibling B will have no recorded revenue from selling in the house price registry. Sibling A's share of the revenue is calculated as the total revenue times the ownership share, from the wealth registry, he/she has for the particular type of real estate. According to 2004's statistics, 23.5% of all sales involved two selling households or more, and the average number of households involved in the sales which involves more than one selling household are 2.31. This means that $23.5\times(1-1/2.31) = 13.3\%$ of all transactions on average go undetected. 128,480 individuals or 1.7% of all individuals sold housing in 2004. Given that the 2004 figures are similar to those in previous years we can expect about $.0.017 \times 13.3 = 0.2\%$ of our sample to have missing income from selling housing during 2000 to 2003.

For coops there is no data such as the housing price registry. This means that the amount of coop housing bought and sold cannot be detected in the same way as with owned housing. For individuals who sell their property, the purchase transaction is fully observed (date, price, etc.). But for individuals who do not, we need to impute these transactions. Note, however, that because we observe the stock of coop wealth in each period, we can restrict the sample of individuals for which these transactions need to be imputed. The purchase transactions in time t are only needed to be imputed for individuals who are observed with some positive stock of co-op wealth in both time t and $t-1$. Individuals whose stock of coop wealth is always equal to zero cannot have been buying at any time in our period of observation, so for them there will be no transaction flows. For individuals whose stock of coop wealth goes from zero to a positive amount at time t, we use the figure reported in the wealth registry. Finally, for individuals whose coop wealth goes from positive to zero at time t, the change in coop wealth is just equal to the amount of the sale which we observe.

To impute purchase transactions for those where the transaction is not observed, we use a statistical model, estimated on the sample for which we observe a purchase. Concretely, we create a sample based on all the purchases observed (i.e. for individuals who did resell). Then we create a balanced panel for years 2005-2007 of all the individuals for which we have a purchase record and who had a positive stock of coop wealth in both time t and $t-1$ (note that individuals will be included in this sample even if their purchase record falls outside the 2005-2007 window). For this sample we estimate a statistical model on the probability of buying a coop:

$$Pr[\text{BUY}_{it}] = X'_{it}\beta + \varepsilon_{it}. \tag{5}$$

The variables used are age-bins, income, household size and change in household size, county or municipality, indicator for moving address in year t (and/or t+1), change in coop wealth (in levels), indicator for selling a coop, indicator for selling an owned house, interaction between selling a coop and change in coop wealth and change in debt and year fixed effects.

Then we estimate a model of the transaction amount for the observed coop purchases (Y_{it}):

$$Y_{it} = X'_{it}\beta + \varepsilon_{it}. \tag{6}$$

The same right-hand side variables as above are used.

Based on these two models, we predict the probability of buying a coop in the sample of individuals observed with positive stock of coop wealth in both time t and $t-1$ in the period 2005-2007. When this probability is large enough we impute the amount of the purchase based on model 6.

Appendix C Comparing Consumption in National Accounts, Survey, and Registry Data

This Appendix reviews how consumption is computed in the Swedish national accounts for the years 2000-2007. It also points out significant differences between what is viewed as consumption in national accounts, in household budget surveys and in a residual consumption measure derived from individual data on income and wealth.

C.1 Computing consumption in national accounts

Consumption in national accounts comprise all consumption by the residents of a country. This means that domestic consumption is only the consumption made by the residents of the country while consumption made by foreigners is left out. Foreigners include both tourist but also asylum seekers who have not yet been granted residence permit. Consumption made by these groups are booked as consumption by foreigners and is a part of the country's exports. Consequently consumption made by the country's residents abroad (as tourists or diplomatic personnel, for instance) is therefore also included in the national account's consumption measure. Consumption in national accounts are meant to include both white market and black market consumption. Black market consumption includes consumption that is not taxed for, when it should, but also activities such as consumption of illegal drugs and money spent on prostitution services (which are illegal to buy in Sweden).

Statistics Sweden use various sources to compute consumption for the national accounts (NA).[41] The household budget survey HUT is one of the main inputs, and the HUT of 1995/1996 was used as a benchmark for NA consumption up until 2004 when the HUT wave of 2003/2004 was meant to replace it. However, HUT, like many other household budget surveys, has experienced higher and higher attrition bias over time and it has been found by Statistics Sweden to severely underestimate various consumption categories. Consequently Statistics Sweden has begun to use more information from business sector to adjust for the decreasing quality of the HUT survey. The perhaps most important data is sales taxes (VAT) and turnover statistics from the retail sector. From 2004 and onward, HUT is, together with surveys to the retail sector, mostly used to spread total consumption expenditure between the different COICOP (classification of individual consumption by purpose) categories in which consumption is divided between. Statistics Sweden also use the vehicle registry to compute consumption of vehicles and special sales statistics from the three major operators on the Swedish grocery market to compute food consumption.

[41] The details in this review is based on two documents provided by Statistics Sweden: "Sweden GNI Inventory", (2016) and "Hushllssektorn i nationalrkenskaperna", (2014). http://www.scb.se/contentassets/c89bb85e14184e92a4d5e4eec5ce4b98/sweden-gni-inventory-2016_public_rev_oct2016.pdf and http://www.scb.se/Statistik/_Publikationer/NR0103_2014A01_BR_X100BR1401.pdf

C.2 Differences between measures of consumption

Consumption is not always the same thing in NA, household budget surveys and in residual consumption expenditures derived from registry data on income and wealth. Before comparing NA consumption to surveyed consumption and residual consumption from registry data there are special types of consumption that are being treated differently in the three measures. The first is imputed rents, the shadow price of occupying one's own home, which is a part of consumption in NA, but is not a part of consumption in a household budget survey. This can however, be made into a consumption object in a residual consumption measure applying the same techniques used when computing the variable in NA.

The second special category is "larger" renovations, where larger is referred to such refurbishing not normally undertaken by tenants. In NA this is viewed as investment and household budget surveys do not normally ask for expenses on larger renovations. In a residual consumption expenditure measure it is not possible to separate expenses on refurbishing from other expenses. Therefore renovation expenses are treated as consumption which, all things being equal, makes registry data consumption larger compared to consumption in NA and surveyed consumption.[42]

The third category is replacement of larger home appliances, such as stoves, dishwashers etc., in individuals' dwellings. This is treated as renovations, i.e. investments, in NA. Consumption of such goods are only incorporated into NA's consumption statistics if a household buys an additional stove or dishwasher.

The fourth category is insurance expenses and expenses on lotteries and gambling. In NA these are measured as the net of what is paid for insurance, lotteries and gambling and what is received by the households in terms of insurance payments or wins. In a residual consumption expenditures measure insurance premiums and lottery tickets are accounted for as they are not considered as saving but insurance payments and wins are in most cases exempt from tax which means that they are not recorded as income. In a sense, insurance payments and wins can in registry data be viewed as unaccounted gifts. Such insurance payments may reduce debts and increase savings but in such as case residual consumption expenditure decreases.

The fifth category is interest payments on liabilities. In NA these are not measured as consumption, rather households consume a service which is the banks' provision of liquidity. The price of this provision is the difference between the borrowing and lending interest rate. Therefore this consumption item is larger in the residual consumption expenditure measure compared to NA. In household budget surveys expenses on interest is a separate category from expenses on housing.

Finally, black market and illegal consumption is estimated by the NA division at Statistics Sweden. This type of consumption is recored in the residual measure as long as the income that is used to pay for it is accounted for. Black market expenses on legal goods is accounted for in household budget surveys but consumption of goods such as illegal drugs and prostitution services

[42] During 2004-2005 there was a special tax deduction for renovation expenses given to individuals who hired construction workers to undertake larger renovations of their properties (the ROT deduction). However, this data is not available to us but is recorded by Statistics Sweden.

are not found in household budget surveys.