Measurement Error in Imputed Consumption

Scott Baker* Lorenz Kueng† Steffen Meyer‡ Michaela Pagel§

May 15, 2018

preliminary and incomplete

Abstract

Because of limitations in survey-based measures of household consumption, a growing literature uses an alternative measure of consumption commonly referred to as 'imputed consumption'. This approach utilizes annual snapshots of household income and wealth from administrative tax registries to calculate household consumption as the residual of the household budget constraint. In this paper we use transaction-level retail investment data to assess the measurement error that can result in imputed consumption due to intra-year changes in asset values and composition. We show that substantial discrepancies between imputed and actual spending can arise due to trading costs, asset distributions, variable trade timing, and volatile asset prices between two annual snapshots. While these errors tend to be quantitatively small and centered around zero on average, we demonstrate that they vary across individuals of different types and income levels and are highly correlated with the business cycle.
JEL Codes: D12, D14, G11, H31.

Keywords: Registry Data, Measurement Error, Retail Investors, Household Portfolios.

*Northwestern University, Kellogg; s-baker@kellogg.northwestern.edu
†Northwestern University, Kellogg, and NBER; l-kueng@northwestern.edu
‡Leibniz Universität Hannover; meyer@gif.uni-hannover.de
§Columbia Business School, NBER, and CPER; mpagel@columbia.edu

1 Introduction

Economists need accurate measures of spending to analyze consumption and saving behavior, to study aggregate fluctuations in consumption, and for constructing measures of economic well-being, such as inequality or poverty. It is therefore important to understand how accurate these measurement of consumer spending really are.

Consumption data without substantial measurement error are difficult to come by, especially when seeking detailed spending data over a long period of time. Consumption surveys use paper or phone interviews to ask stylized questions on spending in a few broad consumption good categories over a particular recall period (e.g., the Interview Survey of the Consumer Expenditure Survey, CEX). Alternatively, households are asked to keep track of recurrent expenditures, such as groceries, for a short period of time in a diary (e.g., the Diary Survey of the CEX). For obvious reasons, survey respondents may have difficulties in recalling past purchases and have little incentive to answer the questions accurately. For instance, respondents may not understand the wording of the questions, may behave differently in practice, or may simply forget some past purchase transactions, or may strategically underreport consumption to avoid more detailed follow-up questions (Parker and Souleles 2017). Moreover, such measurement error or noise in the data generated by surveys that simply ask about past purchases can increase with the length of the recall period (de Nicola and Giné 2014). Additionally, surveys can produce data with systematic biases if respondents have justification bias, concerns about surveyors sharing the information, or stigma about their consumption habits (Karlan and Zinman 2008).[1]

A growing recent literature develops and utilizes an alternative measure of consumption based on highly-detailed administrative data, often referred to as 'imputed consumption', which avoids many of the problems with standard survey-based data.[2] Imputed consumption is constructed as a residual from a household's budget constraint, the part of total income that was not saved or invested. This approach imposes heavy data requirements on the measurement exercise because the researcher needs comprehensive measures of income as well as comprehensive asset holdings and asset price data.

[1] A large existing literature has documented basic problems with survey-based measures of consumption (see e.g., Pistaferri 2015). Ahmed et al. (2006) for example compare two measurements for the same set of households and find that recall food consumption data, which is the basis of a great deal of empirical work, suffers from considerable measurement error, while diaries records are found to be somewhat more accurate. Other work has compared consumption measures across different surveys or across different waves of the same survey (e.g., Pudney 2008, Bound et al. 2001). The measurement error in household-level consumption data, and the difficulty of estimating nonlinear models in the presence of such error, have led some to call for abandoning Euler equation estimation altogether (Carroll 2001).

[2] Important studies include Browning and Leth-Petersen (2003), Koijen et al. (2014), Bach et al. (2015), Fagereng et al. (2016), Sodini et al. (2016), Fagereng et al. (2016), Fagereng and Halvorsen (2017), Fagereng et al. (2017), and Kolsrud et al. (2018).

Sweden, Norway, and Denmark collect most of the required information at an annual level as part of their tax registries (or they collected this data previously in the case of Sweden). In the Swedish case, the tax registry is most comprehensive, containing data on every stock, bond, mutual fund, and bank account each household owns at the end of the year. Furthermore, data on capital gains and dividends is available for some years. Home-ownership and household permanent address can be tracked via the housing registry, and the data also contains information on labor and financial income and transfers. The Norwegian and Danish versions of this data are also highly detailed, though they lack some of this information for parts of the sample periods (for instance, data on actual assets, capital gains, or government transfers); see Appendix A.

The registry-based or imputed consumption approach thus attempts to measure all consumer spending on , services and non-durable and durable goods at an annual frequency. Sweden also runs a standard Household Budget Survey that can be matched with the households in the registry data. This setup allows for the comparison of registry-imputed and survey-based measures of consumption. Koijen et al. (2014) uncover significant discrepancies between registry- and survey-based consumption measures that increase in income and wealth.[3]

In this paper, we use high-frequency transaction-level retail investment data on trades and portfolio holdings for approximately 113,000 retail investors in Germany from 2004 to 2015 to document potential measurement error affecting imputed consumption.[4]

More specifically, discrepancies between imputed and actual consumption may arise whenever investors buy and sell assets at different points within the year, incur trading costs, or are paid dividends or other asset distributions that differ from those of the market at large. That is, miscalculating asset growth experienced by a household will mistakenly attribute changes in asset values to the household consuming more or less than it actually does. As one example, it is well-known that some U.S. retail investors have substantial annual turnover, have significantly different levels of returns relative to the market, and may incur substantial trading costs that can eat up the entire historical equity premium (Barber and Odean 2000). More recently, similar results have been documented for Scandinavian and German investors (e.g., Bach et al. 2015, Koestner et al. 2017).

[3] While the mean and median of the consumption distribution are similar, the survey understates the consumption of wealthy and high-income households, while slightly overstating consumption of the poorest quintile of households. Moreover, Koijen et al. (2014) show that registry-based consumption is sensitive to an accurate imputation of returns that households are earning on their assets. The authors show that incorrectly applying a broad total return measure to a household's financial asset holdings leads to substantial deviations from the properly imputed registry measure and that these discrepancies are increasing in wealth.

[4] Our transaction-based approach to constructing more accurate household-level investment returns mirrors a number of recent papers that take similar approaches using transaction-based spending data to precisely map out household spending and consumption. See for example Gelman et al. (2014), Kuchler and Pagel (2015), Baker (2017), Olafsson and Pagel (2016), or Kueng (2018), for a discussion of the advantages and disadvantages of this type of data.

Our study accomplishes the following goals. First, we demonstrate that imputing consumption from annual portfolio snapshots using a variety of methodologies leads to substantial measurement error, both in absolute terms and relative to household income. Second, we show that these imputation errors, both across households and within household over time, are not purely 'classical' measurement error, i.e., uncorrelated with the outcome variable or the regression error term. In particular, we show that these imputation errors on financial consumption are correlated with household-level financial characteristics like income, wealth, and trading behavior. Third, we show that the measurement error in imputed consumption from investment portfolios is correlated with aggregate economic variables such as home prices, stock market returns, and GDP growth.

Overall, this paper demonstrates that many assumptions in consumption imputation are not innocuous – substantial measurement error and even biased average values can result from an annual snapshot approach – especially for households with higher levels of income, wealth, and financial market trading activity.

2 Transaction-Level Investment Data

To investigate the extent of measurement error in imputed consumption data, we use a unique panel data set that tracks the daily trading of approximately 113,000 private investors in Germany spanning the years 2004 to 2015. The investment data comes from a large German brokerage firm, with a random sample being drawn from the several hundred thousand clients the brokerage serves. With this data, we can precisely measure each trader's daily activity from his logging into an account to every single trade that he makes. We are also able to identify quasi-automatic trades, such as savings plan transactions. Moreover, trading decisions in our sample are not moderated by any influence from third parties, such as financial advisors.

This data set has been used and discussed in detail in previous studies (e.g., Schmittmann et al. 2014). It consists of a monthly asset position file, a daily transactions file, a file containing bookings to cash accounts, and a file containing investor demographics. The monthly asset position file contains identifiers for the investor, the securities as well as the respective volumes and values in euros. The transactions file contains identifiers for the investors and traded securities, transaction volumes, prices, and dates, as well as information on order types (orders with and without limits). Investors may also hold checking, savings and settlement accounts, with transfers between accounts being common.

In addition to portfolio holdings and trades, we also have the time series of checking, savings, and settlement account transactions as well as balances for all investors. The settlement account is used as the vehicle to execute trades into and out of the portfolio. With the data on all transactions in all accounts, we can infer all transfers within the bank as well as all transfers out of the bank. Transactions are automatically categorized and labeled, such as wire transfers, ATM withdrawals, or debit card transactions. We measure wealth as the sum of all assets an investor holds plus his checking, savings, and settlement account balances. The file with cash bookings contains bookings into the investor's cash accounts.

With respect to investor demographics, we observe a customer's age, gender, geographic location (postal code), marital status, employment status, and the length of relationship with the brokerage. We also see self-reported measures of customer income and wealth – which we compare with their corresponding account-based measures – as well as self-reported risk-preferences and whether or not they hold a doctoral degree.

Sample selection We restrict the user sample to individuals who the bank flags as likely using this bank as their only banking institution, i.e., individuals that do not hold other bank relationships. This flag is based on an account being dedicated as containing the tax-free allowance (i.e., Freistellungsauftrag), account activity such as income and spending transactions, and the linking of other accounts via the bank's financial aggregation facility. Nevertheless, our measure of financial/investment consumption may be biased if investors perform offsetting trades in other investment accounts. Therefore, we undertake multiple steps to ensure that we restrict the sample to individuals who primarily utilize only this bank.

It is important to note that the bank, from which we obtained the data, is not only one of the largest retail brokerages in Germany but also a multi-service retail financial institution offering checking and savings account services as well as overdraft facilities, credit cards, mortgage and auto credits, insurance, and retirement savings vehicles. Because cash payments are still prevalent in Germany, the bank also offers a dense network of ATMs. Additionally, in Germany, holding multiple checking accounts is discouraged as this hurts individual credit scores. Needless to say, the bank offers all common online transfer facilities and automatic checking account transactions.

Representativeness Table 1 displays summary statistics for our sample including mean and median checking and savings account balances. The table also shows annual portfolio turnover as defined by Barber and Odean (2000), and the annual sum of all trading fees.

[Table 1 about here. Note: The tables currently only use a 20% random sample of our working data and will be updated shortly.]

Our sample is not representative for the German population as a whole; less than half of Germans are invested in equities, either directly or indirectly. However, it is a relatively representative sample of self-directed retail investors in Germany. We believe that this portion of the population is of particular interest and importance when thinking about measurement error in imputed consumption precisely because these individuals are the most likely to be performing asset trades and seeing the greatest amount of heterogeneity in asset price growth. Households without substantial volatile assets (e.g., equity portfolios) are therefore less likely to experience the types of imputation errors that we describe in our sample.

Our sample does not comprise the entirety of the bank's customer base, but a 10% sample of all customers. The bank did not pick the sample of retail investors by trading frequency but rather chose a random subsample of all bank users who held a brokerage account. In that sense, our sample is representative for individuals in Germany holding an investment portfolio at a major bank.

Table 1 displays demographic statistics as well as self-reported risk aversion and income of our investors. The average age of investors is 44 and 89% of our sample is male. Brokerage clients are generally expected (Cole et al. 2012) and found to be more financially sophisticated than the overall population (Dorn and Huberman 2005). This is also true for our sample: 3% of our investors hold a doctoral degree, which is higher than average in the German population of 1.1% (German Federal Bureau of Statistics 2011).

Investors own portfolios that are worth 64,000 euros, on average. These descriptive statistics are comparable to those reported by household finance studies using U.S. data (Barber and Odean 2000). In addition, we compare average portfolio values to official statistics in Germany. The German central bank reports the average portfolio value of a German stock market investors to be around 48,000 euros (Deutsche Bundesbank 2013). This value seems comparable to the average values we observe in our sample. Additionally, we compare portfolio holdings to self-reported gross annual household incomes for those investors who reported these data. Since income is reported in several ranges, we use the midpoint of each range as a proxy for investor income. The mean ratio of the average portfolio value (over the entire sample period) to annual income is 1.3. For comparison, the ratio of total financial assets to gross household income in the German population is about 1.1 (see German Federal Bureau of Statistics 2008 and Deutsche Bundesbank 2013).

Finally, in Appendix A we discuss annual snapshots based on administrative tax records that are used by a growing literature to impute unobserved annual consumer spending. These records predominantly come from three Scandinavian countries: Sweden, Norway, and Denmark.

3 Calculating Actual and Imputed Financial Consumption

In this section, we discuss our method of computing 'financial consumption' and sources of measurement error that may arise from incomplete information about various aspects of the financial portfolio. By 'financial consumption', we mean a type of active savings or dis-savings flowing to or from investment accounts, defined precisely below. With this formulation, 'financial consumption' need not be actually consumed by the household; these active savings may simply flow from a checking account to an investment account or be observed elsewhere on the balance sheet and would not induce error if measuring consumption directly through spending transactions. However, mismeasurement of 'financial consumption' *will* directly impact imputed consumption derived as a residual of the budget constraint because it implies a mismeasurement of the amount of savings being done.

3.1 Calculating Actual Financial Consumption

3.1.1 The Period Budget Constraint

The goal of this paper is to assess the contribution of a household's financial portfolio to imputed consumption, which is based on administrative records of income taxes and annual snapshots of wealth, and potential measurement error resulting from the need to impute unobserved elements of the household's portfolio.

Financial assets' contribution to imputed consumption play an important role in the residual methodology, especially for the upper tail of the income and wealth distribution where financial asset holdings are a large component of overall wealth. For this group in particular, mismeasurement of financial returns can be potentially large relative to measured income. This is both because financial asset prices and returns can be volatile relative to non-financial asset prices (e.g., home prices) or household income and also because households will buy and sell securities much more often than they will sell homes or change jobs. So, mismeasurement of returns can pose a persistent problem for imputing consumption in every year that an individual or a household is in a sample, even when remaining at the same job and living in the same home.

There are three main approaches to expressing consumption expenditures as a residual of the budget constraint: the flow approach, the stock approach, and the return approach. Denote A_t the stock of financial assets (including cash as well as debt as negative elements) of a household's portfolio at date t (which we define to be the end of period t over which we measure flows) and P_t the corresponding vector of prices (equal to 1 for cash, the numeraire). The flow measure of imputed consumption equates expenditures and revenues. Consumption expenditures C_t during period t can be financed in three ways: (i) with capital income net of capital income and wealth taxes; (ii) with dis-saving (net asset sales) net of capital gains taxes as well as trading and other fees; and (iii) with earned income and transfers net of earned income taxes,

$$C_t = CI_t^{net} - S_t^{net} + E_t^{net}$$

$$= CF_t' A_{t-1} - P_t' \Delta A_t - Fees_t + (E_t - Taxes_t). \tag{1}$$

- $CI_t^{net} = CF_t' A_{t-1} - Tax_t^{CI} - Tax_t^{W}$ is capital income (i.e., cash flows from assets) net of capital income taxes Tax_t^{CI} (if treated differently than earned income taxes) and wealth taxes Tax_t^{W}. CF_t is the vector of cash flows during period t, which include mainly interest payments and dividends.

- $S_t^{net} = P_t' \Delta A_t + Tax_t^{CG} + Fees_t$ is active saving (dissaving if $P_t' \Delta A_t) < 0$) net of capital gains taxes Tax_t^{CG} and fees (both trading fees and other fees, such as annual account fees).

- $E_t^{net} = E_t - Tax_t^{E}$ is earned income E_t, including transfers (government transfers, gifts, etc.), net of earned income taxes Tax_t^{E}. The second line simplifies notation by collecting tax terms, $Taxes_t = Tax_t^{E} + Tax_t^{CI} + Tax_t^{W} + Tax_t^{CG}$.

Budget constraint (1) for period t relates to the different administrative tax records that the residual consumption imputation method can ideally access: personal income tax records and personal wealth tax records. Financial fees on the other hand are typically not observed in tax records.

The second expression of imputed consumption uses a stock approach (see e.g., Eika et al. 2017). Defining financial wealth $W_t = P_t' A_t$, we can decompose changes in financial wealth over period t into active saving and passive capital gains, $\Delta W_t = P_t' \Delta A_t + \Delta P_t' A_{t-1}$. Substituting for active dis-savings $-P_t' \Delta A_t$ in equation (1), we obtain

$$C_t = CF_t' A_{t-1} + \Delta P_t' A_{t-1} - \Delta W_t - Fees_t + (E_t - Taxes_t). \tag{2}$$

The third approach expresses the budget constraint using return notation,

$$C_t = r_t^p \cdot W_{t-1} - \Delta W_t - Fees_t + (E_t - Taxes_t) \tag{3}$$

with portfolio return $r_t^p = \sum_{j \in A_{t-1}} \frac{CF_{jt} + \Delta P_{jt}}{P_{j,t-1}} \frac{P_{j,t-1} A_{j,t-1}}{W_{t-1}}$.

3.1.2 Accounting for Intra-Year Trading

The consumption flow expressed as a residual of the budget constraint above is accurate only between two trading dates but typically not for an entire year. Given that prices fluctuate during the year and individuals trade securities during the year, equations (1)-(3) are incompletely specified. Both realized and unrealized capital gains throughout the year need to be fully taken into account.

Suppose an individual makes N_t trades in year t. Denote t_0 January 1 of year t (which is a trading holiday) and t_{N_t} December 31 of year t when the individual makes a last trade or when the administrative records are consolidated to assess taxes. Hence, there are $N_t + 1$ dates $n = 0, 1, ..., N_t$ that we need to keep track of to accurately impute *annual* consumption expenditure flows starting from period budget constraint (1).

Let X_{t_n} denote the flow of X in period t_n between trading dates t_{n-1} and t_n and $X_t = \sum_{n=1}^{N_t} X_{t_n}$ the annual flow of X from January 1 to December 31 of year t. Using this notation, annual consumption expenditures in year t can be written using equation (1) as

$$C_t = \sum_{n=1}^{N_t} \left[CF'_{t_n} A_{t_{n-1}} - P'_{t_n} \Delta A_{t_n} \right] - Fees_t + (E_t - Taxes_t) \tag{4}$$

The first term, $\sum_n CF'_{t_n} A_{t_{n-1}}$, is the annual flow of dividends and interest. The second term, $\sum_n P'_{t_n} \Delta A_{t_n}$, is the net realized capital gains throughout the year. Note that prices refer not only to the end-of-period security price but the correct price at which the asset was actually bought or sold (e.g., a security contained within A_{t_N} may not have been purchased at price P_{t_N} or $P_{t_0} = P_{t-1_N}$).

Relative to this means of measurement in equation (4), using equation (1) to impute consumption from two annual snapshots will cause measurement error in two ways. First, if P_{t_N} is used to measure the purchase price of a security rather than the true but unobserved purchase prices P_{t_n}, active savings will be measured incorrectly. Second, the annual snapshot approach will entirely miss gains or losses from intermediate trades conducted during the year that are partially netted out, e.g., if $\sum_{n=1}^{N_t} |A_{t_n} - A_{t_{n-1}}| \neq A_{t_{N_t}} - A_{t_0}$. Both of these sources of errors will bias imputed consumption either upwards or downwards, depending on how security prices P_{t_n} vary through the year, unless the research observes active dis-saving flows $-P_{t_n} \Delta A_{t_n}$ directly, say because they are recorded in the capital gains and loss tax registry.[5]

3.1.3 Contribution of Investment Accounts to Imputed Consumption

To focus on the role of financial investment accounts for measures of imputed consumption, we define a financial portfolio's contribution to consumption as

$$FinCons_t \equiv C_t - RealCons_t - CashCons_t - (E_t - Taxes_t) \tag{5}$$

[5] Note that new purchases and hence the initial cost basis of an asset position will not generate a gain or loss and hence will not show up in the capital gains and loss tax registry. The cost basis might be available from the wealth tax registry.

$$= \sum_{n=1}^{N_t} \left[CF'_{t_n} A_{t_{n-1}} - P'_{t_n} \Delta A_{t_n} \right] - Fees_t$$
$$= \sum_{i=1}^{N_t} \left[CF'_{t_n} A_{t_{n-1}} + \Delta P'_{t_n} A_{t_{n-1}} - \Delta W_{t_n} \right] - Fees_t.$$

The three items we exclude from the analysis are the contribution of net earnings $E_t - Taxes_t$ to consumption (which is typically well measured in income tax records) and two types of assets from the vector A, the contribution of real assets, $RealCons_t$ (e.g., owner-occupied housing), and the contribution of cash, $CashCons_t$ (e.g., currency and checking accounts), both of which we do not observe in our data.[6]

While we do not observe these other aspects of household income and consumption, we think that it is relatively unlikely that measurement error in income or real assets would vary negatively with measurement error in financial consumption. Without such a negative correlation, $FinCons_t$ will be linked directly with consumption, so the amount of error we measure in $FinCons_t$ will be translated directly into error that affects the measurement of total consumption.

3.2 Imputing Financial Consumption

3.2.1 Potential Sources of Measurement Error

We now simulate how different issues of missing financial data that often occur with administrative tax records affect imputed consumption. Measurement error in this type of consumption imputation can stem from a number of sources. Compared to a method of constructing a measure of actual financial consumption, these errors range from generally less consequential, like neglecting trading fees or dividend and interest payments, to much more substantial errors based on mismeasuring portfolio composition or neglecting portfolio growth altogether. We seek to understand just how large the potential measurement error might be in each case.

Broadly speaking, these sources of potential error are:

(a) *Missing or incomplete trading fees* – If researchers exclude fees from their calculation of financial consumption ($FinCons_t$), they will *overstate* financial consumption by implicitly assuming that $Fees_t = 0$.

(b) *Missing or incomplete cash flows from assets* – If researchers exclude dividends and interest income from their calculation of financial consumption, they will *understate* financial consumption by implicitly assuming that $CF_{t_n} = 0$. Even when a researcher knows the exact

[6] We abuse notation slightly here as the vectors A, P, and CF now only contain a subset of the elements of the corresponding vectors in the previous equations.

portfolio composition at year-end, if the timing of asset purchases within a year is unknown, researchers may be unable to recover whether an individual would have received a particular cash flow. For example, whether a stock was purchased cum- or ex-dividend.

(c) *Missing intra-year gains or losses* – Researchers with annual snapshots reporting only actual portfolio holdings (and prices) may exclude trading conducted during the year that was 'netted out' in the same year (e.g., buying 10 shares of firm 'X' at $10 in February and selling them 6 months later for $15 each). That is, $A_{t_{N_t}} - A_{t_0}$) will be an accurate portrayal of *net* asset holding changes, but will miss intermediate trades conducted during the year (e.g., if $\sum_{n=1}^{N_t} |A_{t_n} - A_{t_{n-1}}| \neq A_{t_{N_t}} - A_{t_0}$) and thus bias imputed consumption either upwards or downwards depending on how security prices P_{t_n} vary through the year.

(d) *Price errors in intra-year portfolio changes* – Researchers with annual snapshots that only report quantities of securities being held will mis-measure the price at which those securities were purchased or sold. To the extent that securities vary in price throughout the year, this will lead to errors in the measurement of $FinCons_t$. Similarly to above, this means of measuring asset changes (i.e, using $A_{t_{N_t}} - A_{t_0}$ will obscure intra-year gross changes in asset holdings. In addition, P_{t_N} will be mis-measured even when calculating the impact of *net* asset changes on $FinCons_t$

(e) *Incorrect assumptions about portfolio composition and returns* – If researchers only observe the total portfolio value (e.g., they observe $P'A$ rather than P and A separately), they will be unable to determine that actual annual growth of a given portfolio. Researchers may attribute a broad market return to the portfolio, neglecting any variation across households in risk preferences, etc. Again, here both P'_t and $(A_t - A_{t-1})$ will be mis-measured to an even greater extent, depending on how significantly an individual's portfolio differs from a broad market index.

(f) *Neglecting portfolio growth altogether* – Potentially the largest source of measurement error may be to neglect any growth altogether, taking the difference between year-end portfolio balances ΔW_t as being equal to $-FinCons_t$

3.2.2 Illustrative Imputation Methods

For the purposes of this paper, we construct six different measures of imputed financial consumption, denoted $\widehat{FinCons}$ and listed below, corresponding roughly to the sources of measurement error identified above. We then define the imputation error ϵ by comparing measured financial contribution to the portfolio's true contribution to consumption, $FinCons$, which might not be available to the researcher,

$$\epsilon_t = \widehat{FinCons}_t - FinCons_t = \widehat{C}_t - C_t. \tag{6}$$

Each imputation method is not aimed at precisely mimicking procedures taken in any particular paper, but attempt to more broadly span the types of imputation that have been seen in the literature to this point.

Our versions of imputed financial consumption are created as follows:

(i) *No fees* – For this method of imputation, we assume that the researcher has sufficient information on trading behavior to calculate a given portfolio's true growth, purchases, sales, and dividend and interest payments. However, the researcher is unable to observe trading and other portfolio fees and thus *overestimates* the amount of financial consumption obtained by the individual. This is the best case scenario for imputed consumption, assuming that researchers have all relevant information from all tax records. Formally, $\widehat{Fees}_t = 0$.

(ii) *No dividends and interest income* – Again, we assume that the researcher is able to construct the actual financial consumption including all elements (including fees) except cash flows from assets. This scenario might occur if capital income taxes are stored in a separate database that researchers cannot access. Because cash flows are strictly positive, researchers thus *underestimate* the true amount of financial consumption being undertaken by the individual if they ignore cash flow distributions from assets. Formally, $\widehat{CF}_t = 0$.

(iii) *No intra-year gross trades* – With this method of imputation, the researcher can only observe the *net* changes in portfolio composition across annual snapshots and the cost basis of the net purchases. That is, the researcher can observe the individual asset holdings that the portfolio contains and also the price at which they were purchased. However, gross changes in equities that are netted out within a year are unobserved. For example, if an individual purchased 12 shares at 230 euros of Volkswagen on March 1 and sold 10 shares at 160 euros on September 1, the researcher would only observe the net 2 that remained at the end of the year at a cost basis of 230 euros.[7] This scenario might occur if financial institutions directly report to the tax authority (so that the cost basis of each security is observed), but the researcher does not observe realized capital gains and losses, which might be stored in some other administrative tax database (e.g., different tax records for wealth, capital gains, and personal income taxes). Thus, for frequent traders, the researcher will measure trading gains and losses within the year with error. Formally, $\sum_{n=1}^{N_t} |A_{t_n} - A_{t_{n-1}}| \neq A_{t_{N_t}} - A_{t_0}$ and intra-year gross asset positions $\hat{A}_{t_n}$ are not observed.

[7] We assume individuals try to lower tax consequences by using a 'first in, first out' (FIFO) approach.

(iv) *Mid-year prices instead of cost basis* – Similar to the previous method, the researcher can observe year-end individual asset positions held by the individual and year-end prices. However, the researcher is now unaware of the price at which a share was purchased (i.e., the cost basis). Again, any intra-year trading that is netted out (in terms of the *number* of shares, not the *gains* or *losses*) by the end of the year will be unobserved by the researcher. In addition, to the extent that individual securities (stocks, bonds, mutual funds, etc) vary in price throughout the year, the ability to only observe changes in the number of securities will introduce error. This scenario might occur if researchers only have access to wealth tax records, but not the capital gains tax records which also contain the cost basis. Here, a sensible approach is to attribute change in securities held to have occurred mid-year on June 30th at the prevailing price, $P_{t_{6/30}}$, which the researcher can link in from an external source such as Thomson Reuters Datastream. Formally, $\sum_{n=1}^{N_t} \widehat{P'_{t_n}\Delta A_{t_n}} = P'_{t_{6/30}}(A_{t_N} - A_{t_0})$.

(v) *Using market returns, but no individual securities observed* – This method of imputation is simply a less precise version of (iv). Here, the researcher is able to see only aggregate portfolio values at year end, e.g., the value of stocks, bonds, etc. The researcher assumes that the household holds a portfolio with an equity return equal to that seen by the DAX, a standard composite equity index in Germany similar to the S&P 500 Index in the United States. The portfolio is assumed to grow at the same rate as the DAX, with any adjustments to the portfolio being made on June 30th. Thus, the researcher will observe true financial consumption with error if the individual's portfolio differs significantly from broad market holdings. Formally, $\hat{r}_t^p = r_t^{DAX}$ in equation 3.

(vi) *'Raw' portfolio differences* – For the most basic imputation of financial consumption, the researcher looks only at changes in portfolio size between annual snapshots, disregarding any compositional change or growth. That is, the total financial portfolio balance at the end of year t is subtracted from the total financial portfolio balance at the end of year $t-1$ to obtain the imputed financial consumption during year t. Formally, $\widehat{CF'_t A_t} = \Delta \widehat{P'_t A_{t-1}} = 0$ such that $\widehat{FinCons_t} = -\Delta W_t$.

4 Assessing Accuracy of Imputed Financial Consumption

Being able to observe an investor's actual or 'true' financial contribution to imputed consumption, in this section we determine whether and how these different means of imputing financial consumption might impact the interpretation of research using these methods. To begin we compare the imputed consumption measures, by portfolio-year, against the actual financial consumption exhibited by that portfolio-year. Because investors in our sample have dramatically different portfolio sizes and income, we scale the difference in actual and imputed financial consumption by average after-tax income over all investor-years, E_i^{net},

$$\frac{\epsilon_{it}}{E_i^{net}}. \tag{7}$$

We also exclude individuals with portfolio sizes under 1,000 euros so that this scaling is not driven by miniscule portfolios.

[Table 2 about here]

4.1 Distribution of Imputation Errors

Table 2 shows means, standard deviations, and percentiles of this relative error measure for each of the different imputation methods that are enumerated in Section 3. The first row shows the relative error if we only lack information on fees. This is the best case scenario for imputed consumption, assuming that researchers have all relevant information from all tax records. As would be expected, errors stemming from neglecting trading fees are uniformly positive (*overstating* actual financial consumption). In general, these errors are not substantial relative to the individuals' average income, with an interquartile range of 1.5%. However, in the tails, we see that trading fees may be over 20% of income for the heaviest traders.

Previous research shows that the amount of trading has a large person fixed effect (e.g., men tend to trade substantially more; see Figure 3 below). Controlling for individual fixed effects in panel data and individual characteristics such as gender in cross-sectional data could partially alleviate this issue. However, trading might still be substantially correlated with the business cycle, say if investors trade more in downturns, which could lead to biased estimates of the cyclicality of consumption. Researchers might therefore try to construct proxies for the annual number of trades, say by counting the number of asset positions that changed across two annual snapshots.

Errors stemming from omitting portfolio cash flow distributions in the second row are uniformly negative (*understating* actual financial consumption). Similar to the first row, these errors are not substantial relative to the individuals' average income (interquartile range of 2%), but disregarding cash flows understates of financial consumption by 10% of income in the 10th percentile.

For the other imputation methods, errors can be both substantially positive and negative. The third row shows that omitting intra-year trading leads to small errors on average with a mean (median) error of 0.4% (0%) of income, and an interquartile range of only 0.3%. However, the errors in boths tails of the distribution are very large, understating financial consumption by more than 60% of household income in the 10th percentile and overstating it by more than 75% in the 90th percentile.

The fourth row shows that when researchers also lack the cost basis at which additional assets were purchased and instead have to use mid-year prices, the tails of the error distribution become heavier. For instance, the interquartile range of the error increases from 0.3% to 5% of income.

Finally, assuming portfolios simply return the average market return as in the fourth row or neglecting returns entirely as in the fifth row has extremely large effects on imputed financial consumption across the error distribution. These errors are largely driven by a combination of factors, including investors holding portfolios that differ substantially from (efficient) market portfolios and by realized returns over the specific sample period, both for the overall market index and for the the cross-section of individual asset returns. Realized returns can be highly correlated with other measures that researchers want to study in relation to consumption, such as individual income or the business cycle, this measure of imputed consumption could result in significant biases.

[Figure 1 about here]

In Figure 1, we display plots of actual (or 'true') financial consumption on the horizontal axis against financial consumption imputed according to our six different procedures on the vertical axis. Both actual and imputed financial consumption are measured in nominal euros at a portfolio-year level with zero measurement error for a given portfolio-year being represented by a point anywhere along the 45-degree line. The vertical distance between any point and the 45-degree line represents the imputation error, measured in euros. For display purposes in these figures, we censor financial consumption at the 1% and 99% level. With this level of censoring, we find that 'true' annual financial consumption ranges between approximately 150,000 euros and -150,000 euros. That is, households actively are investing or withdrawing up to approximately 150,000 euros.

The top-left panel displays the relationship between actual financial consumption and financial consumption omitting fees, i.e., method (i). As seen in the summary statistics, every deviation from the 45-degree line is above the line, implying that the true level of financial consumption is weakly less than what is imputed while missing any trading fees. In contrast, the top-right panel shows the relationship between true and imputed consumption where the imputation includes fees but excludes cash flows, i.e., method (ii), and all deviations are below the 45-degree line.

In the center and bottom rows, we display the relationship with actual financial consumption for our four other versions of imputation – no intra-year trading (iii), individual securities with wrong prices (iv), using only market returns (v), and 'raw' portfolio differences (vi). These measures have substantially larger levels of error than seen in the top row. In particular, almost every portfolio-year in the final two imputation methods has appreciable levels of error.

One feature to highlight in the bottom row is the tendency for there to be substantial amounts of error even for cases in which there is no actual change in financial consumption. This can be seen as the vertical cluster of data points above and below the origin in several of the graphs. That is, for a large number of households, they exhibit no active saving or dis-saving during the year, solely seeing changes in portfolio value due to passive capital gains and reinvested cash flows. When incorporating only market returns or 'raw' portfolio value changes, we mistakenly attribute these changes to consumption or savings on the part of the household. This particular source of imputation error is greatly diminished in the center row where we are at least taking into account the composition of the portfolio in terms of the number of shares of a given security. Since the main source of error here is in prices of purchases and sales, a household conducting no trades during the year will have a relatively well-measured 0 even when imputing financial consumption.

4.2 Correlation of Imputation Errors with Economics Outcomes

An important question is whether and how these imputation errors correlate with economic outcomes of interest.

[Figure 2 about here]

As a start, Figure 2 shows bin-scatter plots of the imputation error using method (iii) – ignoring intra-year gross trades and assuming that transactions occur at mid-year prices – against investors' average portfolio value and average income. The top row shows that imputation errors (normalized by investors' average income) are systematically related to both financial wealth and income: The relative errors become more negative as investors get richer.

The middle row shows that controlling for household characteristics – including average income when plotting errors against average portfolio balance and average portfolio balance when plotting errors against average income – helps to mitigate the the systematic bias with income but does not help to resolve the systematic relation with portfolio size.[8] The bottom-left panel shows that the relative imputation error is more dispersed for wealthier households, and the bottom-right panel shows that the relative imputation error increases less than proportionally with income.

[Figure 3 about here]

[8] The other controls include age fixed effects, gender, employment status, marital status, and fixed effects for self-reported risk tolerance.

In Figure 3, we show box plots of the distribution of imputation errors (again relative to average income) for a number of subsets of our data. Each panel illustrates relative imputation error for one of our six methods of imputation in Section 3. Note that the y-axis expands in scale for the center and bottom rows, denoting much wider distributions of measurement error for these imputation methods.

Each panel shows how the relative imputation errors are distributed across the full sample and 7 different subsets of our sample. These subsets are: incomes below 60,000 euros, incomes between 60,000 and 100,000 euros, incomes above 100,000 euros; male account-holders, female account-holders; years with positive market returns, and years with negative market returns.

One notable pattern is seen in the measurement error among female account-holders. Across all versions of our consumption imputation, female account-holders tend to have lower levels of measurement error. This is driven primarily by the fact that female account-holders tend to trade less often than do males, thereby generating fewer trading fees, less intra-year trading or pricing measurement error, and perhaps less impulsive selling in down-market years.

Similarly, account-holders with higher income and higher levels of wealth tend to have the highest levels of measurement error associated with them. Again, this seems to be caused by the significantly higher levels of intra-year trading and by deviation of the individual's portfolio return from return of the market portfolio. For instance, we see that intra-year trading measurement error for these individuals, in the left-center panel of Figure 3, exhibits a distribution about 5-10 times wider than the overall distribution.

Finally, the subset of our sample that exhibits the most distinct measurement error distribution in terms of median error is for years in which the market had negative returns. In these years, the entire distribution of measurement error often shifts in the opposite direction as when markets are positive for the year. This fact may be troublesome for researchers interested in investigating the time-series or panel properties of consumption at an individual level. That is, measurement error for consumption shifts significantly in a manner correlated with the business cycle and overall income growth.

[Table 3 about here]

Table 3 displays a set of results regarding some of the time-series properties of two of our measures of imputation errors. The first is the imputation methodology in which we exclude any intra-year trades that were netted out by opposite trades during a given year – method (iii). The second is the imputation methodology in which we assume all households obtain a market rate of return (e.g., the return on the DAX) for a given year – method (v). Errors are measured in three ways: (a) in euros, ϵ_{it}; (b) relative to household labor income, $\epsilon_{it}/\bar{E}_i$; and (c) as the absolute value of (b), $|\epsilon_{it}/\bar{E}_i|$.

For each column, we regress the relevant imputation error on the logged amount of turnover that a portfolio undergoes through the year, including time and household fixed effects. The impact on the average trading error is ambiguous, with some measures increasing and others decreasing as more trading occurs. For instance, a doubling of trading volume decreases the average amount of 'intra-year imputation error' by approximately 153 euros and increases the amount of 'market returns imputation error' by approximately 510 euros.

However, in columns (3) and (6), we see that the absolute amount of error increases unambiguously as more trading is done. A doubling of trading volume tends to increase imputation error relative to income by approximately 1.5-1.7%. Thus, the more trades that an individual or household conducts, the larger the imputation error, either positive or negative, tends to become.

[Table 4 about here]

Table 4 further explores the correlation of the imputation errors with the state of the economy. Each column regresses the measure of imputation error on real GDP growth, the annual return of the market, or the change in home prices in Germany (each independent variable run separately). Because these are all measured on a country-wide basis, we do not include annual fixed effects in any of the regressions.

For almost all macro variables and imputation error measures, we find that increases in the macro variable tends to increase both the average error (measured in euros or relative to income) and the absolute value of the average error. That is, the mean and standard deviation of imputation errors is pro-cyclical by a number of different metrics.

For instance, for every percentage point increase in German home prices, financial consumption tends to be overestimated by approximately 39 euros or 900 euros, depending on the imputation methodology. In addition, for every percentage point increase in German home prices, the absolute value of imputation errors tends to increase by 0.1%-1% of average income.

We do not assert that imputation errors are *driven* by changes in GDP or home prices, but merely note that there exists significant cyclicality in imputation errors that co-varies with important macro trends. Thus, utilizing any sort of aggregate variation in imputed consumption is likely to both be picking up actual changes alongside non-classical measurement error.

[Table 5 about here]

Table 5 demonstrates that these patterns can have much more striking magnitudes across the distribution of wealth. Because households with larger portfolio balances are subject to higher levels of potential financial consumption imputation error, any heterogeneity in error may be similarly magnified for this higher-wealth group. Table 5 performs the same analysis as with market returns in Table 4 but interacting market returns with indicators that indicate which quintile of the portfolio-balance distribution a household is in.

We find that the largest source of errors, both absolute and relative to household income, are found within the households with the largest portfolio balances. These errors are quantitatively large for richer households: if equity markets increase by 50% in a year, the 'market returns' imputation method can overstate consumption by as much as 20% of income for households in the highest portfolio balance quintile. In contrast, the lowest portfolio balance quintile may see an imputation error of only 1.5% of income. This feature of the data is largely repeated across the different imputation methods and error measurements and demonstrate the extent to which imputed consumption can exhibit errors both over the business cycle but also differentially across households with varying financial characteristics.

[Table 6 about here]

Given that many papers turn to imputed consumption data to study how consumption responds to changes in income, Table 6 tests whether changes in household income are directly correlated with imputation error. Here all columns employ the imputation methodology in which we exclude any intra-year trades that were netted out by opposite trades during a given year, which is a conservative method in the sense that it leads to relatively mild imputation errors on average (see Table 2).

Column (1) finds that an increase in household income of 1% yields an increase in mean imputation error of approximately 2.5 euros. In column (2), we find that changes in income do not exhibit a significant relationship with mean imputation error measured as a fraction of household income. However, increases in income do tend to increase the variation in imputation error, as seen in column (3) using the absolute value of imputation error.

Columns (4)-(6) demonstrate that, for all three metrics, households in the highest quintile of average portfolio balance exhibit stronger positive relationships between imputation error and income. This makes intuitive sense, as small portfolios will mechanically have a narrower scope to be measured with substantial error (either measured in euros as in column (4) or relative to income as in column(5)). Column (5) shows that mean imputation error for the lowest three quintiles of portfolio portfolios varies little by income, but there is a significantly positive relationship with income for the highest quintile. Moreover, the variation in imputation errors increases most significantly for the highest quintile, as seen in column (6) (i.e., .00256-.00136 vs. .00256+.00504).

These results suggest that imputation error can be significant for richer households with large asset portfolios and can vary within-household as income varies. Again, there is not necessarily a mechanical linkage that drives this relationship, but households experiencing changes in income may adjust trading strategies to reflect changing beliefs or to better smooth consumption.

Despite the statistically significant differences, the magnitudes in this table are well below those seen in Table 5. This is largely because the imputation errors are primarily driven by equity market behavior and, in our sample period in Germany, equity market returns are virtually uncorrelated with income (correlation = 0.01). In periods or countries with higher levels of correlation, substantial amounts of imputation error on consumption could occur and researchers will be unable to control for this error with simple time effects given the types of heterogeneity we see in our sample.

5 Conclusion

Survey-based measures of consumption are analyzed in a large number of empirical research papers despite suffering from considerable measurement error. For that reason, a large number of recent papers turn to an alternative measure of consumption derived from annual administrative records on income and asset holdings. Commonly referred to as 'imputed consumption', with this approach consumption is calculated as a residual from the household's budget constraint: the part of total income that was not invested or saved.

However, due to incompleteness in asset records, this measure of consumption may suffer from measurement error, as well. In this paper, we use transaction-level data on over 100,000 German households" balances, asset trades, and asset holdings to document the potential shortcomings in using annual snapshots of wealth and income in imputing consumption.

We find that imputing consumption from annual portfolio snapshots leads to substantial measurement error, both in absolute terms and relative to household income. Moreover, these errors are correlated with both household financial characteristics as well as key macroeconomic variables like GDP growth and house price growth. In short, economists should treat annual snapshot-derived imputed consumption with care, since, especially for households with high levels of income and wealth, measurement error can bias or distort the results of common empirical specifications.

The findings in this paper should not be understood as arguing against using this methodology. We certainly believe that using comprehensive administrative tax registries, which include comprehensive data on wealth, income, and changes in financial portfolios, provides an important alternative to survey-based measures of consumption. In many or even most cases, this is probably the prefered approach to measuring consumption. Instead, we encourage researchers to obtain comprehensive access to the different administrative registries to reduce non-classical measurement error in imputed consumption, or to be cautious when imputing consumption for households with large equity portfolios.

A Sources of Annual Administrative Financial Snapshots

A.1 Swedish Administrative Data

The Swedish data (e.g., Koijen et al. 2014, Bach et al. 2015, Sodini et al. 2016, Eika et al. 2017, Kolsrud et al. 2018) is based on the wealth tax that was in place until 2007. The tax registry includes information on the change in bank accounts (checking, savings, certificates of deposit, etc) in 2006 and 2007 if the balance is greater than 10,000 SEK ($1,224) and in prior years if the earned interest exceeds 100 SEK ($12.24). Furthermore, funds, stocks, and bonds are observed including their actual sale prices. Dividend information is either observed at the individual-level or imputed via the International Security Identification Numbers (ISINs) using standard data sources such as Thomson Reuters Datastream.

A.2 Norwegian Administrative Data

The Norwegian data (e.g., Fagereng et al. 2016, Fagereng and Halvorsen 2017, Fagereng et al. 2017) includes broad asset classes of bank deposits, bonds including money market funds, mutual funds, and stocks. Moreover, there is a self-reported category cash holdings that is commonly ignored though. In turn, the authors take the historical annual return of the Oslo Stock Exchange (30%) and MSCI World Index (70%) to impute the returns of mutual funds and the Treasury bill to impute the returns on bonds. Moreover, for the period 2005 to 2011, the portfolio details are observed, allowing for the same measurement as in Swedish data. Dividends are observed via a dividend tax for part of the sample period and otherwise imputed using standard data sources.

A.3 Danish Administrative Data

The Danish data (e.g., Browning and Leth-Petersen 2003, Kreiner et al. 2014) stems from the wealth tax enacted from 1981 to 1996 but does not include the stock of each asset (except for housing). Thus, capital gains as well as contribution to pensions schemes are excluded from the measure of consumption.

References

Ahmed, N., M. Brzozowski, and T. F. Crossley (2006). Measurement errors in recall food consumption data. Technical report, IFS Working Papers, Institute for Fiscal Studies (IFS).

Bach, L., L. E. Calvet, and P. Sodini (2015). Rich pickings? risk, return, and skill in the portfolios of the wealthy.

Baker, S. R. (2017). Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data. *forthcoming at the Journal of Political Economy*.

Barber, B. M. and T. Odean (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance 55*(2), 773–806.

Bound, J., C. Brown, and N. Mathiowetz (2001). Measurement error in survey data. *Handbook of Econometrics 5*, 3705–3843.

Browning, M. and S. Leth-Petersen (2003). Imputing consumption from income and wealth information. *Economic Journal 113*(488).

Carroll, C. D. (2001). A theory of the consumption function, with and without liquidity constraints. *Journal of Economic Perspectives 15*(3), 23–45.

Cole, S., A. Paulson, and G. Shastry (2012). Smart money: The effect of education on financial behavior.

de Nicola, F. and X. Giné (2014). How accurate are recall data? evidence from coastal india. *Journal of Development Economics 106*(0), 52 – 65.

Deutsche Bundesbank (2013). Statistik ueber wertpapierinvestments.

Dorn, D. and G. Huberman (2005). Talk and action: What individual investors say and what they do. *Review of Finance 9*(4), 437–481.

Eika, L., M. Mogstad, and O. Vestad (2017). What can we learn about household consumption from information on income and wealth. Technical report, Mimeo.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2016). Heterogeneity in returns to wealth and the measurement of wealth inequality. *American Economic Review Papers and Proceedings 106*(5), 651–655.

Fagereng, A., L. Guiso, and L. Pistaferri (2017). Portfolio choices, firm shocks, and uninsurable wage risk. *Review of Economic Studies*, rdx023.

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

Fagereng, A., M. B. Holm, and G. J. Natvik (2016). Mpc heterogeneity and household balance sheets.

Gelman, M., S. Kariv, M. D. Shapiro, D. Silverman, and S. Tadelis (2014). Harnessing naturally occurring data to measure the response of spending to income. *Science 345*(6193), 212–215.

German Federal Bureau of Statistics (2008). Einkommens- und verbrauchsstichprobe.

German Federal Bureau of Statistics (2011). Bildungsstand.

Karlan, D. and J. Zinman (2008). Lying about borrowing. *Journal of the European Economic Association 6*(2/3), pp. 510–521.

Koestner, M., B. Loos, S. Meyer, and A. Hackethal (2017). Do individual investors learn from their mistakes? *Journal of Business Economics 87*(5), 669–703.

Koijen, R., S. Van Nieuwerburgh, and R. Vestman (2014). 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*, pp. 308–346. University of Chicago Press.

Kolsrud, J., C. Landais, P. Nilsson, and J. Spinnewijn (2018). The optimal timing of unemployment benefits: Theory and evidence from sweden. *American Economic Review 108*(4-5), 985–1033.

Kreiner, C. T., D. D. Lassen, and S. Leth-Petersen (2014). Measuring the accuracy of survey responses using administrative register data: evidence from denmark. In *Improving the measurement of consumer expenditures*, pp. 289–307. University of Chicago Press.

Kuchler, T. and M. Pagel (2015). Sticking To Your Plan: Empirical Evidence on the Role of Present Bias for Credit Card Debt Paydown. *Working Paper*.

Kueng, L. (2018). Excess Sensitivity of High-Income Consumers. *forthcoming at the Quarterly Journal of Economics*.

Olafsson, A. and M. Pagel (2016). The liquid hand-to-mouth: Evidence from personal finance management software. *Working Paper*.

Parker, J. A. and N. S. Souleles (2017). Reported preference vs. revealed preference: Evidence from the propensity to spend tax rebates.

Pistaferri, L. (2015). Household consumption: Research questions, measurement issues, and data collection strategies. *Journal of Economic and Social Measurement 40*(1-4), 123–149.

Pudney, S. (2008). Heaping and leaping: Survey response behaviour and the dynamics of self-reported consumption expenditure. Technical report, ISER Working Paper Series.

Schmittmann, J. M., J. Pirschel, S. Meyer, and A. Hackethal (2014). The impact of weather on german retail investors. *Review of Finance 19*(3), 1143–1183.

Sodini, P., S. Van Nieuwerburgh, R. Vestman, and U. von Lilienfeld-Toal (2016). Identifying the benefits from home ownership: A swedish experiment. Technical report, National Bureau of Economic Research.

Table 1: Summary Statistics of Retail Investment Accounts

			Percentiles				
	Mean	St.Dev	10th	25th	50th	75th	90th
Portfolio Value	63,958	94,244	2,144	9,359	30,450	76,220	162,788
Annual Trading Fees	121.8	1,145	0	15.4	121.8	493.7	1,489.8
Annual Turnover	67,454	135,245	0	3,075	18,040	66,056	179,424
Income	61,878	55,522	18,785	26,966	42,609	75,451	129,980
Male	.89	.31	0	1	1	1	1
Age	44	13.7	27	33	44	54	64
'Risk Aversion Index'	3.1	1.7	1	1	3	5	5
Number of Individuals	20,557						

Notes: Nominal values are in Euros.

Table 2: Summary Statistics of Imputation Errors (Relative to Average Income, $\epsilon_{it}/\bar{E}_i^{net}$)

				Percentiles				
	N	Mean	St.Dev	10th	25th	50th	75th	90th
(i) No fees	45,307	.0164782	.036686	0	.0005615	.0040838	.0148993	.2015956
(ii) No cash flows	44,267	-.018861	.0200553	-.1009661	-.0252852	-.014062	-.0053282	0
(iii) No intra-year trades	44,481	.003527	.1833919	-.60964	-.0029641	0	.0002488	.7678314
(iv) No cost basis	44,149	.0002777	.2009846	-.8241348	-.0203109	0	.0306379	.6711145
(v) Only market returns	43,706	.0892178	.2194307	-.448988	-.027066	.0710487	.1803856	.8672268
(vi) 'Raw' portfolio diff.	43,630	-.0072111	.278455	-.6719453	-.1565465	-.0449512	.1040751	.8765174

Notes: The six imputation methods (i)-(vi) that we employ in this paper are are as follows: (i) No Fees - disregard any trading fees charged to the household, yielding weakly larger levels of imputed financial consumption; (ii) No cash flows - disregard any dividends paid out to households, yielding weakly smaller levels of imputed financial consumption; (iii) No consideration of intra-year trades which were netted out during the year (e.g., if you bought 12 and sold 10 shares of Volkswagen, only keep track of the net 2 that you were holding at the end of the year) but purchase/sale prices are 'correct' for those net purchases/sales; (iv) Similar to (iii) but using the incorrect prices of purchase/sale for any net changes in securities held (prices assumed to be equal to the June 30th price during the previous year); (v) Market Returns imputation dispenses with actual portfolio holdings and assumes that the household holds the DAX and that any adjustment (i.e., financial consumption) occurs on June 30th at June 30th prices; (v) 'Raw' imputation - disregarding growth altogether and assuming any change in equity portfolio value is the result of active savings or dis-savings.

Table 3: Trading and Imputation Errors

	(iii) No Intra-Year Trades			(v) Only Market Returns		
	ϵ_{it}	$\epsilon_{it}/\bar{E}_i^{net}$	$\lvert\epsilon_{it}/\bar{E}_i^{net}\rvert$	ϵ_{it}	$\epsilon_{it}/\bar{E}_i^{net}$	$\lvert\epsilon_{it}/\bar{E}_i^{net}\rvert$
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Turnover Value)	-153.0***	-0.00268***	0.0159***	510.2***	0.0121***	0.0171***
	(18.24)	(0.000512)	(0.000694)	(32.61)	(0.00102)	(0.000938)
Observations	46,308	22,166	22,166	45,302	21,718	21,718
R^2	0.084	0.098	0.388	0.322	0.336	0.482
Mean of Dep Var	72.9	0.002	0.057	3922	0.083	0.135
Year FE	YES	YES	YES	YES	YES	YES
Individual FE	YES	YES	YES	YES	YES	YES

Notes: Each column denotes the results of regressions examining how individual trading affects the size of imputation errors. Columns (1)-(3) use imputation error derived from our third imputation method, (iii) no intra-year trades. Columns (4)-(6) use imputation error derived from our fifth imputation method, (v) Only Market Returns. Columns (1) and (4) look at levels of imputation errors in euros, ϵ_{it}; columns (2) and (5) look at imputation errors relative to individual investors' average income, $\epsilon_{it}/\bar{E}_i$; and columns (3) and (6) look at the absolute value of imputation errors relative to individual income, $|\epsilon_{it}/\bar{E}_i|$. Trading values are computed as an individual's total gross annual trades multiplied by the price of securities at the time of trade. Standard errors are clustered by investor. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 4: GDP, Market Returns, Home Prices, and Imputation Errors

	(iii) No Intra-Year Trades			(v) Only Market Returns		
	ϵ_{it}	$\epsilon_{it}/\bar{E}_i^{net}$	$\lvert\epsilon_{it}/\bar{E}_i^{net}\rvert$	ϵ_{it}	$\epsilon_{it}/\bar{E}_i^{net}$	$\lvert\epsilon_{it}/\bar{E}_i^{net}\rvert$
	(1)	(2)	(3)	(4)	(5)	(6)
Real GDP Growth	-865.7	-0.0168	0.229***	34,392***	0.628***	0.202***
	(1,153)	(0.0351)	(0.0260)	(1,718)	(0.0509)	(0.0364)
Observations	51,862	24,748	24,748	50,862	24,321	24,321
R^2	0.073	0.088	0.348	0.243	0.260	0.434
Individual FE	YES	YES	YES	YES	YES	YES
Annual DAX Return	1,706***	0.0330***	-0.00771**	6,122***	0.107***	0.00113
	(177.2)	(0.00504)	(0.00345)	(338.3)	(0.00984)	(0.00638)
Observations	50,648	24,169	24,169	50,862	24,321	24,321
R^2	0.078	0.094	0.350	0.248	0.264	0.434
Individual FE	YES	YES	YES	YES	YES	YES
Change in Home Prices	393.7***	0.00999**	0.0125***	9,324***	0.197***	0.0901***
	(138.0)	(0.00430)	(0.00462)	(313.9)	(0.00954)	(0.00776)
Observations	36,640	17,198	17,198	35,906	16,888	16,888
R^2	0.108	0.114	0.417	0.295	0.326	0.495
Individual FE	YES	YES	YES	YES	YES	YES

Notes: Each column and row denotes a regression of imputation errors on annual GDP growth in Germany, the annual return of the DAX (German S&P-500 equivalent), or annual housing price changes in Germany. Columns (1)-(3) use imputation error derived from our third imputation method, (iii) no intra-year trades. Columns (4)-(6) use imputation error derived from our fifth imputation method, (v) Only Market Returns. Columns (1) and (4) look at levels of imputation errors in euros, ϵ_{it}; columns (2) and (5) look at imputation errors relative to individual investors' average income, $\epsilon_{it}/\bar{E}_i$; and columns (3) and (6) look at the absolute value of imputation errors relative to individual income, $|\epsilon_{it}/\bar{E}_i|$. All independent variables are measured on a -1 to 1 scale (i.e., 1% real GDP growth is denoted at 0.01). Standard errors are clustered by investor. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 5: Imputation Errors from Market Returns by Portfolio Balance Quintiles

	(iii) No Intra-Year Trades			(v) Only Market Returns		
	ϵ_{it}	$\epsilon_{it}/\bar{E}_i^{net}$	$\lvert\epsilon_{it}/\bar{E}_i^{net}\rvert$	ϵ_{it}	$\epsilon_{it}/\bar{E}_i^{net}$	$\lvert\epsilon_{it}/\bar{E}_i^{net}\rvert$
	(1)	(2)	(3)	(4)	(5)	(6)
Mkt. Ret.*Eq. Quintile 1	-98.78	-0.00542	-0.0436***	-1,029***	-0.0346***	-0.0735***
	(168.8)	(0.00629)	(0.00564)	(263.2)	(0.00975)	(0.00893)
Mkt. Ret.*Eq. Quintile 2	721.1***	0.0144*	-0.0299***	1,151***	0.00905	-0.0549***
	(261.1)	(0.00835)	(0.00667)	(309.0)	(0.0124)	(0.0101)
Mkt. Ret.*Eq. Quintile 3	2,299***	0.0496***	-0.00537	2,827***	0.0704***	-0.0380***
	(326.8)	(0.00984)	(0.00773)	(480.7)	(0.0153)	(0.0113)
Mkt. Ret.*Eq. Quintile 4	3,337***	0.0604***	0.0137	9,427***	0.189***	0.0303**
	(436.8)	(0.0118)	(0.00887)	(741.0)	(0.0217)	(0.0147)
Mkt. Ret.*Eq. Quintile 5	3,043***	0.0471***	0.0364***	28,179***	0.386***	0.204***
	(652.5)	(0.0145)	(0.0110)	(1,539)	(0.0386)	(0.0261)
Observations	50,648	24,169	24,169	50,862	24,321	24,321
R^2	0.080	0.096	0.353	0.277	0.279	0.442
Individual FE	YES	YES	YES	YES	YES	YES

Notes: Each column and row denotes a regression of imputation errors on the annual return of the DAX (German S&P-500 equivalent) interacted with a household-level indicator denoting which quintile of equity account balance the household is in. Columns (1)-(3) use imputation error derived from our third imputation method, (iii) no intra-year trades. Columns (4)-(6) use imputation error derived from our fifth imputation method, (v) Only Market Returns. Columns (1) and (4) look at levels of imputation errors in euros, ϵ_{it}; columns (2) and (5) look at imputation errors relative to individual investors' average income, $\epsilon_{it}/\bar{E}_i$; and columns (3) and (6) look at the absolute value of imputation errors relative to individual income, $|\epsilon_{it}/\bar{E}_i|$. Market returns are measured in percentage points (i.e., 1% real GDP growth is denoted at 0.01). Standard errors are clustered by investor. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 6: Household Income and Imputation Errors

	(iii) No Intra-Year Trades			(v) Only Market Returns		
	ϵ_{it}	$\epsilon_{it}/\bar{E}_i^{net}$	$\|\epsilon_{it}/\bar{E}_i^{net}\|$	ϵ_{it}	$\epsilon_{it}/\bar{E}_i^{net}$	$\|\epsilon_{it}/\bar{E}_i^{net}\|$
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Income)	258.2**	0.00205	0.00592***			
	(103.7)	(0.00176)	(0.00166)			
ln(Inc)*Eq. Quintile 1				76.66***	0.000950*	-0.00118**
				(27.31)	(0.000524)	(0.000459)
ln(Inc)*Eq. Quintile 2				104.4***	0.00193***	0.000837*
				(28.01)	(0.000554)	(0.000479)
ln(Inc)*Eq. Quintile 3				193.5***	0.00329***	0.00215***
				(30.56)	(0.000596)	(0.000516)
ln(Inc)*Eq. Quintile 4				266.5***	0.00450***	0.00372***
				(33.41)	(0.000652)	(0.000549)
ln(Inc)*Eq. Quintile 5				285.5***	0.00533***	0.00527***
				(36.95)	(0.000708)	(0.000637)
Observations	24,729	24,748	24,748	24,729	24,748	24,748
R^2	0.079	0.091	0.352	0.086	0.097	0.363
Year FE	YES	YES	YES	YES	YES	YES
Individual FE	YES	YES	YES	YES	YES	YES

Notes: Columns (1)-(3) use imputation error derived from our third imputation method, (iii) no intra-year trades. Columns (4)-(6) use imputation error derived from our fifth imputation method, (v) Only Market Returns. Columns (1) and (4) look at levels of imputation errors in euros, ϵ_{it}; columns (2) and (5) look at imputation errors relative to individual investors' average income, $\epsilon_{it}/\bar{E}_i$; and columns (3) and (6) look at the absolute value of imputation errors relative to individual income, $|\epsilon_{it}/\bar{E}_i|$. Standard errors are clustered by investor. ***, **, * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Figure 1: Financial Consumption – Imputed vs. Actual

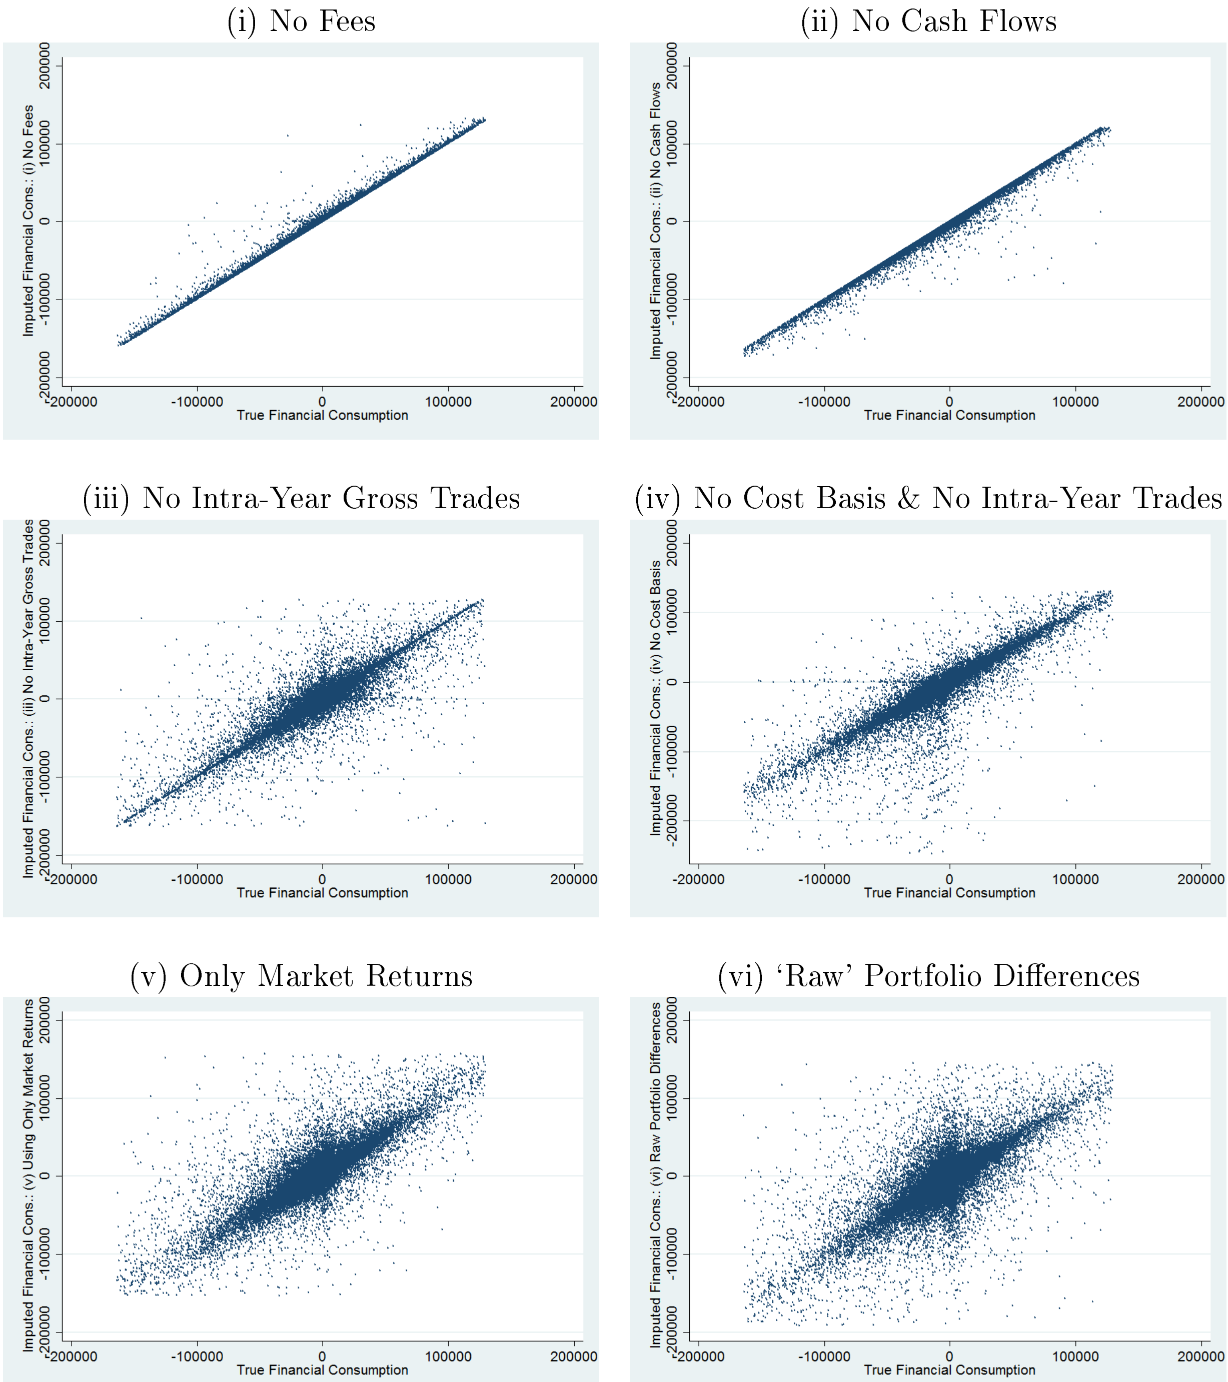

Notes: Each panel plots imputed financial consumption according to one of six imputation methods against true financial consumption. From left to right and top to bottom, the imputation methods follow Section 3: (i) No Fees – disregard any trading fees charged to the household, yielding weakly larger levels of imputed financial consumption; (ii) No Cash Flows – disregard any dividends paid out to households, yielding weakly smaller levels of imputed financial consumption; (iii) No Intra-Year Gross Trades – no consideration of intra-year trades which were netted out during the year (e.g., if you bought 12 and sold 10 shares of Volkswagen, only keep track of the net 2 that you were holding at the end of the year) but purchase/sale prices are 'correct' for those net purchases/sales; (iv) No Cost Basis & No Intra-Year Trades – Similar to (iii) but using the incorrect prices of purchase/sale for any net changes in securities held (prices assumed to be equal to the June 30th price during the previous year); (v) Only Market Returns – imputation dispenses with actual portfolio holdings and assumes that the household holds the DAX Index and that any adjustment (i.e., financial consumption) occurs on June 30th at June 30th prices; (vi) 'Raw' Portfolio Differences – disregarding growth altogether and assuming any change in equity portfolio value is the result of active savings or dis-savings.

Figure 2: Imputation Errors by Portfolio Size and Average Income

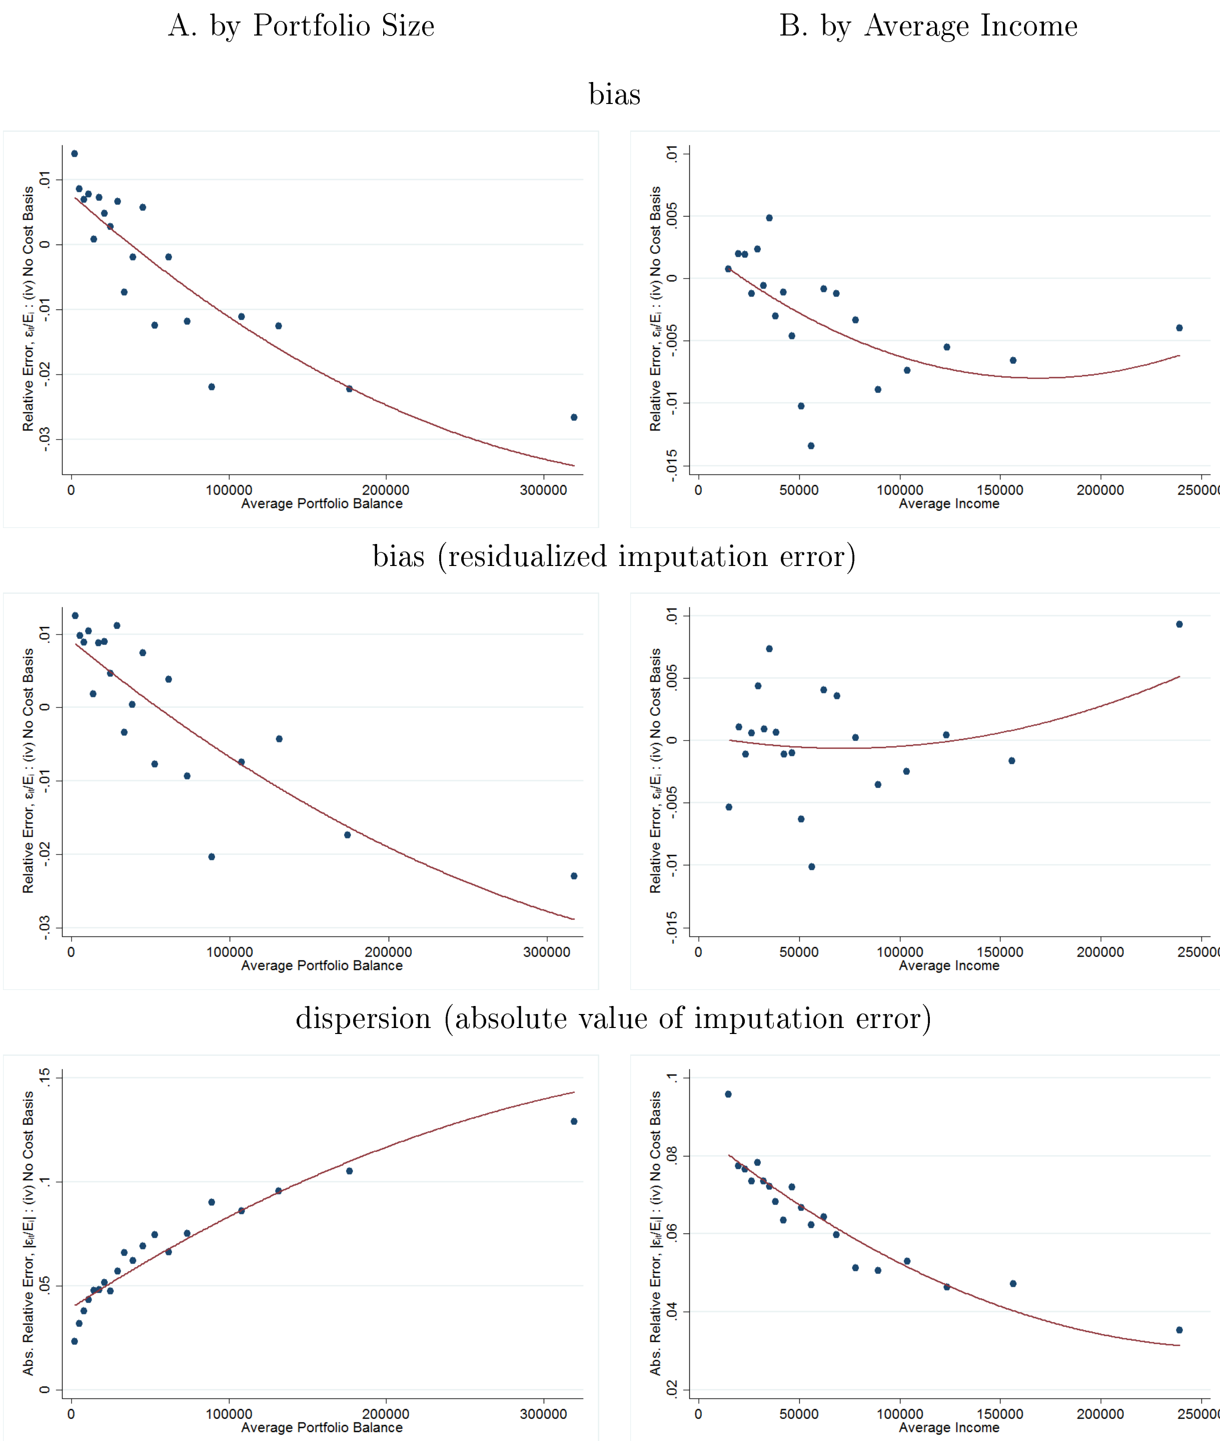

Notes: The left Panel A (Panel B) display bin-scatter plots across 20 quantiles of average financial consumption imputation errors against average individual portfolio balances (average individual income). All imputation errors are measured as in method (iii) of Section 3 and relative to average individual income, $\epsilon_{it}/\bar{E}_i$ (e.g., -0.01 means an annual error 1% the size of average income for that individual that is *underestimating* financial consumption). The middle row first regresses the relative imputation errors on average income in Panel A and on average portfolio balance in Panel B and a set of household characteristics, which include age fixed effects, gender, employment status, marital status, and fixed effects for self-reported risk tolerance. The bottom assesses changes in the dispersion of relative imputation errors using absolute values instead, i.e., $|\epsilon_{it}/\bar{E}_i|$.

Figure 3: Imputation Error by Subset

Notes: Each panel plots imputation errors in financial consumption ($FinCons$) scaled by average individual income according to one of six imputation methods (i)-(vi) of Section 3. Each panel plots the overall distribution of error (left-most bar) as well as the error for 7 different subgroups. These groups are: (1) lowest tercile income individuals; (2) middle tercile income individuals; (3) highest tercile income individuals; (4) males; (5) females; (6) years with positive market returns; (7) years with negative market returns). Bars denote interquartile ranges, with the middle lines being the median error scaled by individual average income.