Saving Behavior Across the Wealth Distribution *

Andreas Fagereng
Statistics Norway

Martin Blomhoff Holm
University of Oslo

Benjamin Moll
Princeton

Gisle Natvik
BI Norwegian Business School

February 28, 2019

Abstract

Do wealthier households save a larger share of their incomes than poorer ones? We use Norwegian administrative panel data on income and wealth to examine how saving rates vary across the wealth distribution. We compare our findings to the prediction of workhorse macro models that saving rates are either independent of or decreasing with wealth. We find that the relation between saving rates and wealth depends on whether saving includes capital gains. Saving rates net of capital gains ("net saving rates") are approximately constant across the wealth distribution, seemingly consistent with workhorse models. However, saving rates including capital gains ("gross saving rates") increase markedly with wealth. Since the predictions of economic theories are about gross saving, our findings challenge workhorse models with approximately constant saving rates. In contrast, our empirical findings are consistent with a theory featuring multiple assets and portfolio adjustment frictions.

*We thank Adrien Auclert, John Campbell, Karen Dynan, Giulio Fella, Adam Guren, Henrik Kleven, Camille Landais, Monika Piazzesi, Luigi Pistaferri, Martin Schneider, Kjetil Storesletten, Olav Syrstad, and seminar participants at various institutions for useful comments.

1 Introduction

Do wealthier households save a larger share of their incomes than poorer ones? The answer to this question has important implications, both for macroeconomics and the study of inequality. In macroeconomics, a prediction of many workhorse models is that saving rates are either independent of or decreasing with wealth (De Nardi and Fella, 2017). These theories then typically also predict that aggregate consumption and saving are approximately independent of the wealth distribution.[1] But whether the individual saving behavior assumed in these models is consistent with the data is an open question. In the study of inequality, how saving rates vary across the wealth distribution is important for understanding wealth inequality and its determinants. In particular, it sheds light on whether saving behavior is a force toward escalating wealth inequality. Better empirical evidence on saving behavior across the wealth distribution is therefore key for making progress on a number of important questions.

We fill this gap by using Norwegian administrative panel data on income and wealth to examine how saving rates out of income vary across the wealth distribution. Because Norway levies both income and wealth taxes on its households, the tax registry data provide a complete account of household income and balance sheets down to the single asset category. We focus on the twelve-year period from 2004 to 2015, for which we combine tax registries with shareholder and housing transactions registries. Taken together, these data contain detailed third-party-reported information on household-level wealth and income, covering the universe of Norwegians from the very bottom to the very top of the wealth distribution.

Our main finding is that how saving rates vary with wealth, crucially depends on whether we include capital gains in the definition of saving. We contrast two notions of saving rates, corresponding to two different ways of writing the budget constraint "consumption plus saving equals income" and, in particular, how capital gains are treated when writing this accounting identity. The first notion of saving, *gross saving*, is the change in a household's net worth from one year to the next, including any revaluation effects due to changing asset prices. In contrast, the second notion, *net saving*, excludes any unrealized capital gains – equivalently, it is the difference between a household's income net of capital gains and its consumption.

[1] As we show in this paper, the prediction that saving rates are either independent of or decreasing with wealth comes out of a wide class of models, in particular standard incomplete-markets models à la Aiyagari (1994) and Huggett (1993) and many of their extensions and generalizations. That these models typically predict that macroeconomic aggregates are approximately independent of the wealth distribution was first noted by Krusell and Smith (1998).

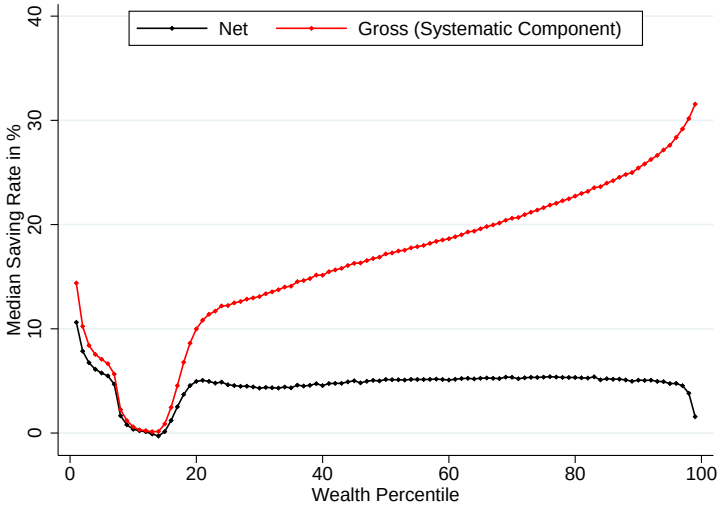


Figure 1: Saving Rates across the Wealth Distribution (Preview).

In our main empirical exercise we ask how the corresponding saving rates vary across the wealth distribution. As is often the case, the resulting findings are easier to communicate visually rather than verbally. Figure 1 therefore plots the corresponding saving rates against percentiles of net worth. To the left of the graph are households with negative net worth. Moving to the right, households become progressively wealthier.

First, consider the net saving rate, i.e. the saving rate excludingf capital gains. For households with negative net worth, the net saving rate declines with wealth. But for the largest part of the wealth distribution it is flat around seven percent. This pattern for net saving rates is seemingly consistent with workhorse macro models which predict exactly such a flat saving rate. But when we instead examine gross saving rates *including capital gains* we find a strikingly different pattern: gross saving rates increase sharply with wealth, from zero at the fifteenth percentile (corresponding to zero net worth) to thirty-five percent for the top one percent of the wealth distribution.[2] The proximate explanation for the diverging behavior of these two notions of saving rates is simple. Wealthier households hold assets like stocks and housing that experience persistent capital gains. Instead of selling off these assets to consume, households hold on to them and therefore have a high gross saving rate. We term this phenomenon "saving by holding."

A simple back-of-the-envelope example helps clarify this point as well as the magnitudes in Figure 1. Assume that the net saving rate is 10% at all points of the wealth distribution

[2]As we explain in more detail below, year-to-year fluctuations in asset prices mean that also the gross saving rate displays large year-to-year fluctuations. The line labelled "gross (systematic component)" therefore plots the systematic component of the gross saving rate (defined in more detail in the main text).

and that capital gains on all assets are 2%. Now compare two individuals at different points of the wealth distribution: the first has an income excluding capital gains of $100,000 and no assets, while the second has the same income but owns assets worth $1,000,000. If the individuals do not consume out of capital gains, their gross savings are $10,000 and $\$10,000 + 2\% \times \$1,000,000 = \$30,000$ respectively. Therefore, the gross saving rate of the first individual is 10% whereas that of the second is $\frac{30,000}{100,000+20,000} = 25\%$. Note that these are roughly the magnitudes we observe in Figure 1 when moving from the 20th to the 95th percentile of the wealth distribution. Also note that even relatively small capital gains (say 2%) can have large effects.

Perhaps surprisingly, both net and gross saving rates are decreasing with wealth among households with negative net worth. The trough of both saving rates is reached where net worth is approximately zero. As detailed below, a challenge for economic models will be to account for both the declining and the increasing segments of the saving rates in Figure 1. Overall, we summarize the main features of our observed saving rates as follows: (1) both net and gross saving rates are decreasing with wealth for households with negative net worth; (2) among households with positive net worth, the net saving rate tends to be flat; (3) among those same households, gross saving rates are sharply increasing with wealth.

To be clear, all our empirical findings (like those in Figure 1) are descriptive and not intended to identify a causal effect of wealth on saving rates. At the same time, we can use our data to better understand the drivers behind these patterns. For instance, workhorse theories suggest that, in the cross-section, saving rates should be correlated with wealth because both are correlated with age, labor income, and innate willingness to save (patience). We show, however, that these factors are not the main drivers of our findings: the relation between saving rates and wealth is qualitatively unchanged when conditioning on observables like age, income and individual savings histories.

Our second contribution is to show that the joint patterns for net and gross saving rates we document have important implications for macroeconomic theory, and in particular challenge many conventional models. To this end, we consider a series of theoretical benchmark models and compare their predictions to our main empirical findings in Figure 1. We start with a particularly simple consumption-saving model that can be solved analytically and then gradually enrich the framework to speak to prevalent features of the data such as changing asset prices and housing. In the simplest version, an infinitely-lived household with homothetic utility saves in a riskless asset and receives a constant stream of labor income. As we show below, this model predicts that a household's saving rate out of income should

be exactly independent of its wealth. We later generalize this model considerably and show that many other benchmark models inherit this property approximately once we condition on appropriate observables like age and labor income.

The benchmark model most relevant for interpreting our results is a version of this pencil-and-paper model with stochastic asset prices. This model is identical to the one just described, with the difference that households buy and sell an asset with a fluctuating price rather than saving in a riskless asset. As is standard, the asset's return is the sum of its dividend yield and capital gains. Because asset prices change over time, the model generates a meaningful distinction between net and gross saving, just like in the data. A key prediction of this model is that the optimal consumption and saving response to capital gains depends crucially on whether these are transitory or persistent. The logic is similar to that of the permanent income hypothesis for labor income shocks. Households experiencing transitory capital gains optimally save these so that such asset price changes translate into changes in their net worth one-for-one. In contrast, households experiencing persistent capital gains optimally consume part of the resulting income flow. This has two implications for saving rates. First, if households experience persistent capital gains, net saving rates should be decreasing with wealth. Second, while fluctuations in asset prices lead to fluctuations in gross saving rates and their dependence on household wealth, gross saving rates purged of transitory capital gains should be independent of wealth. Put differently, the theory predicts that the *systematic* component of the gross saving rate should be independent of wealth.

What we plot in Figure 1 are precisely the empirical counterparts of (i) the net saving rate and (ii) this systematic component of the gross saving rate. Our main takeaway for theory then is that the joint pattern for net and gross saving rates in the data is inconsistent with the benchmark models discussed thus far.

As we show below, housing is the dominant asset in household portfolios for large parts of the wealth distribution. House prices in Norway have appreciated persistently over time and this accounts for part of the patterns in Figure 1. Therefore, our list of benchmark models also includes a model with housing. Housing differs from financial assets such as stocks in two dimensions. First, housing is not only an asset but also a consumption good. Second, housing is indivisible and subject to substantial adjustment costs. We use our model to examine to what extent these two features can help explain our empirical findings, in particular that households seem to hold on to their residences in the face of rising house prices. We show that, by itself, the consumption aspect of housing is not enough to rationalize our findings and that other aspects of housing, like indivisibilities and transaction costs, must be part

of the story. More specifically, we show that a workhorse macro model extended to feature *divisible* housing collapses to a one-asset model with a gross saving rate that is approximately independent of wealth, just like in our simplest benchmark model. In summary, while housing is clearly different from financial assets, the consumption aspect by itself cannot explain our empirical findings.

The importance of housing in the data leads to another, related question: to what extent do our findings in Figure 1 point to a housing-specific story? To answer this question we examine saving rates in financial wealth, that is, saving rates with housing "taken out" of household wealth accumulation. We show that the patterns we document are not exclusive to housing and that households treat capital gains on financial assets similarly to those on housing. Taken together, our findings are consistent with the view that, across different asset classes, households treat capital gains differently from other forms of income and consume very little of these even if they are persistent. Workhorse macro models in which households freely rebalance their portfolios imply that capital gains are just another form of income and there have difficulty generating these patterns.

If standard workhorse models cannot rationalize our empirical findings, what theories can? In a reduced form, what is needed to rationalize Figure 1 is a theory in which households consume very little of any capital gains they experience – that is, they have a saving rate out of capital gains of close to one hundred percent, precisely what we have termed "saving by holding." We argue that a theory featuring multiple assets and portfolio adjustment frictions is promising in this regard. This theory has two key ingredients. First, it takes empirical household balance sheets seriously by incorporating the feature that households hold multiple distinct assets. This contrasts to most workhorse macro models which are either explicitly or effectively one-asset models. In the theory, there is a consumption asset that is used to smooth consumption and an investment asset that experiences capital gains. Second, households are subject to broadly defined portfolio adjustment "frictions", simply meaning some reason for not rebalancing their portfolios most of the time. Potential candidate "frictions" include physical adjustment costs, capital gains taxes generating a "lock-in-effect" (because they are levied on realized rather than accrued gains), inattention to asset price changes, and behavioral factors like mental accounting. We show that this simple model can qualitatively explain our findings and, in particular, generates "saving by holding." We also discuss other candidate explanations, many of which are complementary. These include: asset-demand shifts that affect a majority of households holding a given asset, e.g. a sustained housing preference shift; non-homothetic preferences; and misperceptions

about the stochastic process for asset prices.

As far as we know, ours is the first paper to provide systematic evidence on individual saving rates out of income over the entire wealth distribution.[3] Bach, Calvet and Sodini (2018) perform an exercise close to ours in spirit in that they document patterns in saving behavior across the wealth distribution using administrative data. But our paper differs from theirs in several respects. First and most importantly, their focus is on documenting how the saving rate out of wealth, i.e. the saving-to-wealth ratio or wealth growth rate, varies across the wealth distribution whereas ours is on the saving rate out of income. Given our goal of learning about theories of consumption-saving behavior, the saving rate out of income is the more suitable object to study.[4] Second, given their focus on the saving-to-wealth ratio, they do not study saving behavior at the bottom of the distribution where wealth is zero or negative and hence this ratio is not well defined. In contrast, we study saving behavior across the entire distribution, in particular for the roughly fifteen percent of households with negative net worth, thereby uncovering that saving rates out of income are actually decreasing with wealth in this part of the distribution.

In other related work, Dynan, Skinner and Zeldes (2004) and Straub (2018) document how consumption and saving behavior vary with "permanent income" defined as the permanent component in labor income. This is challenging because permanent income is not directly observable and must be estimated, typically by means of an instrumental variable strategy. We instead focus on how saving behavior varies with wealth which is readily observable.

Fagereng et al. (2016) and Bach, Calvet and Sodini (2016) document that wealthier households obtain higher returns to their wealth whereas we document that they have higher saving rates out of income. This suggests that return heterogeneity and saving rates heterogeneity may both be forces toward diverging wealth inequality and, in fact, reinforce each other. Both findings also suggest that standard macro models that assume away heterogeneity that varies systematically with wealth in both returns to wealth and saving rates miss important

[3] An arguable exception is the handbook chapter by Krueger, Mitman and Perri (2016). In their Table 2, they present some suggestive evidence on consumption rates out of income (i.e. one minus saving rates) using data from the U.S. Panel Study of Income Dynamics (PSID). The much smaller sample size of the PSID forces them to split the wealth distribution into five quintiles only and to compute the consumption rate as total consumption expenditures for a specific wealth quintile divided by total income in that wealth quintile (i.e. not an individual-specific consumption or saving rate like in our paper).

[4] In contrast, as we explain in the main body of the paper, standard consumption-saving models have no clear prediction for the saving-to-wealth ratio except that it should be mechanically decreasing with wealth. This prediction follows because even households with nearly zero wealth, for whom labor income essentially represents their only income source, typically save part of their income, and thus have an extremely high saving-to-wealth ratio.

features of the data and are in need of an overhaul.[5]

Finally, our paper is related to the literature that studies the effect of asset price changes on consumption. While there is much debate about the exact size of such wealth effects on consumer spending, the literature agrees that the marginal propensity to consume capital gains tends to be considerably smaller than the marginal propensity to consume out of other income sources, particularly disposable income.[6] This is consistent with our finding that households "save by holding," meaning that they tend to hold on to assets experiencing persistent capital gains.

Section 2 lays out a series of theoretical benchmarks and shows that workhorse macro models predict that the systematic component of gross saving rates is approximately constant across the wealth distribution. Section 3 describes our data and institutional setting, and discusses some key concepts and definitions regarding the treatment of capital gains. Section 4 presents our empirical results and documents that, inconsistent with the prediction of workhorse macro models, it is net saving rates that are approximately constant across the wealth distribution while gross saving rates increase sharply with wealth. Section 5 theoretically interprets these empirical findings and argues that they are consistent with a theory featuring multiple assets and portfolio adjustment frictions in which the degree of wealth inequality affects macroeconomic aggregates. Section 6 concludes.

2 Theoretical Benchmarks

While our main contribution is empirical, namely to document saving behavior across the wealth distribution, we begin by considering a series of theoretical benchmarks. These will aid the interpretation of our empirical results in Section 4. In particular, we show that

[5] A natural question is whether wealthy households have high saving rates *because* they have high returns. Our findings suggest that the answer is a "partial yes." Fagereng et al. (2016) find that some return heterogeneity is due to allocation of wealth between different asset classes but that there is also return heterogeneity *within* asset class. High-return assets tend to have a large share of their returns in the form of capital gains (e.g. stocks and housing) and so this partly explains our finding that gross saving rates are increasing with wealth. Similarly, to the extent that the return heterogeneity within asset class is explained by heterogeneity of their capital gains component, this would contribute to our findings. On the other hand, as the simple back-of-the-envelope example earlier in the introduction has shown, the gross saving rate can be increasing in wealth even if all households experience the same capital gains as a fraction of their assets, simply because richer individuals hold more assets. Summarizing, while a positive correlation between wealth and returns is not necessary to deliver an increasing gross saving rate, it may magnify the effect.

[6] Poterba (2000) reviews the literature on the consumption effects of changes in stock market wealth and Paiella and Pistaferri (2017) and Christelis, Georgarakos and Jappelli (2015) are two examples of more recent studies. For studies examining, both theoretically and empirically, the effect of house price changes on consumer spending see Berger et al. (2018) and Guren et al. (2018), among others.

workhorse macro models predict that the systematic component of gross saving rates, i.e. saving rates including capital gains, are approximately constant across the wealth distribution. We begin with a particularly simple consumption-saving model that can be solved analytically and then gradually enrich the framework. These models, particularly the richer ones, are designed to speak to important features of the data such as changing asset prices and housing.

2.1 A Pencil-and-Paper Consumption-Saving Model

Constant Asset Prices. We start with a simple consumption-saving model that can be solved by pencil and paper. Households are infinitely lived and have homothetic preferences over utility from consumption c_t:

$$\int_0^\infty e^{-\rho t} u(c_t) dt, \quad u(c) = \frac{c^{1-\gamma}}{1-\gamma}. \tag{1}$$

They receive a constant labor income w and they can save in an asset a_t denominated in units of the consumption good which pays a constant interest rate r. Their budget constraint is

$$c_t + \dot{a}_t = w + r a_t. \tag{2}$$

Finally, they face a natural borrowing constraint $a_t \geq -w/r$. This problem has a very simple analytic solution (see Appendix A.1 for the proof): the optimal saving policy function $\dot{a} = s(a)$ is

$$s(a) = \frac{r-\rho}{\gamma}\left(a + \frac{w}{r}\right). \tag{3}$$

That is, households save (and consume) a constant fraction of their effective wealth $a + w/r$. Importantly, the saving policy function is linear in financial wealth a. Next, the saving rate out of total income $w + ra$ is given by

$$\frac{s}{y} = \frac{s}{w + ra} = \frac{r-\rho}{\gamma r}.$$

That is, households save a constant fraction of their income. Importantly, this fraction is independent of their wealth.[7] We show below that many other benchmark models inherit this property, at least approximately.

[7] Note that this not an artefact of considering infinitesimal time intervals. Instead, it also holds over discrete time intervals: for example, using (3), we have $a_{t+1} - a_t = \int_0^1 \dot{a}_{t+s} ds = \frac{r-\rho}{\gamma} \int_0^1 \left(a_{t+s} + \frac{w}{r}\right) ds = \frac{r-\rho}{\gamma r} \int_0^1 (w + r a_{t+s})\, ds$ and hence saving $a_{t+1} - a_t$ is a constant fraction $\frac{r-\rho}{\gamma r}$ of income $\int_0^1 (w + r a_{t+s})\, ds$.

Changing Asset Prices. Motivated by the fact that changing asset prices turn out to be an important feature of the data, we next extend this simple benchmark model to feature a stochastic asset price. As above, households receive a constant labor income w. In contrast to above, they can now buy and sell an asset k_t at a price p_t. For concreteness we refer to the asset as "stocks." These stocks are the only saving vehicle available to households. Stocks pay dividends and we denote by θ the constant dividend yield. The budget constraint is

$$c_t + p_t \dot{k}_t = w + \theta p_t k_t. \tag{4}$$

The stock price evolves according to

$$\frac{\dot{p}_t}{p_t} = \mu + \varepsilon_t, \tag{5}$$

where $\mu > 0$ is a constant and ε_t is a random variable. We refer to μ as the persistent component of capital gains and to ε_t as the transitory component. A simple special case is when $\varepsilon_t dt = \sigma dW_t$, where dW_t is the innovation to a standard Brownian motion so that the stock price p_t follows a geometric Brownian motion (a continuous-time random walk).

The connection to the setup with constant asset prices is simple. Defining wealth $a_t := p_t k_t$, the budget constraint (4) becomes

$$c_t + \dot{a}_t = w + r_t a_t \quad \text{where} \quad r_t := \theta + \mu + \varepsilon_t. \tag{6}$$

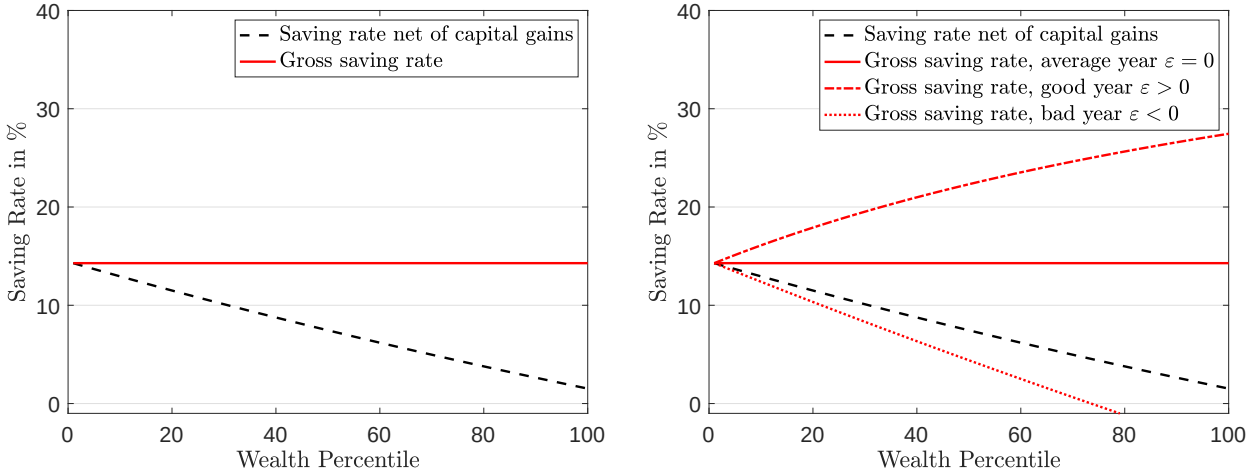
The model with changing asset prices is therefore equivalent to the model with constant asset prices except that the return r_t is time-varying and stochastic. Importantly, the return r_t is the sum of dividends θ and capital gains $\mu + \varepsilon_t$. That this return including capital gains shows up on the right-hand side of the budget constraint (6) is the rationale for the well-known old argument that (unrealized) capital gains are a form of income.[8]

When returns feature a stochastic component $\varepsilon_t \neq 0$ and labor income is positive $w > 0$, it is no longer possible to solve the consumption-saving problem analytically. However, Appendix A.1 derives a useful approximation to the saving policy function under the assumption that return innovations ε_t are Brownian with variance σ^2 and labor income is small $w \approx 0$:

$$\dot{a} \approx \left[w + \bar{r} a - \bar{c} \left(a + \frac{w}{\bar{r} - \gamma \sigma^2} \right) \right] + \varepsilon a \approx \bar{s} \left(a + \frac{w}{\bar{r}} \right) + \varepsilon a, \tag{7}$$

[8] See von Schanz (1896), Haig (1921) and Simons (1938). Also see Eisner (1988) for an argument to include capital gains in the income definition used in the National Income and Product Accounts.

where $\bar{r} := \theta + \mu$, $\bar{c} := \frac{\rho - \bar{r}}{\gamma} + \bar{r} + (1-\gamma)\frac{\sigma^2}{2}$ and $\bar{s} := \bar{r} - \bar{c} = \frac{\bar{r} - \rho}{\gamma} + (\gamma - 1)\frac{\sigma^2}{2}$. The approximation is exact if either $\sigma^2 = 0$ – so that we are back in the case with a deterministic return – or $w = 0$ – in which case the problem is a simplified version of Merton (1969) without portfolio choice. Examination of (7) yields a key result: the optimal consumption and saving response to capital gains depends crucially on whether these are transitory or persistent. The logic is similar to that of the permanent income hypothesis regarding the optimal consumption response to transitory and persistent labor income shocks. When experiencing a transitory capital gain $\varepsilon > 0$, the household optimally saves one hundred percent of that transitory capital gain (see the additive term εa in (7)). Symmetrically, transitory capital losses $\varepsilon < 0$ translate one-for-one into lower wealth. In contrast, when experiencing a persistent capital gain $\mu > 0$, the household optimally consumes part of the resulting income flow.



(a) Only persistent capital gains: $\mu > 0, \varepsilon = 0$

(b) Both types of capital gains: $\mu > 0, \varepsilon \lessgtr 0$

Figure 2: Saving rates across wealth distribution in simple benchmark model

Figure 2 examines the implications of the simple model with changing capital prices for saving rates across the wealth distribution. It plots two measures of the saving rate, corresponding to the two ways of writing the budget constraint: first, the saving rate net of capital gains $p_t\dot{k}_t/(w + \theta p_t k_t)$, corresponding to (4); and second, the saving rate out of total income including capital gains $\dot{a}_t/(w + r_t a_t)$, corresponding to (6). Importantly, the further up in the wealth distribution a household is, the larger are capital gains as a fraction of its income and hence the larger the difference between the two measures. Panel (a) considers the case of only persistent capital gains $\mu > 0$ and no transitory ones $\varepsilon = 0$. In this case, it is easy to see from (7) that the gross saving rate should be independent of wealth. Given that rich households experience larger capital gains, it follows from this that the net saving

rate should be *decreasing* with wealth.[9]

Panel (b) of Figure 2 considers the case where there are transitory capital gains in addition to persistent ones. In the case when $\varepsilon = 0$, meaning that the stock market had an "average year", the figure is the same as in panel (a). As explained above, households save all transitory capital gains meaning that these translate into higher wealth one-for-one. Therefore, when the stock market does well, $\varepsilon > 0$, gross saving is high for wealthy households and the gross saving rate is increasing with wealth. When the stock market does badly, $\varepsilon < 0$, the gross saving rate is instead decreasing with wealth.[10] In all cases, it can be seen from (7), that gross saving rates *purged from transitory capital gains* $(\dot{a}_t - \varepsilon_t a_t)/(w + r_t a_t - \varepsilon_t a_t)$ are independent of wealth. That is, the theory predicts that the *systematic* component of the gross saving rate should be independent of wealth. In contrast, the net saving rate is independent of stock market performance and is always decreasing.[11]

In summary, the simple consumption-saving model in this section predicts that the systematic component of the gross saving rate is independent of wealth. In contrast, if persistent capital gains are positive, the net saving rate should be decreasing with wealth.

2.2 Common Extensions

The above framework was deliberately simplistic so as to introduce our main savings concepts and highlight two properties that largely carry over to richer models typically applied in the macroeconomic literature. These properties are (i) the systematic component of the gross saving rate is approximately flat or slightly decreasing with net worth, and (ii) net saving rates are decreasing with net worth. We now move on to consider such richer model environments and substantiate this claim. In doing so, we will follow much of the macroeconomic literature and ignore changing asset prices. Hence, there is no difference between net and gross saving rates. For versions of these environments with changing asset prices, the

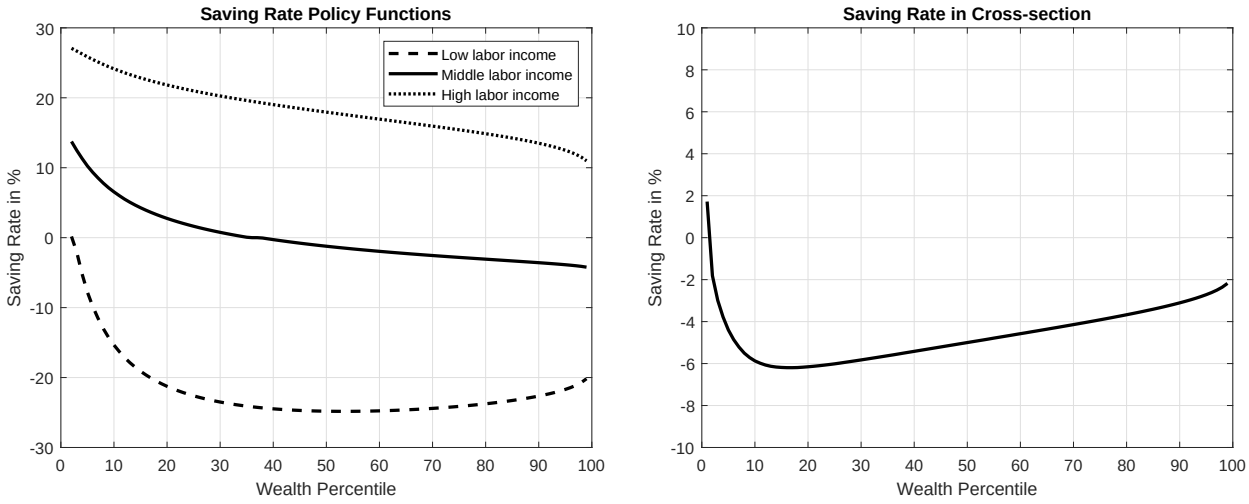
[9] Equation (7) shows that when $\varepsilon = 0$, $\dot{a} = \hat{s}(w + (\theta + \mu)a)$ where $\hat{s} = \bar{s}/\bar{r}$ (recall that with $\varepsilon = 0$ the approximation in (7) is exact). Therefore the gross saving rate is $\frac{\dot{a}}{w+(\theta+\mu)a} = \hat{s}$. Next, net saving is $p\dot{k} = \dot{a} - \mu a = \hat{s}(w + \theta a) - (1 - \hat{s})\mu a$ and therefore the net saving rate is $\frac{p\dot{k}}{w+\theta a} = \hat{s} - (1 - \hat{s})\frac{\mu a}{w+\theta a}$ which is declining in wealth a whenever $\hat{s} < 1, \mu > 0$ and $w > 0$.

[10] Equation (7) shows that $\dot{a} \approx \hat{s}(w + (\theta + \mu)a) + \varepsilon a$ where $\hat{s} = \bar{s}/\bar{r}$. Therefore the gross saving rate is $\frac{\hat{s}(w+(\theta+\mu)a)+\varepsilon a}{w+(\theta+\mu+\varepsilon)a}$ which equals $\hat{s}$ and is independent of wealth when $\varepsilon = 0$, is increasing in wealth a whenever $\varepsilon > 0$ and decreasing whenever $\varepsilon < 0$.

[11] To see this in the most intuitive fashion, note that from (4), net saving equals income excluding capital gains minus consumption, $p\dot{k} = w + \theta pk - c$. From (7) consumption c is independent of transitory capital gains ε and therefore so is net saving. The net saving rate is $p\dot{k}/(w+\theta pk)$ which is therefore also independent of market performance.

statements in this subsection apply to the systematic component of gross saving. The net saving rate would then follow residually, depending on what persistent capital gains happen to be, as illustrated in Figure 2.

A first natural extension is to allow for labor income risk and borrowing constraints, as in Aiyagari (1994) and Huggett (1993). The model and its calibration are described in more detail in Appendix A.2. The model generates a saving policy function $\dot{a} = s(a, w)$ where w is labor income. Figure 3 plots the resultant saving rate out of total income against wealth in this environment. The left panel displays the "saving rate policy function" $s(a, w)/(w + ra)$ for three different levels of labor income w. We see that conditional on labor income, the saving rate is declining in wealth, and more steeply so the closer a household is to the borrowing constraint. For households with different labor income realizations, the policy function simply shifts up or down. The steep decline close to the borrowing constraint reflects two familiar effects. First, households with high enough income realizations save intensively to build a buffer against potentially negative income shocks in the future. Second, households with low income realizations draw down their assets quickly when they are away from the borrowing constraint but eventually hit this constraint and cannot decumulate further (i.e. their saving rate is zero at the constraint but strongly negative above the constraint). The



(a) Policy function

(b) Saving across the wealth distribution

Figure 3: Saving rates under income risk and borrowing constraints

right panel displays saving rates without conditioning on income. That is, it simply plots the cross-sectional relationship between saving rates and wealth in the model's stationary distribution. The relationship is first decreasing, then increasing. This non-linearity reflects two opposing forces. On the one hand, conditional on income, saving rates are decreasing

with wealth, especially close to the constraint. On the other hand, there is an underlying positive correlation between income and wealth: in the stationary distribution, households with high income also have high wealth. And as shown in panel (a) saving rates increase with income. This effect dominates for higher wealth levels and underlies the upward-sloping part of the cross-sectional wealth-saving relationship.

To sum up, there are two takeaways from models with income risk and borrowing constraints: First, conditional on income, saving rates are predicted to decline slightly with wealth. Second, looking in the cross-section without conditioning on income, saving rates may be slightly increasing with wealth. But this increasing relationship is driven by correlation between labor income and wealth.

A natural second extension is to introduce a realistic life-cycle. In particular, it is well-understood that the age-profile for earnings creates very different incentives for saving at different points in the life-cycle. In a standard life-cycle model, households want to save little (or even borrow if possible) when they are young and have low levels of both income and wealth. Thereafter, as their age and income increases, households begin to save and their wealth builds up. Consequently, life-cycle considerations introduce a force toward cross-sectional correlation between saving rates and wealth. For our empirical purposes, we note that this tendency is due to the systematic relationship between age and income. Conditional upon age and income, the underlying forces toward flat or slightly decreasing saving rates, as we highlighted in the simple framework above, are still at work. The observation that conditional on age, saving rates decline with wealth in this class of models, has previously been noted by De Nardi and Fella (2017). Hence, exploring the role of age and income will be important when we turn to the data. In doing so, education levels will also be important, as we would expect education groups to systematically differ in their average life-cycle earnings profiles.

The third model extension that we explore is heterogeneity in individual-specific propensities to save, operationalized as heterogeneity in discount rates. This extension is widespread in applied macroeconomic work, and was emphasized by Krusell and Smith (1998) as a source of wealth dispersion in heterogeneous-agent macro models. Essentially, impatient individuals save little and consequently have low accumulated wealth. An immediate implication is a positive cross-sectional relationship between wealth and saving rates. For our empirical purposes, an important property of this mechanism is that saving rates are "type-dependent" (Gabaix et al., 2016), so that the cross-sectional patterns this mechanism generates should be explained by historical saving behavior at the individual level. We will therefore utilize

the time dimension of our data to gauge the empirical importance of this channel.

2.3 Housing

Our last benchmark model includes housing. Housing is an important asset for many households in the data and housing capital gains are one of the main components of gross saving. Housing differs from financial assets such as stocks in two dimensions. First, housing is not only an asset but also a consumption good. Second, housing is indivisible and subject to substantial adjustment costs. The goal of this section is to build a simple model to examine to what extent these two features can help explain our empirical findings. Economists often emphasize the consumption aspect of housing[12] and intuition suggests that this aspect could help explain why households in our data seem to hold on to their residential wealth in the face of rising house prices. The main result of this section is that, while housing is clearly different from financial assets, the consumption aspect by itself cannot explain our empirical findings. Instead other aspects of housing, like indivisibilities and transaction costs, must be part of the story.

To make this point, we construct a parsimonious model which features the idea that housing is both an asset and a consumption good but that *assumes away* that housing is indivisible and subject to transaction costs. That is, so as to isolate the role of the consumption aspect, we purposely make the counterfactual assumption that housing is fully divisible and not subject to any transaction costs. We then use this model to understand how households treat housing capital gains.

We assume that households receive a constant labor income w and invest in two assets: housing and bonds. The housing asset has two properties: it enters the utility function and its price changes deterministically over time. Households maximize

$$\int_0^\infty e^{-\rho t}\frac{C_t^{1-\gamma}}{1-\gamma}dt \quad \text{with} \quad C=\left(\alpha^{1/\eta}c^{\frac{\eta-1}{\eta}}+(1-\alpha)^{1/\eta}h^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}}, \tag{8}$$

where c is non-durable consumption, h is housing, α is the utility weight on non-durable consumption, and η is the elasticity of substitution between consumption and housing. We

[12] As Glaeser (2000) nicely puts it: "A house is both an asset and a necessary outlay. [...] If house prices double, they then have twice as much wealth, but their cost of living has risen by exactly the same amount. When my house rises in value, that may make me feel wealthier, but since I still need to consume housing there in the future, there is no sense in which I am actually any richer. And because house prices are themselves a major component of the cost of living, one cannot think of changes in housing costs in the same way as changes in the value of a stock market portfolio."

refer to C as "total consumption," i.e. consumption including both non-durable consumption and housing services. The budget constraint is

$$c + \dot{b} + p\dot{h} = w + rb, \tag{9}$$

where b are bonds and p is the house price. We can rewrite this budget constraint in terms of total (non-human) wealth $a := b + ph$ as $c + Rh + \dot{a} = w + ra$ where $R := rp - \dot{p}$ is the "implicit rental cost of living in the house" or the "user cost of housing."

Our goal is to understand theoretically how households treat housing capital gains in this environment which features the consumption-aspect of housing but no indivisibilities or transaction costs. Motivated by our results in Section 2.1, we analyze two cases: (1) a transitory housing capital gain, i.e. an increase in the house price that has no effect on the future price path; (2) a persistent housing capital gain, i.e. a persistent change in the growth rate of house prices.

Conveniently, the model in this section can be solved in closed form (see Appendix A.3).

Proposition 1 *Optimal total consumption and saving are given by*

$$C_t = \frac{m_t}{P_t}\left(\frac{w}{r} + a_t\right), \qquad \dot{a}_t = w + ra_t - m_t\left(\frac{w}{r} + a_t\right), \tag{10}$$

where $P_t := \left(\alpha + (1-\alpha)R_t^{1-\eta}\right)^{\frac{1}{1-\eta}}$, $R_t = rp_t - \dot{p}_t$ *and* $m_t := \left(\int_t^\infty (P_s/P_t)^{1-\frac{1}{\gamma}} e^{-\left(\frac{\rho-r}{\gamma}+r\right)s} ds\right)^{-1}$.

This Proposition underlies two distinct insights regarding how savings are predicted to respond to house price movements. The following corollaries summarize them.

Corollary 1 *Assume that* $r = \rho$ *and* $\dot{p} = 0$. *Then a surprise increase in the house price, i.e. a transitory housing capital gain, has no effect on total consumption and welfare,* $\partial C/\partial p = 0$, *and the housing capital gain translates one-for-one into higher wealth.*

This is Glaeser's point in footnote 12. In the event of a surprise increase in the house price, total consumption and welfare are unchanged, and households save all resulting capital gains. The reason is that although the market value of wealth increases, the price of their consumption basket rises by the exact same amount to offset the initial increase in wealth. By comparing Corollary 1 with the results in Section 2.1, one immediate implication is that households save all transitory capital gains irrespective of whether the capital gains come

from housing wealth or financial assets.[13]

For our second case, a persistent housing capital gain, the results are different.

Corollary 2 *The saving response to persistent housing capital gains is primarily determined by the intertemporal elasticity of substitution (IES)* $1/\gamma$*. When the IES equals one, the model collapses to the one-asset model in Section 2.1 in which the systematic component of the gross saving rate is independent of wealth. In particular, optimal saving in* (10) *becomes* $\dot{a} = \frac{r-\rho}{\gamma}\left(a + \frac{w}{r}\right)$*, which is identical to* (3)*. In this case, persistent housing capital gains have just the same effect on saving rates as persistent capital gains on financial assets or indeed changes in the interest rate. When the IES is close to one, these statements are approximately true.*

Most of these statements can be seen directly by examination of (10). When the IES $1/\gamma = 1$, saving is independent of prices and capital gains. The intuition is similar to how households respond to a change in the real interest rate. When house prices increase, there is an income and a substitution effect and these offset.[14] The elasticity of substitution between non-durable consumption and housing (η) has no separate effect on saving behavior, but only amplifies or weakens intertemporal substitution effects.[15]

Corrollary 2 implies that the consumption aspect of housing will *not*, by itself, result in households keeping housing and consumption unchanged in the face of rising house prices. Why is this the case? The answer is that this intuition ignores that, if housing were divisible and freely adjustable, households would engage in "intertemporal substitution of housing."

[13]To be clear, even though total consumption is fixed, there will be substitution between housing and non-durable consumption. If the elasticity of substitution between housing and non-durable consumption is strictly positive ($\eta > 0$), households increase non-durable consumption in response to higher house prices. They do this by selling housing and buying bonds in proportions to the reallocation between housing and consumption, keeping wealth fixed.

[14]First, a higher future house price implies that future consumption is more expensive relative to income. Hence, the lifetime consumption plan is more expensive and households want to save today in response to the increase in house prices (the income effect). Second, a higher future house price implies that consumption is cheaper today relative to tomorrow. Households therefore want to consume more today to reap the benefit of the low prices (the substitution effect). $1/\gamma$ governs the strength of the substitution effect. If $1/\gamma < 1$, the substitution effect is weaker than the income effect and households save in response to higher prices. If, on the other hand, $1/\gamma > 1$, the substitution effect is stronger than the income effect and households dissave in response to higher prices.

[15]A notable special case is $\eta = 1$ (Cobb-Douglas utility) where changes in relative prices have no effects on the total price of consumption because households adjusts their purchases to exactly offset the changes in relative prices. For low η, changes in house prices have a greater effect on the price index for total consumption because households are unwilling to substitute away from housing into non-durable consumption in response to higher prices. And if total prices change more in response to the future path of capital gains in housing, the intertemporal substitution effects above are also stronger. Conversely, a high η weakens the intertemporal substitution effects.

Figure 4 explains what we mean by this and shows how, in our model, consumption, wealth, and saving rates respond to a persistent housing capital gain. In the Figure, households

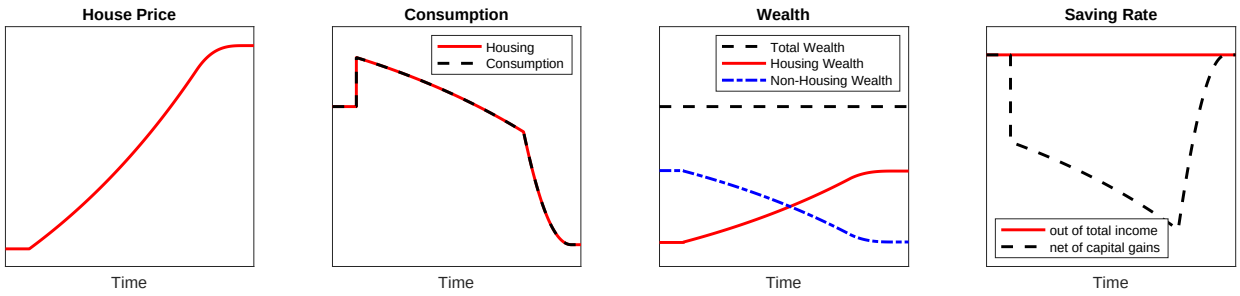


Figure 4: Response of Key Variables to a Persistent Housing Capital Gain.
Note: the Figure assumes $\eta = 0$, IES $= 1$, and $\rho = r = 0.05$.

suddenly learn that house prices will grow by two percent per year for the next twenty years. The figure assumes that the IES equals one and that the elasticity of substitution between housing and consumption η is zero (Leontief utility).[16] Because consumption and housing are perfect complements, they move together. On impact, they both increase because households purchase more housing in anticipation of the future price increases (which lower the user cost $R = rp - \dot{p}$). Along the house price path, households gradually sell off housing to finance its above normal consumption. Since the IES equals one, gross saving is always zero and the market value of wealth is constant. Households only re-balance their portfolio between housing and bonds, but never change the market value of wealth. After twenty years of house price growth, households have the same market value of wealth $b + ph$ with a larger share in the form of housing wealth ph, but less physical units of housing h.

Our results show that even if we internalize that housing is both an asset and a consumption good, it does not follow that saving behavior differs from the benchmark models. Households treat capital gains in housing similar to capital gains in financial assets. Unexpected capital gains in housing are always saved and households spend everything out of persistent capital gains in many standard calibrations (for example $\eta = 1$ or IES $= 1$), similar to how they react to capital gains in a model with only financial wealth. Instead, there are other features of housing that can generate substantially different saving behavior. For example, housing is indivisible and housing transactions entail substantial costs associated with moving, taxes, and brokerage fees. We will explore the role of transactions costs in Section 5.

[16]The figure also assumes that relative weight on non-durable consumption and housing α equals 0.5.

3 Data, Definitions and Institutional Setting

We base our study on Norwegian administrative data. Norway levies both income and wealth taxes on its households and the tax registry data provide a complete account of household income and balance sheets down to the single asset category over time for all Norwegian households. In contrast to other Scandinavian countries, Norway still has a wealth tax in place and the information in our data set is used to compute tax liabilities in practice. Below we describe the data sources in more detail, explain the saving rate measures we construct, and briefly characterize the Norwegian institutional setting.

3.1 Data

We link a set of Norwegian administrative registries, most of which are maintained by Statistics Norway. These data contain unique identifiers at the individual and household level. Our unit of observation is the household.[17] The registers we combine are a rich longitudinal database covering every resident (containing socioeconomic information including sex, age, marital status, family links, educational attainment and geographical identifiers), the individual tax registry, the Norwegian shareholder registry on listed and unlisted stock holdings, balance sheet data for listed and unlisted companies, and registries of housing transactions and ownership. All income flows are (calendar) yearly, and all assets are valued at the end of the year (December 31).

For our purposes, the Norwegian data feature a number of advantages. First, we observe wealth together with income at the household level for the entire population. Neither income nor wealth are top- or bottom-coded. The only sources of attrition are mortality and emigration out of Norway. Second, our data cover a long time period. This is particularly important to understand the role of capital gains in saving behavior. Third, the fact that much of the data is third-party reported limits the scope for tax evasion and other sources of measurement error from self-reporting.

Administrative Wealth and Income Records. Each year, Norwegians must provide complete information about their incomes and wealth holdings to the tax authorities. Wealth taxes are levied on the household as a whole, whereas income taxes are levied on each

[17] Arguably, saving and consumption decisions are made at the household level, which in the data is defined as either a single person household or a married couple or a couple cohabiting with common children. In effect we compute variables at the household level in per capita terms but then weight households by the number of adults in the household.

individual household member separately.[18] Available tax record data are annual and go back to 1993. However, we primarily use data from 2004 to 2015 since the shareholder registry is only available since 2004. Tax information is primarily collected via third parties. Employers submit information on earned labor income to the tax authority. Financial intermediaries (banks, insurance companies, fund brokers, etc.) submit the value of and income from individuals' asset holdings. If an asset is traded, it is the market value that is reported.

We impose a few minor sample selection restrictions in order to reduce errors in the computation of saving rates and wealth. The sample is limited to households with adults above twenty years of age. We drop household-year observations where income is lower than the base amount defined in the Norwegian Social Insurance Scheme (NOK 75,641 in 2010), where the household has immigrated within the last two years, or where the household is either formed or dissolved.[19]

Asset Classes. Our data provide wealth information decomposed by asset class and total debt. An illustration of these data is provided in Figure 5, where portfolio shares are plotted by wealth percentiles. Here, as in the rest of our paper, wealth is defined as the value of all assets minus debt outstanding.

The tax data have separate entries for bank deposits, cash holdings, informal loans, and bond holdings (these include direct holdings of government and corporate bonds, but are mostly consisting of holdings in bond mutual funds, and money market funds[20]). In Figure 5, we collect these together into "safe assets." Vehicles is a separate category containing the estimated value of a household's stock of cars and boats. Their values are calculated based on a valuation schedule depending on list price as newly purchased and the age of the vehicle. Public equity is the directly observed market value of a household's portfolio of listed stocks. The tax records contain the market value of direct stock holdings and the market value of stock funds separately. In Figure 5 we report the sum of these two under the heading "public equity."

"Private business" wealth refers to ownership of firms that are unlisted. Unlike publicly traded shares, no clear-cut market value is observed for equity in unlisted firms, and we

[18]Net wealth above an exemption threshold is taxed at a flat rate which has hovered around 1% during our sample period. The exemption threshold has increased over time. Toward the end of our sample the threshold is approximately NOK 1.5 million for a married couple (and half that for a single person). Importantly for our purposes, assets are reported and recorded even if their total value falls short of the threshold. For several asset classes, tax values are discounted before taxable wealth is computed. We adjust these discounted values back to market values when computing household wealth.

[19]We do not drop households who transition from a status as cohabitants to married.

[20]More disaggregated data from the later years of the sample allow us to see this.

instead apply the "assessed value" which private businesses (by law) must report to the tax authority. The tax authority in turn distributes this value to the shareholders of the firm according to ownership share. This assessed value is closely related to the book value of equity, but omits intangibles (goodwill). As wealth taxation might motivate owners to under-report their firm's true value, the tax authority has control routines to identify under-reporting. Medium- and large-sized firms (with a turnover above about USD 500k) are also required to have their balance sheet audited by an approved auditing entity.

Housing wealth includes the value of a household's principal residence, secondary homes and recreation estates (including "rorbu"). To obtain an up-to date value of households' housing wealth the tax authority have since 2010 updated the value of housing wealth stock with predicted values from area- and housing type specific hedonic regressions of price per square meter of houses transacted during the year on house characteristics. Prior to 2009 the value of housing wealth in the tax records are closely related to the original purchase price, and often far removed from any current valuation. Since housing is such an important asset for the households (Figure 5), obtaining a precise as possible estimate of it is crucial for our purposes. To this end, we combine official transaction data from the Norwegian Mapping Authority (Kartverket), the land registration office, and the population census, which together enable us to identify ownership of every single housing unit and its precise location. We then estimate the price per square meter as a function of house characteristics such as the number of rooms and bathrooms, location, time periods, and their interactions using machine learning tools. More precisely, we utilize an ensemble method combining a random forest algorithm, a regression tree, and LASSO as in Mullainathan and Spiess (2017). A detailed exposition of our house price estimates, which goes beyond the tax authority's own methodology, is described in Fagereng, Holm and Torstensen (2019). We use the predicted values from this procedure as housing values over our sample period.

Debt consists of mortgages, student debt, consumer debt, and personal loans. Figure 5 shows that the ratio of debt to total assets is high across almost the entire wealth distribution.

The portfolio shares in Figure 5 reflect how Norwegians' balance sheets are dominated by housing and debt. They plot the average holdings of different asset classes as a share of total assets (i.e. as opposed to net worth so that the denominator is positive) for different net worth percentiles. To the left, we see households with negative net worth. By definition, these are households whose liabilities exceed their assets which is reflected in a ratio of debt to assets larger than one-hundred percent (the solid green line with dot markers exceeds 100% and is hence out of the figure's range for these households). The least wealthy households

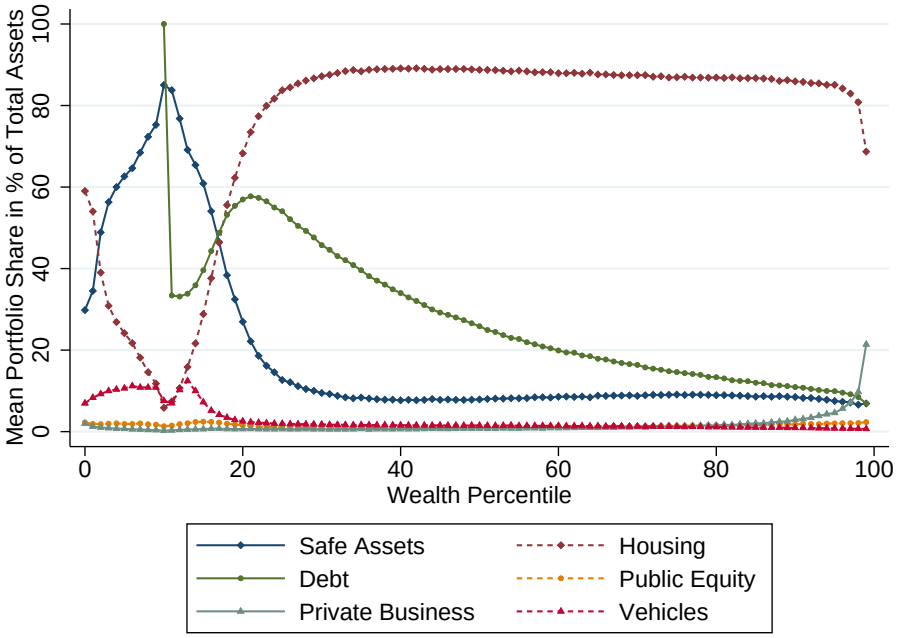


Figure 5: Asset Class Shares in Household Portfolios

own a high ratio of housing wealth to total assets but their debt exceeds their assets. As we move further to the right, the housing share declines steadily until we reach households with approximately zero net worth at about the 15th percentile of the wealth distribution. These households hold approximately no assets and no debt. Thereafter, from the 15th to approximately the 25th wealth percentile, the housing portfolio share grows rapidly. As documented in the Appendix, this increase goes together with an increase in home ownership. Above the 30th wealth percentile home ownership is almost 100% in Norway, and we see in Figure 5 that the portfolio share is almost flat around 85% from the 30th wealth percentile and up to the top 3-4 wealth percentiles.

Besides the dominant role of housing, another striking feature of Figure 5 is the role of safe assets, which predominantly consist of bank deposits. Apart from in the very top of the wealth distribution, safe assets have the second largest portfolio share. Public equity is a small component of Norwegian households' wealth. As shown in the Appendix, only around 25% of Norwegian households hold public equity directly or via stock funds. Finally, we see that private equity becomes an important asset class in the top part of the wealth distribution. The rise in private business wealth among the richest is what drives the housing share down in this part of the wealth distribution.

One reason why the portfolio share of public equity is so low in Figure 5, is that most of household's pension wealth is excluded. As explained in Section 3.3 below, the Norwegian

government provides a relatively generous pension scheme. In addition, some workers have additional private retirement accounts held by their employers, so as to top up the public pension plan which is capped. In extensions of our baseline analysis, we impute each household's public pension entitlement and include it in both its saving and wealth measures. For details on how this affects the plot of portfolio shares against wealth, see the Appendix.

Table 1: **Descriptive Statistics**

	Mean	Std. dev	P10	Median	P90
Age	46.40	15.00	26	46	67
Male	0.51	0.50	0	1	1
Married	0.51	0.50	0	1	1
Safe assets	41,253	137,880	726	12,737	96,921
Housing	441,798	522,941	0	343,893	915,667
Debt	117,426	160,217	0	71,267	282,418
Public equity	7,472	71,192	0	0	12,993
Private equity	21,504	555,650	0	0	0
Vehicles	7,770	33,407	0	2,616	20,735
Net wealth	402,371	891,507	-11,413	270,907	892,641
Gross labor earnings	57,372	50,356	0	57,403	107,585
Change in net worth	22,464	469,244	-70,594	7,950	132,706

Notes: The table reports summary statistics for demographic characteristics of households pooling data for our sample period 2004-15 with a total of 35,806,100 individual-year observations. Values in USD, 2011 prices, using the 2011-average exchange rate between USD and NOK in 2011 (NOK/USD = 5.77).

3.2 Key Concepts and Definitions with Changing Asset Prices

We now discuss in more detail some key conceptual issues due to the presence of changing asset prices. First, we define in more detail what we mean by "net" and "gross" saving and by "net-" and "gross saving rates." We also introduce a concept we term "recurrent" saving which aims to isolate the systematic (rather than cyclical) component of gross saving. Second, we discuss how we operationalize these concepts in the data and discuss our empirical strategy for separating gross saving into net saving and capital gains.

Net, Gross and Recurrent Saving. We find it useful to distinguish between three different notions of saving: net, gross and recurrent saving. In order to define these in the

most transparent manner we make use of the simple model with changing asset prices from Section 2.1. While this simple model features only one asset, all our definitions extend to the case of multiple assets (and liabilities) in the obvious fashion – see Appendix B.1.

Our different notions of saving rates are defined by different ways of writing the budget constraint "consumption plus saving equals income" and, in particular, by how capital gains are treated when writing this accounting identity. Two of these definitions can be gauged from the two budget constraints (4) and (6) which we here restate for convenience (the latter in slightly altered form):

$$c + \underbrace{p\dot{k}}_{\text{net saving}} = \underbrace{w + \theta pk}_{\text{disposable income}}, \tag{11}$$

$$c + \underbrace{p\dot{k} + \dot{p}k}_{\text{gross saving}} = \underbrace{w + (\theta + \dot{p}/p)pk}_{\text{Haig-Simons income}}. \tag{12}$$

The difference between the two accounting identities is simply that the latter adds capital gains $\dot{p}k$ on both sides. Importantly, consumption in the two equations is the same. Then, for the identity "consumption plus saving equals income" to hold, a difference in the saving definition necessarily implies a difference in the income definition. Formulation (11) features disposable income whereas formulation (12) features "Haig-Simons" income which, in particular, counts unrealized capital gains as a form of income (von Schanz, 1896; Haig, 1921; Simons, 1938). Finally, we define the "net saving rate" as the ratio of net saving to disposable income in (11), and the "gross saving rate" as the ratio of gross saving to Haig-Simons income in (12).

Importantly, standard economic theories – and in particular all workhorse one-asset models like those in Section 2 – only have implications for gross saving in (12) and therefore for gross saving rates. However, as we have shown in Section 2.1, these implications differ greatly according to whether capital gains are transitory or persistent: according to theory, households should not consume out of transitory capital gains which should translate one-for-one into higher wealth, but they should consume out of persistent capital gains. An immediate corollary of this is that we expect the gross saving rate to vary strongly with asset market performance: it should be strongly increasing with wealth in good years and strongly decreasing with wealth in bad years (see Figure 2(b)). At the same time, another prediction of our simple theory was that the systematic component of gross saving rates, i.e. the gross saving rate purged of transitory capital gains, should be independent of wealth.

To provide a sharper test for economic theories, we therefore introduce a third notion of

saving (and income) that we term "recurrent saving" and which aims to isolate exactly this systematic component of gross saving. This notion is defined via the budget constraint

$$c + \underbrace{(\dot{k}/k + \mu)pk}_{\text{recurrent saving}} = \underbrace{w + (\theta + \mu)pk}_{\text{recurrent income}}, \tag{13}$$

where $\mu := \overline{\dot{p}/p}$ is a measure of the systematic or persistent component of capital gains (defined in more detail in the next subsection). Recalling the decomposition of capital gains into persistent and transitory components $\dot{p}/p = \mu + \varepsilon$, the key difference to gross saving in (12) is that recurrent saving includes only the persistent component of capital gains μ but not the transitory component ε. Equivalently, recurrent saving is gross saving purged of transitory capital gains, that is $\dot{a} - \varepsilon a$. The definition of recurrent income is implied by the requirement that consumption plus saving equals income.[21] The recurrent saving rate is recurrent saving as a fraction of recurrent income.

The main advantage of working with recurrent saving rates is that the predictions of theory are much clearer than for gross saving rates. In particular, our simple benchmark model with changing asset prices predicts that the recurrent saving rate should be approximately independent of wealth in every year, regardless of stock market performance. When persistent capital gains are positive $\mu > 0$, the situation is depicted in Figure 2(b) by simply relabeling "gross saving rate, average year $\varepsilon = 0$" as "recurrent saving rate."

Owner-Occupied Housing. When the asset under consideration is owner-occupied housing we adjust the definitions of consumption and income in (11), (12) and (13). We adopt a rental equivalence approach that aims to value owner-occupied housing services as the rental income the homeowner could have received if the house had been rented to a tenant. We then add this implicit rental rate to both consumption and income in all three variants of the budget constraint.[22] We follow Eika, Mogstad and Vestad (2018) and distribute the aggregate value of owner-occupied housing services from the National Accounts across

[21] Interestingly, our definition of "recurrent income" arguably coincides with that of Hicks (1939) who defined income as "the maximum value which he can consume during a week, and still *expect* to be as well off at the end of the week as he was at the beginning" (emphasis added).

[22] For example, consider the budget constraint in terms of gross variables (12). Using the notation of Section 2.3 and denoting housing by h, the house price by p, other assets by b, total net worth by $a = b + ph$ and the implicit rental rate of housing by R (for example $Rh = 2.88\% \times ph$ as explained momentarily) the analogue of (12) is

$$\underbrace{c + Rh}_{\text{consumption}} + \underbrace{\dot{p}h + p\dot{h} + \dot{b}}_{\text{gross saving}} = \underbrace{w + rb + Rh + \dot{p}h - \delta ph}_{\text{Haig-Simons income}}.$$

households proportional to the value of their house, which implies a constant rent-to-value ratio across households. We set this ratio equal to 2.88 percent, the average rent-to-value ratio for 1994-2014.

Separating Gross Saving into Net Saving and Capital Gains. As discussed in Section 3.1 our data contain information on the stocks of assets and liabilities at the end of each calendar year. Therefore, we observe the year-to-year change in each household's net worth, that is its gross saving. In order to compute our measures of saving rates we need to separate gross saving into net saving and capital gains. We here briefly summarize our approach. A more detailed description of this important step is in Appendix B.2.

Our approach differs by asset class and we therefore discuss each asset class separately. As already discussed, the most important asset class as a share of total assets is housing (see Figure 5). To separate gross saving in housing into net saving and capital gains, we make use of the housing transaction data we already used to compute house values. For each household we observe whether it engaged in any housing transactions and, if it did, the value of each housing transaction (recorded at the date of the transaction) throughout the year. For households without housing transactions, i.e. households who did not buy or sell a house, we attribute all changes in the value of housing to capital gains. For households with housing transactions, we compute net saving as the net sum of the value of all housing transactions throughout the year and capital gains as the difference between gross and net saving. A precise mathematical version of this decomposition is in equation (A37) in the Appendix.

Next consider public equity which itself consists of directly held stocks and stock funds. For directly held stocks, we make use of the Norwegian shareholder registry which contains information on holdings of individual stocks at the security level. Since all stocks are publicly traded and we therefore observe stock price changes, we can thus compute a measure of capital gains for each individual stock. Net saving for a particular stock is then simply the change in the value of holdings of that stock minus capital gains. Aggregating across all stocks held by an individual, we thus have an individual-level measure of net saving and capital gains in directly held stocks. Unfortunately, information on individual stock holdings at the security-level is not available for stock funds. We therefore measure stock fund capital gains from the Financial Accounts and attribute the same capital gains (in percentage terms) to each individual.

As already discussed, for private equity asset values do not equal market values and

instead are closer to book values. If book values are not marked to market, changes in book value in principle measure net saving. We therefore do not attempt to separate changes in private equity values into net saving and capital gains. Since private equity is an important asset at the very top of the wealth distribution but not elsewhere (Figure 5) this likely results in some measurement error at the very top.

Finally, for bonds we compute capital gains as a weighted average of t-bill rates and medium term bond yields. Decomposed holding data from the Norwegian Fund and Asset Management Association reveal that households holdings in bond and money market fund are concentrated in t-bills with maturity less than 1 year and medium term bonds (2-4 years).[23] All other assets listed in Section 3.1 as well as liabilities do not experience asset price changes so that capital gains are zero and net saving equals gross saving for those assets.[24]

Persistent Capital Gains and Recurrent Saving. We have defined recurrent saving as the systematic component of gross saving or, equivalently, gross saving purged of transitory capital gains. In order to operationalize this concept, we need to split individual capital gains on different assets into persistent and transitory components. That is, to operationalize equation (13), we need to come up with an asset-specific estimate of μ.

Across asset classes, our approach follows three general principles. First, we generally compute the persistent component of capital gains μ as a simple geometric average of year-to-year capital gains over an extended time period.[25] Second, to minimize dependence of μ on episodes of particularly good or bad asset market performance, we generally take as long a time series as possible: ideally from 1950 to 2015, or otherwise as long as the relevant data go back. Third, we compute this persistent component based on asset characteristics rather than computing a household-specific μ.

Table 2 summarizes the results of this exercise by asset class and we now discuss it in more detail. Considering public equity, we compute the persistent component of capital

[23] With an approximate weight of 70% in the short term funds. See https://www.vff.no/about-the-norwegian-fund-and-asset-management-association. The rates are available at https://www.norges-bank.no/en/Statistics/Interest-rates/Treasury-bills-annual/.

[24] Vehicles depreciate over time and we use list price changes to infer the depreciation rates (see Section 3.1). However, we count depreciation of an asset as as decline in the physical units of that asset, that is negative net saving as opposed to a capital loss (i.e. merely a revaluation effect). Therefore, also for vehicles, net saving equals gross saving. Alternatively, we have conducted our main empirical exercise counting vehicle depreciation as a capital loss and results are largely unchanged.

[25] That is, $1+\mu = \left(\prod_{t=0}^{T-1}(p_{t+1}/p_t)\right)^{1/T}$ if the geometric average is computed over T years.

Table 2: Persistent capital gains for key asset classes

Asset	Average capital gain
Directly Held Stocks	3.25 %
Stock Funds	3.18 %
Housing	2.25 %

gains on directly held stocks as the average appreciation of a Norwegian stock index, the OBX Price Index (OBXP), since it came into existence in 1991.[26] Similarly, the persistent component of stock fund capital gains is the geometric average of their yearly counterpart from the national accounts over the period for which this data series is available (since 1995).

For housing, we compute different persistent components for different geographical areas starting in 1950. We use historical house price indices for the Oslo, Bergen, Trondheim and Kristiansand metropolitan areas provided by Norges Bank.[27] The calculated persistent housing capital gains are 2.22% for Oslo, 2.78% for Bergen, 3.16% for Trondheim and 2.82% for Kristiansand, meaning there is some, but not much, geographical variation. In particular, these numbers are close to the 2.25% yearly growth in an aggregate index that ignores this variation (this is the number in column (1) of Table 2). For housing, it is particularly important that we take a long-run perspective and compute persistent capital gains using a long time series going back to 1950. The reason is that, during our sample period from 2004 to 2015, the Norwegian housing market went through a period of abnormally high house price growth. For example, between 2004 and 2015, house prices grew on average by 4.80% and by 6.28% in Oslo. Therefore, for most households in our sample, realized house price growth was considerably above persistent housing capital gains computed as just described. It is easy to see from (12) and (13) that this has the effect that recurrent saving in housing will, on average, be below gross saving. In effect, when calculating recurrent saving rates, we downward-adjust observed housing capital gains, thereby ensuring that our results are not driven by the abnormally high house price growth observed during our sample period. This adjustment mechanically drives a wedge between (time-averaged) gross saving rates and recurrent saving rates and this wedge will be reflected in our main results, in particular Figure 6.

[26] The OBXP Index is a capitalization-weighted pricing index of the largest companies traded on the Oslo Stock Exchange. More information is available at https://www.oslobors.no/ob_eng/markedsaktivitet/#/details/OBXP.OSE.

[27] This series is originally due to Eitrheim and Erlandsen (2004) and is updated annually by Norges Bank and available at https://www.norges-bank.no/en/Statistics/Historical-monetary-statistics/House-price-indices/.

Column (1) of Table 2 summarizes persistent capital gains for these asset classes. In column (2) we also report average rates of return, i.e. adding dividend yields to capital gains. Reassuringly, we find the usual ordering for asset returns: stock returns are higher than housing returns.

3.3 Institutional Setting

Norway is characterized by a generous welfare state that provides relatively rich insurance against income loss, illness and other life events. Services such as child care, education and health care are heavily subsidized and provided at low or no cost. On the income side, this extended welfare state is financed both through relatively high taxes (the ratio of tax income to GDP was 38% in 2016) and through returns from a sovereign wealth fund, Norwegian Pension Fund Global.[28] The main role of this fund is to allow (smooth) public spending in excess of tax revenues, and remove uncertainty regarding future governments' ability to meet public pension entitlements.

The public pension system was reformed in 2010, moving from a defined benefits to a defined contribution system. We provide details in the Appendix, and outline only the main features here. Cohorts born after 1963 are entirely under the new schedule. Here individuals accumulate a pension endowment with 18.3% of their labor income each year, starting at age thirteen. Yearly accumulation is capped at $7.1G$, where G is the base amount defined in the Norwegian Social Insurance Scheme ("Grunnbelopet"). In 2010, $1G$ was NOK 75,641. After retirement, the accumulated endowment is paid out at a flat yearly rate, scaled by an adjustment factor which increases in retirement age (reaching a maximum at age 67) and decreases in cohort-specific life expectancy. For cohorts born before 1954, the schedule is more complicated. The pension endowment is accumulated in proportion to annual income, but the speed of accumulation drops to 1/3 when income exceeds $3G$, before dropping to zero at $12G$ (meaning a maximum of $7G$ can be accumulated per year). The final endowment is then computed based on the individual's twenty best-earning years. In addition, under the old system all public sector employees and approximately half the private sector employees had an additional "service pension". With a service pension, yearly benefits after retirement are either (i) the individual's pension endowment scaled proportionately by her number

[28] See https://www.wider.unu.edu/project/government-revenue-dataset for how the tax-to-GDP ratio has evolved over time. The Norwegian Pension Fund Global is part of Norway's fiscal mechanism for smoothing the use of national oil revenues. The mechanism postulates that the flow government income from oil activity is invested abroad, while an estimated normal return (4.5% previously, 3.5% currently) on the existing fund may be spent each year. Currently, the fund is worth about three times Norwegian GDP.

of working years, or (ii) 66% of their last working year's income, whichever is greatest. Without a service pension, the individual's yearly public pension benefits are computed as in (i). Finally, for cohorts born between 1954 and 1963, benefits were computed as a weighted average of the two systems, with weights depending on birth year. For instance, the 1962 cohort is 9/10 under the new system and 1/10 under the old system.

4 Results

This section presents our main results. We start by presenting descriptive plots of the gross, net, and recurrent saving rates against wealth percentiles. We then move on to control for the main factors that are important for this correlation according to theory: income, age, education, and individual-fixed propensity to save. Thereafter, we address the role of housing, the dominant asset on household balance sheets, in explaining our findings. Finally, we discuss time-aggregation and examine a "consumption-equivalent" saving rate so as to better understand the welfare consequences of our findings.

4.1 Saving Rates across the Wealth Distribution

Figure 6 plots our three saving rates against percentiles of net worth. For every year, we have computed the median saving rate within each of that year's wealth percentiles. The three lines display the average of these within-year medians over our sample period. An alternative approach to time-aggregation is considered in Section **??**. We start with the net saving rate, for which our benchmark theories in Section 2 predicted a declining relationship with wealth. To the left, we see that among the least wealthy households, net savings decline with wealth. Then, after reaching a trough just below the 15th wealth percentile, it rises up to about the 25th wealth percentile, before it flattens out almost completely around $7-8\%$. Note that the trough is reached approximately where net worth is zero (see Figure 5). Hence, the net saving rate declines with wealth for households with negative net worth, increases among households with moderate net worth, and thereafter is flat over most of the wealth distribution.

Consider next the gross saving rate. As capital gains have been positive for most of the period considered, the gross saving rate necessarily lies above the net saving rate. Our interest lies in its shape. To the left, we see that it decreases just like the net saving rate. Then, after reaching a trough where households' wealth is approximately zero, it diverges

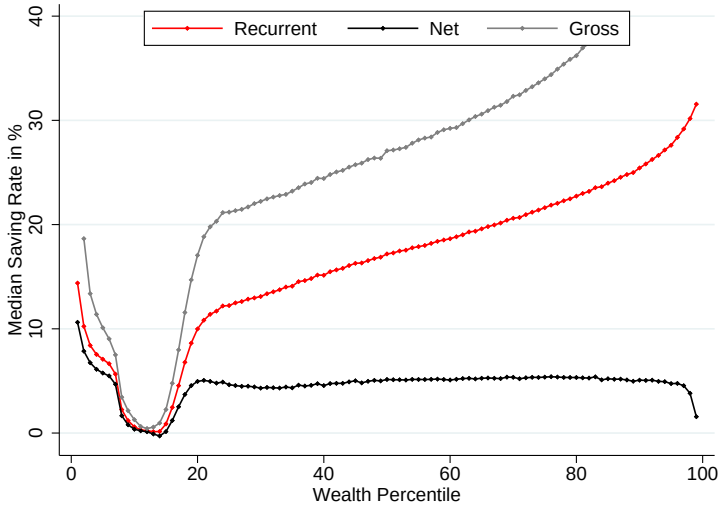


Figure 6: Saving rates across the wealth distribution.

from the net saving rate by rising almost monotonically up to the very wealthiest households. The increase is steepest among households immediately above the savings trough, tapers off somewhat, and then picks up again among the $10-15\%$ wealthiest households. To simplify exposition, we coin this distinct pattern a "swoosh" (in analogy with the logo of a well-known footwear brand).

Our third saving rate, recurrent saving, lies between the net and gross rates because recurrent capital gains are positive, but lower than our sample period's realized capital gains. However, its shape is approximately identical to that of the gross saving rate. Hence, the swoosh shape of gross saving is not driven by abnormally high capital gains in our sample period.

The patterns for gross and recurrent saving rates tell us that wealthier households save a greater fraction of their incomes (including capital gains) than poorer households. Moreover, the divergence between gross or recurrent saving rates and net saving rates tells us something important about the mechanism behind this finding: wealthy households save more by holding on to assets that experience recurrent capital gains.

This mechanism has two logical parts. First, wealthier households hold more assets by definition. Hence, unless they systematically invest into assets with lower capital gains than poorer households, their assets automatically appreciate by a larger dollar amount. Second, households tend to save a higher fraction of these capital gains than of other sources of income. This second part of the mechanism poses a challenge to the workhorse models summarized in Section 2, and is the precise reason why gross saving behaves differently in

the data than in these models.

A more general takeaway from Figure 6 is that asset price changes and how people respond to these, are key to understanding cross-sectional wealth dynamics. In Section 2.1, we moved in this direction by extending a workhorse model of household saving with asset price changes that are either persistent or transitory. As that section highlighted, the implications of these two types of shocks are markedly different. If we compare our results in Figure 6 to the theoretical predictions in Figure 2, we see that the workhorse model with asset price movements actually proposes an explanation of why gross saving rates increase with wealth while net saving rates are flat: if capital gains are purely transitory, households should optimally save them. However, our findings for recurrent saving rates show that this cannot be the explanation. Gross saving rates are not just increasing with wealth for one particular year. Instead, capital gains have a systematic positive component, computed using a long time series, and also the systematic component of gross saving rates is increasing with wealth. A remaining possibility is that households are subject to "expectational errors" and misperceive persistent capital gains as purely transitory. That is, they are repeatedly surprised by the fact that their assets appreciate in line with the recurring historical average. We return to this candidate explanation in Section 5 but argue there that it is unlikely to explain our results in their entirety.

We next turn to a deeper empirical assessment of our findings. In a nutshell, there are three main features of Figure 6 that we wish to understand deeper: (1) all saving rates are decreasing with wealth for households with negative net worth; (2) among households with positive net worth, the net saving rate tends to be flat in wealth; (3) among households with positive net worth, the gross and recurrent saving rates are increasing with wealth.

4.2 The Role of Age, Earnings and Education

As highlighted in Section 2.2, common extensions of our simplest consumption-saving model predict that, in the cross-section, saving rates should be correlated with wealth because both are correlated with labor income and age. In particular, earnings risk may motivate households with high income realizations to save more, thereby inducing a positive correlation between saving rates and wealth. Similarly, life-cycle considerations may lead to a positive relationship between saving rates and wealth if wealth is correlated with age. In this section, we therefore explore to what extent the patterns in Figure 6 can be accounted for by these factors. We focus on the net saving rate as well as the recurrent saving rate, which is the

empirical counterpart to the systematic component of the gross saving rate in theoretical models.

We utilize two approaches: (1) we control parametrically for the covariates in a quantile regression, and (2) we non-parametrically explore saving patterns within groups defined by age, earnings and education. In the first approach, we use the following specification:

$$\frac{s_{it}}{y_{it}} = \phi_1 + \sum_{p=2}^{100} \phi_p D_{it,p} + f(x_{it}) + \tau_t + \epsilon_{it}, \tag{14}$$

where s_{it}/y_{it} is the net or recurrent saving rate, $D_{it,p}$ is a dummy for being in wealth percentile p at the beginning of the year, ϕ_p is the corresponding regression coefficient, x_{it} is the vector of control variables which enters the regression in a flexible fashion as captured by the function $f(\cdot)$, τ_t are time fixed effects, and ϵ_{it} is an error term. In the results presented here, we specify $f(\cdot)$ as a fourth-order polynomial in age, fixed effects for years of education, and a linear function of the natural logarithm of household labor income. The ϕ_p coefficients then represent the object of our interest: the median saving rates across the wealth distribution after controlling for age, earnings, and education. These can then be examined graphically, analogous to the results in Figure 6.

Figure 7 presents the results. It shows the median net and recurrent saving rates after

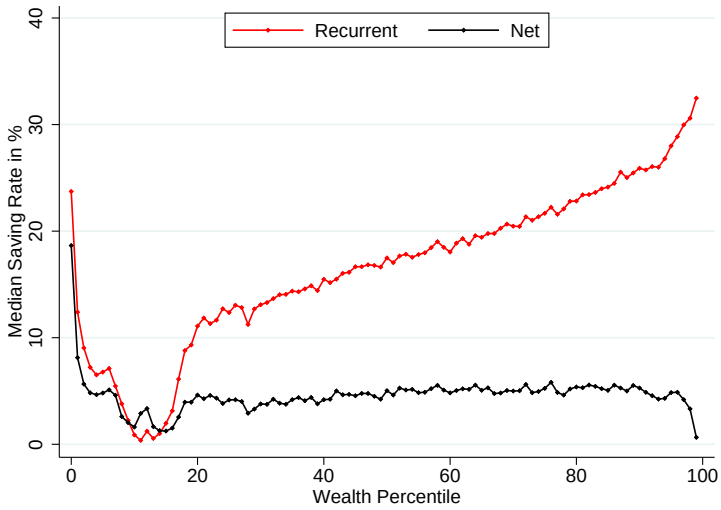


Figure 7: Saving rates, controlling for age, earnings and education

controlling for age, earnings and education through equation (14). Comparison with Figure 6 shows that the resulting saving rates conditional on age, earnings and education are qualitatively very similar to their unconditional analogues. Consistent with the discussion

in Section 2.2, the recurrent saving rate is slightly flatter after controlling for the three variables. This indicates that some of the correlation between saving rates and wealth is due to age, earnings, and education. Still, the main takeaway is that also conditional on these observables, the savings graph maintains its main characteristic: a "swoosh-shaped" recurrent saving rate and a decreasing and then flat net saving rate.

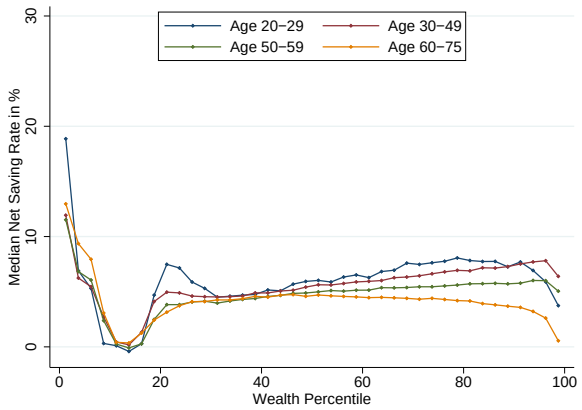
Our second approach is to produce our main graph within groups defined by age, earnings and education. For age we stratify households into four groups (20-29, 30-49, 50-59, and 60-75 years), for earnings we stratify them into earnings deciles, and for education we stratify into three groups (no high school, high school, college). Within each group, we then calculate the median saving rate for different wealth percentiles, just like in Figure 6. Figure 8 shows the median net and recurrent saving rates within groups. Again we observe decreasing and then flat net saving rates and "swooshes" for recurrent saving rates. These patterns tend to hold for all stratifications considered. Conditioning on age, earnings and education has large effects on the *levels* of the saving rates, but only modest effects on how these *covary* with wealth.

4.3 Are rich households rich because they have high saving rates?

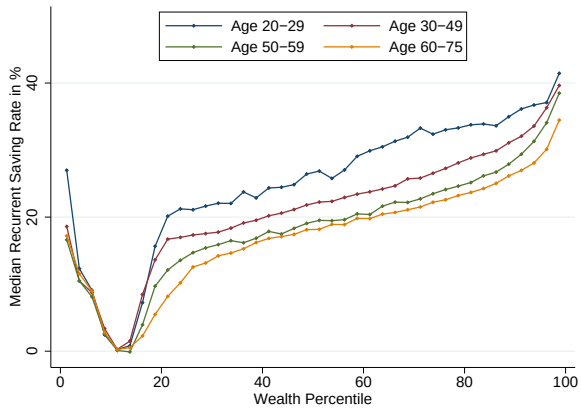
Given that recurrent saving rates are positively correlated with wealth, it is natural to ask: does the saving rate slope upward simply because households who save a lot become rich? Our discussion of benchmark models in Section 2.2 suggests this as a relevant channel. Models with discount rate heterogeneity predict that relatively patient households have higher saving rates and, over time, therefore become wealthier than impatient ones.

We factor out the role of such "type-dependence" in explaining our findings by exploiting the panel dimension of our data. Our approach is as follows. From year five of our sample period and onward, we compute each household's mean saving rate over the previous years (at least four years). This yields one saving rate per household per year. Thereafter, for each year we stratify households by their decile of past saving rates. Within each historical saving rate decile, we compute the median saving rate by wealth percentile. Finally, we construct our main graph by averaging the group-specific medians over years.

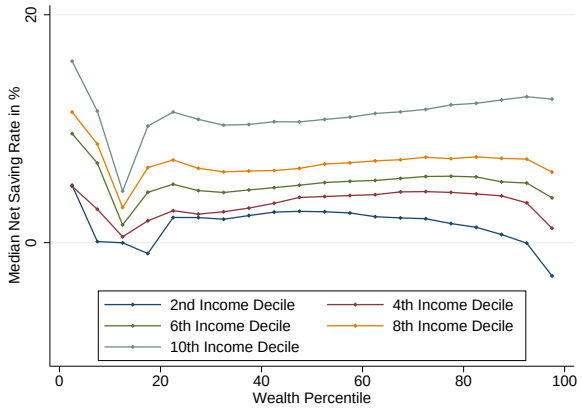
Figure 9 presents the results. We observe that the swoosh-shaped pattern for recurrent saving rates arises even within groups defined by their savings history. Hence, the increasing part of the swoosh shape is not driven by the plausible hypothesis that persistent differences in saving rates cause wealth differences.



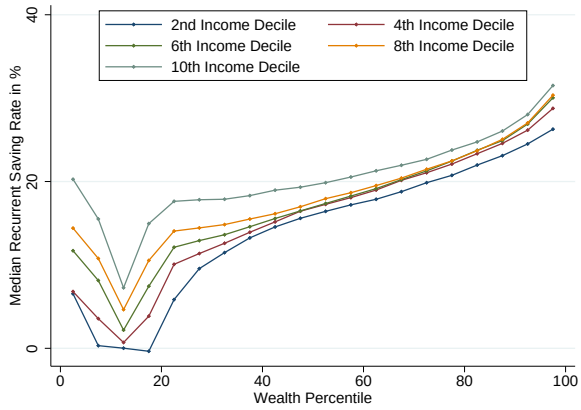
(a) Age, net saving rate



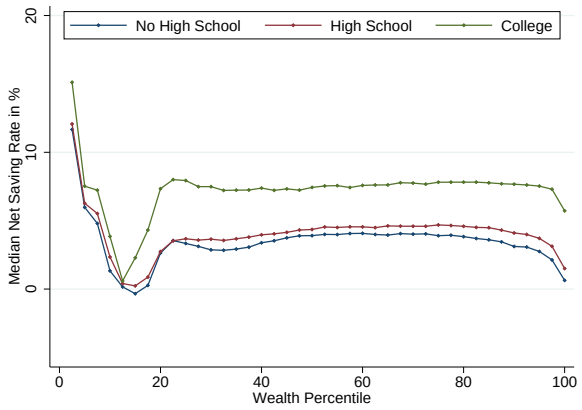
(b) Age, gross saving rate



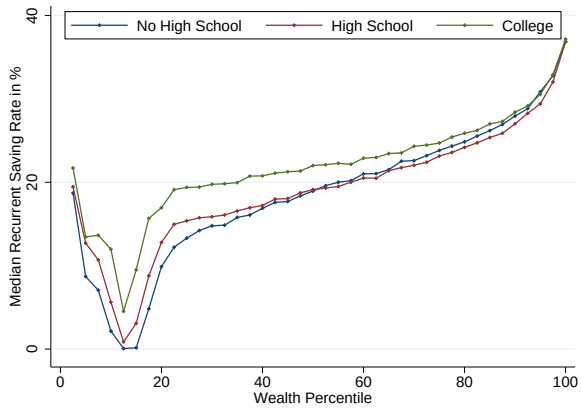
(c) Earnings, net saving rate



(d) Earnings, gross saving rate



(e) Education, net saving rate



(f) Education, gross saving rate

Figure 8: Saving rates across the wealth distribution by age, earnings and education

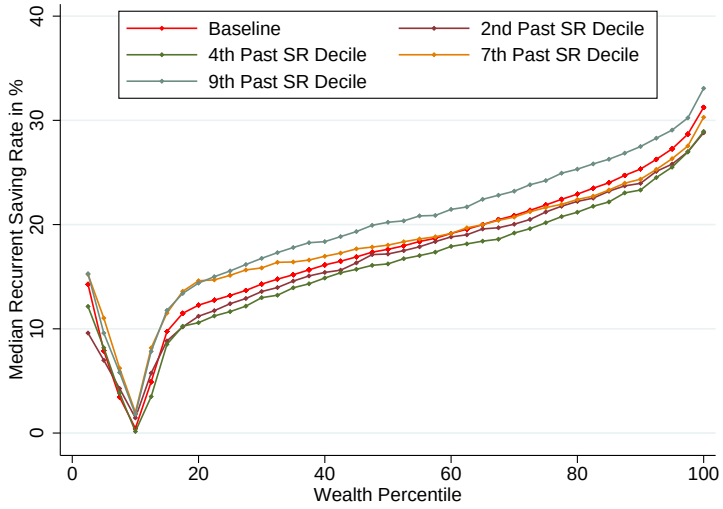


Figure 9: Saving rates across the wealth distribution within deciles of past saving rates

4.4 Is this exclusively a story about housing?

For large parts of the wealth distribution housing is the dominant asset in household portfolios (see Figure 5). This observation naturally leads to the question to what extent our findings in Figure 6 are due to households holding on to their residences in the face of rising house prices. That is, do we see similar patterns for net and gross saving rates once housing is "taken out" of household wealth accumulation? And how do households treat capital gains on assets besides housing? To be clear: this is not a question of robustness of our results, and rather one of interpretation. In particular, the question is whether theoretical explanations of our findings should focus exclusively on housing or whether they should involve other assets as well.

Answering the question how households treat capital gains on assets besides housing is complicated by the fact that Norwegian households hold relatively few other assets that experience asset price changes and hence capital gains (see Figure 5). For the purpose of the current section, we therefore restrict the sample to households that hold at least twenty-five percent of their non-housing assets in the form of stocks or stock funds.[29] In order to "take out" housing, we also focus on what we term "financial wealth" defined like net worth but excluding housing and liabilities, particularly mortgages. As before, we distinguish between gross financial saving, i.e. the year-to-year change in financial wealth, and net financial saving, i.e. gross financial saving net of capital gains. Before proceeding to show the results

[29] Alternatively, we have conducted a related exercise in which we drop all households that own private real estate so that no households in this restricted sample experience housing capital gains by construction.

of this exercise, we already note one caveat: in contrast to directly held stocks, we cannot use the shareholder registry to identify capital gains for stocks that are parts of stock portfolios and we instead use an aggregate index for this asset class (see the discussion in Section 3.2). The split of gross saving in stock funds into net saving and capital gains is therefore subject to some measurement error and our sample restrictions lend more relative importance to this asset class.

Figure 10 plots our three notions of the median financial saving rate versus the percentile of financial wealth for this restricted sample. As before we plot the median net, gross and

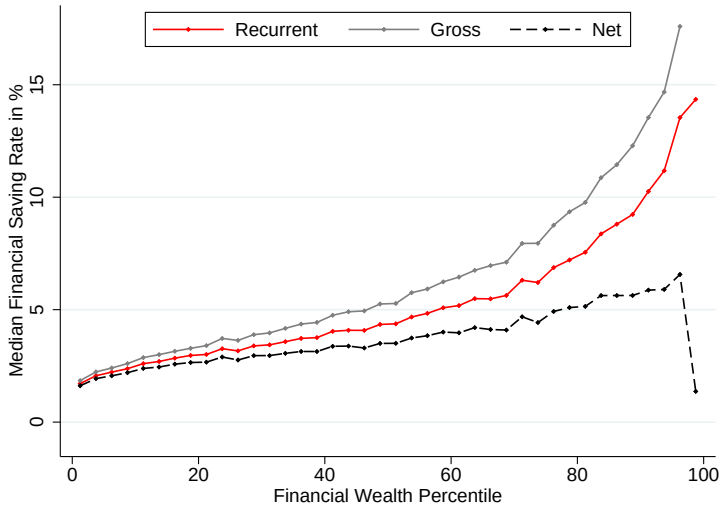


Figure 10: Relation between Financial Saving Rates and Financial Wealth

recurrent saving rates, defined as in Section 4 and analogous to Figure 6. Because the sample is considerably smaller the figure uses twenty-five four-percentile bins rather than one hundred one-percentile binds like in our main results. Like in Figure 6, both gross and recurrent saving rates increase steeply with wealth, ranging from about two percent at the bottom to fourteen percent at the top. Also akin to Figure 6, the net saving rate is considerably flatter than the gross and recurrent saving rates. In contrast to Figure 6 the net saving rate is weakly upward-sloping with financial wealth. Given the measurement concerns about net saving in stock funds discussed in the preceding paragraph, the exact slope of the relation between net financial saving and financial wealth may be subject to measurement error.[30]

[30]In particular, correlation between individual-level stock fund capital gains and financial wealth would bias the relation between the net saving rate and financial wealth. For instance, suppose that true stock fund capital gains are positively correlated with financial wealth, i.e. wealthier individuals experience larger capital gains on stocks held in stock funds. Then, by using an aggregate index, we underestimate stock fund

In summary, the patterns for net, gross and recurrent saving rates are not special to housing and are qualitatively similar for other assets. At the same time, as already noted, the combined capital gains on other assets tend to be smaller than housing capital gains mostly because Norwegian households hold relatively fewer assets that experience substantial capital gains besides housing.

5 Theoretical Interpretation

If standard workhorse model cannot rationalize our empirical findings, what theories can? In a reduced form, what is needed to rationalize Figure 6 is a theory in which households consume very little of any capital gains they experience – that is, they have a saving rate out of capital gains of close to one hundred percent, precisely what we have termed "saving by holding." In this section, we argue that a theory featuring multiple assets and portfolio adjustment frictions is promising in this regard. We also discuss other candidate explanations, many of which are complementary. These include: asset-demand shifts that affect a majority of households holding a given asset, e.g. a sustained housing preference shift; non-homothetic preferences; and misperceptions about the stochastic process for asset prices.

5.1 The Two-Asset Model

Our theory has two key ingredients. First, it takes empirical household balance sheets seriously and, in particular, that households hold multiple distinct assets. This is in contrast to many workhorse macro models in which different assets are perfect substitutes and which then collapse to one-asset models. In the theory, there is a consumption asset that is used to smooth consumption and an investment asset that experiences capital gains. Second, households are subject to portfolio adjustment "frictions" broadly defined, simply meaning some reason for not moving funds from the investment asset to the consumption asset most of the time. The model builds on Kaplan and Violante (2014) and Kaplan, Moll and Violante (2018) who label the two assets as "liquid" and "illiquid."[31] Relative to the models in these papers, the main differences are (i) that households face an asset price that moves over

capital gains for wealthy individuals and hence attribute too large a share of their gross saving in stock funds to net saving rather than capital gains. This would then result in an upward bias of the relationship between the net saving rates and financial wealth. Conversely, a negative correlation between wealth and stock fund capital gains would result in a downward bias.

[31] These papers in turn build on a long literature studying models with multiple assets and adjustment costs generating inaction, going back to Grossman and Laroque (1990).

time and (ii) that we take a broader view of portfolio adjustment frictions rather than just physical transaction costs.

There is a continuum of households indexed by their holdings of a consumption asset b, an investment asset k, and their idiosyncratic labor productivity z. As above, households have standard preferences over utility flows from consumption c_t, discounted at rate ρ. Additionally, they now face a constant hazard of death at rate ϕ so that their effective preferences are given by

$$\mathbb{E}_0 \int_0^\infty e^{-(\rho+\phi)t} u(c_t) dt.$$

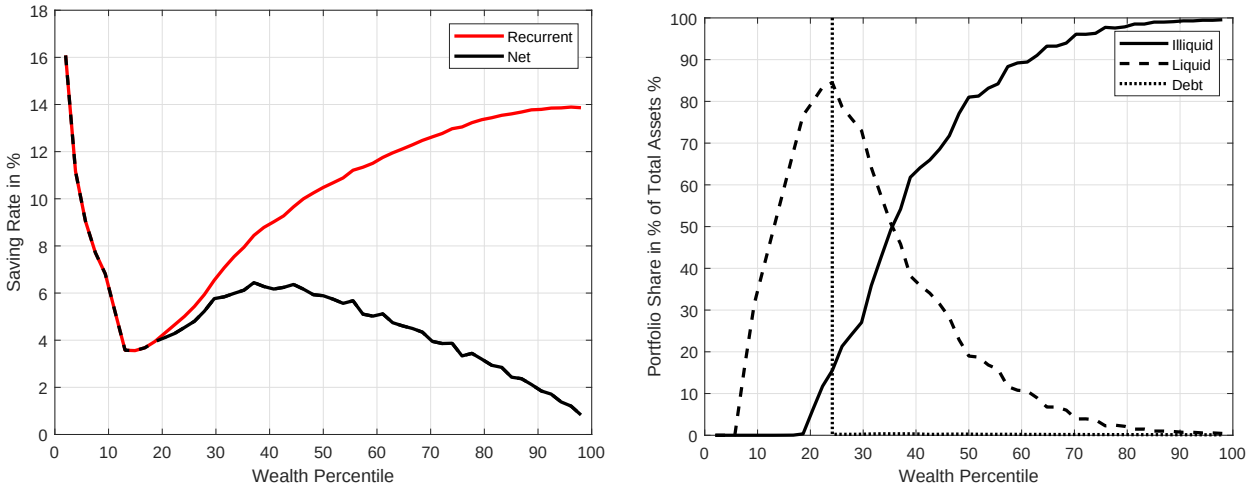
Households can buy and sell investment assets k_t at a price p_t that evolves exogenously according to (5) as in Section 2.1. The investment asset pays dividends, with θ denoting the constant dividend yield. We use d_t to denote a household's purchases of the investment asset (with $d_t < 0$ corresponding to sales). In order to purchase or sell investment assets, households must pay a flow transaction cost $\chi(d_t, p_t k_t)$ which is zero whenever $d_t = 0$. Households can borrow in the consumption asset b up to an exogenous limit $\underline{b}$ at the real interest rate $r^b_- = r^b + \kappa$ where κ is an exogenous wedge between borrowing and lending rates. With a slight abuse of notation, $r^b(b)$ summarizes the full interest rate schedule. A household's assets evolve according to

$$\begin{aligned}
&\dot{b}_t = w z_t + r^b(b_t) b_t + \theta p_t k_t - p_t d_t - \chi(d_t, p_t k_t) - c_t, \\
&\dot{k}_t = d_t, \qquad \frac{\dot{p}_t}{p_t} = \mu + \varepsilon_t, \\
&k_t \geq 0, \qquad b_t \geq \underline{b}
\end{aligned}$$

Following Kaplan, Moll and Violante (2018), we assume the following functional form for the adjustment cost function $\chi(d, a) = \chi_0 |d| + \frac{\chi_1}{2} \left(\frac{d}{a}\right)^2 a$ with $\chi_0, \chi_1 > 0$.[32] The adjustment cost function is symmetric and has a kink at $d = 0$. The two parameters have different implications for behavior: a high χ_0 implies a large inaction region while a high χ_1 gives small deposit rates. The logarithm of idiosyncratic labor productivity z_t follows an Ornstein-Uhlenbeck process (the continuous-time analogue of an AR(1) process) with speed of mean reversion ν. Finally, for computational reasons, we here assume that $\varepsilon_t \equiv 0$, that is, capital gains only have a persistent component μ and the transitory component is assumed to be zero. Transitory, stochastic capital gains could be introduced at the expense of some additional

[32] Because the transaction cost at $a = 0$ is infinite, in computations we replace the term a with $\max\{a, \underline{a}\}$, where the threshold $\underline{a} > 0$ is a small value that guarantees costs remain finite even for households with $a = 0$.

computational complexity.[33]



(a) Saving Rates (b) Portfolio Shares

Figure 11: Saving Behavior Across Wealth Distribution in Two-Asset Model

Figure 11 shows the saving rates across and portfolio shares across the wealth distribution in the two-asset model with the calibration in Appendix D. The model is able to reproduce the swoosh-shape of the recurrent saving rate and the increasing gap between recurrent and net saving rates across the wealth distribution. The swoosh-shaped recurrent saving rate comes from two assumptions. First, we assume a strictly positive interest rate wedge ($r_-^b > r_+^b$) which ensures that the saving rate is high for households with negative net worth. Second, we assume that the return on the investment asset is high, but lower than the borrowing rate ($r_-^b > \mu + \theta > r_+^b$) such that households with negative net worth first repay debt before they start accumulating investment assets. We see this in the portfolio share graph where households accumulate investment assets first when they have repaid most of the debt.

In contrast to the data, the net saving rate is slightly decreasing with net worth for households above the median. As the portfolio share of investment assets becomes high, households receive income primarily from the investment asset. Subsequently, they reduce net saving to consume a greater share of the returns from the investment asset. This prediction is due to our assumption that, while households face a kinked adjustment cost function, they can still sell small percentages of the investment asset (i.e. the investment asset is divisible). We conjecture that the prediction of our model for the net saving rate can be further

[33]Below we want to compute a stationary distribution with a continuum of households. Given this, it is very easy to introduce idiosyncratic asset price risk. Introducing aggregate risk is more challenging, simply because it means that the cross-sectional wealth distribution fluctuates over time.

improved by incorporating indivisibilities and fixed (as opposed to kinked) transaction costs.

5.2 Other Candidate Explanations

Sustained Asset Demand Shifts.

Non-Homothetic Preferences.

Misperceptions about the Stochastic Process for Asset Prices. As alluded to in Section 4.1, a candidate explanation is that households systematically misperceive all capital gains as purely transitory. In terms of the simple model of Section 2.1, they think that persistent capital gains are zero, $\mu = 0$, even though in reality μ is positive. In order for this explanation to rationalize our findings, households must have been repeatedly surprised by the fact that their assets appreciated in line with the recurring historical average. While we cannot rule out these types of systematic expectational errors, they do seem an extreme assumption to maintain. For example, consider the following argument regarding expectations about capital gains on stocks. The equity premium is around six percent in many countries as well as in Norway (see Table 2). Stock returns come in the form of both dividends and capital gains. The dividend yield on stocks is around three percent on average in many countries. If households perceived dividend yields to be the only form of return on holding stocks, they presumably would not hold stocks.

6 Conclusion

Little is known about the distribution of saving rates and how these vary across the wealth distribution. Using Norwegian administrative panel data on income and wealth, we document that how saving rates vary with wealth crucially depends on whether we include capital gains in the definition of saving. Net saving rates are approximately constant across the wealth distribution while gross saving rates increase sharply with wealth. We document that the sharp difference between the two relationships is due to the fact that richer households accumulate wealth through "saving by holding," meaning that they tend to hold on to assets experiencing persistent capital gains.

We show that these patterns are inconsistent with workhorse macro models and discuss theoretical alternatives. In a reduced form, these theoretical explanations have the form

that households have two different saving rates, one out of regular forms of income and one out of capital gains, with the latter being much higher than the former and close to one hundred percent. We show that a model with multiple assets and broadly defined portfolio adjustment frictions has the potential to qualitatively deliver the observed patterns.

Turning to the broader implications of our findings, we think that both the macroeconomics and wealth inequality literatures need to take changing asset prices much more seriously. There are two reasons for this assessment. First, changing asset prices are a prevalent feature of the data and the majority of year-to-year movements in household wealth is due to asset price movements, i.e. capital gains or losses, rather than net saving or dissaving. Second, contrary to the prediction of workhorse macro models, households appear to treat capital gains or losses differently from other forms of income, even if these are persistent. This poses a challenge for the implication of standard models that it does not matter whether the return of an asset comes in the form of dividends or capital gains.

Much of the current macroeconomics literature, in particular the heterogeneous agents literature, as well as the wealth inequality literature instead completely ignore asset price changes and instead treat all assets like money in a checking account.[34] Our findings suggest that these literatures could benefit from incorporating and refining lessons from the household finance literature (Campbell, 2006).

Appendix

A Details on Theoretical Benchmarks

A.1 Simplest Consumption-Saving Model (Section 2.1)

A.1.1 Constant Asset Prices – Derivation of Saving Policy Function in (3)

Households maximize (1) subject to (2). The corresponding HJB equation is

$$\rho v(a) = \max_c \; u(c) + v'(a)(w + ra - c) \tag{A15}$$

[34] See Benhabib and Bisin (2018) for a review of the state-of-the-art of the theoretical wealth inequality literature. To the extent that the literature takes into account asset prices at all, price changes are soaked up into rates of return by appealing to the standard result that an assets rate of return equals its dividend yield plus capital gains. See Gomez (2018) for a contribution that explicitly studies the implications of asset price changes on wealth inequality. The same is true for the public finance literature studying optimal wealth and capital taxation (see e.g. Saez and Stantcheva, 2018).

We solve this equation by using a guess-and-verify strategy: guess

$$v(a) = B\frac{(a + w/r)^{1-\gamma}}{1-\gamma}$$

which implies $v'(a) = B(a + w/r)^{-\gamma}$ and

$$c(a) = v'(a)^{-1/\gamma} = B^{-1/\gamma}(a + w/r) \tag{A16}$$

Substituting into (A15) and dividing by $(a + w/r)^{1-\gamma}$

$$\rho B\frac{1}{1-\gamma} = \frac{1}{1-\gamma}B^{-(1-\gamma)/\gamma} + \frac{B}{r} - BB^{-1/\gamma}$$

Dividing by B and collecting terms we have $B^{-1/\gamma} = \frac{\rho - r}{\gamma} + r$ and hence from (A16) we have

$$c(a) = \left(\frac{\rho - r}{\gamma} + r\right)\left(a + \frac{w}{r}\right) \tag{A17}$$

Since the saving policy function is given by $s(a) = w + ra - c$, this yields (3).

A.1.2 Changing Asset Prices – Derivation of Equation (7)

As mentioned in the text to derive (7), we assume that return innovations are Brownian, i.e. $\varepsilon dt = \sigma dW_t$ where W_t is a Wiener process. In this case (6) becomes

$$da_t = (w + \bar{r}a_t - c_t)dt + \sigma a_t dW_t, \quad \bar{r} := \theta + \mu \tag{A18}$$

Households maximize (1) subject to (A18). To ease notation we drop the bar from $\bar{r}$ and simply write $r = \theta + \mu$ for the remainder of this appendix. The corresponding HJB equation is

$$\rho v(a) = \max_c \; u(c) + v'(a)(w + ra - c) + \frac{\sigma^2 a^2}{2}v''(a). \tag{A19}$$

This problem can no longer be solved in closed form. However, we can use a perturbation method around small w. We first solve (A19) analytically when $w = 0$. We then perturb v that solves (A19) for $w > 0$ around the solution for the case $w = 0$, thus obtaining (7). Our argument makes use of a standard perturbation method as in Fleming (1971), Judd (1996), Anderson, Hansen and Sargent (2012) and Kasa and Lei (2017).

Step 1: Closed form with $w = 0$. For the first step, consider the value function for $w = 0$ which we denote by $v_0(a)$. It solves:

$$\rho v_0(a) = \max_c \; u(c) + v_0'(a)(ra - c) + \frac{\sigma^2 a^2}{2} v_0''(a) \tag{A20}$$

This is the HJB equation for a simplified version without portfolio choice of the problem analyzed by Merton (1969) and it is well-known to have a closed-form solution.

Lemma A1 *The value function and consumption policy function with $w = 0$ are*

$$v_0(a) = B_0^{-\gamma} \frac{a^{1-\gamma}}{1-\gamma}, \quad c_0(a) = B_0 a, \quad B_0 := \frac{\rho - r}{\gamma} + r + (1-\gamma)\frac{\sigma^2}{2} \tag{A21}$$

Proof of Lemma A1: the proof uses a guess-and-verify strategy. Start by guessing that $v_0(a) = B_0^{-\gamma} \frac{a^{1-\gamma}}{1-\gamma}$ for a constant B_0 to be determined. Then $v'(a) = (B_0 a)^{-\gamma}, c(a) = B_0 a$ and $v''(a) = -\gamma B_0^{-\gamma} a^{-\gamma-1}$. Substituting into (A20), we have

$$\rho B_0^{-\gamma} \frac{1}{1-\gamma} = B_0^{1-\gamma}/(1-\gamma) + B_0^{-\gamma} r - B_0^{1-\gamma} - \frac{\sigma^2}{2}\gamma B_0^{-\gamma}$$

Rearranging

$$\rho \frac{1}{1-\gamma} = \gamma B_0/(1-\gamma) + r - \frac{\sigma^2}{2}\gamma.$$

Rearranging again we obtain the expression for B_0 in (A21).□

Step 2: Perturbation around $w = 0$. As already mentioned, it is no longer possible to solve (A19) in closed form. However, we can look for approximate solutions of the form

$$v(a) = v_0(a) + w v_1(a) + O(w^2), \quad c(a) = c_0(a) + w c_1(a) + O(w^2) \tag{A22}$$

where v_0 and c_0 are the value and consumption policy functions from Lemma A1 and where v_1 and c_1 are to be determined.

Proposition A2 *The value and consumption policy functions solving* (A19) *satisfy*

$$\begin{aligned} v(a) &= \bar{c}^{-\gamma}\left(\frac{a^{1-\gamma}}{1-\gamma} + \frac{wa^{-\gamma}}{r-\gamma\sigma^2}\right) + O(w^2) \\ c(a) &= \bar{c}\left(a + \frac{w}{r-\gamma\sigma^2}\right) + O(w^2) \\ \bar{c} &:= \frac{\rho - r}{\gamma} + r + (1-\gamma)\frac{\sigma^2}{2} \end{aligned} \tag{A23}$$

Before proving the Proposition, we first note that it immediately implies the approximate saving policy function (7) in the main text. To see this simply substitute $c(a)$ in (A23) into (A18) to get

$$da_t \approx \left[w + \bar{r}a_t - \bar{c}\left(a + \frac{w}{r-\gamma\sigma^2}\right)\right]dt + \sigma a_t dW_t.$$

Equation (7) is the same equation in terms of $\varepsilon_t dt = \sigma dW_t$.

Proof of Proposition A2: In the proof it is convenient to use slightly different notation than in the statement of the Proposition: for reasons that will become apparent momentarily, we use B_0 in place of the variable $\bar{c}$ used in the statement of the Proposition. As already mentioned we look for solutions of the form (A22). Further, we restrict our attention to solutions such that $v_1'(a) > 0$ for all a which ensures that $v'(a) \approx v_0'(a) + wv_1'(a) > 0$ for all $w > 0$ and therefore consumption $c(a) = (v'(a))^{-1/\gamma}$ is positive. However, we do not make any assumptions about the sign of $v_1(a)$. Substituting (A22) into (A19)

$$\rho(v_0(a) + wv_1(a)) = \max_c \; u(c) + (v_0'(a) + wv_1'(a))(w + ra - c) + (v_0''(a) + wv_1''(a))\frac{\sigma^2}{2}a^2$$

and the first-order condition is $(c_0(a) + wc_1(a))^{-\gamma} = v_0'(a) + wv_1'(a)$. Differentiating both the HJB equation and the first-order condition with respect to w and evaluating at $w = 0$:

$$\rho v_1(a) = v_1'(a)(ra - c_0(a)) + v_1''(a)\frac{\sigma^2}{2}a^2 + v_0'(a), \tag{A24}$$

$$-\gamma c_0(a)^{-\gamma-1}c_1(a) = v_1'(a). \tag{A25}$$

Substituting the expressions for v_0 and c_0 from Lemma A1 into (A24)

$$\rho v_1(a) = v_1'(a)\left(\frac{r-\rho}{\gamma} + (\gamma-1)\frac{\sigma^2}{2}\right)a + v_1''(a)\frac{\sigma^2}{2}a^2 + B_0^{-\gamma}a^{-\gamma} \tag{A26}$$

It remains to find a solution $v_1(a)$ that solves the ODE (A26). We solve it using a guess-and-verify strategy. Guess $v_1(a) = B_1 a^{-\gamma}$. Then $v_1'(a) = -\gamma B_1 a^{-\gamma-1}$ and $v_1''(a) = \gamma(1+\gamma)B_1 a^{-\gamma-2}$. Substituting into (A26)

$$\rho B_1 a^{-\gamma} = -\gamma B_1 a^{-\gamma}\left(\frac{r-\rho}{\gamma} + (\gamma-1)\frac{\sigma^2}{2}\right) + \gamma(1+\gamma)B_1 a^{-\gamma}\frac{\sigma^2}{2} + B_0^{-\gamma}a^{-\gamma}$$

Rearranging we find that $B_1 = \frac{1}{r-\gamma\sigma^2}B_0^{-\gamma}$. Therefore

$$v_1(a) = \frac{1}{r-\gamma\sigma^2}B_0^{-\gamma}a^{-\gamma}$$

Similarly, substituting the expressions for $v_1'(a)$ and $c_0(a) = B_0 a$ into (A25) we find

$$c_1(a) = \frac{1}{r-\gamma\sigma^2}B_0$$

Substituting v_1 and c_1 as well as v_0 and c_0 from Lemma A1 into (A22), we obtain (A23) (recall again that the statement of the Proposition uses $\bar{c}$ to denote B_0).□

A.2 Common Extensions (Section 2.2)

A.3 Housing (Section 2.3)

Recall that households maximize (8) subject to (9). As explained in the text, the budget constraint can be written in terms of total wealth $a = b + ph$ as

$$\dot{a} + c + Rh = w + ra, \quad \text{where} \quad R := rp - \dot{p}$$

is the user cost of housing. It is easy to see that this problem splits into separate inter- and intratemporal problems and can be solved as a two-stage budget problem. The intertemporal problem is

$$\begin{aligned}\max_{\{C_t\}_{t\geq 0}} \int_0^\infty e^{-\rho t}\frac{C_t^{1-\gamma}}{1-\gamma}dt \quad \text{s.t.} \\ \dot{a}_t = w + ra_t - P_t C_t\end{aligned} \tag{A27}$$

where

$$P_t = \left(\alpha + (1-\alpha)R_t^{1-\eta}\right)^{\frac{1}{1-\eta}} \tag{A28}$$

is a CES price index. Given the solution to this intertemporal problem, non-durable consumption and housing can then be found from the intratemporal problem. This problem is just a standard static utility maximization problem with CES utility. The solution is given by

$$c = \frac{\alpha}{R^{1-\eta}(1-\alpha)+\alpha} PC, \qquad h = \frac{R^{-\eta}(1-\alpha)}{R^{1-\eta}(1-\alpha)+\alpha} PC. \tag{A29}$$

A.3.1 Proof of Proposition 1

Problem (A27) is just a consumption-saving problem with a time-varying price level. The solution satisfies

$$C^{-\gamma} = \lambda P, \tag{A30}$$

$$\dot{\lambda} = (\rho - r)\lambda, \tag{A31}$$

$$\dot{a} = w + ra - PC. \tag{A32}$$

The budget constraint can be integrated to yield a lifetime budget constraint

$$\int_0^\infty e^{-rt} P_t C_t dt = \frac{w}{r} + a_0. \tag{A33}$$

Integrating the Euler equation (A31) forward in time, we have

$$\lambda_t = \lambda_0 e^{(\rho - r)t}.$$

Hence from (A30)

$$C_t^{-\gamma}/P_t = C_0^{-\gamma}/P_0 e^{(\rho-r)t}.$$

And in particular

$$P_t C_t / P_0 C_0 = (P_t/P_0)^{1-\frac{1}{\gamma}} e^{-\frac{\rho-r}{\gamma}t}.$$

Substituting into the lifetime budget constraint (A33)

$$P_0 C_0 \int_0^\infty (P_t/P_0)^{-\frac{1}{\gamma}} e^{-\left(\frac{\rho-r}{\gamma}+r\right)t} dt = \frac{w}{r} + a_0,$$

and hence consumption expenditure is

$$P_0 C_0 = \left(\int_0^\infty (P_t/P_0)^{1-\frac{1}{\gamma}} e^{-\left(\frac{\rho-r}{\gamma}+r\right)t} dt \right)^{-1} \left(\frac{w}{r} + a_0 \right),$$

which is the expression for C_t in (10). Since $\dot{a}_0 = w + ra_0 - P_0C_0$, we also get the saving response, namely

$$\dot{a}_0 = w + ra_0 - \left(\int_0^\infty (P_t/P_0)^{1-\frac{1}{\gamma}} e^{-\left(\frac{\rho-r}{\gamma}+r\right)t} dt\right)^{-1} \left(\frac{w}{r} + a_0\right),$$

which is the expression for $\dot{a}_t$ in (10).□

A.3.2 Proof of Corollary 1

When $\dot{p} = 0$ the problem is stationary and we can drop t subscripts from variables. Equation (10) in Proposition 1 becomes

$$PC = \left(\frac{\rho - r}{\gamma} + r\right)\left(a + \frac{w}{r}\right), \qquad \dot{a} = \frac{r-\rho}{\gamma}\left(a + \frac{w}{r}\right) \tag{A34}$$

We are now in a position to derive the marginal effect on total consumption from a surprise and permanent change in the house price. If $r = \rho$, total consumption expenditure in (A34) is simply $PC = rb + rph + w$ where we have used that $a = b + ph$. Totally differentiating this expression with respect to p, we have

$$\begin{aligned}
&PC_p + P_pP = rh \\
&PC_p + PCr\frac{(rp)^{-\eta}(1-\alpha)}{\alpha + (1-\alpha)(rp)^{1-\eta}} = rh \\
&PC_p + rh = rh \\
&\Rightarrow C_p = 0
\end{aligned}$$

where we have used (A29) and (A28). Hence, a surprise permanent increase in the house price has no effect on total consumption and therefore not on welfare. Further, we have that $\dot{a} = 0$ when $r = \rho$ from (A34). This implies that, wealth moves only on impact, immediately after the initial surprise housing capital gain, but not thereafter. That is, the surprise housing capital gain translates into higher wealth one-for-one.□

B Appendix for Section 3

B.1 Net, Gross and Recurrent Saving with Multiple Assets

Section 3.2 defined net, gross and recurrent saving in an environment with one asset. The data obviously features multiple assets and we here generalize these definitions to such an environment. To show that there is nothing special about continuous time and because our data is measured at discrete time intervals, we here present these definitions in discrete time. Finally, we allow for taxes and transfers.

The household receives annual income w_t (labor income and cash transfers), and pays taxes τ_t. There are J assets indexed by $j = 1, ..., J$. Let $k_{j,t-1}$ denote the household's holdings of asset j at the end of period $t-1$. To simplify notation, assume that the household holds each $k_{j,t-1}$ throughout the year and, at the end of the year, sells the assets $k_{j,t-1}$ at prices $p_{j,t}$ and buys a new portfolio $k_{j,t}, j = 1, ..., J$ at the same prices. Throughout the year, asset holdings $k_{j,t-1}$ earn capital income $\theta_{j,t}p_{j,t}k_{j,t-1}$. With this notation, the generalizations of (11) and (12) are

$$\underbrace{c_t + \sum_{j=1}^{J} p_{j,t}(k_{j,t} - k_{j,t-1})}_{\text{net saving}} = \underbrace{w_t - \tau_t + \sum_{j=1}^{J} \theta_{j,t}p_{j,t}k_{j,t-1}}_{\text{disposable income}} \tag{A35}$$

$$\underbrace{c_t + \sum_{j=1}^{J}(p_{j,t}k_{j,t} - p_{j,t-1}k_{j,t-1})}_{\text{gross saving}} = \underbrace{w_t - \tau_t + \sum_{j=1}^{J}\left(\theta_{j,t} + \frac{p_{j,t} - p_{j,t-1}}{p_{j,t-1}}\right)p_{j,t-1}k_{j,t-1}}_{\text{Haig-Simons income}}. \tag{A36}$$

Similarly, the analogue of (13) defining recurrent saving and income is (A36) but with only the persistent component of capital gains added to both sides.

B.2 Separating Gross Saving into Net Saving and Capital Gains

B.2.1 Housing: Using Transaction Data

To explain our approach, it is helpful to introduce some notation. Time is continuous and we consider a household that makes housing transactions at discrete time intervals. We denote by $h(t), p(t)$ and $a_h(t) = p(t)h(t)$ the household's physical number of housing units, the price of housing and the value of the house on January 1 of year t. Throughout a year, i.e. between dates t and $t+1$, the household makes $N \geq 0$ transactions at dates $\tau_n, n = 1, ..., N$

with $t \leq \tau_1 < \tau_2 < ... < \tau_N < t+1$. To simplify notation below, it will be convenient to denote the beginning and end of each year by $\tau_0 = t$ and $\tau_{N+1} = t+1$ (this is just notation and these are not actual transaction dates). The goal is to decompose gross saving, i.e. the change over the year in housing wealth $a_h(t) = p(t)h(t)$ into net saving and capital gains. We use the following decomposition:

$$\underbrace{a_h(t+1) - a_h(t)}_{\text{gross saving}} = \underbrace{\sum_{n=1}^{N+1}(p(\tau_n) - p(\tau_{n-1}))h(\tau_{n-1})}_{\text{capital gains}} + \underbrace{\sum_{n=1}^{N} p(\tau_n)(h(\tau_n) - h(\tau_{n-1}))}_{\text{net saving}} \tag{A37}$$

Note in particular that net saving is only non-zero for households with housing transactions, for example those who move houses. For households without housing transactions throughout the year, net saving is zero by construction and all of gross saving in housing reflects capital gains.

C Public Pension Wealth

This appendix describes how we calculate public pension wealth. We define public pension wealth as the net present value of the future flow of public pension wealth, discounted at the risk free interest rate. The main complication is that there are currently four different pension systems depending on your birth cohort. We describe each of these systems in detail below before we define public pension wealth, savings, and income.

Common notation:

- G_t = basis-amount in the social security system. Adjusted yearly with wage growth.
- pen_t^c = claim to yearly pension flow in units of G at time t
- V_t^c = pension wealth, net present value of future pension income at time t

C.1 The four pension systems

C.1.1 Cohorts born prior to 1944

For these cohorts, we observe actual pension earnings. We therefore calculate pension wealth as the net present value of the flow of pension income. Implicitly, we assume perfect foresight,

but for most of these households, the distance to pension age is maximum 7 years, making it likely that they know their future pension income.

C.1.2 Cohorts born between 1944 and 1953

For these cohorts, there are two parts of the pension system: the social security and the service pension.

Social Security. The social security system is based on a point system. Each year, pensions points were acquired based on the following formula

$$point_t = \begin{cases} \frac{y_t}{G_t} - 1 & \text{if } G_t \leq y_t \leq 6G_t \\ 5 + \frac{y_t - 6G_t}{3G_t} & \text{if } 6G_t \leq y_t \leq 12G_t \\ 7 & \text{if } y_t \geq 12G_t \end{cases}$$

where y_t is gross earnings in year t and G_t is the basis-amount in the social security system.

The persons pension number is the defined as the average of the 20 years with the highest pension points, or the average of the n years that person has earned pension points. If we order the years with positive pension points from largest to smallest with index j, P is defined as

$$P = \begin{cases} \sum_{j=1}^{20} \frac{point_j}{20} & \text{if } n \geq 20 \\ \sum_{j=1}^{n} \frac{point_j}{n} & \text{if } n < 20 \end{cases}$$

The payouts from the social security pension is approximately

$$pen^{SC} = \alpha + \max\left\{\kappa P \frac{min\{n, 40\}}{40}, 1\right\}$$

where $\alpha = 1$ for singles and 0.85 for couples, and κ is a proportionality factor equal to 0.45 if income was accumulated prior to 1992 and 0.42 after 1992.

Service Pension. All public sector employees and about 50 percent of private sector employees have an additional service pension on top of their social security pension. This service pension guarantees an individual a fraction ψ of their final income. This fraction depends on how long the individual has worked for the company/government, but the maximum is $\psi = 0.66$.

The service pension pays the difference between the sum implied by the fraction rule

and the level received from the social security pension. We can therefore approximate the pensions in units of G as

$$pen_t^c = \max\left\{\alpha + \max\left\{\kappa P \frac{min\{n, 40\}}{40}, 1\right\}, \frac{0.66 y^{final}}{G^{final}}\right\}$$

where n is the number of working years, and y^{final} and G^{final} are income and basis-unit in social security in your final working year, respectively.

C.1.3 Cohorts born between 1954 and 1962

Pensions in units of G is a linear combination of the system in Section C.1.2 and C.1.4:

$$pen_t^c = \frac{63 - c}{10} pen_t^{43<c<54} + \frac{c - 53}{10} pen_t^{c>63}$$

C.1.4 Cohorts born after 1963

In 2010, the government simplified the pension system for all earners born after 1964. The new system implies that every year, 18.1 % of your gross income below 7.1 G is added to your pension holdings ("pensjonsbeholdning," P_t).

$$P_t^c = \left(\sum_{\tau=c+13}^{t} \max\left\{0.181 \frac{y_\tau}{G_\tau}, 0.181 \cdot 7.1\right\}\right)$$

P_t^c is, as before, defined in units of G, the basis-amount in the Norwegian pension system.

One complication is that you can start taking out pensions from age 62. However, to simplify the exposition, we assume that all households value their pension as if they would start taking out pensions from age 67. From age 67, your income from pensions is the value of your pension holdings P_t^c in units of G divided by your expected remaining years (e.g. 16.02 for the cohort born in 1964).

$$pen_t^c = \frac{P_t^c}{d_c}$$

C.2 Pension wealth, saving, and income

We define pension wealth as the net present value of pension income from age 67. That income is $pen_t^c G_t$ in each year t when t is greater than 67, where pen_t^c is defined by one of the four systems above depending on your birth cohort. In order to calculate the net present

value, we discount by the risk free real interest rate

$$V_t^c = pen_t^c G_t \, \mathbb{E}_t \left\{ \sum_{\tau=\max\{c+67,t\}}^{\max\{c+67+d_c,t\}} \frac{\prod_{s=t}^{\tau}(1+\pi_{w,s})}{\prod_{s=t}^{\tau}(1+r_s)} \right\} \tag{A38}$$

where $\pi_{w,s}$ is the real growth rate in G in year s and r_s is the real interest rate in year s. The max operator is there because an individual start withdrawing from the pension account after age 67. We have to make two assumptions to calculate pension wealth in the data. First, we assume perfect foresight in the years where we observe r_s and $\pi_{w,s}$. Second, we assume that after 2015, the expected real interest rate and growth rate of G are the observed geometric mean in the years from 1993 to 2015. For example, in order to calculate pension wealth in 2006, we discount by the observed real interest rate from 2006 to 2015, and with the mean real interest rate after 2015.

We define gross pension saving as the change in pension wealth

$$V_{t+1}^c - V_t^c = \underbrace{\underbrace{\text{NPV of pension contributions - pension withdrawals}}_{\text{Net pension saving}} + \underbrace{r_t \frac{V_{t+1}^c}{1+r_t}}_{\text{Capital gains}}}_{\text{Gross pension saving}}, \tag{A39}$$

or more precisely in terms of the variables defined above:

$$V_{t+1}^c - V_t^c = \begin{cases} (pen_{t+1}^c - pen_t^c)\frac{V_t^c}{pen_t^c} & + r_t \frac{V_{t+1}^c}{1+r_t} \text{ if age } < 67 \\ -pen_t^c G_t & + r_t \frac{V_{t+1}^c}{1+r_t} \text{ if age } \geq 67. \end{cases}$$

Changes in pension wealth can be decomposed into two parts: the change in pension wealth due to contributions or withdrawals, and the change in pension wealth due to a revaluation of the net present value of current pension wealth. If an individual works and adds to the pension holdings, we count that as a contribution to pension wealth and is counted as net saving. Similarly, if a person is a retiree, that person receives pension income and reduces the future flow of pension income. We also count such a withdrawal as negative pension saving. We count the revaluation of pension wealth that would occur if there were no contributions or withdrawals as capital gains. In our setting, that is just the revaluation of pension wealth because the retirement age is one year closer.

We define pension income in such a way to ensure that the budget constraint adds up. This implies that net pension income equals net pension saving; and that gross pension

income equals gross pension saving. Furthermore, we define recurrent pension saving and income as the change in pension wealth minus $r_t \frac{V^c_{t+1}}{1+r_{t+1}}$ plus $r \frac{V^c_{t+1}}{1+r}$ where r is the (geometric) average real interest rate from 1993 to 2015.

C.3 Saving rates across the wealth distribution including pensions

[TO BE COMPLETED]

D Details for Two-Asset Model in Section 5

[TO BE COMPLETED]

References

Aiyagari, S Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics*, 109(3): 659–84.

Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent. 2012. "Small noise methods for risk-sensitive/robust economies." *Journal of Economic Dynamics and Control*, 36(4): 468–500.

Bach, Laurent, Laurent Calvet, and Paolo Sodini. 2018. "From Saving Comes Having? Disentangling the Impact of Saving on Wealth Inequality." Swedish House of Finance Research Paper.

Bach, Laurent, Laurent E Calvet, and Paolo Sodini. 2016. "Rich Pickings? Risk, Return, and Skill in the Portfolios of the Wealthy." C.E.P.R. Discussion Papers CEPR Discussion Papers 11734.

Benhabib, Jess, and Alberto Bisin. 2018. "Skewed Wealth Distributions: Theory and Empirics." *Journal of Economic Literature*, 56(4): 1261–91.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2018. "House Prices and Consumer Spending." *Review of Economic Studies*, 85(3): 1502–1542.

Campbell, John Y. 2006. "Household Finance." *Journal of Finance*, 61(4): 1553–1604.

Christelis, Dimitris, Dimitris Georgarakos, and Tullio Jappelli. 2015. "Wealth shocks, unemployment shocks and consumption in the wake of the Great Recession." *Journal of Monetary Economics*, 72(C): 21–41.

De Nardi, Mariacristina, and Giulio Fella. 2017. "Saving and Wealth Inequality." *Review of Economic Dynamics*, 26: 280–300.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy*, 112(2): 397–444.

Eika, Lasse, Magne Mogstad, and Ola Vestad. 2018. "What Can We Learn About Household Consumption from Information on Income and Wealth?" SSB Working Paper.

Eisner, Robert. 1988. "Extended Accounts for National Income and Product." *Journal of Economic Literature*, 26(4): 1611–1684.

Eitrheim, Øyvind, and Solveig K. Erlandsen. 2004. "House Price Indices for Norway 1819-2003." In *Monetary Statistics for Norway 1819-2003. Occasional Paper*, , ed. Øyvind Eitrheim, Jan T. Klovland and Jan F. Qvigstad. Norges Bank.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2016. "Heterogeneity and Persistence in Returns to Wealth." National Bureau of Economic Research, Inc NBER Working Papers 22822.

Fagereng, Andreas, Martin B. Holm, and Kjersti N. Torstensen. 2019. "Housing Wealth in Norway, 1993–2015."

Fleming, Wendell H. 1971. "Stochastic Control for Small Noise Intensities." *SIAM Journal on Control*, 9(3): 473–517.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2016. "The Dynamics of Inequality." *Econometrica*, 84: 2071–2111.

Glaeser, Edward L. 2000. "Comments on "Real Estate and the Macroeconomy"." *Brookings Papers on Economic Activity*, 2000(2): 146–150.

Gomez, Matthieu. 2018. "Asset Prices and Wealth Inequality." Columbia University Working Paper.

Grossman, Sanford J, and Guy Laroque. 1990. "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods." *Econometrica*, 58(1): 25–51.

Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jón Steinsson. 2018. "Housing Wealth Effects: The Long View." National Bureau of Economic Research Working Paper 24729.

Haig, Robert M. 1921. "The Concept of Income – Economic and Legal Aspects." *The Federal Income Tax*, 1–21. New York:Columbia University Press.

Hicks, J. 1939. *Value and Capital: An Inquiry Into Some Fundamental Principles of Economic Theory.* Clarendon Press.

Huggett, Mark. 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of Economic Dynamics and Control*, 17(5-6): 953–969.

Judd, Kenneth L. 1996. "Approximation, perturbation, and projection methods in economic analysis." In *Handbook of Computational Economics*. Vol. 1 of *Handbook of Computational Economics*, , ed. H. M. Amman, D. A. Kendrick and J. Rust, Chapter 12, 509–585. Elsevier.

Kaplan, Greg, and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica*, 82(4): 1199–1239.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review*, 108(3): 697–743.

Kasa, Kenneth, and Xiaowen Lei. 2017. "Risk, Uncertainty, and the Dynamics of Inequality." Department of Economics, Simon Fraser University Discussion Papers dp17-06.

Krueger, D., K. Mitman, and F. Perri. 2016. "Macroeconomics and Household Heterogeneity." Vol. 2 of *Handbook of Macroeconomics*, Chapter 0, 843–921. Elsevier.

Krusell, Per, and Anthony A. Smith. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106(5): 867–896.

Merton, Robert C. 1969. "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case." *The Review of Economics and Statistics*, 51(3): 247–57.

Mullainathan, Sendhil, and Jann Spiess. 2017. "Machine learning: an applied econometric approach." *Journal of Economic Perspectives*, 31(2): 87–106.

Paiella, Monica, and Luigi Pistaferri. 2017. "Decomposing the Wealth Effect on Consumption." *The Review of Economics and Statistics*, 99(4): 710–721.

Poterba, James M. 2000. "Stock Market Wealth and Consumption." *Journal of Economic Perspectives*, 14(2): 99–118.

Saez, Emmanuel, and Stefanie Stantcheva. 2018. "A simpler theory of optimal capital taxation." *Journal of Public Economics*, 162(C): 120–142.

Simons, Henry C. 1938. *Personal Income Taxation: the Definition of Income as a Problem of Fiscal Policy.* University of Chicago Press Chicago, Ill.

Straub, Ludwig. 2018. "Consumption, Savings, and the Distribution of Permanent Income." Harvard Working Paper.

von Schanz, Georg. 1896. "Der Einkommensbegriff und die Einkommensteuergesetze." *FinanzArchiv / Public Finance Analysis*, 13(1): 1–87.