

Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains *

Andreas Fagereng
Statistics Norway

Martin Blomhoff Holm
University of Oslo

Benjamin Moll
Princeton

Gisle Natvik
BI Norwegian Business School

July 14, 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 canonical models of household saving 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 canonical models. However, saving rates including capital gains (“gross saving rates”) increase markedly with wealth. The proximate explanation is that wealthier households own assets that experience persistent capital gains which they hold onto instead of selling them off to consume (“saving by holding”). Since standard theories of wealth accumulation are about gross saving, our findings challenge them. In contrast, our empirical findings are consistent with theories with time-varying discount rates or with multiple assets and portfolio adjustment frictions. Between 1995 and 2015 Norway’s aggregate wealth-to-income ratio rose from approximately 4 to 6.5 and “saving by holding” accounts for between sixty and one hundred percent of this increase.

*We thank Adrien Auclert, John Campbell, Karen Dynan, Giulio Fella, Matthieu Gomez, Adam Guren, Henrik Kleven, Camille Landais, Monika Piazzesi, Luigi Pistaferri, Martin Schneider, Kjetil Storesletten, Olav Syrstad, Martin Weale 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 for a number of fields in economics. In macroeconomics, a prediction of many canonical models of household saving behavior is that saving rates are either independent of or decreasing with wealth.^{1,2} But whether the individual saving behavior predicted by these models is consistent with the data is an open question. Better empirical evidence on saving behavior across the wealth distribution is key for answering such 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 eleven-year period from 2005 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.³

¹As 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 \(1994\)](#), [Bewley \(1986\)](#), [Huggett \(1993\)](#) and [Imrohoroglu \(1989\)](#) and many of their extensions and generalizations. This prediction has previously been emphasized by [De Nardi and Fella \(2017\)](#). The literature also often considers richer versions of such models, particularly lifecycle models. As we show below, these richer versions predict that saving rates are flat or decreasing *conditional* on observables such as age, labor income or education – again, also see [De Nardi and Fella \(2017\)](#), in particular their Figure 1.

²These theories then typically also predict that aggregate consumption and saving are approximately independent of the wealth distribution (see e.g. [Krusell and Smith, 1998](#)).

³Some papers in the literature refer to our “net saving” as “active saving” and our “gross saving” as “passive saving” (e.g. [Dynan, Skinner and Zeldes, 2004](#)). We prefer the terms “net” and “gross” because they are more neutral (e.g. the decision not to sell an asset when its price rises may well be an active decision). But it is important to be clear that this is just different terminology for the same concepts.




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 excluding 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 canonical saving 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.⁴ 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

⁴As 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 labeled “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.

These findings suggest that capital gains are a main source of wealth accumulation at the micro level. But how important is the observed saving of capital gains for wealth accumulation at the macro level? We provide an illustration, by quantifying the contribution of saving-by-holding for wealth accumulation in Norway. From 1995 to 2015, the Norwegian wealth-to-income ratio rose from approximately 4 to 6.5.⁵ By counterfactually imposing that households over this period treated capital gains in the same way as other forms of income, we find that saving by holding accounted for between sixty and one hundred percent of this

⁵As we explain in the main text, these numbers are computed directly from our micro data and feature a smaller increase than some other series, including the national accounts.

increase in aggregate wealth.

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 housing. An infinitely-lived household with homothetic utility receives a constant stream of labor income and saves in an asset with an exogenously given price. We first consider the case of a constant asset price and show that the model predicts that a household's saving rate out of income should be exactly independent of its wealth. We then consider a time-varying and stochastic asset price so that the model generates a meaningful distinction between net and gross saving, just like in the data. As is standard, the asset's return is the sum of its dividend yield and capital gains. 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, 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. Second, if households experience persistent capital gains, net saving rates should be decreasing with 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 or sustained housing demand shifts, must be part of the story.⁶

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.

If canonical 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.”⁷ We propose two candidate explanations for our empirical findings: first, time-varying discount rates or, more generally, asset-demand shifts that affect households throughout the wealth distribution; second, portfolio adjustment frictions that make it costly to liquidate assets experiencing capital gains. Canonical heterogeneous agent models feature neither of these ingredients and instead assume that households have time-invariant discount rates and save in one asset only (as in Aiyagari, 1994; Huggett, 1993). We show that both ingredients can qualitatively explain our findings and, in particular, generate “saving by holding.”⁸ We also briefly discuss other candidate explanations, several of which are complementary. These

⁶More specifically, we show that a canonical consumption-saving 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. Therefore, while housing is clearly different from financial assets, the consumption aspect by itself cannot explain our empirical findings.

⁷Interestingly, already Meade (1964) conjectures that such reduced-form saving rule could be a good approximation to individual saving behavior. See the discussion in footnote 50 in Section 6 of the main text.

⁸There is also *prima facie* evidence for the empirical relevance of both of these ingredients. First, a typical finding in the asset pricing literature is that a large fraction of asset price fluctuations is driven by time-varying discount rates (see e.g. Campbell, 2003; Cochrane, 2005). Second, going back to Grossman and Laroque (1990), many authors have pointed out that some assets are subject to indivisibilities and transaction costs, chief among them housing. Recent incarnations of such theories include Kaplan and Violante (2014) and Kaplan, Moll and Violante (2018) on whose models we build.

include: non-homothetic preferences, misperceptions about the stochastic process for asset prices, and behavioral explanations.

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.⁹ [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. Our paper differs from theirs in two main 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.¹⁰ 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 ill-defined because computing it would require division by zero. 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.^{11,12}

Our findings are relevant for a nascent literature in macroeconomics and the study of inequality that takes portfolio choice and asset prices more seriously. [Hubmer, Krusell and Smith \(2018\)](#) take a useful first step toward developing a quantitative theory of wealth inequality

⁹An 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).

¹⁰In 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.

¹¹[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.

¹²[Fagereng et al. \(2016a,b\)](#) 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. A natural question is whether wealthy households have high saving rates *because* they have high returns. Our findings suggest that this effect may be part of the story but is not at its core. 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. On the other hand, high-return assets tend to have a large share of their returns in the form of capital gains (e.g. stocks and housing) which amplifies the observed divergence between net and gross saving rates.

ity dynamics with these features by modeling them in a reduced form fashion. Also see Gomez (2018). Empirically, Kuhn, Schularick and Steins (2018), Martínez-Toledano (2019) and Feiveson and Sabelhaus (2019) emphasize the importance of portfolio composition and asset price movements for shaping the wealth distribution and wealth accumulation over the lifecycle. We hope that our empirical findings and their implications for theory will be useful building blocks for future contributions to this literature.¹³

Finally, our paper is related to the literature that studies the effect of asset price changes on consumption.¹⁴ The most directly related papers are those that estimate the impact of capital gains on the level of consumption.¹⁵ Baker, Nagel and Wurgler (2007) and Di Maggio, Kermani and Majlesi (2018) argue that marginal propensities to consume (MPCs) out of capital gains are significantly smaller than MPCs out of dividend income. To the extent that their estimates identify the effect of persistent capital gains on consumption, this finding is consistent with our finding that households “save by holding” i.e. they hold on to assets that experience persistent capital gains.

Section 2 lays out a series of theoretical benchmarks and shows that canonical models of wealth accumulation predict that the systematic component of gross saving rates is approximately constant across the wealth distribution. Section 3 describes our data and institutional setting. Section 4 presents our empirical results and documents that, inconsistent with the prediction of canonical models, it is net saving rates that are approximately constant across the wealth distribution while gross saving rates increase sharply with wealth. Section 5 shows that “saving by holding” can account for a large fraction of the rise in Norway’s aggregate wealth-to-income ratio. Section 6 theoretically interprets these empirical findings and argues that they are consistent with theories with time-varying discount rates or with multiple assets and portfolio adjustment frictions. Section 7 concludes.

¹³See Straub (2018) and Cui and Sterk (2018) for examples of how our evidence can be used to discipline macroeconomic theories. On the empirical end, Kuhn, Schularick and Steins and Martínez-Toledano adopt a decomposition of individual wealth dynamics proposed by Saez and Zucman (2016) and use it to conduct counterfactuals. They write $W_{t+1}^i = W_t^i(1 + q_t^i) + S_t^i$ where W_t^i is wealth, q_t^i is capital gains in percentage terms and S_t^i is net saving. Net saving S_t^i is assumed to be independent of capital gains q_t^i which effectively assumes that households always save one hundred percent of their capital gains. While such a specification is inconsistent with standard theories of wealth accumulation, our empirical findings show that it may nevertheless provide a good approximation to individual saving behavior.

¹⁴Poterba (2000) reviews the literature on the consumption effects of changes in stock market wealth and Christelis, Georgarakos and Jappelli (2015), Paiella and Pistaferri (2017) and Chodorow-Reich, Nenov and Simsek (2019) are examples of more recent studies. For studies examining the effect of house price changes on consumer spending, both theoretically and empirically, see Berger et al. (2018) and Guren et al. (2018) among others.

¹⁵As opposed to the impact of the *level* of asset prices on the *level* of consumption or, equivalently, *changes* in asset prices on *changes* in consumption, like in many of the studies in footnote 14.

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 canonical models of individual wealth accumulation predict that the systematic component of gross saving rates, i.e. saving rates including capital gains, are approximately constant across the wealth distribution. The theory also guides our definition of different saving concepts that we use in our empirical analysis (“net”, “gross” and “recurrent”). 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 :

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

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

$$c_t + \dot{a}_t = w + ra_t. \quad (2)$$

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

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

That is, households save (and consume) a constant fraction of their effective wealth $a + w/r$. 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.¹⁶ Importantly, this fraction is independent of their wealth. We show below that many other benchmark models inherit this property, at least approximately.

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

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

The stock price evolves according to

$$\frac{\dot{p}_t}{p_t} = \mu + \varepsilon_t, \quad (5)$$

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

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

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 $\mu + \varepsilon_t$.

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, Ap-

¹⁶Intuitively, households save a constant fraction of their effective wealth and their income is the constant return to that effective wealth. Hence saving is also a constant fraction of their income. Finally note that a constant saving rate is not an artifact of considering infinitesimal time intervals. Instead, it also holds over discrete time intervals: for example, using (3), we have $a_{t+1} - a_t = \int_0^1 \dot{a}_{t+s} ds = \frac{r-\rho}{\gamma} \int_0^1 (a_{t+s} + \frac{w}{r}) ds = \frac{r-\rho}{\gamma r} \int_0^1 (w + r a_{t+s}) ds$ and hence saving $a_{t+1} - a_t$ is a constant fraction $\frac{r-\rho}{\gamma r}$ of income $\int_0^1 (w + r a_{t+s}) ds$.

pendix 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 $w \approx 0$:

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

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.

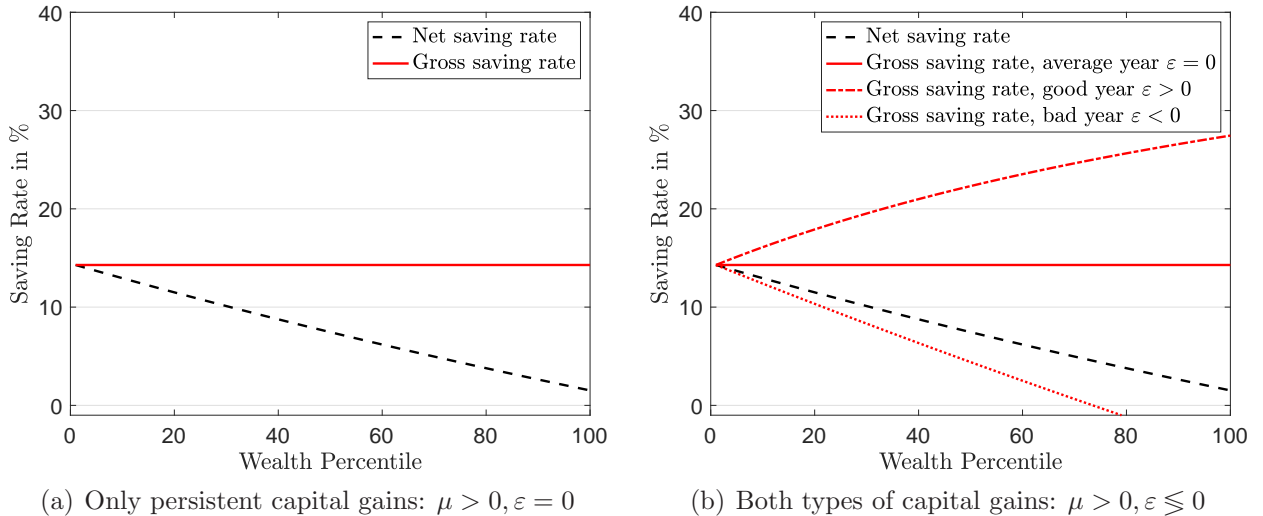


Figure 2: Saving rates across wealth distribution in simple benchmark model.

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 (4) and (6): first, the saving rate net of capital gains $p_t \dot{k}_t / (w + \theta p_t k_t)$, or “net saving rate,” corresponding to (4); and second, the saving rate out of total income including capital gains $\dot{a}_t / (w + r_t a_t)$, or “gross saving rate,” 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 that the net saving rate should be *decreasing* with wealth.¹⁷

Panel (b) of Figure 2 considers the case where there are both transitory and persistent capital gains. In the case when $\varepsilon = 0$, meaning that the stock market had an “average year”, the figure is identical to 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.¹⁸ 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.¹⁹

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.

Key Concepts: Net, Gross and Recurrent Saving. Motivated by these theoretical results we now define some key concepts that we will later use in our empirical analysis. We briefly revisit two notions of saving we have already defined informally, namely “net” and “gross” saving. We thereafter introduce a concept we term “recurrent” saving which aims to isolate the systematic (rather than cyclical) component of gross saving.

All three notions of saving rates are defined by alternative ways of writing the budget constraint “consumption plus saving equals income”, in particular by how capital gains are

¹⁷When $\varepsilon = 0$ the approximation in (7) is exact and so $\dot{a} = \hat{s}(w + (\theta + \mu)a)$ where $\hat{s} = \bar{s}/\bar{r}$, i.e. the gross saving rate is constant at \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 when $\mu > 0$ and $w > 0$.

¹⁸Equation (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 $\varepsilon < 0$.

¹⁹To 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.

treated when writing this accounting identity. As already discussed, the definition of “net” and gross” saving can be gauged from the two budget constraints (4) and (6) which we here restate for convenience (the latter in slightly altered form):

$$c + \underbrace{\dot{p}k}_{\text{net saving}} = \underbrace{w + \theta pk}_{\text{disposable income}}, \quad (8)$$

$$c + \underbrace{\dot{p}k + \dot{p}k}_{\text{gross saving}} = \underbrace{w + (\theta + \dot{p}/p)pk}_{\text{Haig-Simons income}}. \quad (9)$$

The difference between these 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 (8) features disposable income, whereas formulation (9) features “Haig-Simons” income which includes unrealized capital gains as a form of income (von Schanz, 1896; Haig, 1921; Simons, 1938).²⁰ Finally, we define the “net saving rate” as the ratio of net saving to disposable income in (8), and the “gross saving rate” as the ratio of gross saving to Haig-Simons income in (9).

Typical macro models of wealth accumulation are models of gross saving in (9), and it therefore seems natural to compare gross saving rates generated by such models to those in the data. However, recall from the simple model we just analyzed that we should expect the gross saving rate to vary strongly with asset market performance as in Figure 2(b). Fortunately, another theoretical prediction was that the systematic component of gross saving rates should be independent of wealth. To provide a sharper test for economic theories, we therefore introduce a third notion of saving that we term “recurrent saving”, meant to isolate exactly this systematic component. This notion is defined via the budget constraint

$$c + \underbrace{(\dot{k}/k + \mu)pk}_{\text{recurrent saving}} = \underbrace{w + (\theta + \mu)pk}_{\text{recurrent income}}, \quad (10)$$

where μ is a measure of the systematic, or persistent, component of capital gains. Recalling the decomposition of capital gains into persistent and transitory components, $\dot{p}/p = \mu + \varepsilon$, the key difference from gross saving in (9) is that recurrent saving includes only the persistent component of capital gains, μ , but not the transitory component, ε . Equivalently, recurrent

²⁰That unrealized capital gains are a form of income and should therefore be taxed is a well-known argument in the public finance literature dating back to the contributions of these authors. Also see Eisner (1988) and Robbins (2019) for arguments to include capital gains in the income definition used in the National Income and Product Accounts.

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 considerably 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.”

2.2 Housing

Our next benchmark model includes housing. Housing is an important asset for most 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 housing²² 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 clearly is 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 housing as a consumption good but that counterfactually assumes that housing is fully divisible and can be freely transacted. Households receive a constant labor income w and invest in two assets: housing and bonds. Housing enters the utility function and its price changes deterministically

²¹Arguably, our definition of “recurrent income” 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).

²²As [Glaeser \(2000\)](#) 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.”

over time. Households maximize

$$\int_0^\infty e^{-\rho t} \frac{C_t^{1-\gamma}}{1-\gamma} dt \quad \text{with} \quad C = \left(\alpha^{1/\eta} c^{\frac{\eta-1}{\eta}} + (1-\alpha)^{1/\eta} h^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \quad (11)$$

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

$$c + \dot{b} + p\dot{h} = w + rb, \quad (12)$$

where b are bonds and p is the house price. We can rewrite this budget constraint in terms of total (non-human) wealth, $a := b + ph$. The rewritten constraint is $c + Rh + \dot{a} = w + ra$ 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. Conveniently, our model with housing can be solved in closed form.

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

$$C_t = \frac{m_t}{P_t} \left(\frac{w}{r} + a_t \right), \quad \dot{a}_t = w + ra_t - m_t \left(\frac{w}{r} + a_t \right), \quad (13)$$

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

In particular, note the presence of the price index for total consumption P_t , which depends on the user cost of housing R_t , thereby capturing in a precise fashion the intuition that housing is both an asset and a consumption good (just like in Glaeser’s quote in footnote 22.)

Motivated by our results in Section 2.1 we again distinguish between transitory and persistent housing capital gains. Consider first the transitory case, i.e. a one-time increase in the house price that has no effect on the future price path. Recall that in the simple model above, the household optimally saved one hundred percent of any transitory capital gain. As we show in the Appendix, it follows from Proposition 1 that the same is true with housing: a one-time, unexpected house price increase translates one-for-one into higher net worth while total consumption is unchanged.²³

²³Because total consumption C is unchanged in this case, so is welfare. This is Glaeser’s point in footnote

The more interesting case is that of a persistent housing capital gain, i.e. an anticipated change in the growth rate of house prices.

Corollary 1 *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 recurrent saving rate is independent of wealth. In particular, optimal saving in (13) becomes $\dot{a} = \frac{r-\rho}{\gamma} (a + \frac{w}{r})$, 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 (13). When the IES $1/\gamma = 1$, $m_t = r + \frac{\rho-r}{\gamma}$ and saving is independent of house 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.²⁴ 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.²⁵

Corollary 1 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.” Figure 3 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

22. To be clear, even though total consumption is unchanged, 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.

²⁴The income effect is that a higher future house price makes future total consumption more expensive and hence households consume less today. The substitution effect is that a higher future house price makes total consumption cheap today relative to tomorrow and hence households consume more today. $1/\gamma$ 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.

²⁵A 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 adjust 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 3: Response of key variables to a change in persistent housing capital gain.

Notes: the Figure assumes $\eta = 0$, $\alpha = 0.5$, $IES = 1/\gamma = 1$, 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).²⁶ 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 fewer physical units of housing h .

Our results show that even if we take into account that housing is not just an asset but also a consumption good, households should treat capital gains in housing similar to capital gains in financial assets. Households should save transitory housing capital gains while they should partly consume out of persistent ones. 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 transaction costs in Section 6.

2.3 Common Extensions

The framework in Section 2.1 was deliberately simplistic so as to introduce our main saving concepts and highlight two properties that largely carry over to many richer models from the macroeconomics literature. These properties are (i) the recurrent saving rate, i.e. the systematic component of the gross saving rate, is approximately flat or slightly decreasing

²⁶The figure also assumes that relative weight on non-durable consumption and housing α equals 0.5.

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, gross and recurrent saving rates. For versions of these environments with changing asset prices, the statements in this subsection apply to recurrent 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.3. The model generates a saving policy function $\dot{a} = s(a, w)$ where w is labor income. Figure 4 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 . 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.²⁷ The right panel displays saving rates without conditioning on income and instead

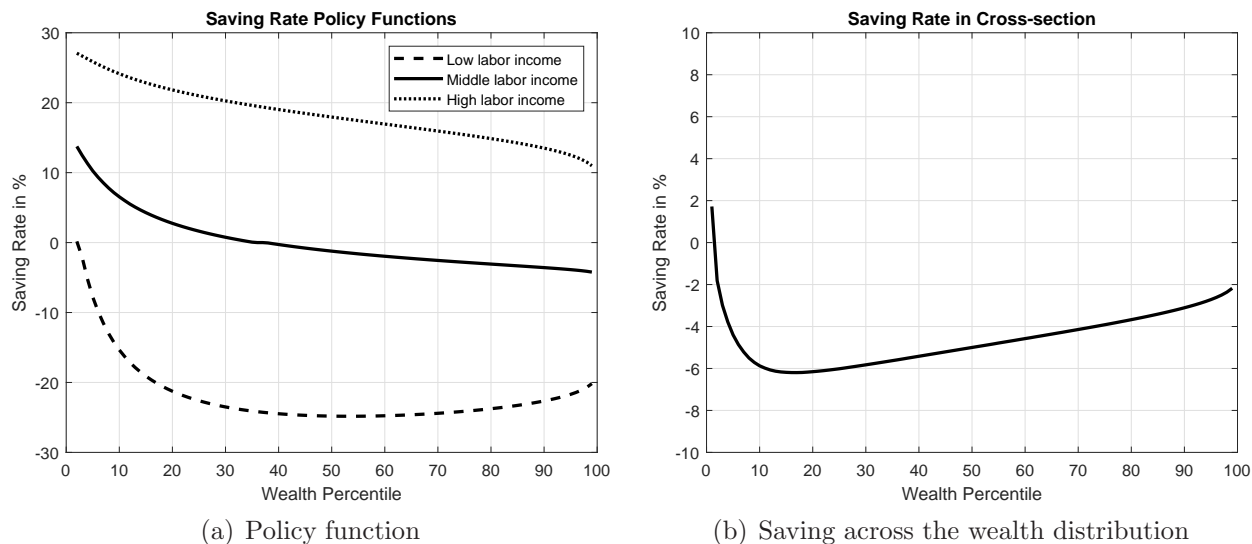


Figure 4: Saving rates under income risk and borrowing constraints.

²⁷It is straightforward to show that the flat-saving-rate result from Section 2.1 now applies as wealth becomes large. More precisely, for all w , $s(a, w)/(w + ra) \rightarrow (r - \rho)/(\gamma r)$ as $a \rightarrow \infty$ meaning that the saving rate policy function even converges to the same value as in Section 2.1. 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).

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 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 at 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.²⁸ Consequently, life-cycle considerations introduce a cross-sectional correlation between saving rates and wealth, simply because both are correlated with age and income. Conditional on age and income, saving rates are slightly decreasing with wealth, an observation that has previously been made by [De Nardi and Fella \(2017\)](#). Hence, we will conduct empirical exercises in which we condition on age, income as well as education.

The third model extension that we explore is heterogeneity in individual-specific propensities to save, operationalized as heterogeneity in discount rates. Such “heterogeneous betas” are a popular device for generating wealth dispersion in economic models (e.g. [Krusell and Smith, 1998](#)). Patient individuals save a lot and consequently have high accumulated wealth. An immediate implication is a positive cross-sectional relationship between wealth and saving rates. For our empirical purposes, saving rates in this case are “type-dependent” ([Gabaix et al., 2016](#)), so that this cross-sectional correlation should be explained by historical saving behavior at the individual level. We will therefore present results conditioning on past saving behavior.

²⁸In 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.

3 Data, Definitions and Institutional Setting

We base our study on Norwegian administrative data. In contrast to other Scandinavian countries, Norway still levies a wealth tax (in addition to income taxes) on its households. In the process of collecting these taxes, all households are every year obliged to provide a complete account of household income and balance sheet components down to the single asset category. These data are reported to the tax authorities by third parties, and are the foundation of our empirical analyses. Through unique identifiers for individuals and households we can link these data to further data sources to address our research question. Below we describe these data 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.²⁹ We combine 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). Details on these sources and data sets are provided in the Appendix.

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 reported electronically by third-parties limits the scope for tax evasion and other sources of measurement error from self-reporting.

²⁹Arguably, saving and consumption decisions are made at the household level, which in the data is defined either as a single-person household, or a married couple, or a cohabiting couple with common children. In effect we compute variables at the household level in per-capita terms, and then weight households by the number of adults in the household.

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 individually.³⁰ Available tax record data are annual and go back to 1993. However, we primarily use data from 2005 to 2015 since the shareholder registry is only available from 2004 (and therefore capital gains from 2005). Tax information is primarily collected from 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 at the end of the year 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 disposable income is lower than the base amount defined in the Norwegian Social Insurance Scheme (in 2011 equal to NOK 79,216, or about USD 13,700), where the household has immigrated within the last two years, or where the household is either formed or dissolved.³¹

Asset Classes. We observe wealth information 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 total assets net of outstanding debt, alternatively termed net worth.

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 primarily consist of holdings in bond mutual funds and money market funds³²). In Figure 5, we collect these 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 with a valuation schedule based on list price as newly purchased and vehicle age. Public equity is the sum of stock funds and listed stocks held directly or indirectly via private businesses, all at market values (how we unpack stocks that are indirectly held via private businesses is explained in

³⁰Taxable wealth above an exemption threshold is taxed at a flat rate which has been around 1% during our sample period. The exemption threshold has increased over time. Toward the end of our sample the threshold is approximately USD 260,000 (NOK 1.5 million) for a married couple (and half of that for a single person). Importantly for our purposes, assets are reported and recorded even if the household’s total wealth falls short of the threshold. For several asset classes, values are discounted when measuring taxable wealth. We adjust these discounted values back to market values when computing household wealth.

³¹We do not drop households who transition from a status as cohabitants to married.

³²More disaggregated data from the later years of the sample allow us to see this.

Appendix B.4).

“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 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 derived from the book value, 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. Note that when computing private business wealth, we have extracted listed stocks and debt and placed them, respectively, in the public equity and debt categories on the owners’ balance sheets (see Appendix B.4).

Housing wealth includes the value of a household’s principal residence, secondary homes and recreation estates. To obtain an up-to date value of housing wealth, the tax authority have since 2010 utilized predicted values from area- and housing-type-specific hedonic regressions. Here the price per square meter of houses transacted during the year are regressed on house characteristics. Prior to 2009 the housing wealth values in the tax records are related to the original purchase prices, which could deviate substantially from a current valuation. To alleviate these discrepancies we obtain a corrected value of housing wealth for our purposes. To do so, 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 techniques. 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 estimation, 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 to calculate housing wealth over our entire sample period.

The portfolio shares in Figure 5 reflect how Norwegians’ balance sheets are dominated by housing and debt, where debt is the sum of mortgages, student loans, consumer debt, personal loans in the household and debt by private businesses. Figure 5 plots the average holdings of different asset classes relative to total assets for different wealth percentiles. To




Figure 5: Asset class shares in household portfolios.

Notes: The figure displays the mean portfolio share in percent of total asset across the wealth distribution. Safe assets is the sum of deposits, bonds, and informal loans. Debt is the sum of private debt and debt held indirectly via private firms. Public equity is the sum of directly-held stocks, stock funds, and stocks held indirectly via private firms. Private business is the book value of private firms, taking out public stocks and debt.

the left, we see households with negative net worth. By definition, these are households whose liabilities exceed their assets, reflected in a ratio of debt to assets above one-hundred percent (the solid green line with dot markers exceeds 100% and hence escapes the figure’s range for these households). We also note that the least wealthy households hold a high ratio of housing wealth to total assets. As we move further to the right, the housing share declines steadily until we reach households with approximately zero net worth around the 12th percentile of the wealth distribution. These households hold approximately no assets and no debt. Thereafter, from the 12th to approximately the 25th wealth percentile, the housing portfolio share grows rapidly.

Another striking feature of Figure 5 is the portfolio share of safe assets (predominantly bank deposits) which is high across the whole wealth distribution. In contrast, public equity is a relatively small component of Norwegian households’ wealth. 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.

Table 1 provides main descriptive statistics of our sample.³³ The first panel in the table displays demographics. The second panel describes the wealth components. While the portfolio shares of public equity above appear low, participation rates in public equity or mutual funds are high relative to comparable EU-countries (ECB, 2016). In Appendix B.4 we verify the consistency between our micro-level data on stockholding and the aggregate holdings by households at the Norwegian stock exchange. The third panel in Table 1 provides statistics on gross labor earnings and yearly changes in net worth. We see how the mean change in net worth is almost half as high as labor earnings, and how the variation in net worth growth well exceeds the variation in labor earnings.

On top of the asset classes displayed here comes pensions, which we return to in Section 3.3 and impute in extensions that are detailedly explained in the Appendix.

Table 1: Descriptive statistics.

	Mean	SD	P10	Median	P90	Part. Rate
<i>Demographics:</i>						
Age	49.90	17.65	27	48	75	
Male	0.49	0.50	0	0	1	
Years of education	0.31	0.46	0	0	1	
Less than high school	12.15	3.20	8	12	16	
High school	0.30	0.46	0	0	1	
College education	0.39	0.49	0	0	1	
<i>Asset values:</i>						
Safe assets	43,466	113,421	850	14,733	1071,03	0.99
Housing	423,416	394,561	0	345,802	925,513	0.80
Debt	142,873	1,387,289	0	63,224	298,969	0.85
Public equity	10,040	596,421	0	0	13,473	0.38
Private business	53,374	1,639,997	0	0	4,964	0.14
Vehicles	6,884	103,678	0	1,389	19,236	0.57
Net wealth	394,306	829,754	-7,589	286,830	915,238	
Gross labor earnings	52,207	49,448	0	52,178	105,150	
Change in net worth	18,225	466,655	-71,843	7,364	126,472	

Notes: The table summarizes demographic characteristics and asset holdings of households pooling data for our sample period 2005-15 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).

³³Throughout the paper NOK figures are in 2011 prices, converted to USD, using the 2011 NOK/USD exchange rate of 5.77.

3.2 Implementing Different Saving Rate Concepts

Net, Gross and Recurrent Saving with Multiple Assets. As discussed in Section 2.1, we distinguish between three different notions of saving: net, gross and recurrent saving. We there defined these concepts in the context of a simple model with changing asset prices. While this simple model featured only one asset, all our definitions extend to the case of multiple assets (and liabilities) in the obvious fashion – see Appendix B.1. We now discuss how we operationalize these concepts in the data and discuss our empirical strategy for separating gross saving into net saving and capital gains.

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. We therefore directly 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 utilize the same housing transactions data that we use 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 provided in equation (A42) in the Appendix.

For housing we additionally adjust the definition of income. 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 let out. We then add this implicit rental rate to income.³⁴ We follow Eika, Mogstad and Vestad (2018) and distribute

³⁴For example, consider the budget constraint in terms of gross variables (9). Using the notation of Section 2.2 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

the aggregate value of owner-occupied housing services from the national accounts across households in proportion to the value of their house, which implies a constant rent-to-value ratio across households equal to 2.88 percent.

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, we directly 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, the values we observe for private businesses are not market values. Instead, we use assessed tax values, that are related to book values. If book values are not marked to market, changes in book value in principle measure net saving. Year-to-year changes in private business values are therefore likely a better approximation to net saving than to gross saving. Apart from extracting publicly listed stocks from their balance sheets (as explained in Appendix B.4), we therefore also do not attempt to separate changes in private equity values into net saving and capital gains. Since private businesses are an important asset at the very top of the wealth distribution but not elsewhere (Figure 5), this likely results in an underestimate of capital gains and gross saving among the very wealthiest.

Finally, consider bonds and money market funds. Decomposed holding data from the Norwegian Fund and Asset Management Association reveal that households holdings in bond and money market funds are concentrated in government bonds with maturity below one year and medium-term bonds (two to four years).³⁵ These short-term bonds are unlikely to experience substantial price changes and bond holdings are a very small fraction of households' asset holdings, with an aggregate portfolio share that averages just 0.4% over our

of (9) is

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

³⁵Among these short- and medium-term bonds, the weight on short-term bonds is approximately seventy percent.

sample period. We therefore make the simplifying assumption that bonds do not experience capital gains.

All other assets listed in Section 3.1, as well as liabilities, do not experience asset price changes. Hence, capital gains are zero, and net saving equals gross saving for those assets.³⁶

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 (10), we need 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.³⁷ 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 .

For public equity, we compute the persistent component of capital gains on directly held stocks as the average appreciation of the Norwegian ex-dividend stock price index, the OBX Price Index (OBXP), as far as it goes back (1991).³⁸ 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 (four of Norway’s main cities) provided by Norges Bank.³⁹ The calculated persistent housing capital gains are 2.22% for Oslo, 2.78% for Bergen,

³⁶Vehicles 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 a decline in the physical units of that asset, as opposed to merely a revaluation effect. That is, depreciation leads to negative net saving as opposed to a capital loss. Therefore, also for vehicles, net saving equals gross saving. Alternatively, we have conducted our main empirical exercise counting vehicle depreciation as capital loss and results are largely unchanged.

³⁷That 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.

³⁸The 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.

³⁹This 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/>.

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. For housing, it is particularly important that we take a long-run perspective and compute persistent capital gains using a long time series. The reason is that during our sample period from 2005 to 2015, the Norwegian housing market went through a period of abnormally high house price growth. For example, between 2005 and 2015 house prices grew on average by 4.80% nationwide, 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 then follows from (9) and (10) 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.

Summarizing, we use the following numbers for persistent capital gains: 3.25% for directly held stocks, 3.18% for stock funds, and 2.26% for housing. To put these numbers into perspective, the corresponding average rates of return (persistent capital gains plus average dividend yield) are 6.03% for directly held stocks and 5.14% for housing.⁴⁰ In particular, we recover the typical finding that stocks yield a higher average return than housing.

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.⁴¹ The main role of this fund is to allow (smooth) public spending in

⁴⁰For stocks, we compute the total return from the OBX Total Return Index (https://www.oslobors.no/ob_eng/markedsaktivitet/#!/details/OBX.OSE) which is the analogue of the OBX Price Index we used to compute persistent capital gains, but including dividends. For housing, the “dividend yield” is simply the implicit rental rate that we set at 2.88% as discussed above. Hence the return to housing is $2.26\% + 2.88\% = 5.14\%$.

⁴¹See <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

excess of tax revenues and remove uncertainty regarding future governments' ability to meet public pension entitlements.

As the Appendix explains in detail, 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 reported in Section 4.5 and the Appendix, we impute each household's public pension entitlement and include pensions in measures of both saving and wealth.

4 Results

This section provides 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 theory suggests are important for this correlation: 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. Then we discuss a number of additional exercises to better understand the robustness of our findings. These include the role of pensions and time-aggregation issues.

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 4.5 which also shows saving rates separately for each year.

We start with the net saving rate, for which our benchmark theories in Section 2 predicted a declining relationship with wealth. To the left in Figure 6, we see that among the least wealthy households, the net saving rate does indeed decline with wealth. Then, after reaching a trough just below the 15th wealth percentile, it rises up to about the 20th wealth percentile, before flattening out almost completely around 6 – 8%. Note that the trough is reached approximately where net worth is zero (see Figure 5). Hence, the upshot is that net saving

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.




Figure 6: Saving rates across the wealth distribution.

Notes: The figure displays the median saving rates within wealth percentile and year, averaged across all years (2005-2015). The net saving rate is defined in equation (8) as net saving divided by disposable income. The recurrent saving rate is defined in equation (10) as recurrent saving divided by recurrent income. The gross saving rate is defined in equation (9) as gross saving divided by gross income.

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. The two saving rates follow each other down to the trough around zero net worth. Thereafter the gross diverges 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 do.

Moreover, the divergence between recurrent saving rates and net saving rates tells us the mechanism behind this finding: wealthy households save more because they have more assets that experience capital gains and they hold on to these assets.

The mechanism has two logical components. First, wealthier households hold more assets. Unless they systematically invest in assets with lower capital gains than poorer households do, their total asset holdings automatically appreciates by a larger dollar amount. Figure 7 reveals that this is true in our data: wealthier households hold more assets and there is no systematic variation in capital gains rates to these assets across the wealth distribution. 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 canonical models summarized in Section 2, and is the reason why gross saving behaves differently in the data than in these models.

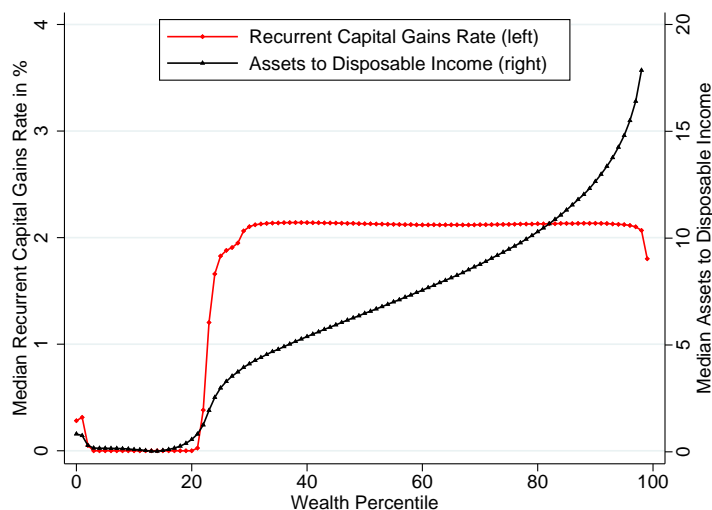


Figure 7: Average capital gains and asset-to-income ratio.

Notes: This figure displays the median recurrent capital gains rate (recurrent capital gains as a fraction of total assets) and the median total assets as a share of disposable income across the wealth distribution. All variables are computed as the median within wealth percentile and year, averaged across all years (2005-2015).

A more general takeaway from Figure 6 is that asset price changes and how people respond to these are key to understanding saving behavior across the wealth distribution. 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 benchmark 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, optimizing households should 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.

4.2 The Role of Age, Earnings, and Education

As highlighted in Section 2.3, 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:

$$\frac{s_{it}}{y_{it}} = \sum_{p=1}^{100} \phi_p D_{it,p} + f(x_{it}) + \tau_t + \epsilon_{it}, \quad (14)$$

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 year t , ϕ_p is the corresponding regression coefficients, x_{it} is the vector of control variables, 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 represent the objects of our interest: the median saving rates across the wealth distribution after controlling for age, earnings, and education. These can be displayed graphically, analogous to the results in Figure 6.

Figure 8 presents the results. It shows the median net and recurrent saving rates after

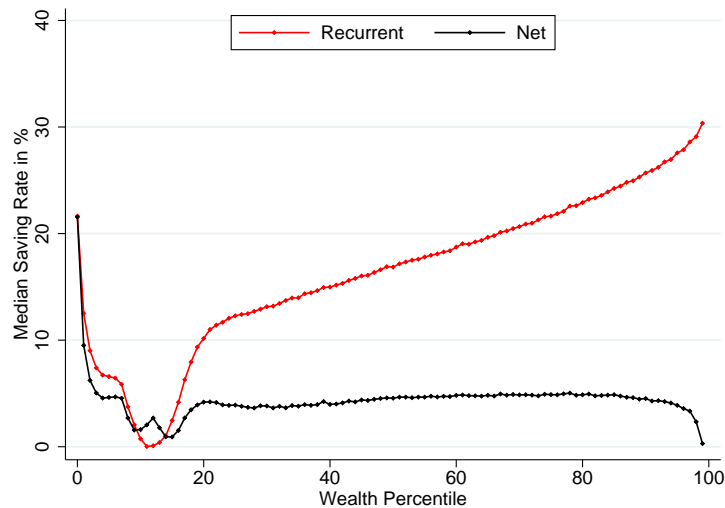



Figure 8: Saving rates, controlling for age, earnings and education

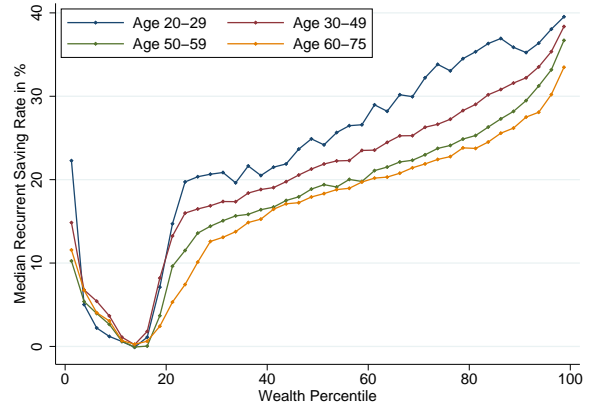
Notes: This figure displays the coefficients on wealth percentiles (ϕ_p) from estimating equation (14) as a quantile regression for net and gross saving rates, respectively.

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 counterparts. Consistent with the discussion in Section 2.3, 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 even 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 compute the median saving rate for different wealth percentiles, just like in Figure 6. Figure 9 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 affect the *level* of the saving rates, but only modest effects on how these *covary* with wealth.




(a) Age, net saving rate




(b) Age, recurrent saving rate




(c) Earnings, net saving rate



(d) Earnings, recurrent saving rate



(e) Education, net saving rate



(f) Education, recurrent saving rate

Figure 9: Saving rates across the wealth distribution by age, earnings and education
Notes: The figures display the median net saving rates (left column) and median recurrent saving rates (right column) within age, labor earnings, and education groups. All saving rates are computed as the median saving rate within wealth percentile and year, averaged across all years (2005-2015).

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 2.3 suggests that persistent individual variation in saving behavior might result in a positive correlation between saving rates and wealth. For example, 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” (Gabaix et al., 2016) 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 10 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.

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 begs the question if our findings in Figure 6 are only 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 other assets than housing? To be clear: this is not a question of our results’ robustness, but one of interpretation. In particular, the question is whether theoretical explanations of our findings should focus exclusively on housing or whether they should apply to other assets as well.

Answering how households treat capital gains on other assets than housing is complicated because the vast majority of Norwegian households hold relatively few other assets that experience asset price changes and hence capital gains (see Figure 5). For the purpose of the

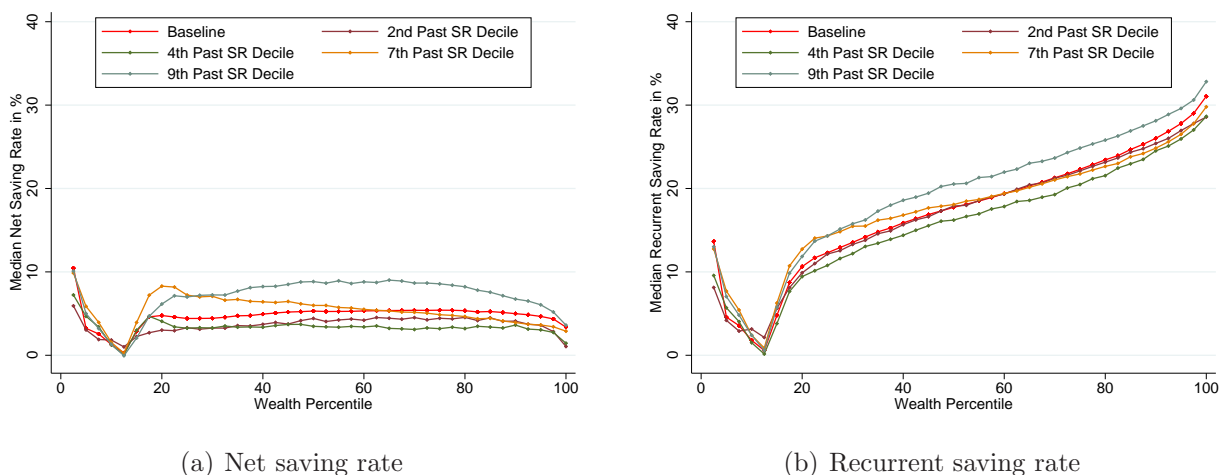
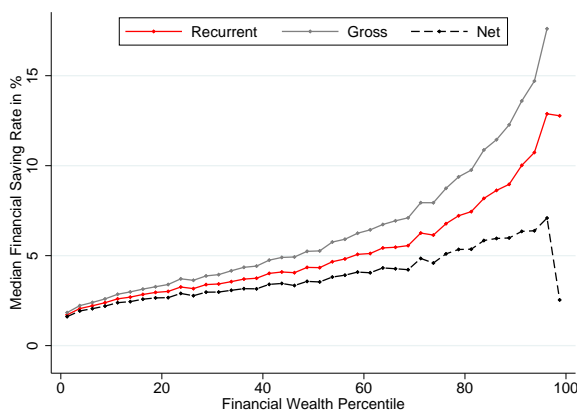


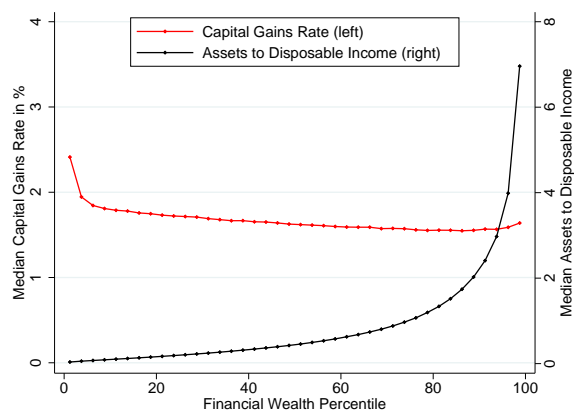
Figure 10: Saving rates across the wealth distribution within deciles of past saving rates
Notes: The figures display the median net saving rates (left) and median recurrent saving rates (right) within percentiles of past saving rates. Conditional upon observing a household for at least 4 prior years, we compute each household’s past recurrent saving rate for every year, and thereafter stratify each household-year observation by average past recurrent saving rate. All variables are computed as the median within wealth percentile and year, averaged across all years (2005-2015).

current section, we therefore restrict our sample to households that hold at least twenty-five percent of their non-housing assets in the form of stocks or stock funds. In order to exclude housing, we also focus on what we term “financial wealth” defined as net worth excluding housing and liabilities. 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 the results of this exercise, we already note one caveat: in contrast to directly held stocks, for stocks that are held in stock funds we cannot use the shareholder registry to identify capital gains for individual stocks at the security level and we instead use an aggregate index for this asset class (see the discussion in Section 3.2). For stock funds, the split of gross saving 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 11 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 recurrent saving rates, defined as in Section 4 and analogous to Figure 6. Because the sample is 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



(a) Financial saving rates



(b) Capital gains and financial assets to income

Figure 11: Relation between financial saving rates and financial wealth.

Notes: These figures display the median financial saving rates, capital gains rate, and financial assets to income ratio for a sample of individuals that hold at least USD 1,000 in public equity and at least 25% of their financial wealth in public equity (10.4% of benchmark sample). All variables are computed as the median within wealth percentile and year, averaged across all years (2005-2015).

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.⁴²

In summary, the patterns for net, gross and recurrent saving rates are not limited to accumulation of housing wealth, but are qualitatively similar for wealth accumulation in other assets too. At the same time, as already noted, the combined capital gains on other assets tend to be substantially lower than housing capital gains because Norwegian households hold relatively few assets that experience capital gains besides housing.

⁴²In 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 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.

4.5 Additional exercises


We have explored a range of robustness checks and extensions of our main analysis. In this section, we present the set of additional exercises that we find particularly interesting.

Public pensions. Our first exercise involves including public pensions in the definition of saving. Norway has a public pension system with full coverage of all citizens. We compute public pension saving as the change in the discounted value of future pension benefits, taking into account discounting, expected wage growth, and expected life-time of the household. We describe the details of the system and how we compute pension saving in Appendix B.5.

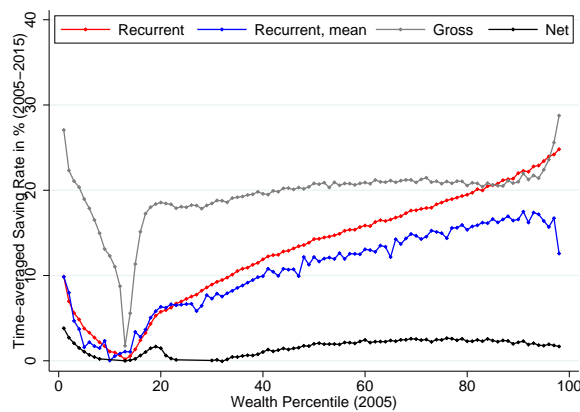
Figure 12(a) presents our main graph where we include pension saving in the definition of income and saving plotted against the original wealth distribution for comparison. By including the pension saving, we lift the saving rate of all households, but the cross-sectional variation remains the same: the net saving rate is approximately flat for the positive part of the wealth distribution, and the recurrent and gross saving rates increase with wealth.

Saving rates with alternative time averaging. Our main graph was constructed by first computing median saving rates within each percentile for every year, and then plotting the average of these yearly median saving rates. A concern with this approach is that if households transact irregularly, for example if they buy homes every fifth year, the median households might always be non-transactors who therefore save a relatively large fraction of capital gains. To address this concern, we here present an alternative way of time-averaging saving rates. We first compute the average saving rate for each household over the entire period (2005-2015). Thereafter we stratify households by their percentile in the 2005 wealth distribution. Then we compute the within-percentile median and mean of households' time-averaged saving rate. This way, the saving rate we compute also includes years in which households transact. Figure 12(b) presents the median net, gross, and recurrent saving rates, in addition to the mean recurrent saving rate, against the 2005 wealth distribution on the horizontal axis. We see that also with this definition our main result holds qualitatively: the net saving rate is approximately flat across the wealth distribution, while the recurrent saving rate increases with wealth.


Saving rates by year. Instead of averaging across time, we could simply plot saving against wealth for each year separately. Figures 12(c) and (d) do this for the net and gross saving rates. Each lighter gray line is constructed for a specific year. The thicker black lines




(a) Including Public Pensions



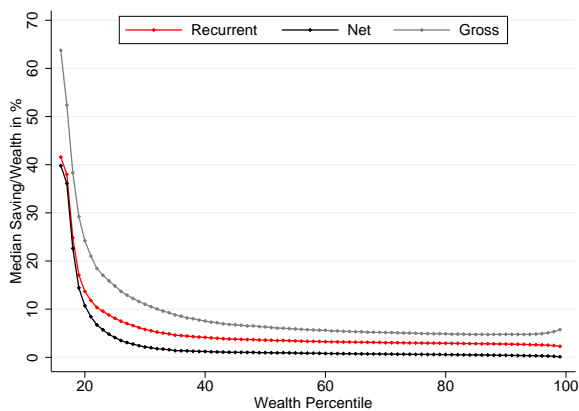
(b) Alternative Time-averaging



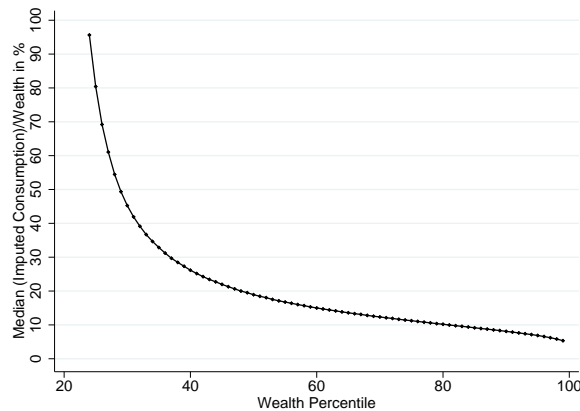
(c) Net saving rates across years



(d) Gross saving rates across years



(e) Saving rates as fraction of wealth



(f) Imputed consumption as fraction of wealth

Figure 12: Additional exercises.

Notes: (a) displays the median saving rates when we include pension saving and income from the public pension system. Appendix B.5 describes the public pension system and how we compute pension wealth, income, and saving. (b) displays the median of the average saving rates for the period 2005-2015, plotted against the wealth percentile at the beginning of 2005. (c) and (d) display the median net and gross saving rates for each year in our sample. (e) and (f) display the median saving and consumption as a fraction of wealth across the wealth distribution. We compute imputed consumption as disposable income minus net saving.

are the averages in our baseline exercises. Notice that the vertical axes are different in the two figures. Together, the two plots give a stark illustration of saving by holding. In all years, including the two where house prices fell, the plotted net saving rate has the same shape. In particular, it is approximately flat from the 25th wealth percentile and out. In contrast, the shape of the gross saving rate varies greatly across years. The gross saving rate increases most distinctly with wealth in good years where asset prices increased the most, and increases the least with wealth in bad years where asset prices increased the least. These large year-to-year fluctuations in gross saving rates are consistent with those predicted by our benchmark model (Figure 2). But what is inconsistent with it is the flat net saving rate and the increasing recurrent saving rate.

Saving rates as a fraction of wealth. Because our objective is to test the saving behavior implied by economic theory, our object of interest is defined as saving relative to income. However, other definitions of saving rates are interesting for other purposes. In particular, [Bach, Calvet and Sodini \(2018\)](#) investigate the role of saving as a fraction of wealth, i.e. the growth rate of wealth, across the wealth distribution to discuss the implications of saving behavior for the dynamics of wealth inequality. To ease the comparison with their work, we present saving as a fraction of wealth across the wealth distribution in Figure 12(e).⁴³ The figure reveals a similar pattern across the wealth distribution as in Figure 4 in Bach et al.. Hence, our results are complementary to theirs, not contradictory.⁴⁴ For comparison, Figure 12(f) also shows imputed consumption (defined as disposable income minus net saving – see equation (8)) as a fraction of wealth. Just like saving as a fraction of wealth, also consumption as a fraction of wealth declines strongly with wealth. This is consistent with the common intuition that very rich households cannot possibly consume their entire wealth. At the same time, the declining pattern in Figures 12(e) and (f) is somewhat mechanical. Part of it is simply driven by the fact that even households with little wealth typically still have some labor income out of which they save and consume so that the ratio of these quantities to wealth blows up as wealth becomes small.⁴⁵

⁴³We follow [Bach, Calvet and Sodini \(2018\)](#) and cut the figure at the bottom of the distribution, namely at the percentile below which net worth is zero or negative so that the ratio of saving to wealth is ill-defined (because computing it would require division by zero).

⁴⁴Note that our gross saving over wealth is the same as their total saving over wealth while the other definitions (our recurrent and net, and their active saving) are not comparable.

⁴⁵For instance consider the saving and consumption policy functions in our benchmark model which are given by (3) and $c(a) = \left(\rho - \frac{r-\rho}{\gamma}\right) \left(a + \frac{w}{r}\right)$. We have $s(a)/a = \frac{r-\rho}{r} \left(1 + \frac{w}{ra}\right)$ and $c(a)/a = \left(\rho - \frac{r-\rho}{\gamma}\right) \left(1 + \frac{w}{ra}\right)$, both of which are declining with wealth and blow up as $a \rightarrow 0$, just like in panels (e) and (f).

Zooming in on the right tail of the wealth distribution. We next zoom in on the right tail of the wealth distribution, in Figure A15 in the Appendix. This subset of the population is particularly interesting because they hold a disproportionately high share of the economy-wide stock of wealth, and because their asset portfolio is considerably less tilted toward housing. Figure A15(a) visualizes the latter point. A distinguishing feature of the right tail of the wealth distribution is that the wealthiest households hold more than 50% of their assets in private businesses. Furthermore, their portfolio is also more leveraged than is the case for most other households, with a debt share of almost 40% of total assets.

When it comes to our study’s main graphs on net and recurrent saving, Figure A15(b) shows that zooming in on the tight tail yields one substantial difference from our main plots. Within the top percentile, the gross saving rate first increases similarly to what it does across the main part of the population, but drops as we move into the top 0.1 percent. The explanation is found in Figure A15(c), which plots capital-gains-to-income ratios across the wealth-to-income distribution. Relative to income, observed capital gains drop markedly as we step into the top 0.1 percent group. Hence, the saving pattern observed in Figure A15(b) is driven by the fact that households at the very top hold relatively less wealth in assets that experience yearly capital gains, and is consistent with saving by holding.⁴⁶

5 Implications for Aggregate Wealth Accumulation

Our results suggest that households accumulate wealth by holding on to assets that experience persistent capital gains. Hence, in contrast to what canonical models of household saving would suggest, persistent capital gains constitute a main source of micro-level wealth accumulation. A natural follow-up question then becomes, to what extent does saving-by-holding matter at the aggregate level? We now turn to this question by quantifying the contribution of saving-by-holding for aggregate wealth accumulation in Norway.

To this end we proceed by means of counterfactual experiments motivated by the benchmark models in Section 2. For each household, we impose a counterfactual saving rate and then compute how the household’s wealth would have evolved under the imposed saving behavior. Thus, apart from the base year in which we start (1995), our procedure yields a

⁴⁶Part of the low measured capital gains for the top 0.1% is likely due to the fact that tax values of private businesses are related to book values rather than market values – see the discussion in Section 3.1 – and therefore private business wealth does not generate capital gains. Since private business wealth accounts for a large share of the assets held by the top 0.1% (see panel (a)), capital gains for this group are likely an underestimate of their capital gains if private businesses were valued at market values.

sequence of counterfactual wealth levels \hat{a}_{it} for each year t for every household i . Starting from the budget constraint “consumption plus recurrent saving equals recurrent income” (equation (10)), we construct counterfactual wealth levels from one year to the next as follows:

$$\hat{a}_{i,t+1} = \hat{a}_{i,t} + \hat{s}_{i,t} (w_{i,t} + \theta_{i,t}\hat{a}_{i,t} + \mu_{i,t}\hat{a}_{i,t}) + \varepsilon_{i,t}\hat{a}_{i,t}. \quad (15)$$

Here $\hat{s}_{i,t}$ is the imposed counterfactual saving rate out of recurrent income, $w_{i,t}$ is labor income, $\theta_{i,t}$ is the capital income rate, $\mu_{i,t}$ are persistent capital gains, and $\varepsilon_{i,t}$ are transitory capital gains. Note that variables with a hat are counterfactually constructed and variables without a hat are unchanged from our baseline empirical exercise. For each individual household, its actual saving behavior corresponds to a particular split of recurrent income into consumption and recurrent saving (again see (10)), and our counterfactuals simply explore different assumptions about this split by varying the recurrent saving rates $\hat{s}_{i,t}$.

We consider two counterfactual scenarios, differentiated by our choice of $\hat{s}_{i,t}$.⁴⁷

1. **Flat recurrent saving rate.** We adjust the recurrent saving rate $\hat{s}_{i,t}$ so that its within-percentile mean is independent of wealth (i.e. flat with wealth) and equal to the unweighted cross-sectional average, while the within-percentile dispersion in saving rates is unchanged. More precisely, we set $\hat{s}_{i,t} = s_{i,t} - \bar{s}_{p,t} + \bar{s}_t$, where $s_{i,t}$ is the recurrent saving rate for household i , $\bar{s}_{p,t}$ is the average recurrent saving rate for percentile p , and \bar{s}_t is the recurrent saving rate, all in year t .
2. **Recurrent saving rate equal to net saving rate.** For every year, we compute how much wealth each household would accumulate if their recurrent saving rate were at the level of the household’s observed net saving rate. Hence, given our main empirical finding in Figure 6 above, this exercise reduces the average recurrent saving rate and shows how wealth would have evolved if households had treated persistent capital gains in the same way as they treat net income.

The two scenarios thus exemplify how aggregate wealth would have evolved if the saving rate were independent of wealth, in the first scenario at the level of average recurrent saving rates and in the second scenario at the level of the net saving rate. In terms of an economic

⁴⁷We also conducted a third counterfactual that starts from the budget constraint “consumption plus gross saving equals Haig-Simons income” in equation (9), namely $\hat{a}_{i,t+1} = \hat{a}_{i,t} + \hat{s}_{i,t}^g (w_{i,t} + \theta_{i,t}\hat{a}_{i,t} + \mu_{i,t}\hat{a}_{i,t} + \varepsilon_{i,t}\hat{a}_{i,t})$ and then assumes a gross saving rate that is independent of wealth. The results are similar to those of the first counterfactual exercise below.

model, the scenarios represent how wealth-to-income ratio would have evolved if one used the model in Section 2.1 calibrated either to the mean recurrent saving rate (scenario 1) or the mean net saving rate (scenario 2). We depict these counterfactual scenarios in Figure

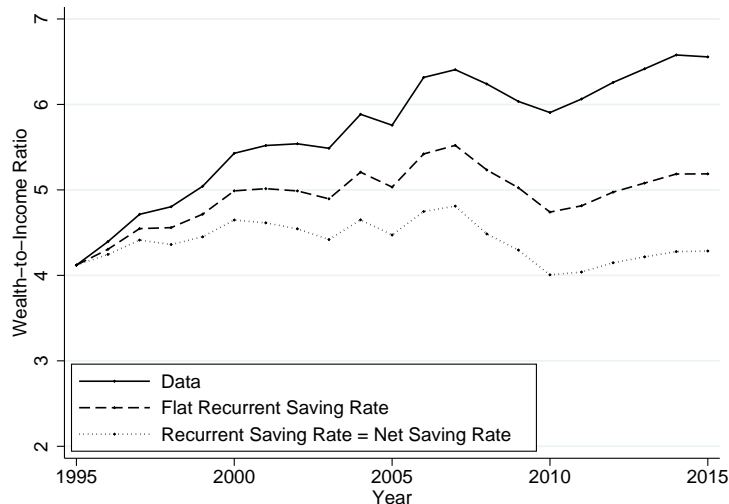


Figure 13: Actual and counterfactual wealth-to-income ratios.

Notes: This figure displays the evolution of the ratio of net wealth to disposable income in the data and two counterfactual scenarios. See the text for an explanation of the exact procedure. We adjust individual’s capital gains and capital income in both counterfactual experiments such that they correspond to that individual’s counterfactual wealth, assuming fixed portfolio shares. In both scenarios, we assume a 100% saving rate out of transitory capital gains.

13. The solid line is the wealth-to-income ratio across all households in our sample. The wealth-to-income ratio increased markedly over our sample period, from approximately 4 in 1995 to more than 6.5 in 2015.⁴⁸

This development would have been very different if households behaved according to our counterfactual scenarios. In scenario 1 (flat recurrent saving rate), the economy-wide accumulation of wealth relative income would have been approximately halved. This effect comes only from the fact that when the rich save less and the poor save more, the former household’s influence dominates as their total income is (far) greater than the latter household’s total income. In scenario 2 (recurrent saving rate = net saving rate), aggregate wealth accumulation drops further, resulting in approximately the same counterfactual wealth-to-

⁴⁸These numbers are computed directly from our micro data and feature a smaller increase than some other series, including the national accounts. For example, the World Inequality Database (<https://wid.world/country/norway/>) reports an increase in the Norwegian wealth-to-income ratio from approximately 3 to 6 over the same time period, i.e. an increase by 3 rather than 2.5 as in our data. The levels in our series are higher, mostly because we use our own higher-quality measure of housing values (see Section 3.1).

income ratio in 2015 as the actual ratio in 1995. The additional reduction comes from the fact that our counterfactual now assumes people on average consume considerably more out of their persistent capital gains.⁴⁹

Our counterfactual scenarios share the main properties of the simple models reviewed in Section 2. These properties are common to structural frameworks that have been widely applied in macroeconomics over the past decades. Although crude, by for instance ignoring general equilibrium feedback from wealth to saving behavior, we believe our counterfactuals suggest that the empirical saving behavior we have documented is crucial to carry in mind when studying phenomena related to aggregate wealth accumulation.

6 Theoretical Interpretation

If canonical models cannot rationalize our empirical findings, what theories can? In a reduced form, what is needed to rationalize Figure 6 is a theory where households consume very little out of any capital gains they experience – that is, they have a saving rate out of capital gains close to one hundred percent, precisely what we have termed “saving by holding.”⁵⁰ In this section, we propose two candidate explanations for our empirical findings: first, time-varying discount rates or, more generally, asset-demand shifts that affect households throughout the wealth distribution, e.g. a sustained housing preference shift; second, a theory featuring multiple assets and portfolio adjustment frictions. We also briefly discuss other candidate explanations, several of which are complementary. These include: non-homothetic preferences, misperceptions about the stochastic process for asset prices, and behavioral explanations.

⁴⁹Of course, both counterfactuals ignore general equilibrium effects on asset prices. For example, consider scenario 2 in which households consume a large fraction of their persistent capital gains, i.e. they sell off assets that experience these persistent capital gains. A natural conjecture is that such a scenario would result in downward pressure on the prices of these assets, i.e. lower capital gains.

⁵⁰ Interestingly, that such a reduced-form saving rule is a useful approximation to individual saving behavior was already conjectured by Meade (1964). In a discussion of individual wealth dynamics and wealth inequality, he discusses the potential importance of capital gains and remarks (see endnote 14): “The influence of capital gains could be even more marked than is implied in the text. Suppose that property owners regard as their income only the income paid out on their property and save a fraction of this, but in addition automatically accumulate 100 per cent of any capital gain not paid out in dividend or rent or interest. Then the formula for [the growth rate of wealth] k becomes $k = S\frac{E}{K} + SV + V'$ where [S is the saving rate, E is earnings, K is wealth,] V is the paid-out rate of return on capital and V' is the rate of return from capital gains.” In a reduced form, this simple saving rule is precisely what is needed to explain our empirical findings in Figure 6.

6.1 Time-Varying Discount Rates or Asset Demand Shifts

We now show that theories with time-varying discount rates have the potential to explain our empirical findings. More generally, any time variation in preferences and beliefs that affects stochastic discount factors and hence asset demand falls into this category. A typical finding in the asset pricing literature is that a large fraction of asset price fluctuations is driven by time-varying discount rates (see e.g. Campbell, 2003; Cochrane, 2005). We here argue that they also have the potential to help explain our findings.

Individuals have preferences

$$\int_0^\infty e^{-\int_0^t \rho_{is} ds} u(c_{it}) dt, \quad u(c) = \frac{c^{1-\gamma}}{1-\gamma} \quad (16)$$

where the discount rate $\{\rho_{it}\}_{t \geq 0}$ may be both individual-specific (i) and time-varying (t). Such discount rate variation could, for example, be due to intertemporal preference shifts.⁵¹

Apart from more general discounting, there are two additional differences from the model with changing asset prices in Section 2.1. First, we assume away any form of uncertainty because this complication is inessential for the points we want to make. Second, one goal of this section is to think more carefully about the sources of asset-price fluctuations. We therefore write equation (4) slightly differently as $c_{it} + p_t \dot{k}_{it} = w + \Theta_t k_{it}$, where $\{\Theta_t\}_{t \geq 0}$ is the asset's dividend which follows an exogenous time path. We also adopt the approach of the asset-pricing literature to treat the required asset return $\{r_t\}_{t \geq 0}$ as a primitive and the asset price as an outcome.⁵²

$$p_t = \int_t^\infty e^{-\int_t^s r_\tau d\tau} \Theta_s ds. \quad (17)$$

As before, the optimal policy functions can be solved analytically.

Proposition 2 *Let $\chi_{it} := \left[\int_t^\infty \exp\left(-\int_t^s \left(r_\tau + \frac{\rho_{i\tau} - r_\tau}{\gamma}\right) d\tau\right) ds \right]^{-1}$. Then optimal consumption is given by*

$$c_{it} = \chi_{it} \left(p_t k_{it} + \int_t^\infty e^{-\int_t^s r_\tau d\tau} w ds \right). \quad (18)$$

How persistent capital gains affect consumption depends on the source of these capital gains.

⁵¹The preferences in (16) are equivalent to $\int_0^\infty e^{-\rho t} \alpha_{it} u(c_{it}) dt$ where α_{it} is a preference shifter satisfying $\alpha_{i0} = 1$ for all i . The implied discount rate is then given by $\rho_{it} := \rho - \dot{\alpha}_{it}/\alpha_{it}$.

⁵²In contrast, in Section 2.1, we treated the asset price path $\{p_t\}_{t \geq 0}$ as the primitive and the return $\{r_t\}_{t \geq 0}$ as the outcome. For now, we do not take a stand where the required rate of return $\{r_t\}_{t \geq 0}$ in (17) comes from. One intuitive reason is that individuals can save in another asset (e.g. bonds) that pays a return $\{r_t\}_{t \geq 0}$. Arbitrage then requires $r_t = (\Theta_t + \dot{p}_t)/p_t$ which implies (17). Alternatively, $\{r_t\}_{t \geq 0}$ could be pinned down from preferences in general equilibrium (see the discussion in Special Case I below).

From (17), growing asset prices can only be due to one of two factors: dividend growth or declining returns. We consider these in turn. We argue that if capital gains are due to dividend growth, then individuals should consume out of the resulting capital gains so that this version of the theory cannot explain our empirical findings. However, if capital gains are due to time-varying returns then the theory can explain Figure 6 under one additional assumption: individuals' discount rates shift in line with these returns throughout the wealth distribution. This happens, for example, in a closed economy with homogeneous preferences.

Special Case I: Capital Gains due to Dividend Growth. First consider the case where the asset return is constant, $r_t = r$ for all t , and the only source of capital gains is dividend growth. For simplicity assume that dividends grow at a constant rate $\mu > 0$, $\Theta_t = \Theta_0 e^{\mu t}$. Then from (17), the asset price is given by the Gordon formula $p_t = \Theta_t / (r - \mu)$ and the price grows at $\dot{p}_t / p_t = \mu$. Note that these results correspond exactly to the assumptions in Section 2.1: defining $\theta := \Theta_t / p_t$, the budget constraint is equivalent to (4) and (5). If we further assume that $\rho_{it} = \rho$ for all i and t , we obtain the same expressions for saving policy functions and saving rates as in Section 2.1.⁵³ Individuals should therefore consume out of persistent capital gains $\dot{p}_t / p_t = \mu$. Because the setup is equivalent to that in Section 2.1, the gross saving rate is flat and the net saving rate is decreasing with wealth. Consequently, the theory in this section cannot rationalize our findings in Figure 6 if dividend growth is the source of capital gains. Note that canonical macro models of household saving behavior implicitly assume that either there are no capital gains or that all capital gains are due to dividend growth.

Special Case II: Capital Gains due to Time-Varying Returns. Next consider the opposite case where dividends are constant, $\Theta_t = \Theta$ for all t , but required returns $\{r_t\}_{t \geq 0}$ are declining over time and this is the source of a rising asset price in (17). The following corollary to Proposition 2 states this section's main result.

Corollary 2 *Assume that dividends are constant over time, $\Theta_t = \Theta$ for all t , and that returns $\{r_t\}_{t \geq 0}$ are declining over time so that $\dot{p}_t > 0$. Consider an individual whose time-varying discount rate equals the return to wealth at every point in time, $\rho_{it} = r_t$ for all t .*

⁵³Under these assumptions (18) implies that $c_{it} = \left(r + \frac{r-\rho}{\gamma}\right) \left(a_{it} + \frac{w}{r}\right)$ with $a_{it} := p_t k_{it}$ and hence gross saving is $\dot{a}_{it} = \frac{r-\rho}{\gamma} \left(a_{it} + \frac{w}{r}\right)$. This is the same expression as (3) and also the special case of (7) with $\sigma^2 = 0$.

Then this individual's optimal consumption, net and gross saving are

$$c_{it} = w + \Theta k_{it}, \quad p_t \dot{k}_{it} = 0, \quad \dot{a}_{it} = \frac{\dot{p}_t}{p_t} a_{it}. \quad (19)$$

The net saving rate is $p_t \dot{k}_{it}/(w + \Theta k_{it}) = 0$ which is independent of wealth, and the gross saving rates is $\dot{a}_{it}/(w + r_t a_{it}) = (\dot{p}_t/p_t)/(w/a_{it} + r_t)$ which is increasing with wealth a_{it} .

The proof is simple and instructive: first, with constant dividends the asset price from (17) is $p_t = \Theta \int_t^\infty e^{-\int_t^s r_\tau d\tau} ds$; second, with $\rho_{it} = r_t$ all t , $\chi_{it} := [\int_t^\infty e^{-\int_t^s r_\tau d\tau} ds]^{-1}$. Plugging into (18) yields $c_{it} = w + \Theta k_{it}$ in (19). All other results follow immediately. \square

Intuitively, changes in the required asset return and the corresponding capital gains have income and substitution effects. The assumption that dividends are constant over time eliminates the income effect. The assumption that $\rho_{it} = r_t$ eliminates the substitution effect. It follows that consumption is independent of capital gains as in (19).

To see that constant dividends eliminate the income effect, consider the case where the intertemporal elasticity of substitution $1/\gamma$ equals zero. Consumption is then given by (19) with $\chi_{it} := [\int_t^\infty e^{-\int_t^s r_\tau d\tau} ds]^{-1}$, independently of the time path of ρ_{it} . The income effect of capital gains is thus captured by the ‘‘annuity value of financial wealth’’ $[\int_t^\infty e^{-\int_t^s r_\tau d\tau} ds]^{-1} p_t k_{it}$. As already noted, with constant dividends, $p_t = \Theta \int_t^\infty e^{-\int_t^s r_\tau d\tau} ds$ and hence the annuity value is Θk_{it} which is independent of capital gains.⁵⁴ As already noted, the additional assumption $\rho_{it} = r_t$ eliminates intertemporal substitution effects of return changes. This is easy to see from individuals' Euler equations $\dot{c}_{it}/c_{it} = (r_t - \rho_{it})/\gamma$.⁵⁵

Takeaways. Summarizing, the simple model in this section can rationalize our findings if all asset price changes are due to changes in required returns $\{r_t\}_{t \geq 0}$ and, additionally, individuals' discount rates equal that return, $\rho_{it} = r_t$ for all t . At this point, the reader may wonder: why would it be the case that the asset return and discount rates are equalized even though both are moving over time? One answer is that this is the outcome in a

⁵⁴In contrast, in special case I above with constant returns $r = \theta + \mu$ we have $[\int_t^\infty e^{-\int_t^s r_\tau d\tau} ds]^{-1} = r$ and hence the annuity value of financial wealth is $r p_t k_{it} = (\theta + \mu) p_t k_{it}$ which *does* include capital gains. The formalism of [Auclert \(2019\)](#) provides a somewhat more succinct explanation for the absence of an income effect in special case II with constant dividends: changes in r_t would, in general, have an income effect but this income effect is zero because the household's ‘‘unhedged interest exposure’’ (URE) is zero.

⁵⁵If r_t were declining while ρ_{it} were constant so that eventually $r_t < \rho_{it}$, individuals would sell off some of their assets to consume, a standard intertemporal substitution effect. As this discussion suggests, Corollary 2 holds for any ρ_{it} not necessarily equal to r_t under the extreme alternative assumption that the intertemporal elasticity of substitution $1/\gamma$ equals zero.

closed economy in which discount rates are common across individuals $\rho_{it} = \rho_t$ and decline exogenously over time: equilibrium then requires that $r_t = \rho_t$. Put differently, if *the reason* why asset prices are increasing is that everyone’s demand for the asset has increased, then those individuals will not sell off their asset holdings in the face of rising asset prices. Note that this narrative requires that preferences shift *throughout the wealth distribution* (a slightly weaker version of the assumption that $\rho_{it} = \rho_t$ for all i). For example, to rationalize Figure 6, demand needs to have shifted at every wealth percentile.

An example of such preference shifts that may be relevant in our Norwegian empirical setting are sustained housing demand shifts. Recall from Section 4 that a substantial part of our main results in Figure 6 are due to housing capital gains. If, over time, Norwegians at all percentiles of the wealth distribution put increasingly higher value on living in the capital, Oslo, this will drive up Oslo house prices. And those individuals will not downgrade their houses in the face of rising house prices – after all, they were responsible for the house price boom in the first place.

6.2 Portfolio Adjustment Frictions

We next show that theories with portfolio adjustment frictions have the potential to explain our findings. Our model has two key ingredients. First, motivated by empirical household balance sheets, households are assumed to hold multiple distinct assets.⁵⁶ Second, households are subject to portfolio adjustment “frictions” broadly defined.

We assume that households hold two assets. A “consumption asset” b is used to smooth consumption. An “investment asset” k experiences capital gains. The portfolio adjustment friction is a stand-in for several alternative reasons why households only infrequently move funds from the investment asset to the consumption asset. Candidates include physical adjustment costs and capital gains taxes that generate a “lock-in-effect” (because they are levied on realized rather than accrued gains). The model builds on [Kaplan and Violante \(2014\)](#) and [Kaplan, Moll and Violante \(2018\)](#) who label the two assets “liquid” and “illiquid.”⁵⁷ The main differences from the models in those papers, are that in our framework: (i) households face an asset price that moves over time, and (ii) we take a broader view of portfolio adjustment frictions rather than just physical transaction costs.

⁵⁶This contrasts to canonical models of household saving behavior where different assets are perfect substitutes, so that these frameworks essentially collapse to one-asset models

⁵⁷These papers in turn build on a long literature studying models with multiple assets and adjustment costs generating inaction, going back to [Grossman and Laroque \(1990\)](#).

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 . Their expected lifetime utility is

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

Households can buy and sell investment assets k_t at a price p_t that grows exogenously according to (5) as in Section 2.1, but again without transitory capital gains: $\varepsilon_t \equiv 0$.⁵⁸ The investment asset pays dividends, with θ denoting the constant dividend yield.⁵⁹ We use d_t to denote a household's purchases of the investment asset (with $d_t < 0$ corresponding to sales). In order to buy 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

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

with the borrowing constraints $b_t \geq \underline{b}$ and $k_t \geq 0$, and where z_t denotes an individual's idiosyncratic labor productivity.⁶⁰ Following Kaplan, Moll and Violante (2018), we assume a kinked functional form for the adjustment cost function which generates an inaction region.⁶¹ The adjustment cost function is symmetric and has a kink at $d = 0$. The model's calibration is given in Appendix D Figure 14 shows the saving rates and portfolio shares across the wealth distribution in the two-asset model. The model is able to reproduce some key qualitative characteristics of our empirical findings in Figure 6: The recurrent saving rate first drops

⁵⁸That is, capital gains only have a persistent component μ and the transitory component is assumed to be zero. We make this assumption for computational reasons. Transitory, stochastic capital gains could be introduced at the expense of some additional computational complexity. 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.

⁵⁹In terms of the discussion in Section 6.1 about the sources of capital gains, this corresponds to Special Case I where capital gains are due to dividend growth.

⁶⁰The logarithm of z_t follows an Ornstein-Uhlenbeck process (the continuous-time analogue of an AR(1) process) with speed of mean reversion ν .

⁶¹More precisely, we assume $\chi(d, a) = \chi_0 |d| + \frac{\chi_1}{2} \left(\frac{d}{a}\right)^2 a$, with $\chi_0, \chi_1 > 0$. 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. 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 $a = 0$.

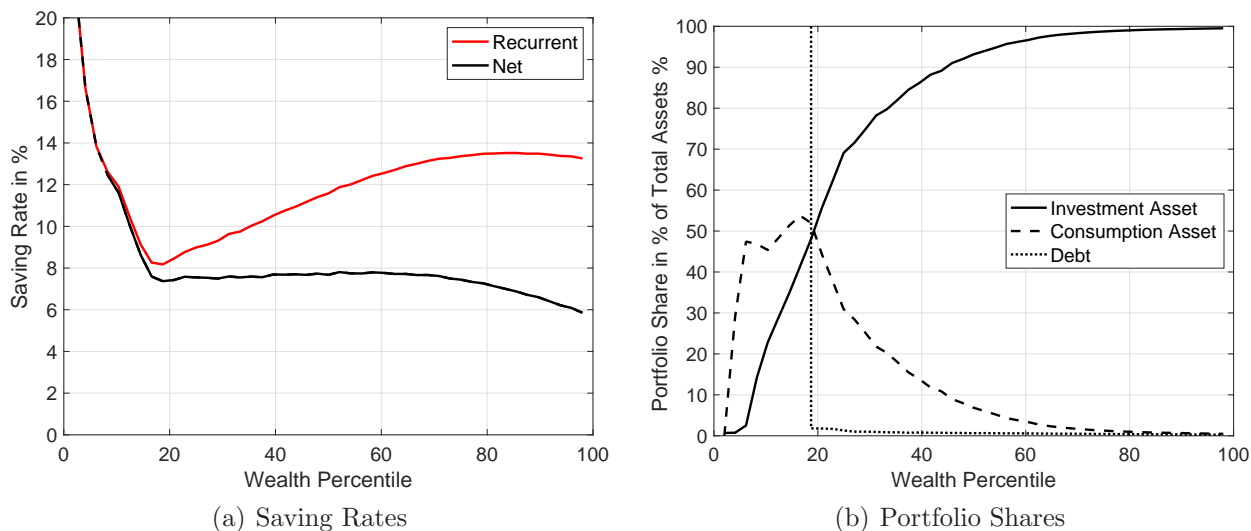


Figure 14: Saving behavior in two-asset model with portfolio adjustment friction.

and then increases, while net saving rates first decrease and then flatten out.⁶²

In contrast to the data, net saving rates decrease slightly with wealth for households above the 70th percentile of the wealth distribution. The reason is that as the portfolio share of investment assets becomes high, they become households’ primary source of income. Subsequently, the wealthiest households reduce net saving to consume a greater share of the returns from the investment asset. This outcome 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 our model’s prediction for the net saving rate can be further improved by incorporating indivisibilities and fixed (as opposed to kinked) transaction costs.

6.3 Other Candidate Explanations

Non-Homothetic Preferences. The canonical models’ prediction that the systematic part of the gross saving rate is flat, derives from the assumption of homothetic preferences as in (1). Since [Atkinson \(1971\)](#) it is well understood that saving can increase more strongly

⁶²The 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 ($r_-^b > \mu + \theta > r_+^b$) so that households with negative net worth primarily repay debt before they start accumulating investment assets. We see this in the portfolio share graph [Figure 6\(b\)](#) where households accumulate investment assets in a large scale after they have repaid most of their debt.

with wealth if preferences instead feature some form of non-homotheticity.⁶³ Intuitively, relative to Figure 2(a), non-homotheticities rotate saving rates counter-clockwise. Hence, such theories can readily generate *gross* saving rates that are increasing with wealth. However, non-homotheticities do not generally predict the striking flatness of the *net* saving rate in Figure 6. Instead, a flat net saving rate will only result as a knife-edge case, for one specific degree of non-homotheticity.⁶⁴

Misperceptions about the Stochastic Process for Asset Prices. As alluded to in Section 4.1, a candidate explanation for increasing recurrent saving rates is that households systematically misperceive all capital gains as purely transitory. In terms of the simple model of Section 2.1, they might think that persistent capital gains are zero, $\mu = 0$, even though in reality they are positive, $\mu > 0$. In order for this explanation to rationalize our findings, households in our sample 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.⁶⁵

Inattention and Behavioral Explanations. Another possibility is that households do not pay attention to asset price changes (Sims, 2003; Gabaix, 2017). Alternatively, it is possible that they perceive capital gains as a distinct form of income in the budgeting process (Shefrin and Thaler, 1988; Baker, Nagel and Wurgler, 2007; Di Maggio, Kermani and Majlesi, 2018) or simply perceive them as “paper gains.”

⁶³See e.g. Carroll (1998), De Nardi (2004), Straub (2018), Saez and Stantcheva (2018) and other papers cited in the review by Benhabib and Bisin (2018). Like Atkinson (1971), De Nardi (2004) assumes a warm-glow bequest motif with a different curvature than the period utility while alive, a particular form of non-homotheticity.

⁶⁴Consistent with this discussion, Straub (2018) shows that a calibrated version of his non-homothetic model (extended to include capital gains) can produce a U-shaped gross saving rate, just like in Figure 6. On the other hand, while his model’s implied net saving rate does not rise with wealth in a systematic fashion, it features considerably more variation across wealth percentiles than that in Figure 6.

⁶⁵For example, consider the following argument regarding capital gains on stocks. The equity premium is around six percent in many countries as well as in Norway. 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 them in the first place.

7 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 our definition of saving. Net saving rates are approximately constant across the wealth distribution, while gross saving rates increase sharply with wealth. We document that these distinct relationships are present because richer households accumulate wealth through “saving by holding,” meaning that they tend to hold on to assets experiencing persistent capital gains.

As we show, these patterns are inconsistent with canonical models of household saving behavior and we discuss theoretical alternatives. In a reduced form, these theoretical explanations have the property 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 propose two candidate explanations for our empirical findings: first, time-varying discount rates or, more generally, asset-demand shifts that affect households throughout the wealth distribution; second, portfolio adjustment frictions that make it costly to liquidate assets experiencing capital gains.

Turning to the broader implications of our findings, we believe that both the macroeconomics and wealth inequality literatures need to take changing asset prices more carefully into consideration. There are two reasons for this assessment. First, changing asset prices are a prevalent feature of the data. 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 canonical saving models, households appear to treat capital gains or losses differently from other forms of income, even if they are persistent. This poses a challenge for the standard models which typically carry the implication 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 treat all assets like money in a checking account.⁶⁶ Our findings suggest that

⁶⁶See [Benhabib and Bisin \(2018\)](#) for a review of the state-of-the-art in the theoretical wealth inequality literature. To the extent that the literature takes into account asset prices at all (see e.g. [Gomez, 2018](#), for a notable exception), price changes are soaked up into rates of return by appealing to the standard result that an asset’s rate of return equals its dividend yield plus capital gains. The same is true for the public finance literature studying optimal wealth and capital taxation (see e.g. [Saez and Stantcheva, 2018](#)).

these literatures could benefit from incorporating and refining lessons from the household finance literature (Campbell, 2006).

References

- AIYAGARI, S. R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *The Quarterly Journal of Economics*, 109(3), 659–84.
- ALSTADSÆTER, A., M. JACOB, W. KOPCZUK, AND K. TELLE (2016): “Accounting for business income in measuring top income shares: Integrated accrual approach using individual and firm data from Norway,” Discussion paper, National Bureau of Economic Research.
- ALSTADSÆTER, A., N. JOHANNESSEN, AND G. ZUCMAN (2017): “Tax evasion and inequality,” Discussion paper, National Bureau of Economic Research.
- ANDERSON, E. W., L. P. HANSEN, AND T. J. SARGENT (2012): “Small noise methods for risk-sensitive/robust economies,” *Journal of Economic Dynamics and Control*, 36(4), 468–500.
- ATKINSON, A. B. (1971): “Capital Taxes, the Redistribution of Wealth and Individual Savings,” *Review of Economic Studies*, 38(2), 209–227.
- AUCLERT, A. (2019): “Monetary Policy and the Redistribution Channel,” *American Economic Review*, 109(6), 2333–2367.
- BACH, L., L. CALVET, AND P. SODINI (2018): “From Saving Comes Having? Disentangling the Impact of Saving on Wealth Inequality,” Research paper, Swedish House of Finance.
- BACH, L., L. E. CALVET, AND P. SODINI (2016): “Rich Pickings? Risk, Return, and Skill in the Portfolios of the Wealthy,” CEPR Discussion Papers 11734, C.E.P.R. Discussion Papers.
- BAKER, M., S. NAGEL, AND J. WURGLER (2007): “The Effect of Dividends on Consumption,” *Brookings Papers on Economic Activity*, 38(1), 231–292.
- BENHABIB, J., AND A. BISIN (2018): “Skewed Wealth Distributions: Theory and Empirics,” *Journal of Economic Literature*, 56(4), 1261–91.
- BERGER, D., V. GUERRIERI, G. LORENZONI, AND J. VAVRA (2018): “House Prices and Consumer Spending,” *Review of Economic Studies*, 85(3), 1502–1542.
- BEWLEY, T. (1986): “Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers,” in *Contributions to Mathematical Economics in Honor of Gerard Debreu*, ed. by W. Hildenbrand, and A. Mas-Collel. North-Holland, Amsterdam.
- CAMPBELL, J. Y. (2003): “Consumption-based asset pricing,” in *Handbook of the Economics of Finance*, ed. by G. Constantinides, M. Harris, and R. M. Stulz, vol. 1 of *Handbook of the Economics of Finance*, chap. 13, pp. 803–887. Elsevier.
- (2006): “Household Finance,” *Journal of Finance*, 61(4), 1553–1604.
- CARROLL, C. D. (1998): “Why Do the Rich Save So Much?,” NBER Working Papers 6549, National Bureau of Economic Research, Inc.
- CHODOROW-REICH, G., P. T. NENOV, AND A. SIMSEK (2019): “Stock Market Wealth and the Real Economy: A Local Labor Market Approach,” NBER Working Papers 25959, National Bureau of Economic Research, Inc.
- CHRISTELIS, D., D. GEORGARAKOS, AND T. JAPPELLI (2015): “Wealth shocks, unemployment shocks and consumption in the wake of the Great Recession,” *Journal of Monetary Economics*, 72(C), 21–41.
- COCHRANE, J. (2005): *Asset Pricing*. Princeton University Press.
- CUI, W., AND V. STERK (2018): “Quantitative Easing,” CEPR Discussion Papers 13322, C.E.P.R. Discussion Papers.

- DE NARDI, M. (2004): “Wealth Inequality and Intergenerational Links,” *Review of Economic Studies*, 71(3), 743–768.
- DE NARDI, M., AND G. FELLA (2017): “Saving and Wealth Inequality,” *Review of Economic Dynamics*, 26, 280–300.
- DI MAGGIO, M., A. KERMANI, AND K. MAJLESI (2018): “Stock Market Returns and Consumption,” NBER Working Papers 24262, National Bureau of Economic Research, Inc.
- DYNAN, K. E., J. SKINNER, AND S. P. ZELDES (2004): “Do the Rich Save More?,” *Journal of Political Economy*, 112(2), 397–444.
- EIKA, L., M. MOGSTAD, AND O. VESTAD (2018): “What Can We Learn About Household Consumption from Information on Income and Wealth?,” Working paper, SSB.
- EISNER, R. (1988): “Extended Accounts for National Income and Product,” *Journal of Economic Literature*, 26(4), 1611–1684.
- EITRHEIM, Ø., AND S. K. ERLANDSEN (2004): “House Price Indices for Norway 1819-2003,” in *Monetary Statistics for Norway 1819-2003*, ed. by Ø. Eitrheim, J. T. Klovland, and J. F. Qvigstad, Occasional Paper. Norges Bank.
- EUROPEAN CENTRAL BANK (2016): “The Household Finance and Consumption Survey: results from the second wave,” Discussion paper, ECB Statistics Paper, No. 18.
- FAGERENG, A., L. GUISO, D. MALACRINO, AND L. PISTAFERRI (2016a): “Heterogeneity and Persistence in Returns to Wealth,” NBER Working Papers 22822, National Bureau of Economic Research, Inc.
- (2016b): “Heterogeneity in Returns to Wealth and the Measurement of Wealth Inequality,” *American Economic Review: Papers & Proceedings*, 106(5), 651–655.
- FAGERENG, A., M. B. HOLM, AND K. N. TORSTENSEN (2019): “Housing Wealth in Norway, 1993–2015,” .
- FEIVESON, L., AND J. SABELHAUS (2019): “Lifecycle Patterns of Saving and Wealth Accumulation,” Finance and Economics Discussion Series 2019-010, Board of Governors of the Federal Reserve System (US).
- FLEMING, W. H. (1971): “Stochastic Control for Small Noise Intensities,” *SIAM Journal on Control*, 9(3), 473–517.
- GABAIX, X. (2017): “Behavioral Inattention,” NBER Working Papers 24096, National Bureau of Economic Research, Inc.
- GABAIX, X., J.-M. LASRY, P.-L. LIONS, AND B. MOLL (2016): “The Dynamics of Inequality,” *Econometrica*, 84, 2071–2111.
- GLAESER, E. L. (2000): “Comments on “Real Estate and the Macroeconomy”,” *Brookings Papers on Economic Activity*, 2000(2), 146–150.
- GOMEZ, M. (2018): “Asset Prices and Wealth Inequality,” Working paper, Columbia University.
- GROSSMAN, S. J., AND G. LAROQUE (1990): “Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods,” *Econometrica*, 58(1), 25–51.
- GUREN, A. M., A. MCKAY, E. NAKAMURA, AND J. STEINSSON (2018): “Housing Wealth Effects: The Long View,” Working Paper 24729, National Bureau of Economic Research.
- HAIG, R. M. (1921): *The Concept of Income – Economic and Legal Aspects* pp. 1–21. Columbia University Press, New York.
- HICKS, J. (1939): *Value and Capital: An Inquiry Into Some Fundamental Principles of Economic Theory*. Clarendon Press.
- HUBMER, J., P. KRUSELL, AND T. SMITH (2018): “A Comprehensive Quantitative Theory of the U.S. Wealth Distribution,” Working paper.
- HUGGETT, M. (1993): “The risk-free rate in heterogeneous-agent incomplete-insurance economies,”

- Journal of Economic Dynamics and Control*, 17(5-6), 953–969.
- IMROHOROĞLU, A. (1989): “Cost of Business Cycles with Indivisibilities and Liquidity Constraints,” *Journal of Political Economy*, 97(6), 1364–83.
- JUDD, K. L. (1996): “Approximation, perturbation, and projection methods in economic analysis,” in *Handbook of Computational Economics*, ed. by H. M. Amman, D. A. Kendrick, and J. Rust, vol. 1 of *Handbook of Computational Economics*, chap. 12, pp. 509–585. Elsevier.
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): “Monetary Policy According to HANK,” *American Economic Review*, 108(3), 697–743.
- KAPLAN, G., AND G. L. VIOLANTE (2014): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82(4), 1199–1239.
- KASA, K., AND X. LEI (2017): “Risk, Uncertainty, and the Dynamics of Inequality,” Discussion Papers dp17-06, Department of Economics, Simon Fraser University.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2016): *Macroeconomics and Household Heterogeneity* vol. 2 of *Handbook of Macroeconomics*, chap. 0, pp. 843–921. Elsevier.
- KRUSELL, P., AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.
- KUHN, M., M. SCHULARICK, AND U. I. STEINS (2018): “Income and Wealth Inequality in America, 1949-2016,” Working Papers 9, Federal Reserve Bank of Minneapolis, Opportunity and Inclusive Growth Institute.
- MARTÍNEZ-TOLEDANO, C. (2019): “Housing Bubbles and Wealth Inequality,” Working paper, Paris School of Economics.
- MEADE, J. (1964): *Equity, Efficiency and the Ownership of Property*. Allen and Unwin.
- MERTON, R. C. (1969): “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case,” *The Review of Economics and Statistics*, 51(3), 247–57.
- MULLAINATHAN, S., AND J. SPIESS (2017): “Machine learning: an applied econometric approach,” *Journal of Economic Perspectives*, 31(2), 87–106.
- PAIELLA, M., AND L. PISTAFERRI (2017): “Decomposing the Wealth Effect on Consumption,” *The Review of Economics and Statistics*, 99(4), 710–721.
- POTERBA, J. M. (2000): “Stock Market Wealth and Consumption,” *Journal of Economic Perspectives*, 14(2), 99–118.
- ROBBINS, J. (2019): “Capital Gains and the Distribution of Income in the United States,” Working paper, Brown University.
- SAEZ, E., AND S. STANTCHEVA (2018): “A Simpler Theory of Optimal Capital Taxation,” *Journal of Public Economics*, 162(C), 120–142.
- SAEZ, E., AND G. ZUCMAN (2016): “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data,” *The Quarterly Journal of Economics*, 131(2), 519–578.
- SHEFRIN, H. M., AND R. H. THALER (1988): “The Behavioral Life-Cycle Hypothesis,” *Economic Inquiry*, 26(4), 609–643.
- SIMONS, H. C. (1938): *Personal Income Taxation: the Definition of Income as a Problem of Fiscal Policy*. University of Chicago Press Chicago, Ill.
- SIMS, C. A. (2003): “Implications of rational inattention,” *Journal of Monetary Economics*, 50(3), 665–690.
- STRAUB, L. (2018): “Consumption, Savings, and the Distribution of Permanent Income,” Working paper, Harvard.
- VON SCHANZ, G. (1896): “Der Einkommensbegriff und die Einkommensteuergesetze,” *FinanzArchiv / Public Finance Analysis*, 13(1), 1–87.