

NBER WORKING PAPER SERIES

SAVING BEHAVIOR ACROSS THE WEALTH DISTRIBUTION:
THE IMPORTANCE OF CAPITAL GAINS

Andreas Fagereng
Martin Blomhoff Holm
Benjamin Moll
Gisle Natvik

Working Paper 26588
<http://www.nber.org/papers/w26588>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2019

We thank Adrien Auclert, John Campbell, Karen Dynan, Giulio Fella, Matthieu Gomez, Adam Guren, Henrik Kleven, Claus Thustrup Kreiner, Camille Landais, Soren Leth-Petersen, Stefan Nagel, Monika Piazzesi, Luigi Pistaferri, Martin Schneider, Antoinette Schoar, Kjetil Storesletten, Ludwig Straub, Amir Sufi, Olav Syrstad, Martin Weale and seminar participants at various institutions for useful comments. This paper is part of a research project at Statistics Norway generously funded by the Norwegian Research Council (#287720). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Andreas Fagereng, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains
Andreas Fagereng, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik
NBER Working Paper No. 26588
December 2019
JEL No. D14,E21,E40

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 answer this empirical question. The relation between saving rates and wealth crucially depends on whether saving includes capital gains. Saving rates net of capital gains ("net saving rates") are approximately constant across the wealth distribution. 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"). These joint patterns for net and gross saving rates challenge canonical models of household wealth accumulation. They are instead consistent with theories in which time-varying discount rates or portfolio adjustment frictions keep households from realizing capital gains. Between 1995 and 2015 Norway's aggregate wealth-to-income ratio rose from approximately 4 to 7. "Saving by holding" accounts for up to 80 percent of this increase.

Andreas Fagereng
Research Department
Statistics Norway
Pb 8131
0033 Oslo, Norway
afagereng@gmail.com

Benjamin Moll
Department of Economics
Princeton University
192 Julis Romo Rabinowitz Building
Princeton, NJ 08544
and NBER
moll@princeton.edu

Martin Blomhoff Holm
Department of Economics
University of Oslo
Moltke Moes Vei 31
0851 Oslo, Norway
martin.b.holm@outlook.com

Gisle Natvik
Department of Economics
BI Norwegian Business School
Nydalsveien 37
0484 Oslo, Norway
gisle.j.natvik@bi.no

A data appendix is available at <http://www.nber.org/data-appendix/w26588>

1 Introduction

Do wealthier households save a larger share of their incomes than poorer ones? When this issue comes up in the popular debate, participants frequently assert that the answer is “yes.”¹ In contrast, economists’ canonical models of household wealth accumulation predict that saving rates are either independent of or *decreasing* with wealth.² Unfortunately, precise evidence is missing, leaving us with an open empirical question.

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 that correspond to two alternative ways of writing the budget constraint “consumption plus saving equals income” and differ by how capital gains are treated when writing this accounting identity. The first notion of saving, *active* or *net saving*, is the change in a household’s net worth from one year to the next *holding asset prices constant* – the difference between a household’s income (excluding capital gains) and its consumption. In contrast, the second notion, *gross saving*, is simply the change in a household’s net worth, including any revaluation effects due to changing asset prices.³

In our main empirical exercise we ask how the corresponding saving rates vary across the wealth distribution. Our findings are easiest to communicate visually. Figure 1 therefore

¹For example: “The rich save more than the poor, and the more they have, the more they’ll save. Money that’s being saved isn’t being spent [...] In turn, that means less work for everybody and a lethargic economy” (Ehrenfreud, 2015). Dynan, Skinner and Zeldes (2004) feature similar media quotes from the 1990s.

²As 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 pointed out by De Nardi and Fella (2017, see Figure 1).

³Some papers in the literature refer to our “gross saving” as “passive saving” (e.g. Dynan, Skinner and Zeldes, 2004). We prefer to use “gross” for two reasons. First, we found that the “passive” nomenclature can lead to confusion, for example, because active saving is then part of passive saving (the other part being unrealized capital gains). Second, the term “gross saving” avoids taking a stand on individuals’ intentions and information sets, e.g. the decision not to sell an asset when its price rises may well be an active decision.




Figure 1: Saving rates across the wealth distribution (preview).

previews our main findings by plotting the saving rates against percentiles of net worth. To the far left are households with negative net worth. Households with approximately zero wealth are located around the 15th percentile. And as we move to the right, households become progressively wealthier.

We see that among households with low or negative net worth, it does not matter much if capital gains are included or excluded from saving: net and gross saving rates lie almost on top of each other and decline with wealth. However, among the vast majority of households who have positive wealth, matters are very different. Here the net saving rate (excluding capital gains) is remarkably flat around seven percent, while the gross saving rate (including capital gains) increases sharply from around zero for households with zero net worth, to thirty-five percent for the top one percent of the wealth distribution.⁴ Returning to our opening question whether wealthier households have higher saving rates, we therefore obtain the following answer: No, they do not have higher saving rates in the traditional sense of net or active saving; but, yes, they still accumulate more wealth through capital gains.

The proximate explanation for the divergence of the two 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

⁴As we explain in more detail below, year-to-year fluctuations in asset prices also generate large year-to-year fluctuations in the gross saving rate. The line labeled “gross (systematic component)” therefore plots the systematic component of the gross saving rate (defined in more detail in the main text).

the net saving rate is 10% at all points of the wealth distribution 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.

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. But are there any obvious correlates that account for 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 find, however, that these factors are not the main drivers of our observed saving profiles: the relation between saving rates and wealth is qualitatively unchanged when conditioning on observables like age, income and individual saving histories.

Our results 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 7. By counterfactually imposing that households over this period treated capital gains in the same way as other forms of capital income, we find that saving by holding accounted for up to 80 percent of this increase in aggregate wealth.

Our second contribution is to view net and gross saving rates through the lens of macroeconomic theory and to show that our observations challenge canonical models of household wealth accumulation. 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 address prevalent features of the data such as housing. The simple benchmark features a household with homothetic utility that receives a constant stream of labor income and saves in an asset with an exogenously given price. When this asset price is constant, the model predicts that a household's saving rate

out of income should be exactly independent of its wealth. When the asset price varies stochastically over time, so that the model generates a meaningful distinction between net and gross saving, the optimal consumption and saving response to capital gains depends crucially on whether these are transitory or persistent. The logic follows from the permanent income hypothesis. Households experiencing transitory capital gains optimally save these so that such asset price changes translate one-for-one into changes in their net worth. 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 that depend on household wealth, gross saving rates purged of transitory capital gains should be independent of wealth. That is, 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) the 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 across 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 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 transaction costs or sustained housing demand shifts must be part of the story.⁵

The prominence of housing in our data leads to another, related question: to what extent do the patterns 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

⁵More 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 without housing. Therefore, while housing is clearly different from financial assets, the consumption aspect by itself cannot explain our empirical findings.

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 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, Figure 1 calls for 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 close to one hundred percent, precisely what we have termed “saving by holding.” 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 models of wealth accumulation (e.g. Aiyagari, 1994; Huggett, 1993) feature neither of these ingredients and instead assume that households have time-invariant discount rates and save in one asset only. We show that both ingredients can qualitatively explain our findings and, in particular, generate “saving by holding.”⁶ 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 various behavioral explanations.

As far as we know, we are the first to provide systematic evidence on individual saving rates out of income over the entire wealth distribution.⁷ Bach, Calvet and Sodini (2018) also examine saving behavior across the wealth distribution using administrative data but with a different focus. Complementary to this paper, they examine how the saving rate out of *wealth*, i.e. the saving-to-wealth ratio or wealth growth rate, varies across the wealth distribution whereas we focus 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

⁶There is strong precedent for both of these ingredients in the existing literature. First, a well-established argument holds that asset price fluctuations are primarily driven by time-varying discount rates (Campbell, 2003; Cochrane, 2005). Second, going back to Grossman and Laroque (1990), economists have pointed to indivisibilities and transaction costs as important features of some assets, chief among them housing.

⁷Arguable exceptions are Krueger, Mitman and Perri (2016) and Saez and Zucman (2016). However, neither paper provides evidence on *individual*-specific saving rates like we do. Krueger, Mitman and Perri present suggestive evidence on consumption rates out of income (i.e. one minus saving rates) computed as total consumption expenditures for a specific wealth quintile divided by total income in that wealth quintile (they work with quintiles rather than percentiles like we do due to the small sample size of their dataset, the U.S. Panel Study of Income Dynamics). Saez and Zucman provide evidence on “synthetic saving rates” that are computed by following percentile groups, rather than individuals, over time.

informative object to study.⁸ Given their focus on the saving-to-wealth ratio, they also do not study saving behavior at the bottom of the wealth distribution where this ratio is ill-defined because the denominator is zero or negative. In contrast, we study saving behavior across the entire distribution, including 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.⁹

Our findings are relevant for a nascent literature in macroeconomics and the study of inequality that emphasizes portfolio choice and asset price changes (Hubmer, Krusell and Smith, 2018; Gomez, 2018; Kuhn, Schularick and Steins, 2019; Feiveson and Sabelhaus, 2019; Martínez-Toledano, 2019). We hope that our empirical findings and their theoretical implications will be useful building blocks for future contributions to this literature.¹⁰ Finally, our paper is related to the literature on the consumption effects of asset price changes, in particular papers that estimate how capital gains affect consumption.¹¹ Broadly consistent with our findings, Baker, Nagel and Wurgler (2007) and Di Maggio, Kermani and Majlesi (2019) argue that marginal propensities to consume (MPCs) out of capital gains are significantly smaller than MPCs out of dividend income. More generally, our findings are consistent with a household finance literature that finds substantial inertia in households' financial decisions (e.g. Brunnermeier and Nagel, 2008; Calvet, Campbell and Sodini, 2009).

Section 1 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 2 describes our data and institutional setting. Section 3 presents our empirical results and documents that net saving rates are approximately flat while gross saving rates increase sharply with wealth. Section 4 shows that “saving by holding” can account for a large fraction of the rise in Norway’s aggregate wealth-to-income ratio. Section 5 theoretically interprets our empirical findings, and Section 6 concludes.

⁸In 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.

⁹Dynan, Skinner and Zeldes (2004), Alan, Atalay and Crossley (2015) and Straub (2018) document how consumption and saving behavior vary with “permanent income” defined as the permanent component in labor income. Permanent income is not directly observable and must be estimated, typically by means of an instrumental variable strategy. We instead focus on wealth which is readily observable.

¹⁰For example, Straub (2018) and Cui and Sterk (2018) use our evidence to discipline their macro models.

¹¹As 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. Poterba (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.

2 Theoretical Benchmarks

While our main contribution is empirical, we begin by considering a series of theoretical benchmarks. These will guide our empirical analysis in Section 3. 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. Theory also motivates 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 to address important features of the data such as 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 with pencil and paper. Households are infinitely lived and have homothetic utility

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

where c_t is consumption. They receive a constant labor income w and can save in an asset a_t paying a constant interest rate r . Their budget constraint is $\dot{a}_t = w + ra_t - c_t$ and they face a natural borrowing constraint $a_t \geq -w/r$. Utility maximization yields a simple analytic solution for the optimal saving policy function $\dot{a} = s(a)$ (see Appendix A.1 for the proof):

$$s(a) = \frac{r-\rho}{\gamma} \left(a + \frac{w}{r} \right). \quad (2)$$

That is, households save (and consume) a constant fraction of their effective wealth $a + w/r$. It follows that the saving rate out of total income $w + ra$ is constant too:

$$\frac{s}{y} = \frac{s}{w + ra} = \frac{r-\rho}{\gamma r}. \quad (3)$$

Hence, the saving rate out of income is independent of wealth.¹² We show below that many other benchmark models inherit this property, at least approximately.

¹²Households 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 income. Note that the saving rate is also constant over discrete time intervals (not just infinitesimal ones): using (2), we have $a_{t+1} - a_t = \int_0^1 \dot{a}_{t+s} ds = \frac{r-\rho}{\gamma} \int_0^1 (a_{t+s} + \frac{w}{r}) ds$ and hence saving $a_{t+1} - a_t$ is a constant fraction $\frac{r-\rho}{\gamma r}$ of income $\int_0^1 (w + ra_{t+s}) ds$.

Changing asset prices. We next extend the 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.” Stocks are the only saving vehicle available to households. They 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 $\dot{p}_t/p_t = \mu + \varepsilon_t$, 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 ε_t as the transitory component. A simple special case is when $\varepsilon_t dt = \sigma dW_t$, where W_t is a standard Brownian motion, so that p_t follows a random walk.

Defining wealth by $a_t := p_t k_t$, the budget constraint (4) becomes

$$c_t + \dot{a}_t = w + r_t a_t \quad \text{where} \quad r_t := \theta + \mu + \varepsilon_t. \quad (5)$$

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, r_t is the sum of dividends θ and capital gains $\dot{p}_t/p_t = \mu + \varepsilon_t$.

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, Appendix A.1 derives a useful approximation to the saving policy function under the assumptions that return innovations ε_t are Brownian with variance σ^2 and labor income is small, $w \approx 0$:

$$\dot{a} \approx \bar{s} \left(a + \frac{w}{\bar{r}} \right) + \varepsilon a, \quad \text{where} \quad \bar{r} := \theta + \mu \quad (6)$$

and $\bar{s} := \frac{\bar{r} - \rho}{\gamma} + (\gamma - 1) \frac{\sigma^2}{2}$. 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 (6) yields a key result: the optimal 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 all of it (see the additive term εa in (6)). 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.

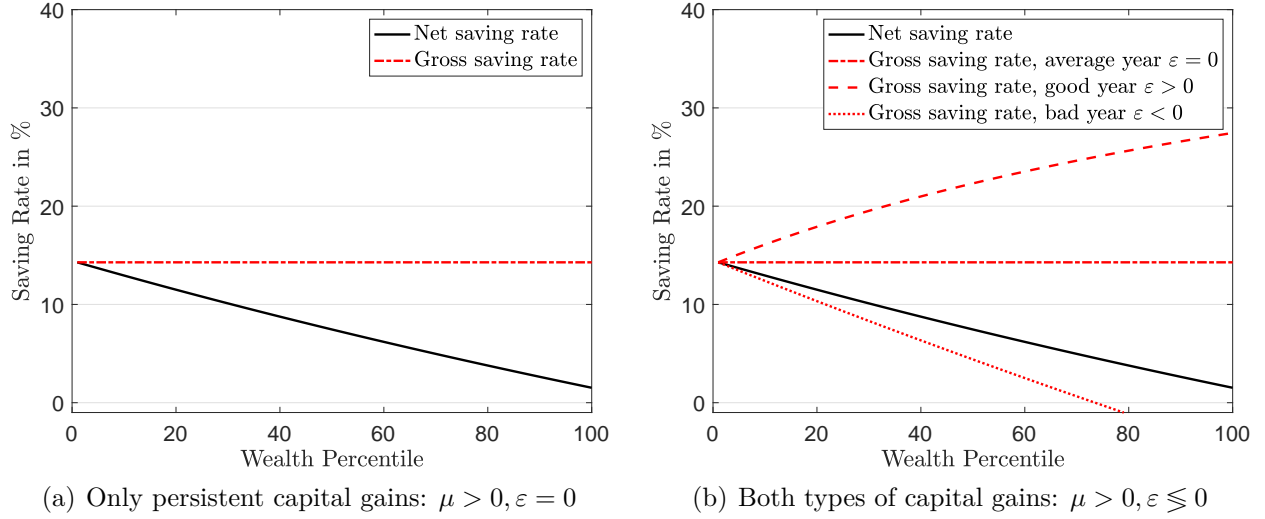


Figure 2: Saving rates across wealth distribution in simple benchmark model.

Figure 2 displays saving rates across the wealth distribution from the model with changing asset prices. It plots two measures of the saving rate, corresponding to the two ways of writing the budget constraint. The “net saving rate” follows from (4) and is the saving rate net of capital gains $p_t \dot{k}_t / (w + \theta p_t k_t)$. The “gross saving rate” follows from (5) and is the saving rate out of total income including capital gains $\dot{a}_t / (w + r_t a_t)$. Importantly, the wealthier a household is, the larger are capital gains relative to income and the more do the two measures differ. Panel (a) considers the case of only persistent and no transitory capital gains, $\mu > 0$ and $\varepsilon = 0$. In this case (6) implies that the gross saving rate should be independent of wealth. Given that rich households experience larger capital gains, the net saving rate should be *decreasing* with wealth.¹³

Panel (b) of Figure 2 considers the case with both transitory and persistent capital gains. When $\varepsilon = 0$, meaning that the stock market had an “average year,” the figure is identical to panel (a). In contrast, when the stock market does well, $\varepsilon > 0$, gross saving is increasing with wealth. When the stock market does badly, $\varepsilon < 0$, the gross saving rate is instead decreasing with wealth.¹⁴ The reason is that households save all transitory capital gains, so these impact wealth one-for-one.

¹³When $\varepsilon = 0$ the approximation in (6) 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 $w > 0$.

¹⁴Equation (6) 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 $\varepsilon < 0$.

In all cases it can be seen from (6) that gross saving rates *purged from transitory capital gains* $(\dot{a}_t - \varepsilon_t a_t)/(w + (r_t - \varepsilon_t)a_t)$ are independent of wealth. That is, the simple consumption-saving model in this section predicts that the *systematic* component of the gross saving rate is independent of wealth. In contrast, the net saving rate is independent of stock market performance and predicted to be decreasing with wealth.¹⁵

Key concepts: net, gross and recurrent saving. Motivated by our theoretical results we now define key concepts that we will use in our empirical analysis. These are “net” and “gross” saving, which we have discussed already, and “recurrent” saving which aims to isolate the systematic (rather than cyclical) component of gross saving.

All three saving rates result from alternative ways of writing the budget constraint “consumption plus saving equals income.” They differ by how capital gains are treated when writing this accounting identity. As already mentioned, the definition of “net” and “gross” saving can be gauged from the two budget constraints (4) and (5) respectively:

$$c + \underbrace{\dot{p}k}_{\text{net saving}} = \underbrace{w + \theta pk}_{\text{disposable income}} \quad , \quad (7)$$

$$c + \underbrace{\dot{p}k + \dot{p}k}_{\text{gross saving}} = \underbrace{w + (\theta + \dot{p}/p)pk}_{\text{Haig-Simons income}} \quad . \quad (8)$$

The difference between these two accounting identities is simply that the latter adds capital gains $\dot{p}k$ on both sides. Intuitively, since consumption in the two equations is the same, a difference in the saving definition necessarily implies a difference in the income definition. Formulation (7) features disposable income, whereas formulation (8) features “Haig-Simons” income which includes unrealized capital gains (von Schanz, 1896; Haig, 1921; Simons, 1938).¹⁶ Finally, we define the “net saving rate” as the ratio of net saving to disposable income, and the “gross saving rate” as the ratio of gross saving to Haig-Simons income.

Canonical macro models of wealth accumulation are models of gross saving in (8), and it therefore seems natural to compare gross saving rates generated by such models to those in the data. However, as seen in Figure 2(b), the simple model above implied that gross saving rates should vary strongly with asset market performance. Fortunately, another theoretical

¹⁵From (6) consumption c is independent of transitory capital gains ε . Then from (4) so is net saving $\dot{p}k$.

¹⁶This income definition forms the basis for the argument in the public finance literature that capital gains should be taxed on accrual rather than realization. Related, Eisner (1988) and Robbins (2019) argue for including capital gains in the income definition used in the National Income and Product Accounts.

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

$$c + \underbrace{(\dot{k}/k + \mu)pk}_{\text{recurrent saving}} = \underbrace{w + (\theta + \mu)pk}_{\text{recurrent income}}, \quad (9)$$

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 (8) is that recurrent saving includes only the persistent component of capital gains, μ , and not the transitory component, ε . Equivalently, recurrent saving is gross saving purged of transitory capital gains, that is $\dot{a} - \varepsilon a$. The definition of recurrent income follows from the requirement that consumption plus saving equals income. The recurrent saving rate then is recurrent saving divided by 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 is 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

Housing is a prominent asset for most households in our data and housing capital gains are an important component 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 examine how these two features shape model-based predictions for saving rates. Economists often emphasize the consumption aspect of housing, and intuition suggests that this aspect could make households hold on to their residential wealth in the face of rising house prices.¹⁷ The main result of this section is that while housing is different from financial

¹⁷As Glaeser (2000) puts it: “A house is both an asset and a necessary outlay. [...] 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.”

assets, the consumption aspect *by itself* does not qualitatively change the relation between saving rates and wealth. Instead other aspects of housing, like indivisibilities and transaction costs, are needed to generate such changes.

We build on Berger et al. (2018) and construct a parsimonious model in which housing is a durable consumption good but counterfactually assume that it is fully divisible and freely adjustable. Households receive a constant labor income w and invest in two assets: housing and bonds. House prices change deterministically over time. Households maximize

$$\int_0^\infty e^{-\rho t} \frac{C_t^{1-\gamma}}{1-\gamma} dt \quad \text{with} \quad C = \left(\alpha^{1/\eta} c^{\frac{\eta-1}{\eta}} + (1-\alpha)^{1/\eta} h^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \quad (10)$$

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. The budget constraint is

$$c + \dot{b} + p\dot{h} = w + rb, \quad (11)$$

where b are bonds and p is the house price. We can rewrite this budget constraint in terms of total wealth, $a := b + ph$ as $c + Rh + \dot{a} = w + ra$ where $R := rp - \dot{p}$ is the “implicit rental cost” or “user cost” of housing. Our goal is to understand how households treat housing capital gains in this environment. Conveniently, the model can be solved in closed form.

Proposition 1 *Optimal total consumption C_t and saving \dot{a}_t are given by*

$$C_t = \frac{m_t}{P_t} \left(\frac{w}{r} + a_t \right), \quad \dot{a}_t = w + ra_t - m_t \left(\frac{w}{r} + a_t \right), \quad (12)$$

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

The presence of the user cost of housing R_t in the price index P_t captures precisely the intuition that housing is both an asset and a consumption good (just like in footnote 17).

Analogous to Section 1.1, we distinguish between transitory and persistent housing capital gains. For the transitory case, we show in Appendix A.2 that an unexpected house price increase translates one-for-one into higher net worth, analogous to the result in our simple benchmark model without housing.¹⁸ The more relevant case for us is that of persistent housing capital gains, meaning an anticipated change in the growth rate of house prices \dot{p}/p .

¹⁸One can also show that total consumption C is unchanged in the face of higher house prices and therefore so is welfare (as stated in footnote 17). To be clear, even though *total* consumption is unchanged, households will sell off housing to increase their *non-durable* consumption as long as $\eta > 0$, i.e. there is a positive housing wealth effect on non-durable consumption (Berger et al., 2018).

Corollary 1 *The saving response to persistent housing capital gains depends on the intertemporal elasticity of substitution (IES) $1/\gamma$. When the IES equals one, $m_t = \rho$ and optimal saving in (12) becomes $\dot{a} = \frac{r-\rho}{\gamma} \left(a + \frac{w}{r}\right)$, meaning that the model with housing collapses to the one-asset model in Section 1.1. Housing capital gains then affect saving behavior like any other asset returns. When the IES is close to one, this is approximately true.*

In general, a persistent house price appreciation affects optimal saving not only by raising the return to housing wealth, but also by affecting the price index of the total consumption basket P_t . When the IES $1/\gamma = 1$, this effect disappears. Intuitively, changes in the price index have both intertemporal substitution and income effects and when the IES equals one, these exactly offset.¹⁹ Corollary 1 implies that the consumption aspect of housing will *not*, by itself, result in households keeping their portfolios unchanged in the face of rising house prices. Why is this the case? The answer is that if housing were divisible and freely adjustable, households would engage in “intertemporal substitution of housing.” To understand what this means, consider a scenario where households suddenly learn that house prices will grow by two percent per year for the next twenty years. Figure A13 in Appendix A.2 plots the optimal response to such a persistent housing capital gain, assuming that the IES equals one and that the elasticity of substitution between housing and consumption η is zero (perfect complements). On impact, households increase non-durable consumption and purchase more housing in anticipation of the future price increases (which lowers the user cost R). Afterwards, they gradually sell off housing. 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.

Our results show that the consumption aspect of housing does not by itself change how households are predicted to respond to capital gains. Transitory housing capital gains should be saved, while persistent gains should be partly consumed. Still, there are other features of housing that can qualitatively affect saving behavior. In particular, housing is indivisible and entails substantial transaction costs. We explore these aspects in Section 5.

2.3 Common Extensions

The framework in Section 1.1 was deliberately simplistic so as to introduce our main saving concepts and highlight two properties: (i) the recurrent saving rate, i.e. the systematic

¹⁹Higher future house prices implies (i) that future total consumption is more expensive (income effect) and (ii) that total consumption is cheaper today relative to tomorrow (substitution effect). As in standard consumption theory, $1/\gamma$ governs the relative strength of the substitution and income effects.

component of the gross saving rate, is approximately flat with wealth, and (ii) the net saving rate is decreasing with wealth. We here show that these properties carry over to many richer models of household wealth accumulation from the macroeconomics literature. As in much of this literature, we will here ignore changing asset prices so that there is no difference between net, gross and recurrent saving. With changing asset prices, the statements we make here would apply to recurrent saving rates. The net saving rate would then follow residually, depending on what persistent capital gains happen to be (see 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 3 plots the resultant saving rate out of total income against wealth in this environment. The left panel displays the “saving rate policy function” $s(a, w)/(w + ra)$ for three different levels of labor income w . 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.²⁰ The right panel displays saving rates without conditioning on income. It

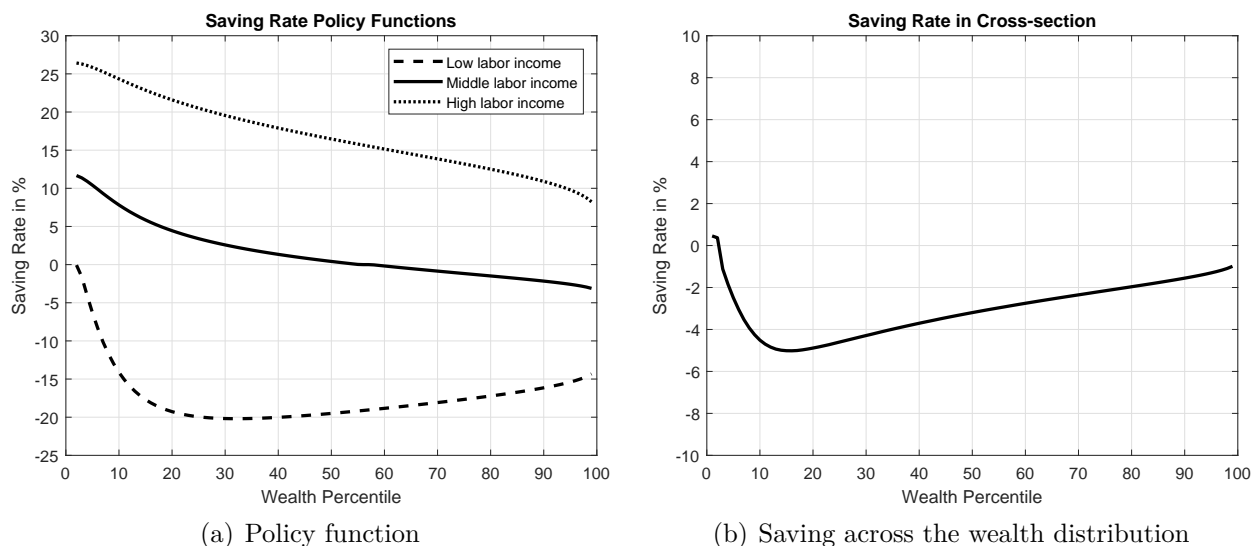


Figure 3: Saving rates with income risk and borrowing constraints.

simply plots the cross-sectional relationship between saving rates and wealth in the model’s

²⁰It is straightforward to show that the flat-saving-rate result from Section 1.1 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 1.1. The steep decline close to the borrowing constraint reflects two familiar effects. First, precautionary saving of high-income, low-wealth households. Second, low-income households rapidly decumulating wealth and then hitting the constraint (i.e. their saving rate is zero at the constraint but strongly negative above the constraint).

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. On the other hand, saving rates increase with labor income (see panel (a)), and labor income and wealth are positively correlated in the stationary distribution. Summarizing, while saving rates conditional on income should decline slightly with wealth, they may be increasing with wealth when not conditioning on income. But this increasing relationship is driven by a cross-sectional correlation between labor income and wealth.

A natural second extension is to introduce the realistic assumption that labor income grows over the life cycle, creating different saving incentives at different ages.²¹ Consequently, life-cycle considerations introduce a cross-sectional correlation between saving rates and wealth because both are correlated with age and income. Conditional on age and income, saving rates are slightly decreasing with wealth, as noted by [De Nardi and Fella \(2017\)](#). Hence, we will conduct empirical exercises in which we condition on age and labor income.

The third model extension we explore is heterogeneity in individual-specific propensities to save, operationalized as heterogeneity in discount rates. Such heterogeneity is 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](#)), and the cross-sectional wealth-saving correlation should be explained by historical saving behavior at the individual level. We will therefore present results conditioning on past saving behavior.

3 Data and Definitions

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. Below we describe these data in more detail and explain the saving rate measures we construct.

²¹In a standard life-cycle model, households save little (or borrow, if possible) when they are young and have low income and wealth. As both age and income rise, they begin to save and accumulate wealth.

Like other Northern European countries, Norway has 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 subsidized and provided at low or no cost. On the income side, this welfare state is financed both through taxes (the ratio of tax income to GDP was 38% in 2016) and through returns from a sovereign wealth fund, the Norwegian Pension Fund Global.²² The main role of this fund is to smooth public spending and support future governments' ability to meet public pension entitlements.

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, household and firm level. Our unit of observation is the household.²³ 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 Appendix B.

For our purposes, the Norwegian data have a number of advantages. First, we observe wealth together with income at the household level for the entire population. Neither income nor wealth is top- or bottom-coded. The only sources of attrition are mortality and emigration out of Norway. Second, our data cover a long time period. 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.

Administrative wealth and income records. Each year, Norwegians must provide complete information about their incomes and wealth holdings to the tax authorities. Wealth

²²See <https://www.wider.unu.edu/project/government-revenue-dataset> for how the tax-to-GDP ratio has evolved over time. The Norwegian Pension Fund Global is part of Norway's fiscal mechanism for smoothing the use of national oil revenues. The mechanism postulates that the flow government income from oil activity is invested abroad, while an estimated normal return (3% currently, 4% previously) on the existing fund may be spent each year. Currently, the fund is worth about three times Norwegian GDP.

²³Arguably, 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.

taxes are levied on the household as a whole, whereas income taxes are levied individually.²⁴ 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.

For the majority of our analyses we use data for the period 2004 to 2015. 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.

Household balance sheets. The tax records contain 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 funds). The sum of these is labeled “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 (from the stockholder registry)²⁵ or indirectly via private businesses (next paragraph), all at market values.

“Private business” wealth refers to ownership of firms that are not publicly traded. Unlike publicly traded shares, no 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

²⁴Taxable 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 revert these discounted values back to market values when computing household wealth.

²⁵In Appendix B.5 we verify the consistency between our micro-level data on stock holdings and aggregate household stock holdings at the Norwegian stock exchange.

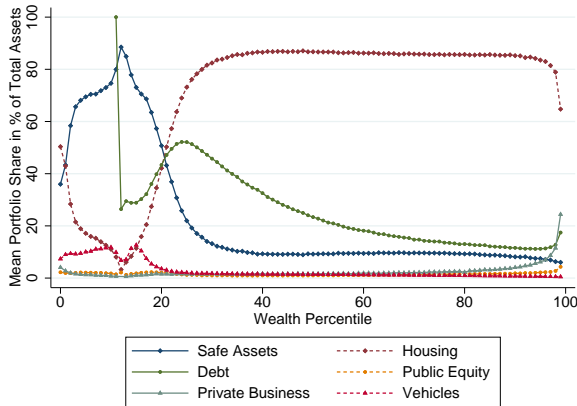
audited by an approved auditing entity. As noted, when computing private business wealth, we have extracted listed stocks and debt and placed them, respectively, in the public equity and debt categories on the owners’ balance sheets (see Appendix B.5).

“Housing wealth” includes the value of a household’s principal residence, secondary homes and recreational estates. Traditionally, housing values in the tax records have been linked to original purchase prices often deviating substantially from current market values. To improve on this, we use transaction level data from the Norwegian Mapping Authority combined with the land registration office and the population census to regress price per square meter house on house characteristics. The predicted values from this procedure are then used as housing wealth over our sample period.²⁶ The corresponding mortgages, together with student loans, consumer debt and personal debt, and debt in private businesses owned by the individual constitute the household “debt.”

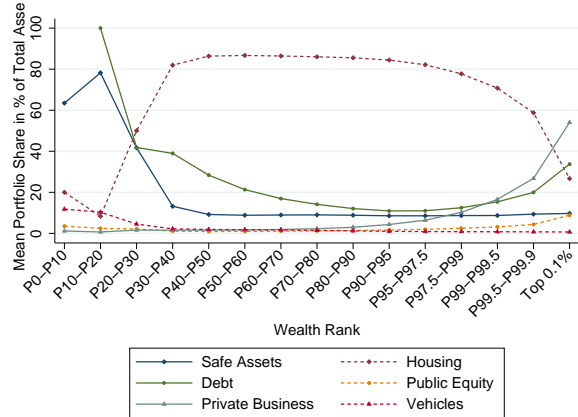
In addition to the balance sheet components discussed here, households also have pension claims. The Norwegian government provides a relatively generous pension scheme. Some workers additionally have private retirement accounts held by their employers, so as to top up the public pension plan which is capped. In extensions of our baseline analysis, we impute each household’s public pension entitlement and include pensions in measures of income, saving, and wealth in Section 3.5 and Appendix B.6.

Descriptive statistics. Figure 4 plots the average holdings of different asset classes and debt relative to total assets within each wealth percentile. Here, as in the rest of the paper “wealth” is defined as the value of total assets net of debt, “net worth.” Panel (a) in Figure 4 displays portfolio shares for the entire population, panel (b) focuses on the wealthiest households. To the left on the axes are households with negative net worth. By definition their debt-to-asset ratio exceeds 100 percent. We also note that the least wealthy households hold a high ratio of housing wealth relative 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 15th wealth percentile. These households hold almost no assets and no debt. Thereafter, from the 15th to approximately the 25th wealth percentile, the housing portfolio share grows rapidly. The portfolio share of safe assets (predominantly bank

²⁶The price per square meter is estimated as a function of house characteristics such as the number of rooms and bathrooms, location, time periods, and their interactions using machine learning techniques. We utilize an ensemble method combining a random forest algorithm, a regression tree, and LASSO as in Mullainathan and Spiess (2017). Fagereng, Holm and Torstensen (2019) evaluate the predictive power of the method in a holdout sample of actual transactions, showing that 85% of the predicted housing values are within $\pm 10\%$ of the transaction value.



(a) Mean portfolio shares by wealth percentile



(b) Mean portfolio shares by wealth rank

Figure 4: Asset class shares in household portfolios.

Notes: The figure displays the mean portfolio share in percent of total assets across the wealth distribution, by percentile group. 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.

deposits) is high across the entire wealth distribution. In contrast, public equity is a relatively small component of Norwegian households' wealth. Finally, we see that private business ownership becomes an important asset class in the top part of the wealth distribution, especially for the top 0.5% and 0.1%.

Table 1 provides the main descriptive statistics of our sample. The first panel in the table displays demographics. The second panel describes the wealth components. The third panel provides statistics on income and savings. Household disposable income is defined as the sum of labor income, business income, capital income, transfers, and housing service flows, minus taxes.²⁷

²⁷To compute the housing service flow, 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. For example, consider the budget constraint in terms of gross variables (8). Using the notation of Section 1.2 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 $Rh = 2.04\% \times ph$ as explained momentarily) the analogue of (8) is

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

We follow Eika, Mogstad and Vestad (2018) and distribute 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 rent-to-value ratio which is decreasing from about 2.04% in 2005 to 1.78% in 2015.

Table 1: Descriptive statistics.

	Mean	SD	P10	Median	P90	Part. Rate
Age	49.91	17.65	27	48	75	
Male	0.49	0.50	0	0	1	
Years of education	12.15	3.20	8	12	16	
Less than high school	0.30	0.46	0	0	1	
High school	0.39	0.49	0	0	1	
College education	0.31	0.46	0	0	1	
Safe assets	43,427	113,105	854	14,744	107,041	0.99
Housing	422,914	394,045	0	345,499	924,637	0.80
Debt	142,485	1,387,280	0	63,154	298,423	0.85
Public equity	9,991	595,108	0	0	13,461	0.38
Private business	53,050	1,639,922	0	0	4,830	0.14
Vehicles	6,881	103,783	0	1,417	19,233	0.57
Net wealth	393,778	827,460	-7,624	286,566	914,446	
Disposable income	58,828	204,020	29,525	51,411	84,987	
Haig-Simons income	75,276	508,691	4,690	57,106	184,346	
Net saving	1,814	632,071	-30,164	2,021	29,653	
Gross saving	18,262	466,676	-71,564	7,366	126,272	

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 are in USD, 2011 prices - using the 2011-average exchange rate between USD and NOK in 2011 (NOK/USD = 5.77), throughout the paper.

3.2 Implementing Different Saving Rate Concepts

As discussed in Section 1.1, we distinguish between three notions of saving: *net*, *gross*, and *recurrent saving*, and we 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.2. We now discuss how we operationalize these concepts in the data and our empirical strategy for separating gross saving into net saving and capital gains.

Separating gross saving into net saving and capital gains. In the data we directly observe the year-to-year change in each household’s net worth, that is its gross saving. To compute our measures of saving rates we need to separate gross saving into net saving and capital gains. Our approach differs by asset class and we therefore discuss each asset class separately (for details see Appendix B.3).

The largest balance sheet component of most households is housing. To separate gross saving in housing into net saving and capital gains, we utilize the same housing transactions data that we use to compute housing wealth. 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 month 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 (see equation (A40) for the details of this decomposition).

Next we consider public equity, consisting of directly held stocks and stock funds. For directly held stocks, we use the Norwegian shareholder registry, which contains year-end information on holdings of individual stocks at the security level. Since all stocks are publicly traded, we observe stock price changes. We can thus compute a measure of capital gains for each individual stock. Net saving for a particular stock becomes the change in the value of holdings of that stock minus capital gains, adjusted for net purchases of that stock.²⁸ 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 explained above, for private businesses our data provides 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.5), 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 4), this likely results in an underestimate of capital gains and gross saving among the very wealthiest.

Lastly, consider bonds and money market funds. Decomposed holding data from the Norwegian Fund and Asset Management Association reveal that households' holdings in

²⁸Stock holdings are observed at the end of the year. We assume net purchases in any given stock happen in the middle of the year when computing net saving. Alternative assumptions (such as all purchases happen at the first or last day of the year, or throughout the year) yield similar results.

bond and money market funds are concentrated in government bonds with maturity below one year and medium-term bonds (two to four years).²⁹ These short-term bonds are unlikely to experience substantial price changes and bond holdings are a very small fraction of households' asset holdings (an aggregate portfolio share of 0.4% over the sample period). We therefore make the simplifying assumption that bonds do not experience capital gains.

Other assets listed in Section 2.1, as well as liabilities, do not experience asset price changes. Hence, capital gains are zero, and net saving equals gross saving for those assets.³⁰

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. To operationalize this concept, we need to split individual capital gains on different assets into persistent and transitory components. That is, to operationalize equation (9), 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. Second, to minimize the 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 .^{31,32} An important limitation of our procedure is the relatively short time span of our data (between 21 and 66 years depending on the asset class). This means that our estimates of μ come with large standard errors and this should be kept in mind when interpreting our estimates.³³

²⁹Among short- and medium-term bonds, the weight on short-term bonds is approximately 70 percent.

³⁰Vehicles depreciate over time and we use list price changes to infer the depreciation rates (see Section 2.1). 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.

³¹An alternative is to construct an estimate of μ that is not only asset-specific but also depends on each household's wealth percentile. We present such an exercise in Figure A15(b).

³²Note that our procedure effectively assumes that expected capital gains are constant within an asset class, implying that total expected returns vary with dividend yields. This consequence is consistent with the empirical finding that expected returns can be predicted by valuation ratios such as the price-dividend ratio (the inverse of the dividend yield), as argued by for instance Campbell (2018, ch.5). An alternative to our procedure would be to assume that expected total returns are constant.

³³To illustrate the issue, assume that asset prices are a random walk so that capital gains are white noise. The standard error of the mean of such a white noise series is σ/\sqrt{T} , where σ is the standard deviation and T the sample length. Thus if the standard deviation of yearly capital gains is say $\sigma = 15\%$, the standard error of μ over $T = 25$ years is $15\%/5 = 3\%$ which is large relative to the μ 's we estimate below (e.g. 3.25%).

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).³⁴ 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 Oslo, Bergen, Trondheim and Kristiansand (four of Norway’s main cities), in addition to a nationwide index for the remainder of the country.³⁵ For housing, it is particularly important to 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 (8) and (9) 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 because housing capital gains in our sample period have been higher than its historical average, 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 5.

Summarizing, we use the following numbers for persistent capital gains: 3.25% for directly held stocks, 3.18% for stock funds, and the area specific gains 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 4.30% for housing.³⁶

³⁴The 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.

³⁵The values for persistence capital gains are 2.22% for Oslo, 2.78% for Bergen, 3.16% for Trondheim and 2.82% for Kristiansand, and 2.26% for the rest of the country. These series are originally due to Eitrheim and Erlandsen (2004) and are updated annually by Norges Bank and available at <https://www.norges-bank.no/en/Statistics/Historical-monetary-statistics/House-price-indices/>.

³⁶For 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 is time-varying and we set to 2.04% in 2005 as discussed above. Hence the return to housing (using the non-metropolitan nationwide persistent capital gains) in 2005 was $2.26\% + 2.04\% = 4.30\%$.

In particular, we recover the typical finding that stocks yield a higher average return than housing.

4 Results


This section provides our main results. We start by presenting descriptive plots of the net, gross, 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 5(a) plots our three saving rates against percentiles of net worth. It depicts the main finding of our paper, and the net and recurrent saving rates here are the same as we saw in our introductory Figure 1. For every year we have computed the median saving rate within each year-specific wealth percentile. The three lines plot the average of these within-year medians over our sample period. An alternative approach to time-aggregation is discussed in Section 3.5 which also shows saving rates separately for each year.

To better understand the different patterns for the three saving rate concepts in panel (a) and in particular the role of capital gains, panel (b) shows how these capital gains vary across the wealth distribution. We plot both observed capital gains (i.e. the difference between gross and net saving) as well as their persistent component. In addition, we plot households' asset holdings relative to income. All three curves use the same time-aggregation approach as in panel (a).

In panel (a), we start with the *net* saving rate, for which our benchmark theories in Section 1 predicted a declining relationship with wealth (see Figure 2). To the left, we see that among the least wealthy households, the net saving rate does indeed follow this prediction. Then, after reaching a trough around the 15th wealth percentile, it rises up to the 20th wealth percentile. Note that the trough lies approximately where net worth is zero (see Figure 4). In panel (a), the curve flattens out around 6 – 8% from the 20th wealth percentile and up.



(a) Saving rates by wealth percentiles.

(b) Average capital gains and asset-to-income ratio.

Figure 5: Saving rates across the wealth distribution.

Notes: (a) displays the median saving rates within wealth percentile and year, averaged across all years (2005-2015). (b) displays households' capital gains as a fraction of their assets, the persistent component of these capital gains, and median total assets as a share of disposable income. All variables are computed as the median within-wealth percentile and year, averaged across all years (2005-2015).

Consider next the gross saving rate. As capital gains have been positive on average over 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 net and gross saving rates follow each other down to the trough around zero net worth. Thereafter the gross saving rate diverges 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, the recurrent saving rate, lies between the net and gross rates because recurrent capital gains are positive, but lower than the average of observed capital gains over our sample period. However, its shape is approximately the same as that of the gross saving rate. Hence, the fact that the gross saving rate is increasing with wealth is not simply driven by abnormally high capital gains in our sample period.

In the beginning of this paper we asked whether wealthier households have higher saving rates. Figure 5 suggests a nuanced answer to this question. On the one hand, net saving rates are almost completely flat across the wealth distribution – which is to say that, no, wealthier households do *not* have higher saving rates in the traditional sense. On the other hand, the gross and recurrent saving rates increase sharply with wealth – that is, even though the net

saving rate is flat, wealthier households do nevertheless accumulate more wealth through capital gains.

The proximate explanation for the diverging behavior of these notions of saving rates is simple: wealthy households hold assets that experience persistent capital gains. And instead of selling off these assets to consume, they hold on to them in the face of rising asset prices. The mechanism has two logical components. First, wealthier households hold more assets. Hence, their total asset holdings automatically appreciate by a larger dollar amount, unless they systematically invest in assets with lower capital gains than poorer households do. In our data we can observe this component of the mechanism: panel (b) of Figure 5 visualizes how wealthier households hold more assets relative to their income, and capital gains rates do not systematically decline with wealth. Second, households tend to save a higher fraction out of these capital gains than out of other income sources. This second part of the mechanism poses a challenge to the canonical models summarized in Section 1, and is the reason why gross saving behaves differently in the data than in these models.

A more general takeaway from Figure 5 is that asset price changes and how people respond to these are key to understanding saving behavior across the wealth distribution. In Section 1.1, 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 5 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, the pattern for capital gains in Figure 5(b) challenges this explanation. Capital gains have a systematic positive component, computed using a long time series, and so the systematic component of gross saving rates is increasing with wealth.

4.2 The Role of Age, Earnings, and Education

Common extensions of our simplest consumption-saving model (Section 1.3) 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.³⁷ In this section, we therefore explore to what extent the patterns in Figure 5 can be accounted for by these factors. We focus on the net saving rate as well as the recurrent saving rate, which theory has sharper predictions for than the gross saving rate (see Section 1.1).

We utilize three approaches: (1) we compute median saving rates by *within-cohort* wealth percentiles thereby eliminating the cross-sectional correlation between age and wealth, (2) we control parametrically for covariates in a quantile regression, and (3) we plot saving by wealth within age, earnings, and education groups.

Figure 6(a) displays the results from approach (1). It breaks the mechanical link between wealth and age: because we rank households according to the wealth percentile for their respective cohort, every percentile now consists of an equal number of households from all age groups. As before, we computed the median saving rate within each percentile for each year, and then averaged these medians over time. The resulting pattern is very similar to that in Figure 5. The one notable difference is that the net saving rate now increases slightly with wealth, constituting an even bigger challenge for canonical models.


In approach (2), the quantile regression, we use the following specification

$$\frac{s_{it}}{y_{it}} = \sum_{p=1}^{100} \phi_p D_{it,p} + f(x_{it}) + \tau_t + \epsilon_{it}, \quad (13)$$

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 coefficient for percentile p , x_{it} is a vector of control variables, τ_t are time fixed effects, and ϵ_{it} is an error term. In what follows, we specify $f(\cdot)$ as a fourth-order polynomial in age and household labor income with fixed effects for years of education. We have considered a multitude of alternatives here and the results are insensitive to the exact specification. 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 are displayed graphically in Figure 6(b).

The figure shows that conditional on age, earnings and education, the relation between wealth and saving rates is qualitatively similar to its unconditional counterpart in Figure 5. Consistent with the discussion in Section 1.3, the recurrent saving rate is slightly flatter after controlling for the three variables. This indicates that some of the correlation between saving

³⁷An alternative mechanism with similar implications is that personal experience with macroeconomic events shapes expectations. Malmendier and Nagel (2011) find evidence in this direction, documenting that individuals' experienced returns shape their portfolio choice and optimism. If wealth correlates with age, this could lead to a systematic relationship between expected capital gains and wealth.



(a) Saving rates by within-cohort wealth percentiles. (b) Controlling for age, earnings and education.

Figure 6: Saving rates adjusted for age, cohort, earnings and education.

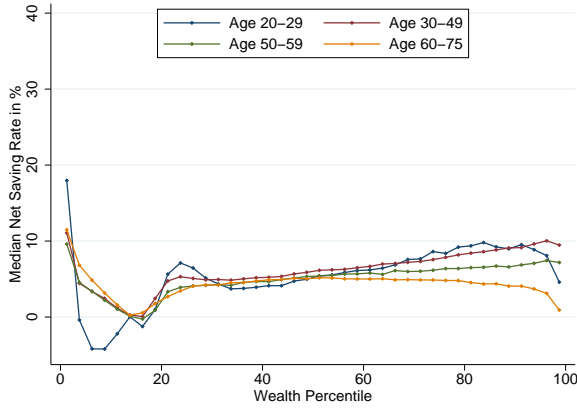
Notes: (a) displays the median saving rates by within-cohort wealth percentile (with cohorts defined as birth-years) and by year, averaged across all years (2005-2015). The net saving rate is defined in equation (7) as net saving divided by disposable income. The recurrent saving rate is defined in equation (9) as recurrent saving divided by recurrent income. The gross saving rate is defined in equation (8) as gross saving divided by gross income. (b) displays the coefficients on wealth percentiles (ϕ_p) from estimating equation (13) as a quantile regression for net and recurrent saving rates, respectively.

rates and wealth is due to age, earnings, and education. Still, the main takeaway is that even conditional on these observables, the saving graph maintains its main characteristic: a “swoosh-shaped” recurrent saving rate, and a first decreasing and then flat net saving rate.


With approach (3) we produce our main graph within groups defined by age, earnings and education. For age we divide households into four strata (20-29, 30-49, 50-59, and 60-75 years), for earnings we stratify them by earnings decile, 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 5. Figure 7 shows the median net and recurrent saving rates within groups (the results for education are in Appendix figure A14). Again we observe decreasing and then flat net saving rates, and “swooshes” for recurrent saving rates. Conditioning on age, earnings and education affects the *level* of saving rates, but only modestly affects how these *covary* with wealth.

4.3 Do Our Results Arise Because Saving Causes Wealth?


Our discussion of benchmark models in Section 1.3 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 house-




(a) Net saving rate by age group



(b) Recurrent saving rate by age group



(c) Net saving rate by earnings group



(d) Recurrent saving rate by earnings group

Figure 7: Saving rates across the wealth distribution by age and earnings

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).

holds have higher saving rates and, over time, therefore become wealthier than impatient ones. Hence, a natural question is: do the gross and recurrent saving rates tend to slope upward simply because households who save a lot become rich?

We exploit the panel dimension of our data to factor out the role of such “type dependence” (Gabaix et al., 2016). 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.

To conserve on space, we relegate the resultant graph to the appendix, Figure A14. In short, our main results remain the same also within groups defined by saving history. Hence, our results do not seem to be due to high saving rates causing high wealth.

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 4). This fact naturally begs the question whether our findings in Figure 5 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. The issue is whether theoretical explanations of our findings should focus exclusively on housing or whether they should apply to other assets as well. We perform two exercises to shed light on this issue. We first look at the financial saving rates among households who hold a significant portion of their portfolio in public equity. Thereafter we zoom in on the wealthiest part of the population.

Financial saving rates among households who hold public equity. Answering how households treat capital gains from a specific asset as part of a larger portfolio is complicated. Furthermore, the vast majority of Norwegian households hold relatively few assets other than housing that experience asset price changes (see Figure 4). For the purpose of this section, we therefore restrict our sample to households who 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 focus on financial wealth defined as net worth excluding housing and liabilities. 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.³⁸ Before proceeding to the results of this exercise, we already note one caveat: in contrast to directly held stocks, for stock fund holdings we cannot use the shareholder registry to identify capital gains for individual stocks at the security level. Instead, we use an aggregate index for this asset class (see the discussion in Section 2.2). Hence, for stock funds, which our sample restrictions now lend relatively more importance too than in our baseline analysis above, the split of gross saving into net saving and capital gains becomes subject to some measurement error.

³⁸To be clear: The net financial saving rate is the change in financial wealth minus capital gains, divided by disposable income excluding housing income. Recurrent and gross financial saving rates are defined similarly.

Figure 8(a) plots the “financial versions” of our three saving rates defined in Section 3, against financial wealth for this restricted sample. The graphs display within-percentile medians, as in our baseline Figure 5. Because the sample now is relatively small, the figure uses two-and-a-half percentile bins rather than the one-percentile bins in our main results.

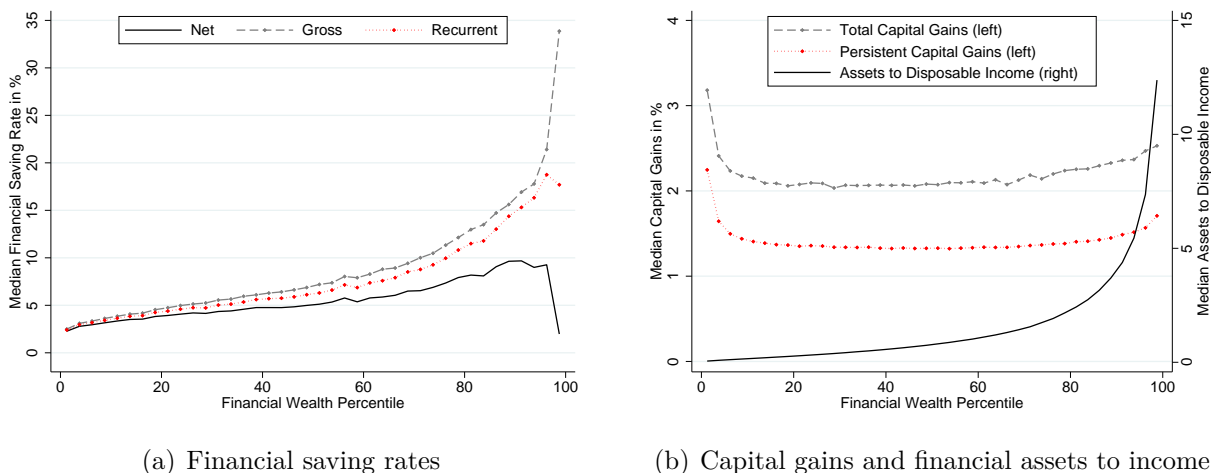


Figure 8: Relation between financial saving rates and financial wealth.

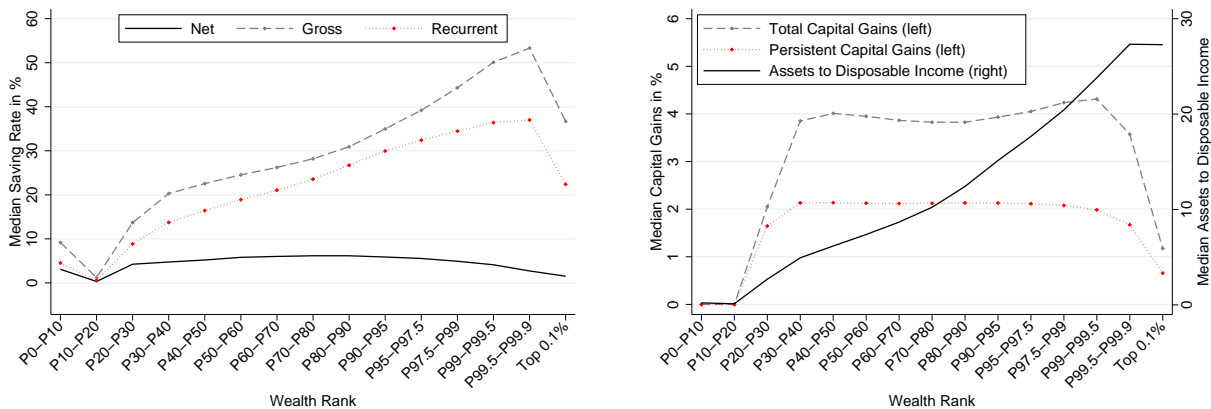
Notes: These figures display the median financial saving rates, capital gains in percent of financial assets, and financial assets-to-income ratio for the sample of households who hold at least USD 1,000 in public equity and at least 25% of their financial wealth in public equity (10.4% of benchmark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

In Figure 8(a), we see the familiar pattern that both gross and recurrent saving rates increase steeply with wealth, ranging from about two percent at the bottom to between twenty and thirty percent at the top. Also akin to Figure 5, the net saving rate is considerably flatter than the gross and recurrent saving rates. In contrast to Figure 5, the net saving rate is weakly upward-sloping with financial wealth. However, the concerns about measuring net saving in stock funds discussed in the preceding paragraph call for caution in emphasizing the exact quantitative relationship between net financial saving and financial wealth.³⁹

In sum, the patterns we observed for net, gross and recurrent saving rates do not seem limited to accumulation of housing wealth. As far as we can tell, under the limitation

³⁹In particular, a non-zero 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.

that Norwegian households hold relatively few assets with capital gains besides housing, the patterns are qualitatively similar for wealth accumulation in other assets too.



(a) Saving rates across the wealth distribution

(b) Capital gains and asset-to-income ratio

Figure 9: Saving behavior in the right tail of the wealth distribution.

Notes: The figures display the median saving rates, capital gains in percent of assets, and assets to income ratio. All variables are computed as the median within wealth percentile and year, averaged across all years (2005-2015).


Zooming in on the right tail of the wealth distribution. We next zoom in on the right tail of the wealth distribution in Figure 9. This subset of the population is particularly interesting because they hold a disproportionately high share of aggregate wealth, and because their asset portfolio is considerably less tilted toward housing.

Figure 9(a) 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 continues to increase 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 9(b), which plots capital gains relative to income. This ratio drops markedly as we step into the top 0.1 percent group. Hence, the saving pattern observed in Figure 9(a) 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.⁴⁰


⁴⁰The low measured capital gains for the top 0.1% are likely in large part due to the fact that tax values of private businesses are related to book values rather than market values and therefore measured capital gains on private businesses are zero, see the discussion in Sections 2.1 and 2.2. 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.

4.5 Additional Exercises


We have explored a range of robustness checks and extensions of our main analysis. In this section, we present those that we find particularly informative.



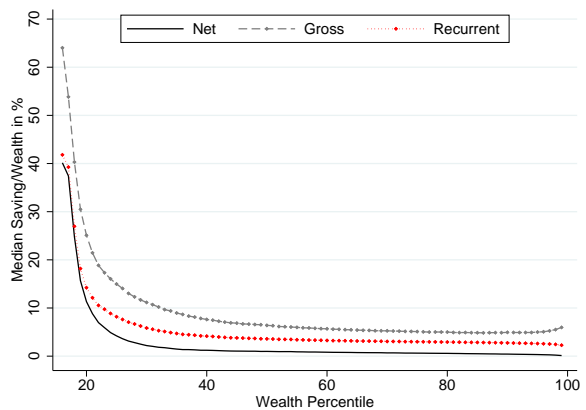
(a) Net saving rates across years



(b) Gross saving rates across years



(c) Including public pensions



(d) Saving rates as fraction of wealth

Figure 10: Additional exercises.

Notes: (a) and (b) display the median net and gross saving rates for each year in our sample. (c) displays the median saving rates when we include pension saving and income from the public pension system. Appendix B.6 describes the public pension system and how we compute pension wealth, income, and saving. (d) displays the median saving rate as a fraction of wealth.

Saving rates by year. Instead of averaging across time, our first exercise in Figures 10(a) and (b) simply plots saving against wealth for each year separately. Each lighter gray line is constructed for a specific year. The thicker black lines are our baseline counterparts. Notice that the vertical axes are different in the two panels.

Together, the two plots give a stark illustration of saving by holding. In all years, including the two where asset prices fell for large parts of the wealth distribution (2008 and 2009), 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 with large capital gains, and actually *decreases* with wealth in bad years in which large parts of the distribution experienced capital losses.

These large year-to-year fluctuations in gross saving rates are consistent with the ones predicted by our benchmark model (Figure 2). But what remains inconsistent with our benchmark model, is the flatness of net saving rates and the recurring tendency for gross saving rates to increase with wealth.

Public pensions. Our second 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.6.

By including the pension saving in the definition of income Figure 10(c), 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, while the gross and recurrent saving rates increase with wealth.

Saving rates as a fraction of wealth. Because our objective is to test the saving behaviour implied by economic theory, our objects of interest are saving relative to income. However, alternative 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, with the objective of addressing how saving behavior affects the dynamics of wealth inequality. To ease comparison with their work, we present saving as a fraction of wealth across the wealth distribution in Figure 10(d).⁴¹ The figure reveals a similar pattern as Figure 4 in Bach et al.⁴² Appendix Figure A15(h) also shows imputed consumption (defined as disposable income minus net saving –

⁴¹We 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. The rationale is that the ratio of saving to wealth is ill-defined when wealth is zero or negative.

⁴²Note 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.

see equation (7)) 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, part of the downward sloping nature of both graphs is somewhat mechanical: even low-wealth households 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.⁴³

Saving rates by income. Appendix Figure A15(g) additionally plots our three saving rates measures across the distribution of *income*, as opposed to *wealth* like in our main Figure 5. The main takeaway is that not just gross saving rates but also net saving rates are strongly increasing with income. There are a number of theoretical reasons why this is the qualitative finding one should expect, for example the permanent income hypothesis suggests that households should save high transitory income realizations.

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.

Over our sample period, the Norwegian aggregate wealth-to-income ratio increased markedly from approximately 4 in 1995 to more than 7 in 2015 (see the solid line in Figure 11).⁴⁴ We here conduct a counterfactual exercise which asks: what would the evolution of the wealth-to-income ratio have been if households behaved as in our benchmark models in Section 1,

⁴³For instance consider the saving and consumption policy functions in our benchmark model, namely (2) 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 decrease with wealth and blow up as $a \rightarrow 0$, just like in panels (e) and (f).

⁴⁴These 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. The levels in our series are higher, mostly because we use our own higher-quality measure of housing values (see Section 2.1). For the exercise in this section, we extend our dataset back to 1995. The main limitation in the data prior to 2005 is that we do not observe individual stock ownership at the security level. To extend the sample period for this exercise we therefore approximate capital gains on publicly-traded stocks prior to 2005 using the OBXP stock price index. For other assets, we continue to compute capital gains with the methodology of Section 2.2 also prior to 2005.

i.e. if they did not save by holding? We start each household with their observed wealth in 1995. We then construct its counterfactual wealth level for all subsequent years by assuming a counterfactual saving behavior, namely that each household’s saving rate out of persistent capital gains equals its observed net saving rate. In each year, we adjust capital income and capital gains to correspond with the counterfactual wealth level for each household.

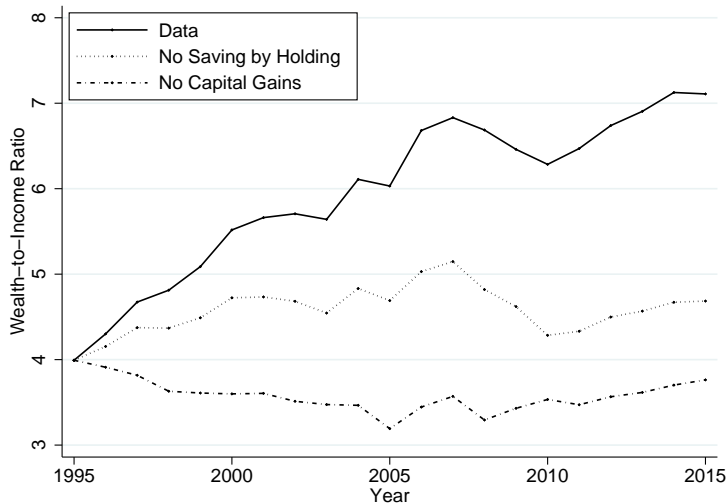


Figure 11: Actual and counterfactual wealth-to-income ratios.

Notes: This figure displays the evolution of the ratio of net wealth to disposable income in the data and in the counterfactual scenario. We adjust each individual’s capital gains and capital income such that they correspond to that individual’s counterfactual wealth, assuming fixed portfolio shares.

Figure 11 plots the counterfactual scenario along with the actual Norwegian wealth-to-income ratio. While this ratio increased from 4 to 7 in the data, it increased from 4 to only 4.7 in the “no saving by holding” counterfactual, meaning that saving by holding accounts for around 80 percent of the observed increase.⁴⁵ The reduced wealth accumulation is due to households in the counterfactual consuming considerably more out of their persistent capital gains. For comparison, the figure also plots a second scenario in which we simply set all capital gains to zero. The “no saving by holding” counterfactual lies in between this scenario and the data. Although crude, by for instance ignoring general equilibrium feedback from saving behavior to asset prices, our counterfactual suggests that the empirical saving behavior we have documented is crucial for understanding aggregate wealth accumulation.⁴⁶

⁴⁵In addition we have also computed an alternative counterfactual in which the recurrent saving rate is independent of wealth but at the mean of the observed recurrent saving rate (rather than being equal to the net saving rate). Even though the average recurrent saving rate is unchanged by construction, this counterfactual nevertheless accounts for around 60 percent of the increase in the wealth-to-income ratio.

⁴⁶In our counterfactual households sell off a large fraction of assets that experience persistent capital gains. One would expect this to lead to downward pressure on the prices of these assets, i.e. lower capital gains.

6 Theoretical Interpretation

If canonical models cannot rationalize our empirical findings, what theories can? In a reduced form, Figure 5 calls for a theory where households have a saving rate out of capital gains close to one hundred percent, what we have termed “saving by holding.”⁴⁷ 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; second, a theory featuring multiple assets and portfolio adjustment frictions or, more broadly, sources of inertia. 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.

6.1 Time-Varying Discount Rates or Asset Demand Shifts

A typical finding in the asset pricing literature is that a large fraction of high-frequency asset price fluctuations is driven by time-varying discount rates (see e.g. [Campbell, 2003](#); [Cochrane, 2005](#)). We here argue that there may also be analogous lower-frequency discount rate changes and that these have the potential to help explain our findings. Individuals have preferences

$$\int_0^\infty e^{-\int_0^t \rho_{is} ds} u(c_{it}) dt, \quad u(c) = \frac{c^{1-\gamma}}{1-\gamma} \quad (14)$$

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 demand shifts.⁴⁸

Apart from more general discounting, there are two additional differences from the model with changing asset prices in Section 1.1. 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

⁴⁷That such a reduced-form saving rule is a useful approximation to individual saving behavior was already conjectured by [Meade \(1964\)](#). In a discussion of wealth dynamics and inequality, he remarks: “The influence of capital gains could be even more marked [... if] 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” ([Meade, 1964](#), endnote 14).

⁴⁸The preferences in (14) 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 $\rho_{it} := \rho - \dot{\alpha}_{it}/\alpha_{it}$.

asset price as an outcome:⁴⁹

$$p_t = \int_t^\infty e^{-\int_t^s r_\tau d\tau} \Theta_s ds. \quad (15)$$

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*

$$c_{it} = \chi_{it} \left(p_t k_{it} + \int_t^\infty e^{-\int_t^s r_\tau d\tau} w ds \right). \quad (16)$$

How persistent capital gains affect consumption depends on the source of these capital gains. From (15), 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 5 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 (15), the asset price is given by the Gordon formula $p_t = \Theta_t / (r - \mu)$ and the price grows at $\dot{p}_t / p_t = \mu$. This case therefore corresponds exactly to the assumptions in Section 1.1: defining $\theta := \Theta_t / p_t$, the budget constraint is equivalent to (4). 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 1.1. Individuals should therefore consume out of persistent capital gains $\dot{p}_t / p_t = \mu$. Because the setup is equivalent to that in Section 1.1, the gross saving rate is flat and the net saving rate is decreasing with wealth. Therefore it cannot rationalize our findings in Figure 5. Note that canonical macro models of household wealth accumulation (e.g. Aiyagari, 1994; Huggett, 1993) implicitly assume that either there are no capital gains or that all capital gains are due to dividend growth.

⁴⁹In contrast, in Section 1.1, 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 (15) 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 (15). Alternatively, $\{r_t\}_{t \geq 0}$ could be pinned down from preferences in general equilibrium (see the discussion below).

Special case II: capital gains due to time-varying returns. Next consider the opposite case where dividends are constant and the asset price (15) rises only due to a declining required return. 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 declines in tandem, $\rho_{it} = r_t$ for all t . Then this individual’s optimal consumption, net and gross saving are*

$$c_{it} = w + \Theta k_{it}, \quad p_t \dot{k}_{it} = 0, \quad \dot{a}_{it} = \frac{\dot{p}_t}{p_t} a_{it}. \quad (17)$$

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 (15) is $p_t = \Theta \int_t^\infty e^{-\int_t^s r_\tau d\tau} ds$; second, with $\rho_{it} = r_t$ all t , $\chi_{it} := \left[\int_t^\infty e^{-\int_t^s r_\tau d\tau} ds \right]^{-1}$. Plugging into (16) yields $c_{it} = w + \Theta k_{it}$ in (17). The other results follow immediately. ■

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 (17).⁵⁰

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 . But why would asset returns and discount rates be equal even though both are moving over time? One answer is that this is the outcome in a 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 an asset price is rising is that everyone demands more of the

⁵⁰To 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 (17) with $\chi_{it} := \left[\int_t^\infty e^{-\int_t^s r_\tau d\tau} ds \right]^{-1}$, independently of the time path of ρ_{it} . The income effect of capital gains is thus captured by the “annuity value of financial wealth” $\left[\int_t^\infty e^{-\int_t^s r_\tau d\tau} ds \right]^{-1} p_t k_{it}$. 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. (Auclert (2019) provides a somewhat more precise explanation for the absence of an income effect in the special case 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” is zero.) The additional assumption $\rho_{it} = r_t$ eliminates intertemporal substitution effects of return changes. This is easy to see from individuals’ Euler equations $\dot{c}_{it}/c_{it} = (r_t - \rho_{it})/\gamma$.

asset, then individuals will not sell off their asset holdings in the face of a rising price. Note that this narrative requires that preferences shift *throughout the wealth distribution*. In particular, to rationalize Figure 5, demand needs to have shifted at every wealth percentile.

6.2 Portfolio Adjustment Frictions and Inertia

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.⁵¹ Second, households are subject to portfolio adjustment “frictions” broadly defined. These frictions inhibit withdrawals to finance consumption and give rise to inertia in response to asset price changes. Related types of inertia are an important idea in the household finance literature (e.g. Brunnermeier and Nagel, 2008; Calvet, Campbell and Sodini, 2009).

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 move funds from the investment asset to the consumption asset only infrequently. 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 Grossman and Laroque (1990) and, more recently, Kaplan and Violante (2014) and Kaplan, Moll and Violante (2018) who label the two assets “liquid” and “illiquid.” Our framework differs from theirs in two dimensions: (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. If the investment asset is interpreted as housing, the model is a version of what Guren et al. (2018) call the “new canonical model of housing” except that we do not model long-term mortgages.

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 $\mathbb{E}_0 \int_0^\infty e^{-(\rho+\phi)t} u(c_t) dt$. The investment asset k pays a constant dividend yield θ and its price p_t grows exogenously according to $\dot{p}_t/p_t = \mu$, i.e. as in Section 1.1 but without transitory capital gains. 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

⁵¹In contrast, canonical models of household wealth accumulation typically collapse to one-asset models.

interest rate $r_-^b = r^b + \kappa$, where κ is an exogenous wedge between borrowing and lending rates. We denote by $r^b(b)$ the full interest rate schedule. A household's assets evolve according to

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

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, the logarithm of which follows an Ornstein-Uhlenbeck process (the continuous-time analogue of an AR(1) process). Following [Kaplan, Moll and Violante \(2018\)](#), we assume that the adjustment cost function is symmetric and has a kink at $d = 0$ which generates an inaction region.⁵² We describe the model's calibration in Appendix D.

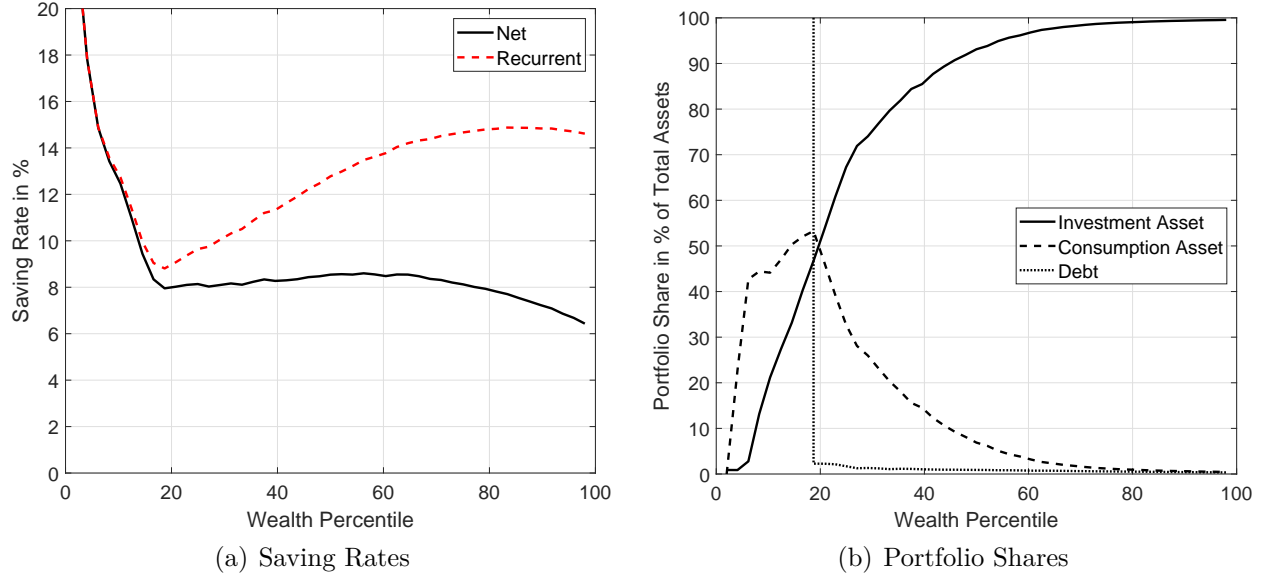


Figure 12: Saving behavior in two-asset model with portfolio adjustment friction.

Figure 12 shows the saving rates and portfolio shares across the wealth distribution in our model with portfolio adjustment frictions. The model is able to reproduce some key qualitative characteristics of our empirical findings in Figure 5: The recurrent saving rate first drops and then increases, while the net saving rate first decreases and then flattens out.⁵³ In contrast to the data, net saving rates decrease slightly with wealth for households

⁵²We assume $\chi(d, a) = \chi_0 |d| + \frac{\chi_1}{2} \left(\frac{d}{a}\right)^2 a$, with $\chi_0, \chi_1 > 0$. Because this function is infinite at $a = 0$, in computations we replace the term a with $\max\{a, \underline{a}\}$, where $\underline{a} > 0$ is a small value guaranteeing a finite cost.

⁵³Two parametric assumptions are responsible for generating the swoosh-shaped recurrent saving rate: a positive interest rate wedge, $\kappa > 0$; and that the return on the investment asset is high, but lower than the borrowing rate $r_-^b > \mu + \theta > r^b$. The second assumption ensures that households with negative net worth primarily repay debt before accumulating investment assets. See the portfolio share graph Figure 5(b).

above the 70th percentile of the wealth distribution, meaning that the wealthiest households reduce net saving to consume some of their investment assets. This is due to our assumption that, while households face a kinked adjustment cost function, the investment asset is still 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 with wealth if preferences instead feature some form of non-homotheticity.⁵⁴ Intuitively, relative to [Figure 2\(a\)](#), non-homothetic preferences rotate saving rates counter-clockwise. Hence, such theories can readily generate *gross* saving rates that are increasing with wealth. However, they do not generally predict the striking flatness of the *net* saving rate in [Figure 5](#). Instead, a flat net saving rate will only result for a knife-edge degree of non-homotheticity.⁵⁵

Returns to wealth are increasing with wealth. [Fagereng et al. \(2016, 2019\)](#) and [Bach, Calvet and Sodini \(2016\)](#) document that wealthier households obtain higher returns to their wealth whereas we document that they have higher gross 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 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. Finally, through the lens of many canonical models, returns that are increasing with wealth should rotate saving rates counter-clockwise just like non-homothetic preferences, and would only predict a flat net saving rate in a knife-edge case.

⁵⁴See 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, a particular form of non-homotheticity.

⁵⁵Consistent 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 5](#). On the other hand, while the net saving rate implied by his model does not rise with wealth in a systematic fashion, it features considerably more variation across wealth percentiles than that in [Figure 5](#).

Misperceptions about the stochastic process for asset prices. As alluded to in Section 3.1, 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 1.1, 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.

Inattention and behavioral explanations. Another possibility is that households do not pay attention to asset price changes (Sims, 2003; Gabaix, 2019). Alternatively, it is possible that they perceive capital gains as a distinct form of income in the budgeting process, e.g. due to “mental accounting” (Shefrin and Thaler, 1988; Baker, Nagel and Wurgler, 2007; Di Maggio, Kermani and Majlesi, 2019) or simply perceive them as “paper gains.”

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 saving includes capital gains. Perhaps surprisingly, net or active saving rates are approximately constant across the wealth distribution, meaning that wealthier households do *not* have higher saving rates than poorer ones in the traditional sense. On the other hand, gross saving rates increase sharply with wealth so that wealthier households nevertheless accumulate more wealth through capital gains. These distinct relationships are present because richer households “save by holding,” meaning that they tend to hold on to assets experiencing persistent capital gains.

These patterns are inconsistent with canonical models of household wealth accumulation and we therefore 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 asset-demand shifts that affect households throughout the wealth distribution; second, portfolio adjustment frictions or

other sources of inertia 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 into consideration more carefully. 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 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.

Fortunately, a nascent literature in macroeconomics and the study of inequality is starting to emphasize portfolio choice and asset price changes (Hubmer, Krusell and Smith, 2018; Gomez, 2018; Kuhn, Schularick and Steins, 2019; Martínez-Toledano, 2019; Feiveson and Sabelhaus, 2019). We hope that our empirical findings and discussion of potential theoretical explanations will be useful building blocks for future contributions to this literature. More broadly, a productive avenue for future research in the macro and wealth inequality literatures is to incorporate and refine lessons from household finance (Campbell, 2006).

References

- AIYAGARI, S. R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *The Quarterly Journal of Economics*, 109(3), 659–84.
- ALAN, S., K. ATALAY, AND T. F. CROSSLEY (2015): “Do the Rich Save More? Evidence from Canada,” *Review of Income and Wealth*, 61(4), 739–758.
- 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 (2019): “Tax evasion and inequality,” *American Economic Review*, 109(6), 2073–2103.
- 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.
- BRUNNERMEIER, M. K., AND S. NAGEL (2008): “Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-evidence on Individuals,” *American Economic Review*, 98(3), 713–736.
- CALVET, L. E., J. Y. CAMPBELL, AND P. SODINI (2009): “Fight or Flight? Portfolio Rebalancing by Individual Investors,” *The Quarterly Journal of Economics*, 124(1), 301–348.
- 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.
- (2018): *Financial Decisions and Markets*. Princeton University Press.
- 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 (2019): “Stock Market Returns and Consumption,” *Journal of Finance*, forthcoming.
- DYNAN, K. E., J. SKINNER, AND S. P. ZELDES (2004): “Do the Rich Save More?,” *Journal of Political Economy*, 112(2), 397–444.
- EHRENFREUD, M. (2015): “Where the poor and rich really spend their money,” *Washington Post*.
- 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.
- FAGERENG, A., L. GUISO, D. MALACRINO, AND L. PISTAFERRI (2016): “Heterogeneity in Returns to Wealth and the Measurement of Wealth Inequality,” *American Economic Review: Papers &*

- Proceedings*, 106(5), 651–655.
- (2019): “Heterogeneity and Persistence in Returns to Wealth,” *Econometrica*, *forthcoming*.
- FAGERENG, A., M. B. HOLM, AND K. N. TORSTENSEN (2019): “Housing Wealth in Norway, 1993–2015,” Discussion paper.
- 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. (2019): “Behavioral Inattention,” in *Handbook of Behavioral Economics: Applications and Foundations 1*, ed. by B. D. Bernheim, S. DellaVigna, and D. Laibson, vol. 2, chap. 4, pp. 261–343. Elsevier.
- 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.
- 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 (2018): “Risk, uncertainty, and the dynamics of inequality,” *Journal of Monetary Economics*, 94, 60–78.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2016): *Macroeconomics and Household Heterogeneity*vol. 2A of *Handbook of Macroeconomics*, chap. 11, pp. 843–921. Elsevier (Amsterdam).
- 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 (2019): “Income and Wealth Inequality in America, 1949-2016,” *Journal of Political Economy*, *forthcoming*, (9).
- MALMENDIER, U., AND S. NAGEL (2011): “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?,” *The Quarterly Journal of Economics*, 126(1), 373–416.
- 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.