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

Andreas Fagereng
BI Norwegian Business School

Martin Blomhoff Holm
University of Oslo

Benjamin Moll
London School of Economics

Gisle Natvik
BI Norwegian Business School

May 2021

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 and interpret our findings through the lens of economic theory. We find that saving rates net of capital gains are approximately flat across the wealth distribution, i.e., the rich do *not* actively save a larger share of their incomes than the poor. At the same time, saving rates including capital gains increase strongly with wealth because wealthier households own assets that experience capital gains which they hold on to instead of selling them off to consume. We show that these findings are consistent with standard models of household wealth accumulation with homothetic preferences under one additional assumption: rising asset prices are accompanied by declining asset returns rather than rising dividends (cashflows). We conclude that saving rate heterogeneity across wealth groups is not likely to be an important contributor to changes in aggregate saving and the wealth distribution, but capital gains are.

*We thank Adrien Auclert, John Campbell, Karen Dynan, Giulio Fella, Matthieu Gomez, Adam Guren, Henrik Kleven, Claus Thustrup Kreiner, Camille Landais, Søren 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 project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (#851891) and the Norwegian Research Council (#287720).

1 Introduction

A large and growing literature in macroeconomics studies the determinants of secular trends in income and wealth inequality and how such distributional shifts feed back to macroeconomic aggregates like the economy's saving rate and equilibrium interest rates, or how they affect the transmission of macroeconomic policy. A key ingredient in many of the theories in this literature is how individuals' saving behavior varies across the wealth distribution, in particular whether the rich save a larger share of their incomes than the poor.[1] Empirically disciplining the proposed theoretical mechanisms requires evidence on how saving rates vary with wealth. Unfortunately, such empirical evidence is largely lacking.[2]

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, and by interpreting our findings through the lens of economic theory.

Because Norway levies both income and wealth taxes on 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 study highlights that the relation between wealth and saving rates crucially depends on whether saving includes capital gains. We distinguish between two saving concepts which correspond to two alternative ways of writing the household budget constraint and differ by how capital gains are treated when writing this accounting identity. *Net saving*, or *active 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. *Gross saving* is simply the total change in a household's net worth, including any revaluation effects due to changing asset prices.[3]

[1]For such theories of secular inequality trends, see for example Boerma and Karabarbounis (2021); Greenwald et al. (2021); Gomez and Gouin-Bonenfant (2020); Kaymak, Leung and Poschke (2020); Benhabib and Bisin (2018); Hubmer, Krusell and Smith (2020); De Nardi and Fella (2017); Gabaix et al. (2016); Kaymak and Poschke (2016); De Nardi (2004); Carroll (1998). For theories of macro aggregates and policy, see for example Mian, Straub and Sufi (2021); Melcangi and Sterk (2020); Rachel and Summers (2019); Straub (2018); Auclert and Rognlie (2016); Krueger, Mitman and Perri (2016); Kumhof, Rancière and Winant (2015); Krusell and Smith (1998).

[2]A few existing studies do provide related evidence, e.g., on "synthetic saving rates." See the related literature section at the end of this introduction for a more detailed discussion, particularly footnote 5.

[3]The literature has used a number of other names for the same concepts, for example "change-in-wealth

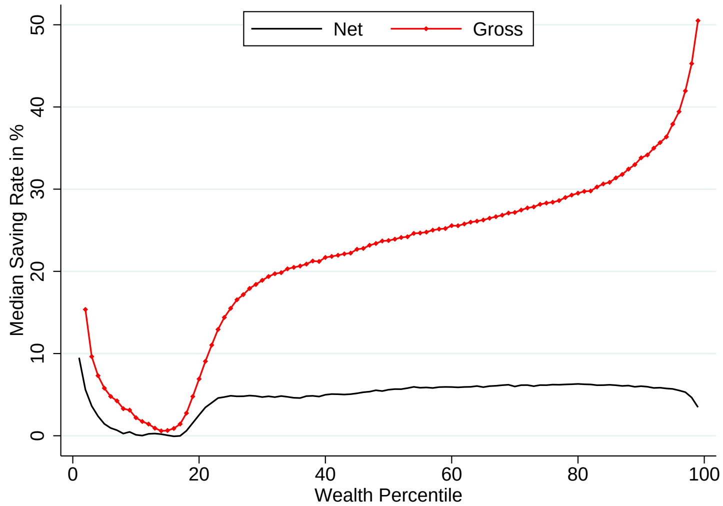

Figure 1: Saving rates across the wealth distribution.

Our main finding is that among households with positive net worth, net or active saving rates are remarkably flat across the wealth distribution. Gross saving rates instead increase markedly with wealth. Hence, the answer to how saving rates vary with wealth, crucially depends on whether capital gains are included in the definition of saving. Our second contribution is to provide a theoretical interpretation. We show that the empirical finding is consistent with standard models of household wealth accumulation with homothetic preferences under one additional assumption: rising asset prices are accompanied by declining asset returns rather than rising dividends (cashflows).

The empirical relationships between wealth and saving rates are easiest to communicate graphically. Figure 1 therefore plots saving rates against percentiles of net worth. To the left are households with negative net worth, while zero-wealth households are located around the 15th percentile. We see that among households with low or negative net worth, it is unimportant if capital gains are included or excluded from saving: net and gross saving rates track each other and decline with wealth. However, among the majority of households who have positive wealth, matters are very different. While the gross saving rate (including capital gains) increases sharply up to fifty percent for the top one percent of the wealth distribution, the net saving rate (excluding capital gains) is remarkably stable around seven percent. Moreover, these characteristics of the observed saving rates still hold when we control for the main determinants of heterogeneous saving behavior in economic models, namely age, income and inclination to save (e.g., due to

saving" in place of "gross saving" and "passive saving" in place of capital gains. Whereas we say that gross saving is the sum of net saving and capital gains, these studies would say that change-in-wealth saving is the sum of active and passive saving. The two statements are equivalent.

patience). They also hold when we examine saving rates in financial wealth, that is, saving rates with housing "taken out" of household wealth accumulation. Our data thus provide a nuanced answer to how saving rates vary with wealth: The rich do not have particularly high net saving rates, but they still accumulate more wealth than others through capital gains.

The decline in saving rates to the left in Figure 1 is consistent with standard models with borrowing constraints. The remarkable fact is how flat the net saving rate is across the rest of the wealth distribution. Indeed, once one appreciates the flat net saving rate in Figure 1, the increasing gross saving rate is a simple corollary. In particular, note that the diverging *gap* between the gross and net saving rates (though not their respective level or slope) has a mechanical explanation: wealthier households hold more assets like stocks and housing whose prices increase over time. Therefore, once we have found that the net saving rate is flat across the wealth distribution, the gross saving rate must increase with wealth if capital gains are positive. A flat net saving rate means that, even though wealthier households hold more assets experiencing capital gains, they do not sell off these assets to consume, but instead predominately hold on to them. They 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 and illuminates the magnitude of divergence between the two saving rates in Figure 1. Assume that 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 with different wealth: 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 neither individual consumes 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 even relatively small capital gains (say 2%) can induce a sizable divergence between net and gross saving.

These findings beg the questions: Why are net saving rates flat? And why do they remain flat in the face of changing asset prices? Our second contribution is to interpret our findings through the lens of economic theory. To this end, we consider a series of theoretical benchmark models and compare their predictions for net and gross saving rates to Figure 1.

We start with a particularly simple consumption-saving model that can be solved analytically and thereafter enrich the framework to address prevalent features of the data. 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 independent of wealth.

When the asset price varies 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 they come with increased future cashflows (dividends) or not. When they do, the household optimally consumes part of these future cashflows (an income effect) by selling some of its assets. As a consequence, the net saving rate will be decreasing with wealth. In contrast, when asset prices rise but cashflows do not – equivalently, when asset returns decline – the household does not consume out of these capital gains and, with homothetic utility, the net saving rate is independent of wealth. Intuitively, constant cashflows eliminate the income effect we just described and the falling asset return induces primarily a substitution effect. But in contrast to the income effects from dividend growth, this substitution effect does not systematically vary with household wealth. Together with homothetic utility, the implication is that the net saving rate is constant across the wealth distribution.

In summary, our findings are consistent with standard models of household wealth accumulation with homothetic preferences, borrowing constraints and one additional assumption: rising asset prices are accompanied by declining asset returns rather than rising dividends (cashflows).[4] While declining asset returns are our preferred explanation for the observed saving behavior, we also briefly discuss other candidate mechanisms that have the potential to generate this behavior even if rising asset prices are instead accounted for by dividend growth. These explanations include portfolio adjustment frictions, non-homothetic preferences, and various behavioral explanations. All have in common that they moderate the tendency for wealthier households to consume out of future cashflows, but none of them generate flat net saving rates except as a knife-edge case.

How important are the saving patterns we uncover at the micro level for outcomes at the macro level? As an illustration, we quantify how (i) net saving rate heterogeneity across the wealth distribution and (ii) capital gains shaped aggregate saving and the wealth distribution in Norway from 1995 to 2015. Over this period the aggregate wealth-to-income ratio rose from approximately 4 to 7. If we counterfactually impose that every wealth percentile had the same median net saving rate, the implied wealth-to-income trajectory remains indistinguishable from the actual one. In contrast, if we turn off capital gains and impose that they were zero every year, the trajectory flattens out and the wealth-

[4]Note that this explanation of our empirically observed saving behavior is consistent with the dominant view in the finance literature that asset price fluctuations are primarily driven by time-varying discount rates (Cochrane, 2005; Campbell, 2003), not dividends. For a recent study relating long-run stock price growth to declining interest rates, see van Binsbergen (2020).

to-income ratio in 2015 remains at 4. Moreover, imposing that each wealth percentile had the same net saving rate casts similarly negligible influence across the wealth distribution. In contrast, imposing zero capital gains has marked distributional consequences. For instance, the evolution of wealth at the 10th percentile was flat in the data and would have been flat without capital gains. Median wealth more than doubled in the data, but would have remained almost constant without capital gains. Wealth at the 90th percentile more than tripled in the data, but would only have doubled without capital gains. To sum up, net saving rate heterogeneity across wealth groups has been unimportant for changes in aggregate saving and the wealth distribution in our data, while capital gains have been key for both.

Related Literature. We hope that our findings will be useful building blocks for the large and growing literature on macroeconomic implications of micro-level saving behavior cited in the first paragraph. Many of the studies in this literature assume homothetic preferences and we find that this feature may suffice for explaining observed saving behavior when coupled with the additional assumption that rising asset prices are accompanied by declining asset returns. Notably, this additional assumption is not typically made in the literature. Instead, canonical models typically either assume away capital gains completely by imposing constant asset prices, or they impose a constant asset return which implicitly assumes that all capital gains are driven by dividend growth. Important exceptions that model rising asset prices due to declining returns are Greenwald et al. (2021) and Gomez and Gouin-Bonenfant (2020). Besides this theoretical literature, our findings are also relevant for a nascent empirical literature in macroeconomics and the study of inequality that emphasizes portfolio choice and asset price changes (Feiveson and Sabelhaus, 2019; Kuhn, Schularick and Steins, 2019; Martínez-Toledano, 2019).

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.[5] Bach, Calvet and Sodini (2018)

[5]Arguable exceptions are Krueger, Mitman and Perri (2016), Saez and Zucman (2016), and Smith, Zidar and Zwick (2020). However, none of these papers provide evidence on *individual*-specific saving rates like we do. Krueger, Mitman and Perri document 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 due to the small sample size of their dataset, the U.S. Panel Study of Income Dynamics). Saez and Zucman and Smith, Zidar and Zwick provide evidence on "synthetic saving rates" that are computed by following percentile groups, rather than individuals, over time. Interestingly and in line with our results, Smith, Zidar and Zwick find that using their preferred capital gains estimates considerably attenuates the saving rate disparities of Saez and Zucman and increases the importance of asset price growth for understanding wealth growth.

Straub (2018), Alan, Atalay and Crossley (2015), and Dynan, Skinner and Zeldes (2004) 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 also examine saving behavior across the wealth distribution using administrative data but with a different focus. Complementary to our 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 informative object to study.[6] Moreover, given their focus on the saving-to-wealth ratio, they do not study households 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.

Our paper is also related to the literature on the consumption effects of asset price changes, in particular papers that estimate how capital gains affect consumption.[7] Broadly consistent with our findings, Di Maggio, Kermani and Majlesi (2019) and Baker, Nagel and Wurgler (2007) argue that marginal propensities to consume (MPCs) out of capital gains are significantly smaller than MPCs out of dividend income.[8] Our evidence differs substantively though, as the object of our interest is not households' marginal responses to changes in available resources, but their average propensity to save out of income and capital gains. As pinpointed by Aguiar, Bils and Boar (2020), theory may have very different predictions for the two objects.

Roadmap. Section 2 presents what a series of theoretical benchmarks predict for saving rates across the wealth distribution. Section 3 describes our data and how we measure saving rates. Section 4 presents our main empirical results on net and gross saving rates. Section 5 presents additional evidence on saving rates by capital gains and briefly discusses alternative explanations why net saving rates are flat. Section 6 discusses the macroeconomic and distributional implications of our findings. Section 7 concludes.

an instrumental variable strategy. We instead focus on wealth which is readily observable.

[6] In contrast, as we explain in the main body of the paper, standard consumption-saving models have no clear prediction for the saving-to-wealth ratio except that it should be mechanically decreasing with wealth.

[7] As 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 Chodorow-Reich, Nenov and Simsek (2021), Paiella and Pistaferri (2017), and Christelis, Georgarakos and Jappelli (2015) 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. (2021) among others.

[8] Our findings are also consistent with a household finance literature that finds substantial inertia in households' financial decisions (e.g., Calvet, Campbell and Sodini, 2009; Brunnermeier and Nagel, 2008).

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 4. In particular, we show that standard models of individual wealth accumulation predict that net saving rates, i.e., saving rates excluding capital gains, are approximately constant across the wealth distribution under two key assumptions: (i) homothetic preferences and (ii) rising asset prices are accompanied by declining asset returns rather than rising dividends (cashflows). Theory also motivates our definition of different saving concepts that we use in our empirical analysis ("net" and "gross"). We begin with two simple consumption-saving models that can be solved with pencil and paper: first, a model with a constant asset price as in most existing work; second, a model with a changing asset price as in the data. Thereafter we address other important features of the data such as housing and risk.

2.1 Saving Decisions with Constant Asset Price

Households are infinitely lived and have homothetic utility

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

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):

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

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:

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

Hence, the saving rate out of income is independent of wealth.[9] We show below that many other benchmark models inherit this property, at least approximately.

2.2 Saving Decisions with Changing Asset Prices

Setup. We next extend the simple benchmark model to feature a time-varying asset price. As above, households have homothetic preferences (1) and 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. This asset pays a dividend D_t and is the only saving vehicle available to households. Both the price p_t and the dividend D_t evolve according to exogenous and deterministic time paths. The budget constraint is

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

Households maximize (1) subject to (4). Note that we assume away any form of uncertainty because this complication is inessential for the points we want to make. We briefly discuss the role of uncertainty in Section 2.3.

The interest rate relevant for households' investment decisions is the asset's return which consists of both dividend payments and capital gains:

$$r_t := \frac{D_t + \dot{p}_t}{p_t}. \tag{5}$$

In particular, the budget constraint (4) can be written in terms of the market value of wealth $a_t := p_t k_t$ as $\dot{a}_t = w + r_t a_t - c_t$. Relative to the model in Section 2.1 the return r_t is now potentially time-varying.

Rather than taking the asset price $\{p_t\}_{t\geq 0}$ as given as we have just done, we can equivalently take the perspective of the asset pricing literature to treat the required asset return $\{r_t\}_{t\geq 0}$ as a primitive and the price as an outcome.[10] Integrating (5) forward in time and assuming a no-bubble condition, the asset price equals the present discounted value of

[9] Households save a constant fraction of their effective wealth and their income is the constant return to that effective wealth. Hence saving is also a constant fraction of income. Note that the saving rate is also constant over discrete time intervals (not just infinitesimal ones): using (2), we have $a_{t+1} - a_t = \int_0^1 \dot{a}_{t+s} ds = \frac{r-\rho}{\gamma} \int_0^1 \left(a_{t+s} + \frac{w}{r}\right) ds$ and hence saving $a_{t+1} - a_t$ is a constant fraction $\frac{r-\rho}{\gamma r}$ of income $\int_0^1 (w + r a_{t+s})\, ds$.

[10] For now, we do not take a stand where the required rate of return $\{r_t\}_{t\geq 0}$ in (6) comes from. One intuitive reason is that individuals can save in another asset (for example, a bond) that pays a return $\{r_t\}_{t\geq 0}$. Arbitrage then requires (5) which implies (6). Alternatively, $\{r_t\}_{t\geq 0}$ could be pinned down from preferences in general equilibrium.

future dividend streams:[11]

$$p_t = \int_t^{\infty} e^{-\int_t^s r_\tau d\tau} D_s ds. \tag{6}$$

The goal of this section is to understand how households' optimal consumption and saving decisions are affected by rising asset prices. From (6), a growing asset price can only be due to one of two factors: dividend growth or declining returns. That is, either the asset price rises and so do cashflows (dividends); or the asset price rises even though cashflows do not, i.e., the return declines. We show below that these two cases – whether growing asset prices come with increased cashflows or not – have drastically different implications for households' optimal consumption and saving decisions.

Key Concepts: Net and Gross Saving. Before turning to the characterization of households' saving decisions, we define key concepts that we will use in our empirical analysis. These are "net" and "gross" saving and the corresponding net and gross saving rates. The different definitions follow from different ways of writing the budget constraint "consumption plus saving equals income." They differ by how capital gains are treated when writing this accounting identity. We have:

$$c + \overbrace{p\dot{k}}^{\text{net saving}} = \overbrace{w + Dk}^{\text{disposable income}}, \tag{7}$$

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

The difference between these two accounting identities is 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 (Simons, 1938; von Schanz, 1896; Haig, 1921).[12] 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.

Optimal Consumption and Saving Decisions with Changing Asset Prices. With these definitions in hand, Proposition 1 characterizes households' optimal choices in the face

[11] The no bubble condition is $\lim_{T\to\infty} e^{-\int_0^T r_s ds} p_T = 0$.

[12] This 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, Robbins (2019) and Eisner (1988) argue for including capital gains in the income definition used in the National Income and Product Accounts.

of changing asset prices. As we explain momentarily, Figure 2 illustrates some of the proposition's main findings.

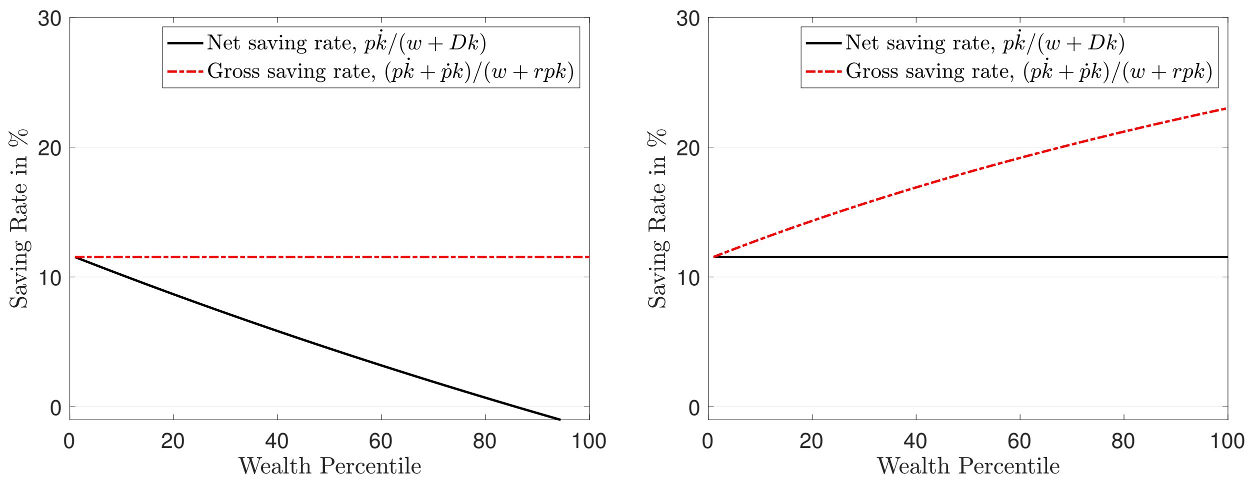

(a) Case 2: growing dividend (constant return) (b) Case 3: constant dividend (declining return)

Figure 2: Saving rates across wealth distribution with growing asset price (Proposition 1).

Proposition 1. *Consider a household with current asset holdings k_t who maximizes the homothetic utility function (1) subject to the budget constraint (4) with perfect foresight over $\{p_s, D_s, r_s\}_{s\geq t}$ with $\{r_s\}_{s\geq t}$ defined in (5). Its optimal consumption and net saving are*[13]

$$c_t = \xi_t \int_t^\infty e^{-\int_t^s r_\tau d\tau}(w + D_s k_t)ds, \tag{9}$$

$$p_t \dot{k}_t = \phi_t(w + D_t k_t) - \xi_t k_t \int_t^\infty e^{-\int_t^s r_\tau d\tau}(D_s - D_t)ds, \tag{10}$$

where $\xi_t = \xi_t\left(\frac{1}{\gamma}, \rho, \{r_s\}_{s\geq t}\right)$ and $\phi_t = \phi_t\left(\frac{1}{\gamma}, \rho, \{r_s\}_{s\geq t}\right)$ are defined in (A4) and are independent of k_t and w. When the intertemporal elasticity of substitution is zero, $1/\gamma = 0$, then $\phi_t\left(0, \rho, \{r_s\}_{s\geq t}\right) = 0$. We further have the following special cases:

1. ***Constant asset price, dividend and return:** When $p_t = p$ and $D_t = D$, all t so that $r_t = \frac{D}{p} =: r$ for all t, then $\xi_t = r - \frac{r-\rho}{\gamma}$ and $\phi_t = \frac{r-\rho}{\gamma r}$ and from (10) net saving (which also*

[13] Equivalently, we can write consumption and gross saving in terms of wealth $a_t := p_t k_t$:

$$c_t = \xi_t \left(a_t + \int_t^\infty e^{-\int_t^s r_\tau d\tau} w ds\right), \qquad \dot{a}_t = \phi_t(w + r_t a_t) - \xi_t a_t \int_t^\infty e^{-\int_t^s r_\tau d\tau}(r_s - r_t)ds.$$

equals gross saving) is $p\dot{k}_t = \frac{r-\rho}{\gamma r}(w + rpk_t)$, i.e., we recover (2) with $a_t := pk_t$. Both net and gross saving rates are independent of wealth.

2. ***Price growth with dividend growth (constant return):*** *When p_t and D_t increase over time while the asset return is constant, $r_t = r$ for all t, then $\xi_t = r - \frac{r-\rho}{\gamma}$, $\phi_t = \frac{r-\rho}{\gamma r}$, and $\int_t^\infty e^{-\int_t^s r_\tau d\tau}(D_s - D_t)ds = \dot{p}_t/r$. Therefore net and gross saving are*

$$p_t\dot{k}_t = \frac{r-\rho}{\gamma r}(w + D_t k_t) - \left(1 - \frac{r-\rho}{\gamma r}\right)\dot{p}_t k_t, \tag{11}$$

$$p_t\dot{k}_t + \dot{p}_t k_t = \frac{r-\rho}{\gamma r}(w + rp_t k_t). \tag{12}$$

Therefore the gross saving rate is independent of wealth and equal to $\frac{r-\rho}{\gamma r}$ and the net saving rate decreases with wealth as in Figure 2(a).

3. ***Price growth without dividend growth (falling return):*** *When p_t increases over time while the dividend is constant, $D_t = D$ all t, and the return r_t declines over time, then*

$$p_t\dot{k}_t = \phi_t(w + Dk_t), \tag{13}$$

$$p_t\dot{k}_t + \dot{p}_t k_t = \phi_t(w + r_t p_t k_t) + (1 - \phi_t)\dot{p}_t k_t. \tag{14}$$

Therefore the net saving rate is independent of wealth and equal to ϕ_t and the gross saving rate increases with wealth as in Figure 2(b). Wealthier households "save by holding."

First consider optimal consumption in (9). With homothetic CRRA utility (1), households consume a fraction ξ_t of the present value of their future income stream $\int_t^\infty e^{-\int_t^s r_\tau d\tau}(w + D_s k_t)ds$. Net saving in (10) is best understood by considering the three special cases in the Proposition. As expected, when the asset price, dividend, and return are all constant as in special case 1, we recover the characterization of optimal saving behavior from Section 2.1. The more interesting cases are special cases 2 and 3, in which the asset price increases over time as in the data.

In special case 2, this asset price growth is entirely accounted for by dividend growth; equivalently the asset return is constant over time, $r_t = r$ for all t. As already noted, the model can be written in terms of market wealth $a_t := p_t k_t$, and with a constant asset return, the model is exactly isomorphic to the one in Section 2.1. It therefore makes sense that the expression for gross saving in (12) again recovers (2) from Section 2.1. To understand the expression for net saving (11), it is helpful to consider consumption in (9) becomes $c_t = \left(r - \frac{r-\rho}{\gamma}\right)\int_t^\infty e^{-r(s-t)}(w + D_s k_t)ds$. Therefore, when the growing asset price comes with additional future cashflows $D_s > D_t, s > t$, the household optimally consumes part of this future income – a standard income effect. The household achieves this by selling some of its assets and more so the larger are its asset holdings, hence the second term in (11) which becomes more and more negative the larger are assets k_t. The implications for saving *rates* are then straightforward and as illustrated in Figure 2(a): when the rising asset price comes with future cashflows and the asset return is constant over time, our simple theory predicts that the net saving rate should be decreasing with wealth and the gross saving rate independent of wealth.

In special case 3, the asset price grows but cashflows do not; equivalently, the asset return r_t declines (see (5) and (6)). In this case, there no longer is an income effect from growing dividends as in special case 2. The only effects that remain are due to the changing asset return r_t, in particular there is a substitution effect.[14] But this substitution effect does not systematically vary with household wealth. Therefore, while asset price changes do affect the net saving rate (ϕ_t depends on $\{r_t\}_{t\geq 0}$, capturing the substitution effect), they do not affect it in a way that varies systematically with wealth – see (10). As a result, wealthier households who experience larger capital gains passively save most of these, i.e., they "save by holding." The implications for saving *rates* are illustrated in Figure 2(b): when the asset price rises but cashflows do not, our simple theory with homothetic utility predicts that the net saving rate is independent of wealth and the gross saving rate is increasing with wealth.

More generally, rising asset prices may be accompanied by both dividend growth and declining returns. In this case, the key determinant of how the net saving rate varies with wealth is whether the combination of rising D_t and falling r_t triggers any income effects that differ systematically with wealth. Whether this happens can be seen by examining the special case of (10) when the intertemporal elasticity of substitution is zero, $1/\gamma = 0$ so that $\phi_t = 0$:

$$p_t \dot{k}_t = -\xi_t k_t \int_t^{\infty} e^{-\int_t^s r_\tau d\tau}(D_s - D_t)ds,$$

with $\xi_t = 1/\int_t^{\infty} e^{-\int_t^s r_\tau d\tau} ds$. This expression shows that, also in the more general case, such

[14] The changing asset return r_t also has an income effect. However, in special case 3 with a constant dividend income flow D, this income effect is either zero or does not systematically differ across the wealth distribution. To see this, consider the case where additionally the intertemporal elasticity of substitution (IES), $1/\gamma = 0$ so that there is no substitution effect and the household wants to perfectly smooth its consumption. With both constant dividend and labor income, the household can achieve this by simply consuming the constant overall income flow $w + Dk_t$ each period. In the language of Auclert (2019), the household has an "unhedged interest rate exposure" of zero. Under less stringent assumptions, a changing return r_t does have an income effect but it does not systematically differ across the wealth distribution. For example, one can break the assumption that the household's labor income w is constant over time. In this case (13) becomes $p_t \dot{k}_t = \phi_t(w_t + Dk_t) - \xi_t \int_t^{\infty} e^{-\int_t^s r_\tau d\tau}(w_s - w_t)ds$, where the second term captures the income effect due to changing r_t. But this term does not systematically vary with wealth k_t.

income effects are small – and hence the net saving rate is approximately flat – unless rising asset prices are accompanied by fast dividend growth.

To summarize, our main theoretical result is that the net saving rate is approximately constant across the wealth distribution under two key assumptions: (i) households have homothetic preferences and (ii) rising asset prices are accompanied by declining asset returns rather than rising dividends (cashflows). See case 3 of Proposition 1. It is worth noting that, while canonical models of wealth accumulation (e.g., Aiyagari, 1994) almost always make the first assumption, homothetic preferences, they do not typically make the second. Instead, canonical models typically either assume away capital gains by imposing constant asset prices (special case 1 of Proposition 1) or impose a constant asset return which implicitly assumes that all capital gains are driven by dividend growth (special case 2).

2.3 Extensions

The above framework was deliberately simplistic so as to introduce our main saving concepts and highlight our main theoretical predictions in Proposition 1. We next show how these predictions carry over to richer models of household wealth accumulation, covering the main extensions in the macroeconomics literature.

The extensions we consider are the consumption aspect of housing, asset-price risk, income risk, correlation between age and earnings, and heterogeneity in patience (innate inclination to save). Of these extensions, only the first two explicitly model changing asset prices. For the other extensions we therefore do not distinguish between net and gross saving, and our statements should be interpreted as statements about the net saving rate. In versions of these theories with rising asset prices accompanied by falling returns, the gross saving rate would follow residually, depending on what capital gains happen to be. The models corresponding to these extensions are described in Appendix B. For our purposes, the important conclusions are summarized as follows:

1. **Housing as a consumption good.** The fact that housing is a consumption good as well as an asset, does not alter the theoretical prediction that net saving rates are approximately flat across the wealth distribution when asset (house) price changes are accompanied by falling returns rather than rising dividends (rents or implied rents).

2. **Asset-price risk.** In models with a stochastic asset price, the net saving rate continues to be flat across the wealth distribution as long as persistent asset price movements

are primarily accounted for by discount factor movements rather than cashflows (in contrast, higher-frequency price changes can also be due to cashflow shocks).

3. **Income risk and borrowing constraints.** In models with income risk, cross-sectional correlation between labor income and wealth can induce an upward sloping relation between saving rates and wealth. However, conditional on households' income and age, the saving rate is approximately flat. Borrowing constraints result in elevated saving rates for indebted households.

4. **Life-cycle earnings profile.** In models where income varies systematically over the life cycle, cross-sectional correlation between income, age and wealth can induce an upward sloping relation between saving rates and wealth. However, conditional on households' income and age, the saving rate is approximately flat.

5. **Heterogeneous inclinations to save.** In models where households differ in their innate inclination to save (patience or returns), cross-sectional correlation between saving inclination and wealth can induce an upward sloping relation between saving rates and wealth. However, conditional on households' saving inclination, the saving rate is approximately flat.

The first statement follows from extending our benchmark model to feature housing. We start with this extension because intuition might suggest that the consumption aspect of housing would be enough to make households hold on to their residential wealth in the face of rising asset prices, and that this could explain the observed flat net saving rate regardless of what happens to asset returns.[15] In Appendix B.1 we build on Berger et al. (2018) and study an environment where households invest in two assets, housing and bonds, the former of which is also a durable consumption good. We show that, just like in our benchmark model in Section 2.2, the key determinant of differences in saving rates across the wealth distribution is whether a rising house price is accompanied by a falling return or rising (implied) rents. When the house price rises in tandem with rents such that the housing return is constant, the theory predicts that households optimally increase their total consumption expenditure (consisting of both non-durable consumption and housing) by selling off parts of their assets, and therefore the net saving rate is decreasing

[15] As 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."

with wealth (as in case 2 of Proposition 1).[16] When the house price instead rises even though (implied) rents do not, i.e., the housing return falls, the net saving rate is constant across the wealth distribution (as in case 3 of Proposition 1). These results show that the key ingredient needed to generate a flat net saving rate is not the consumption aspect of housing but instead that rising house prices are accompanied by falling returns.

The second statement is based on an extension of the model in Section 2.2 in which the asset price is stochastic and driven by dividend- and stochastic discount factor shocks. As in Proposition 1, the net saving rate continues to be flat across the wealth distribution when asset price changes are accounted for by discount factor shocks. When asset price changes are instead accounted for by changing cashflows, an additional subtlety emerges: the optimal consumption and saving response to dividend shocks depends on whether these are transitory or persistent. The logic follows from the permanent income hypothesis. Households experiencing persistent dividend shocks optimally consume part of the resulting income flow so that the net saving rate is decreasing with wealth. In contrast, households experiencing transitory dividend shocks optimally save these, which in turn generates an increasing gross saving rate and flat net saving rate. Thus, also with asset-price risk, the net saving rate continues to be flat across the wealth distribution as long as persistent asset price movements are primarily accounted for by discount factor movements.

The third statement applies to models with labor income risk and borrowing constraints, as in Aiyagari (1994) and Huggett (1993). We analyze such a model in Appendix B.3 and show that it generates a policy function for saving which conditional on income is declining in wealth. The slope is distinct for households close to the borrowing constraint and almost flat for richer households. For households with higher labor income realizations, the policy function simply shifts up. If we instead consider the cross-sectional relationship between saving rates and wealth in the model's stationary distribution, without conditioning on income, the relationship is first decreasing and then increasing. This reflects two opposing forces. On the one hand, conditional on income, saving rates decrease with wealth. On the other hand, saving rates increase with labor income which in turn is positively correlated with wealth. Hence, in our empirical exercises we will condition on labor income.

In 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

[16] As we show in Appendix B.1 this is true even in an extreme case in which the household's utility function over non-durable consumption and housing is Leontief. The key observation is that Leontief preferences do *not* restrict the household to keep its consumption of non-durables and housing services unchanged; instead such preferences only restrict these to move in fixed proportions.

begin to save and accumulate wealth. Consequently, life-cycle considerations introduce a cross-sectional correlation between saving rates and wealth because both are correlated with age and income. Hence, we will conduct empirical exercises in which we condition on age and labor income. Conditional on age and income, models tend to predict that saving rates are slightly decreasing with wealth, as noted by De Nardi and Fella (2017).

Households might differ in their innate inclination 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 will be explained by historical saving behavior at the individual level. Hence, we will conduct empirical exercises in which we condition on past saving behavior.

3 Data and Definitions

Our study uses Norwegian administrative data. In contrast to other Scandinavian countries, Norway still levies a wealth tax on households (in addition to income taxes). In the process of collecting these taxes, all households are obliged to provide a complete account of household income and balance sheet components down to the single asset category every year. These data are reported to the tax authorities by third parties and are the foundation of our empirical analysis. Below we describe these data in more detail and explain the saving rate measures we construct.

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, designed to smooth public spending and support future governments' ability to meet public pension entitlements.[17]

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

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.[18] 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 C.

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 taxes are levied on the household as a whole, whereas income taxes are levied individually.[19] 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. For traded assets, 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 2005 to 2015. We impose a few minor sample restrictions in order to reduce errors in the computation of saving rates

[18] Arguably, 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.

[19] Taxable 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.

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)[20] 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 in proportion 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 500,000) are also required to have their balance sheet audited by an approved auditing entity. As noted in the pre- and proceeding paragraphs, 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 C.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 which often deviate 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 the price per square meter on house characteristics. The predicted values from this

[20] In Appendix C.5 we verify the consistency between our micro-level data on stock holdings and aggregate household stock holdings at the Norwegian stock exchange.

procedure are then used as housing wealth over our sample period.[21] 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 – see Section 4.5 and Appendix C.6.

Descriptive Statistics. Figure 3 plots the average holdings of different asset classes and debt relative to total assets across the wealth distribution. 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 3 displays portfolio shares within each wealth percentile, panel (b) repeats the exercise plotting the portfolio shares within groups that give more weight to the top of the wealth distribution. 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 deposits) is high across the entire wealth distribution. In contrast, public equity is a relatively small component of Norwegian households' wealth. Finally, we see (in panel (b)) that private business ownership becomes important 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.[22]

[21] The 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 (2020) evaluate the predictive power of the method in a holdout sample of actual transactions, showing that 85% of the predicted housing values

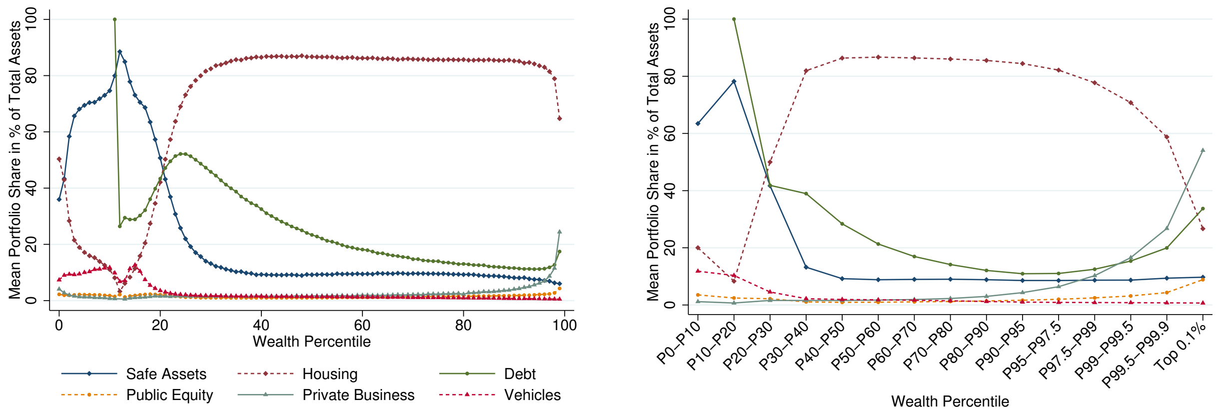

(a) Mean portfolio shares by wealth percentiles

(b) Mean portfolio shares by wealth percentiles

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.

Figure 3: Asset class shares in household portfolios.

3.2 Measuring Net and Gross Saving Rates

As discussed in Section 2.2, we distinguish between two notions of saving: *net* and *gross 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 C.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.

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

are within ±10% of the transaction value.

[22] To 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). Denoting housing by h, the house price by p, other assets by b, total net worth by $a = b + ph$, housing depreciation by δ, and the implicit rental rate of housing by R (for example $Rh = 2.04\% \times ph$ as explained momentarily) the analogue of (8) is

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

We follow Eika, Mogstad and Vestad (2020) 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.

	Mean	SD	P10	Median	P90	Participation 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 (net) income	58,828	204,020	29,525	51,411	84,987	
Haig-Simons (gross) 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.

Table 1: Descriptive statistics.

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 C.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 (A60) 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.[23] 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 C.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 3), this likely results in an underestimate of capital gains and gross saving among the very wealthiest.

Households' holdings in bond and money market funds are concentrated in government bonds with maturity below one year and medium-term bonds (two to four years), according to holding data from the Norwegian Fund and Asset Management Association. Among short- and medium-term bonds, the share of short-term bonds is approximately 70 percent. Because such bonds are unlikely to experience substantial price changes and bond holdings are a very small fraction of households' asset holdings (0.4% of the total portfolio over the sample period), we make the simplifying assumption that bonds do not experience capital gains.

The remaining assets and liabilities listed in Section 3.1 do not experience asset price changes.[24] Hence, with the above measures of net saving and capital gains within asset classes, we have all the necessary components to compute the net and gross saving rates (from Section 2.2) in the data.

[23]Stock 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.

[24]Vehicles 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.

4 Results

This section provides our main results. We start by plotting the net and gross saving rates for each wealth percentile. We next control for the main factors theory suggests are important for this correlation: income, age, education, and individual-fixed inclination to save. Thereafter, we address the role of housing, the dominant asset on household balance sheets, in explaining our findings. Finally we discuss a number of additional exercises to better understand the robustness of our findings, including the role of pensions and issues related to time-aggregation.

4.1 Saving Rates across the Wealth Distribution

Figure 4(a) plots net and gross saving rates against percentiles of net worth. For every year we have computed the median saving rate within each year-specific wealth percentile. The two lines plot the average of these within-year medians over our sample period.[25] The shaded areas reflect 95% confidence intervals, computed from quantile regressions.[26]

To unveil the importance of capital gains for the pattern in panel (a), panel (b) shows how capital gains vary across the wealth distribution. In addition, we plot households' asset holdings relative to income. Panel (b) uses the same time-aggregation approach as panel (a).

In panel (a), we start with the *net* saving rate. To the left, we see that the net saving rate is decreasing with wealth among the least wealthy households. Then, after reaching a trough around the 15th wealth percentile, it rises up to the 25th wealth percentile. Note that the trough lies where net worth is around zero (see Figure 3). The net saving rate flattens out at 6 – 8% from the 25th wealth percentile and up.

Consider next the gross saving rate. As capital gains have been positive on average over the sample period, 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

[25] We present results from an alternative time-aggregation approach in Section 4.5.

[26] The confidence bands in Figure 4(a) are computed as follows. For every year t we run a quantile regression of saving rates on indicators of wealth percentiles. Based on the resulting robust variance-covariance matrices, we compute the following time-averaged standard error for each percentile: $\widehat{\sigma}_p = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\widehat{\sigma}_{pt}^2}$, where T is the number of years and $\widehat{\sigma}_{pt}$ is the standard error of the regression coefficient on the indicator for percentile p in year t.

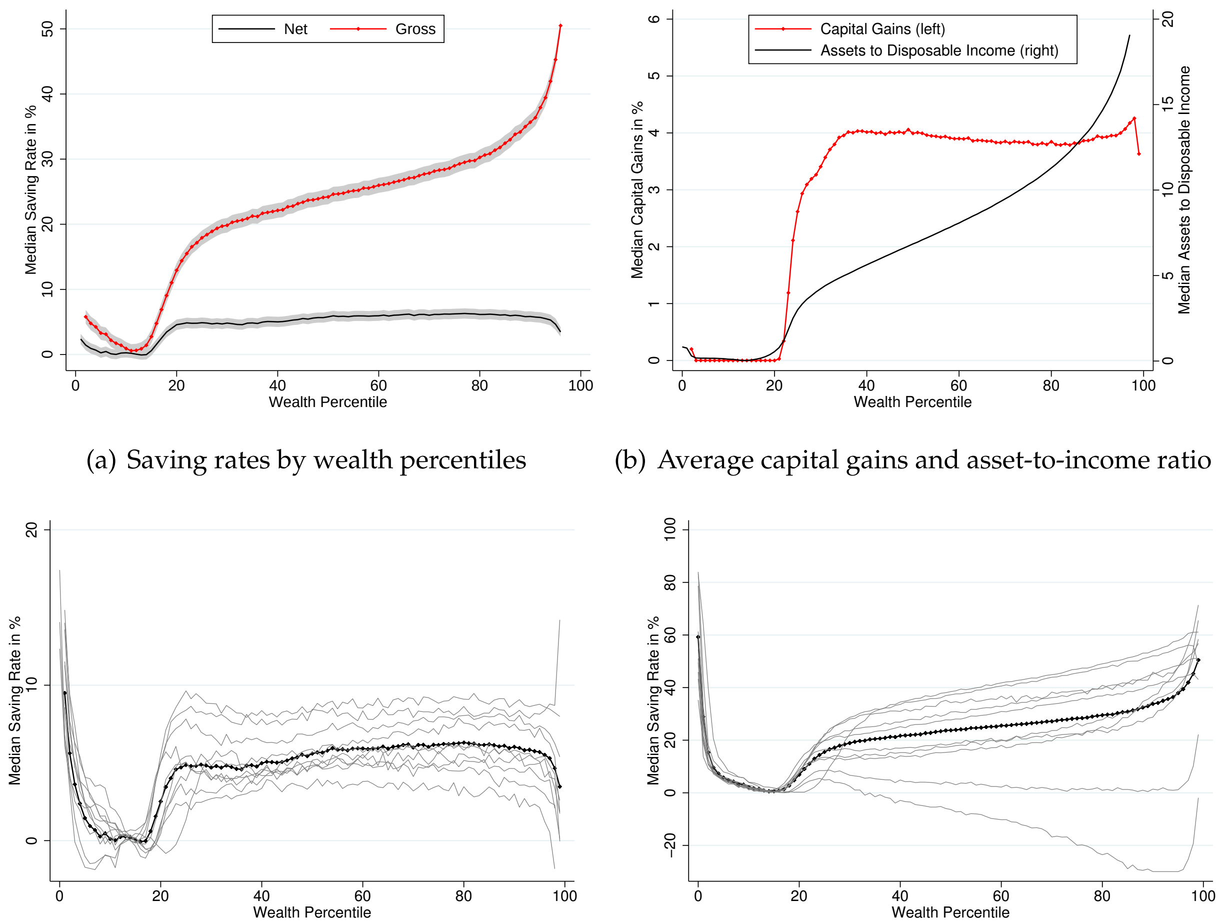

(a) Saving rates by wealth percentiles

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

(c) Net saving rates in different years

(d) Gross saving rates in different years

Notes: (a) displays the median saving rates within wealth percentile and year, averaged across all years (2005-2015). The grey areas show 95% confidence intervals computed from the variance-covariance matrix from quantile regressions of saving rates on indicators of wealth percentiles each year, averaged across all years (see footnote 26). (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 in (a) and (b) are computed as the median within-wealth percentile and year, averaged across all years (2005-2015). (c) and (d) display median saving rates for each year of the sample (grey) together with the average across all years (black).

Figure 4: Saving rates across the wealth distribution.

above zero wealth, 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).

Instead of averaging across time, Figures 4(c) and (d) plot saving against wealth for each year separately. The lighter gray lines are constructed for each specific year. The thicker black lines are our baseline time-averaged counterparts. Notice that the scale on the vertical axes differs between the two panels.

Together, Figures 4(c) and (d) 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 relationship between net saving rates and wealth 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 considerably from year to year. It increases most distinctly with wealth in good years with high capital gains, and then *de*creases with wealth in bad years where large parts of the distribution experienced capital losses.

In the beginning of this paper we asked whether wealthier households have higher saving rates. Figure 4 suggests a nuanced answer to this question. On the one hand, net saving rates are almost flat across the wealth distribution, especially among households with some positive net worth – which is to say that, no, wealthier households do *not* have higher saving rates in the traditional sense. On the other hand, gross saving rates increase sharply with wealth – that is, even though the net saving rate is flat, wealthier households still accumulate more wealth through capital gains.

The proximate explanation for our observed patterns of saving rates 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 4 visualizes how wealthier households hold more assets relative to their income, and capital gains rates do not systematically decline with wealth. Second, households do not systematically adjust their saving in response to asset price changes. In the face of rising asset prices, households do not sell off assets to consume part of their capital gains. Conversely, in the face of falling asset prices, households also do not purchase more assets to compensate for the paper loss in wealth. Instead, households tend to hold on to assets as their price fluctuates, leading to flat net saving rates every year and gross saving rates that increase or decrease depending on the sign of capital gains.

Having documented our main empirical finding in Figure 4, we can now link it back to the theoretical discussion in Section 2. What explains the observed patterns for net and gross saving rates? In particular, what explains (i) the remarkably flat net saving rate for households with positive net worth; and (ii) the elevated net and gross saving rates for households with debt at the left of the figure?

Comparing our results in Figure 4 to the theoretical predictions in Figure 2 suggests a simple explanation for the flat net saving rate: households behave according to a completely standard theory of household wealth accumulation with homothetic preferences but one additional twist, namely that rising asset prices are accompanied by declining asset returns rather than rising cashflows. See case 3 of Proposition 1 and Figure 2(b). Intuitively, when asset prices rise even though cashflows do not, richer households do not experience a larger income effect than poorer ones and therefore consume approximately the same fraction of their disposable incomes. Rich households instead optimally "save by holding" on to most of their assets experiencing capital gains. Section 2 also contains a simple explanation of the elevated net and gross saving rates for households with debt: borrowing constraints and high-interest debt. Intuitively, households facing borrowing constraints or borrowing at high interest rates, accumulate wealth more quickly so as to repay this debt and distance themselves from these constraints (see Section 2.3). We flesh out these theoretical implications further in Section 5 where we provide some additional, more direct evidence consistent with the declining-returns interpretation of our findings and discuss alternative explanations.

4.2 Controlling for Age, Earnings, and Education

As summarized in Section 2.3, common extensions of consumption-saving models imply that, in the cross-section, saving rates might correlate with wealth because both saving and wealth are related to labor income and age. Earnings risk may motivate households with high income realizations to save more, thereby inducing a positive correlation between saving rates and wealth. Life-cycle considerations may lead to a positive relationship between saving rates and wealth because both are correlated with age.[27] In this section, we therefore explore to what extent the patterns in Figure 4 can be accounted for by these factors.

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 5(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

[27] An 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.

for each year, and then averaged these medians over time. The resulting pattern is very similar to that in Figure 4. The one notable difference is that the net saving rate now increases slightly with wealth.

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

$$\frac{s_{it}}{y_{it}} = \sum_{p=1}^{100} \phi_p D_{it,p} + f(x_{it}) + \tau_t + \epsilon_{it}, \tag{15}$$

where s_{it}/y_{it} is the net or gross 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 $f(\cdot)$ to include a full set of age, education (no high school, high school, college), and household labor income (in quintiles) indicator variables. 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 5(b).

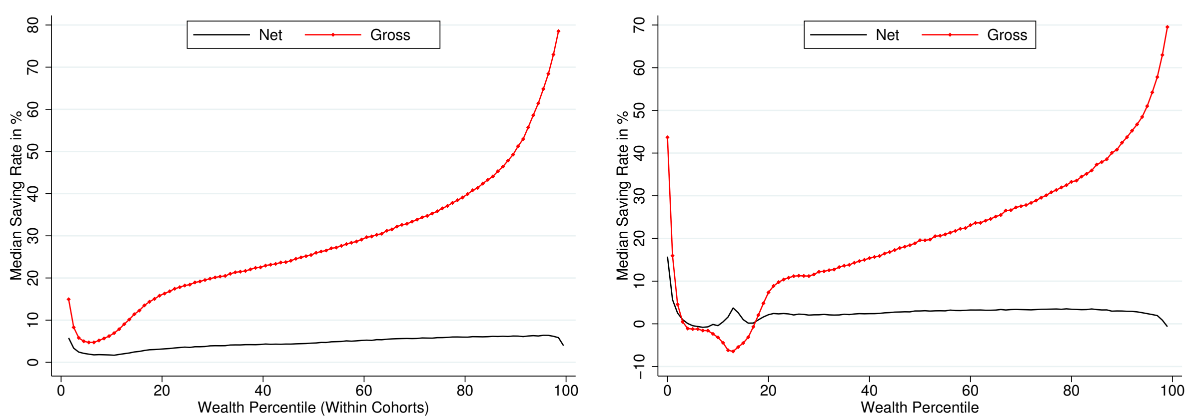

(a) Saving rates by within-cohort wealth percentiles (b) Controlling for age, earnings and education

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 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 (15) as a quantile regression for net and gross saving rates, respectively.

Figure 5: Saving rates adjusted for age, cohort, earnings and education.

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 4. Consistent with the discussion in Section 2.3, the gross saving rate is slightly flatter after controlling for the three variables. This indicates that some of the correlation between saving rates and wealth is due to age, earnings, and education. Still, the main takeaway is that even conditional on these observables, the saving graph maintains its main characteristics: a "swoosh-shaped" gross saving rate, and a remarkably flat net saving rate among households with positive net worth.

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 4. Figure 6 shows the median net and gross saving rates within groups (the results for education are in Appendix figure A15). Again we observe decreasing and then flat net saving rates, and "swooshes" for gross 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 2.3 suggests that persistent individual variation in saving rates might result in a positive correlation between (gross) saving rates and wealth. For example, models with discount rate heterogeneity predict that relatively patient households have higher saving rates and, over time, therefore become wealthier than impatient ones. Hence, a natural question is: do the gross 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 A15. 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.

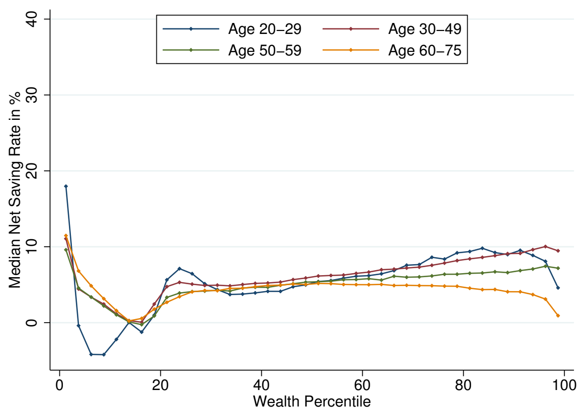

(a) Net saving rate by age group

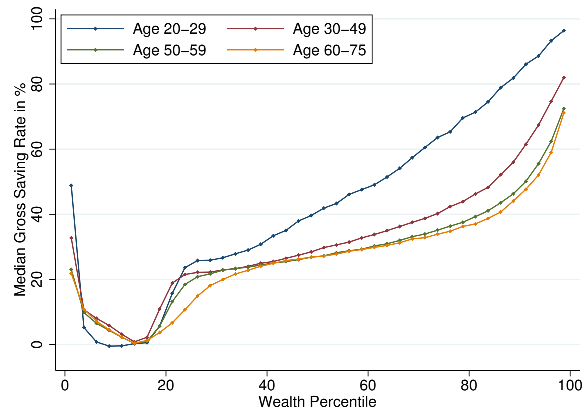

(b) Gross saving rate by age group

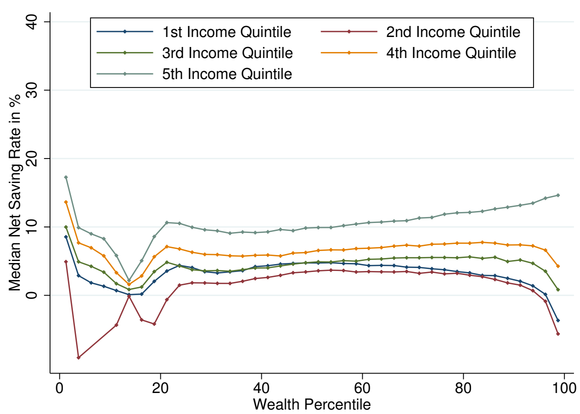

(c) Net saving rate by earnings group

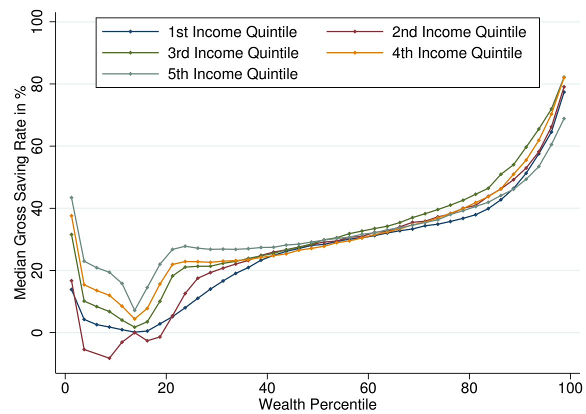

(d) Gross saving rate by earnings group

Notes: The figures display the median net saving rates (left column) and median gross 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).

Figure 6: Saving rates across the wealth distribution by age and earnings.

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 3). This fact naturally begs the question whether our findings in Figure 4 are mostly 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 3). For the purpose of this section, we therefore restrict attention 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.[28] 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 3.2). Hence, for stock funds, which our sample restrictions now lend relatively more importance to than in our baseline analysis above, the split of gross saving into net saving and capital gains becomes subject to some potential measurement error.

Figure 7(a) plots the "financial versions" of our two saving rates, against financial wealth for this restricted sample. The graphs display within-percentile medians, as in our baseline Figure 4. 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.

In Figure 7(a), we see the familiar pattern that the gross saving rates increase steeply with wealth, ranging from about two percent at the bottom to more than thirty percent at the top. Also akin to Figure 4, the net saving rate is considerably flatter than the gross saving rates. In contrast to Figure 4, 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.[29]

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

[29]In 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

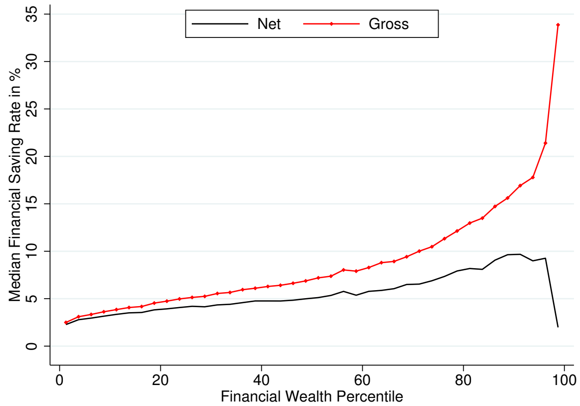

(a) Financial saving rates

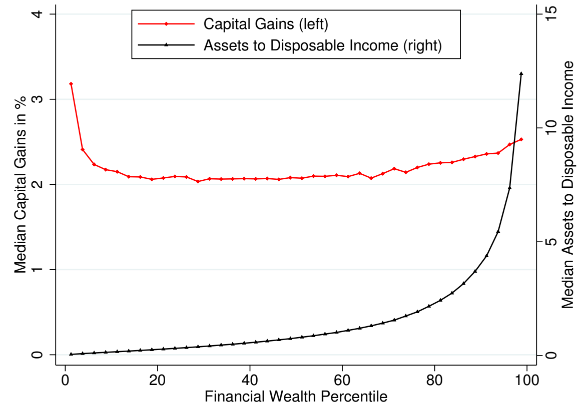

(b) Capital gains and financial assets to income

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).

Figure 7: Relation between financial saving rates and financial wealth.

In sum, the patterns we observed for net and gross saving rates do not seem limited to accumulation of housing wealth. As far as we can tell, under the limitation that Norwegian households hold relatively few assets with capital gains besides housing, the patterns are qualitatively similar for wealth accumulation in other assets too.

Zooming in on the Right Tail of the Wealth Distribution. We next turn our attention to the right tail of the wealth distribution. This subset of the population is particularly interesting because they hold a disproportionately large share of aggregate wealth, and because their asset portfolio is considerably less tilted toward housing.

Figure 8(a) shows that zooming in on the right 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 8(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 8(a) is driven by the fact that households at the very top hold relatively less wealth in assets that experience yearly

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.

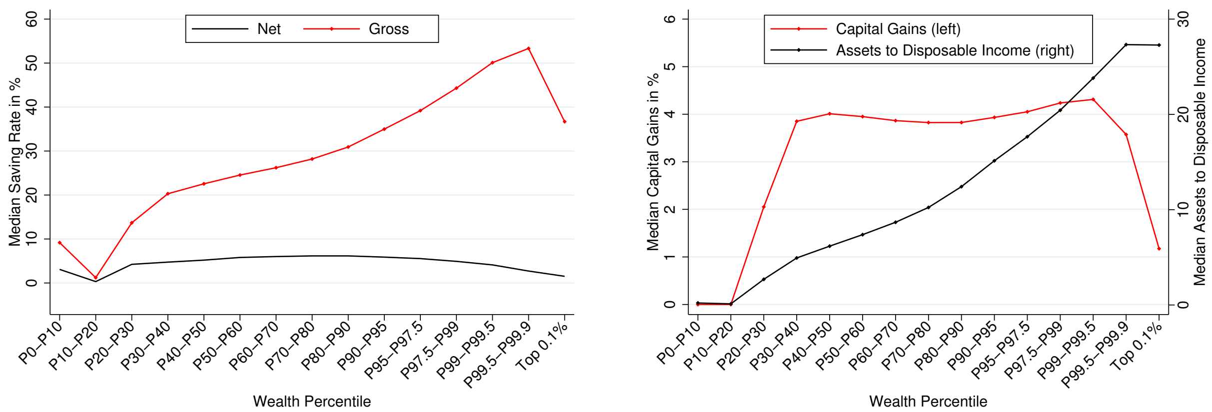

(a) Saving rates across the wealth distribution

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

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).

Figure 8: Saving behavior in the right tail of the wealth distribution.

capital gains and is consistent with saving by holding.[30]

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.

Public Pensions. Our first exercise is to include accumulation of 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 C.6. By including pension saving in Figure 9(a), 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 saving rates increase with wealth.

[30]The 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 3.1. Since private business wealth accounts for a large share of the assets held by the top 0.1% (see panel (a)), the observed capital gains in this group are likely an underestimate of the capital gains if private businesses were valued at market values.

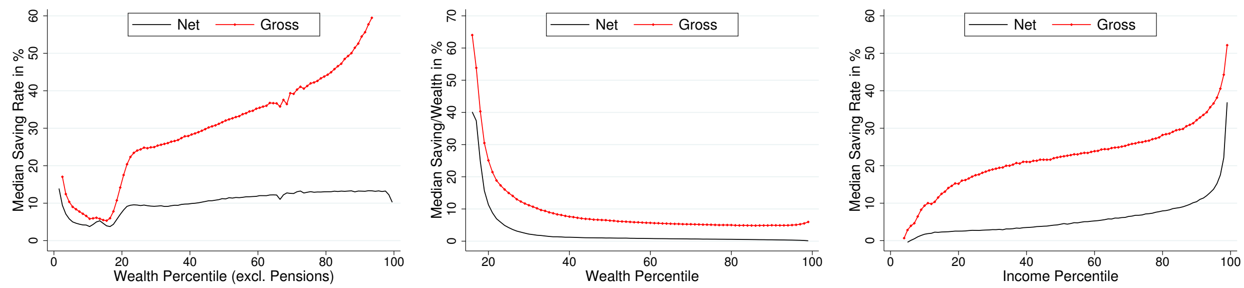

(a) Including public pensions (b) Saving as fraction of wealth (c) Saving rates by income

Notes: (a) displays the median saving rates when we include pension saving and income from the public pension system. Appendix C.6 describes the public pension system and how we compute pension wealth, income, and saving. (b) displays the median saving rate as a fraction of wealth. (c) displays the median saving rate by income. All variables are computed as the median within capital gains percentile and year, averaged across all years (2005-2015).

Figure 9: Additional exercises.

Saving Rates as a Fraction of Wealth. Because our objective is to test the saving behavior 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 9(b).[31] The figure reveals a similar pattern as Bach et al's Figure 4.[32] In results not reported in the paper, we have also plotted imputed consumption (defined as disposable income minus net saving – 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.[33]

[31] We 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.

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

[33] For 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 $a \to 0$, just like in Figure 9(b).

Saving Rates by Income. Figure 9(c) plots our two saving rates measures across the distribution of *income*, as opposed to *wealth* like in our main Figure 4. 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 relationship to expect, for example the permanent income hypothesis suggests that households will save high transitory income realizations.

5 Saving by Holding: Further Evidence and Alternative Explanations

As shown in Proposition 1, our empirical finding of a flat net saving rate and increasing gross saving rate is consistent with a model where households have homothetic preferences and capital gains are primarily accounted for by declining asset returns. Intuitively, when asset prices rise even though cashflows do not, richer households do not experience a larger income effect than poorer ones and therefore consume approximately the same fraction of their disposable incomes. Rich households instead optimally "save by holding" on to most of their assets experiencing capital gains. In this section we provide some additional evidence on this mechanism and thereafter review alternative explanations.

5.1 Further Evidence on Saving by Holding

Saving Rates by Capital Gains. Rather than plotting saving rates across the wealth distribution as we did in our main Figures 1 and 4, we here directly examine how saving rates vary with capital gains, i.e., to what extent households "save by holding." Our main Figures can then be viewed as a simple corollary since richer households hold more assets experiencing capital gains.

To this end, we examine how saving rates vary with capital gains as a share of Haig-Simons income (i.e., gross income including capital gains as defined in (8)). Besides having a number of convenient features (e.g., that it is bounded above by 100%), we choose this variable because Proposition 1 provides distinct predictions for how saving will vary with it. When rising asset prices are due to falling returns (case 3), we can rewrite (13) as

$$\frac{p_t\dot{k}_t}{w + r_t p_t k_t} = \phi_t - \phi_t \frac{\dot{p}_t k_t}{w + r_t p_t k_t} \tag{16}$$

for a number ϕ_t close to zero (ϕ_t is the net saving rate which we expect to take a number

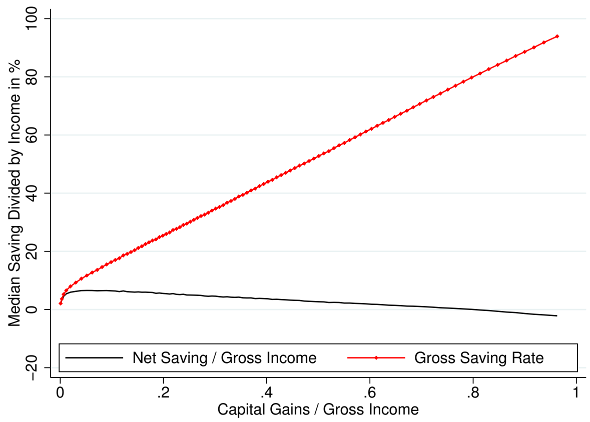

Notes: The figure displays the median saving rate by capital gains relative to gross income. All variables are computed as the median within capital gains percentiles and year, averaged across all years (2005-2015). The sample is restricted to households with positive capital gains to avoid issues with small or negative gross income.

Figure 10: Saving rates by capital gains.

Dependent variable: $\frac{\text{Net saving}}{\text{Gross income}}$		
	(1)	(2)
Constant	0.070 (0.00008)	
$\frac{\text{Capital gains}}{\text{Gross income}}$	-0.088 (0.00013)	-0.090 (0.00013)
Controls		X
Observations	19,090,221	19,090,221

Notes: The table displays coefficients from quantile regressions of net saving relative to gross income on capital gains relative to gross income for 2005-2015. The sample is restricted to households with positive capital gains to avoid issues with small or negative gross income. Controls include indicators for age, year, education, and income quintiles. Robust standard errors are reported in parentheses.

Table 2: Saving rates by capital gains.

like 5 or 10%). Our theory thus implies that net saving as a share of gross income should decrease linearly with capital gains as a share of gross income with a slightly negative slope. Equivalently, gross saving as a share of gross income (the gross saving rate) should increase linearly with capital gains as a share of gross income with a slope close to one.[34] In contrast, if rising asset prices instead came with rising dividends (case 2), the net saving rate should decrease strongly with capital gains (with a slope of minus one) and the gross saving rate should be flat.[35] We now explore these relationships directly in the data.

Figure 10 plots the empirical relations suggested by (16) (black line), applying the same procedure as in our main graphs earlier. The curve for net saving over gross income is approximately linearly decreasing in capital gains with a shallow slope, consistent with equation (16). Equivalently, the gross saving rate increases approximately linearly with capital gains, again consistent with our theory.

[34] To see this, rewrite (16) in terms of gross saving: $\frac{p_t \dot{k}_t + \dot{p}_t k_t}{w + r_t p_t k_t} = \phi_t + (1 - \phi_t)\frac{\dot{p}_t k_t}{w + r_t p_t k_t}$. Note that estimating this relationship rather than (16) by construction yields the same results (except that the slope is now $1 - \phi_t$ rather than $-\phi_t$). Results for regressions of gross saving relative to gross income on capital gains relative gross income corresponding to Table 2 and A6 are available upon request.

[35] To see this rewrite (11) and (12) as $\frac{p_t \dot{k}_t}{w + r p_t k_t} = \frac{r - \rho}{\gamma r} - \frac{\dot{p}_t k_t}{w + r p_t k_t}$ and $\frac{p_t \dot{k}_t + \dot{p}_t k_t}{w + r p_t k_t} = \frac{r - \rho}{\gamma r}$.

Next, we run the regressions suggested by (16), namely

$$\left(\frac{\text{Net saving}}{\text{Gross income}}\right)_{i,t} = \alpha_t + \beta_t \left(\frac{\text{Capital gains}}{\text{Gross income}}\right)_{i,t} + \epsilon_{i,t} \tag{17}$$

Besides the theoretical predictions we already discussed, note that (16) generated by our theory with declining returns makes the precise prediction that $\beta_t = -\alpha_t$. If returns were instead associated with dividend growth, the prediction from Proposition 1 is instead that $\beta_t = -1$ (see the discussion above, in particular footnote 35).

Table 2 shows the results from this regression when all data are pooled across years. Appendix Table A6 reports results from running (17) year-by-year. Column (1) in Table 2 shows that on average across years, the empirical relationship is close to the theoretical prediction when capital gains are associated with lower returns. The estimated value of α is 7% while the estimate for β is -8.8%. In columns (2) we add controls, so as to check if the estimate of β is driven by omitted variables. The estimate remains practically unchanged.[36] The year-by-year results in Table A6 give a similar picture. The estimated value of β_t deviates only slightly from $-\alpha_t$ each year. Moreover, the two parameter estimates move systematically together over time: when α_t is high, $-\beta_t$ is high too.

In sum, the observed relationships between net saving and capital gains align with the theoretical prediction from Proposition 1, in the case where asset price changes are associated with changes in returns.

Sources of Capital Gains. Our preferred theoretical interpretation postulates that capital gains are induced by falling returns rather than rising dividends. Therefore a natural next question is whether, empirically, this was the case for rising asset prices experienced by the households in our data. To examine this question, we consider various measures of price-to-rent ratios for housing, price-dividend ratios for public equity, as well as interest rates. Between 2005 and 2021, the risk-free interest rate declined from 2.17% to 0.52%, and simple benchmark values of the price-to-rent ratio for housing and the price-dividend ratio for stocks rose by 53% and 18% (as shown in Figure 3, housing is the much more important of the two asset classes).[37] These trends are consistent with the view that rising

[36] Note that when we add controls, the prediction $\beta = -\alpha$ from (16) no longer holds as the constant has a different interpretation. We therefore do not report the estimated values of the constant term when we add controls.

[37] The rationale for considering the period from the beginning of our sample (2005) to the present (2021) is that asset prices are forward-looking. The measure of the risk-free interest rate is the annualized one-month NIBOR interest rate. The price-to-rent ratio is calculated using the rental price index (from the consumer price index) and the aggregate house price index from Statistics Norway. The price-dividend ratio for stocks is the inverse of the dividend yield for the MSCI World stock index computed as the dividends accrued

asset prices experienced by households in our sample were primarily accounted for by declining returns – as required for our theory to be able to generate our empirical findings.

The idea that rising asset prices are primarily accounted for by declining returns might not come as a surprise for readers familiar with the finance literature. In particular, extensive research dating back to at least Campbell and Shiller (1988) has documented the importance of movements in discount rates for understanding asset-price movements.[38] While much of this research focuses on price variation at higher frequencies than our sample period, some studies also consider longer-run price growth. For instance, van Binsbergen (2020) studies international stock price growth over the past decades and argues it can be accounted for by declining interest rates.

5.2 Alternative Explanations of Saving by Holding

Portfolio Adjustment Frictions. Portfolio adjustment frictions are likely important for housing which is indivisible and subject to transaction costs (e.g. Grossman and Laroque, 1990; Kaplan and Violante, 2014; Kaplan, Moll and Violante, 2018; Guren et al., 2021). They may also be relevant for other asset classes. For instance, Fagereng, Gottlieb and Guiso (2017) rely on stock market adjustment costs to explain life-cycle saving patterns in Norway. Portfolio adjustment frictions can inhibit asset sales, thereby giving rise to inertia in response to asset price changes. Hence, they will contribute to saving-by-holding behavior. However, if rising asset prices were accompanied by higher cashflows, households with large asset holdings would still want consume out of these rising cashflows (the income effect we discussed in Section 2.2). Thus, even in models with transaction costs we will find that net saving tends to decline with asset holdings if capital gains are positive and associated with higher cashflows. This class of explanations can therefore explain our findings only under extreme parametric assumptions, such as high portfolio adjustment frictions. Appendix E uncovers these points further by applying such a model.

Non-homothetic Preferences. The canonical models' prediction that the systematic part of the gross saving rate is flat, follows 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.[39] Intuitively,

over the past 12 months divided by the average price in a month. The rationale for using the MSCI World stock index rather than a Norway-specific index is that Norwegians also hold non-Norwegian stocks.

[38] Also see Campbell and Cochrane (1999) and Bansal and Yaron (2004), for example.

[39] See e.g., Mian, Straub and Sufi (2021), Saez and Stantcheva (2018), Straub (2018), De Nardi (2004), Carroll (1998), and other papers cited in the review by Benhabib and Bisin (2018). Like Atkinson (1971),

relative to Figure 2(a), non-homothetic preferences rotate saving rates counter-clockwise. Hence, such theories can generate *gross* saving rates that are increasing with wealth even when asset price growth comes with rising cashflows. However, non-homotheticity does not generally predict that *net* saving rates are flat, as in Figure 4. Instead, a flat net saving rate will only result as a knife-edge case, for a specific degree of non-homotheticity.[40]

Returns to Wealth are Increasing with Wealth. Correlation between wealth and returns could in principle also play a role. Bach, Calvet and Sodini (2020) and Fagereng et al. (2016, 2020) document that wealthier households obtain higher returns to their wealth whereas we document that they have higher gross saving rates. It is natural to wonder if perhaps wealthy households have high gross saving rates because their access to higher returns gives stronger incentives for intertemporal substitution, a mechanism which holds irrespective of whether returns are associated with cash flows or not. Again, this effect may be part of the story but is unlikely to lie at its core. First, saving-by-holding implies that the gross saving rate will 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. Second, our exercise where we control for historical saving rates directly addresses this type of explanation, and shows that our main findings are not driven by heterogeneous innate inclination to save, for instance due to personal access to high returns. Third, and more generally, returns that are increasing with wealth will rotate model-implied saving rates counter-clockwise just like non-homothetic preferences, but would only predict a flat net saving rate as a knife-edge case.

Behavorial Explanations. It follows from our baseline theory in Proposition 1 that even if rising asset prices were associated with rising cash flows, households' net saving rates might still be flat if households do not perceive that this association exists. That is, if households believe asset price hikes are associated with lower returns, they will save as our theory predicts, irrespective of whether capital gains actually are driven by cash flows. Alternatively, households might not pay attention to asset prices altogether. (Sims, 2003; Gabaix, 2019), or 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

De Nardi (2004) assumes a warm-glow bequest motif with a different curvature than the period utility, a particular form of non-homotheticity.

[40]Consistent 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 4. 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 4.

Wurgler, 2007; Di Maggio, Kermani and Majlesi, 2019). These behavioral features are plausible, but do not, to our knowledge, generally predict that net saving rates are flat and capital gains are saved.

To sum up, alternative explanations of our findings exist, but they typically predict a flat net saving rate only as a knife-edge case.

6 Implications

Our results suggest that asset price changes associated with lower returns are important for understanding saving behavior across the wealth distribution. But how important is the observed saving behavior for aggregate outcomes, distribution, and ultimately welfare? We turn to these questions now.

6.1 Implications for Aggregate Saving and the Wealth Distribution

We here quantify how the following two factors have contributed to wealth accumulation in Norway: (i) heterogeneity in net saving rates and (ii) capital gains. First, we consider the evolution of aggregate wealth relative to income. Thereafter, we consider the evolution of the wealth distribution.

For both exercises we extend our dataset back to 1995.[41] We then consider two counterfactual scenarios. In the first counterfactual, we impose that every year, each percentile of the wealth distribution has the same median net saving rate instead of their actual net saving rate, i.e. we eliminate the (small) net-saving-rate heterogeneity in Figure 4.[42] This exercise isolates how important heterogeneity in net saving rates across the wealth distribution actually was from 1995 to 2015. In the other counterfactual exercise, we impose that capital gains are zero each year. This exercise isolates how important capital gains were. For all exercises, we start each household with their observed wealth in 1995.

[41] The main limitation in the data prior to 2005 is that we do not observe individual stock ownership at the security level and therefore use market returns. To extend the sample period, we 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 3.2 also prior to 2005.

[42] In each year, we rescale every household's net saving rate such that the within-percentile median saving rate is uniform across wealth percentiles, i.e. our counterfactual eliminates saving rate heterogeneity *across* percentiles (the heterogeneity in Figure 4) but not the heterogeneity *within* percentiles. More precisely, we first compute the within-percentile median net saving rate in our actual data, then compute the average of these within-percentile medians, and then eliminate the difference between each within-percentile median and this average by subtracting it from every household's saving rate.

Figure 11 plots the actual wealth-to-income ratio together with what it would have been in our two counterfactual scenarios. Over our sample period, the Norwegian aggregate wealth-to-income ratio increased from approximately 4 in 1995 to above 7 in 2015.[43] If we impose uniform net saving rates across the wealth distribution, we see that the trajectory remains almost unchanged. If we instead impose zero capital gains, we see that the trajectory changes completely. The counterfactual exercise suggests that without capital gains, the wealth-to-income ratio in 2015 would have been nearly the same as in 1995.

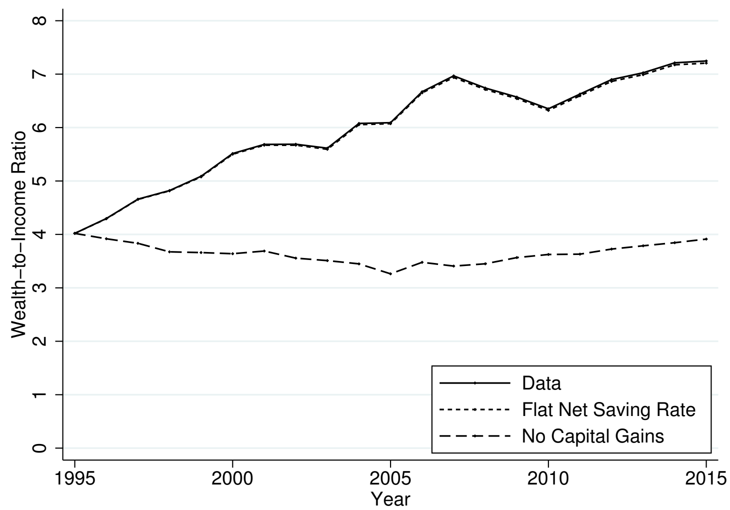

Notes: This figure displays the evolution of the ratio of net wealth to disposable income in the data and in the counterfactual scenario. We adjust each individual's capital gains and capital income such that they correspond to that individual's counterfactual wealth, assuming fixed portfolio shares.

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

We next turn to the wealth distribution. Figure 12 plots how three specific percentiles of the wealth distribution, the 10th, 50th or the 90th percentiles, actually evolved together with our two counterfactual scenarios. Panel (a) shows there are negligible effects of imposing that net saving rates are uniform across the wealth distribution. Hence, heterogeneity in net saving rates has not been important for the evolution of the wealth distribution. Panel (b) shows that matters are very different when it comes to capital gains. These have had minor effects on wealth levels at the 10th percentile of the wealth distribution, reflecting that households in this percentile have very low asset holdings. For

[43]These 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 3.1).

the median wealth level, capital gains are highly important. In the data, median wealth nearly tripled, but median wealth would have remained almost unchanged without capital gains. At the 90th percentile of the wealth distribution, wealth more than tripled in the data but would have less than doubled without capital gains.

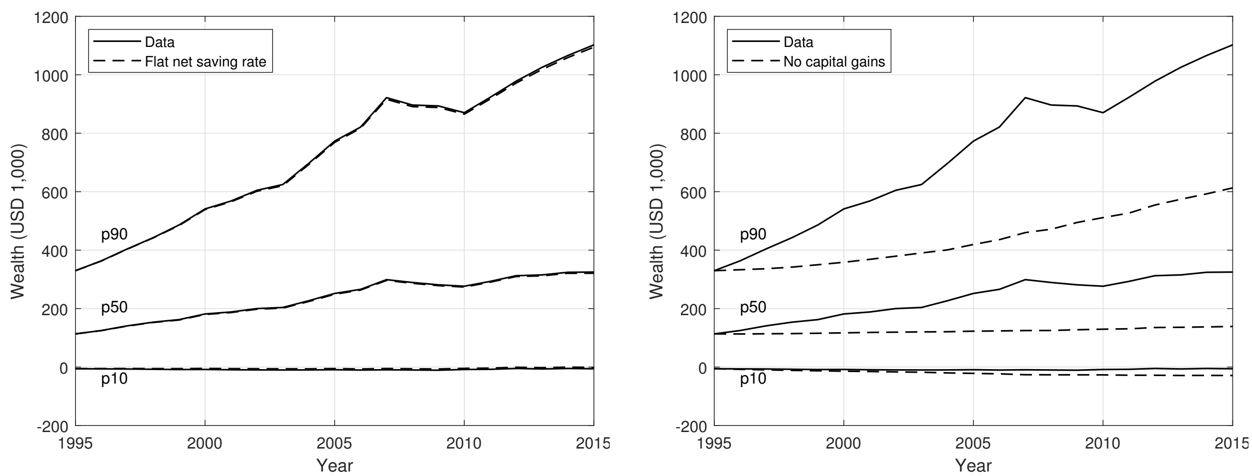

(a) Flat net saving rate (b) No capital gains

Notes: Panel (a) and (b) display the evolution of percentile thresholds in the wealth distribution over time. (a) shows the counterfactual where we set net saving rates equal across wealth percentiles. (b) displays the counterfactual where we set capital gains to zero.

Figure 12: Actual and counterfactual wealth percentiles.

To more precisely unveil how capital gains have affected the wealth distribution, Table 3 reports the actual wealth levels at several different percentiles for 1995 and 2015. In between those columns, we report what these percentiles would have been in 2015 without capital gains, according to our counterfactual scenario. We see that capital gains have elevated wealth levels throughout the distribution. At the bottom, households would have become poorer if asset prices had not increased. Further up, wealth levels would have increased also without capital gains, but not nearly to the same extent as they actually did. The contribution from capital gains to wealth levels is particularly high from the median and up to the 90th percentile, whereas it is more muted at the very top of the wealth distribution as illustrated by the top 1 percent in the table (though see the measurement issue concerning private business wealth discussed in footnote 30).

In terms of wealth inequality, the contribution from capital gains is subtle. If we consider measures of *relative* wealth, such as the percentile ratios in Table 3, they all peak in our counterfactual scenario without capital gains. This implies that asset price growth

	1995 Data	2015 No capital gains	2015 Data
Wealth Percentile Thresholds ($1,000)			
P1	-82	-173	-150
P10	-6	-29	-5
P25	29	12	78
P50	113	139	317
P75	207	325	635
P90	330	613	1,081
P99	686	2,040	2,418
Wealth Ratios			
P99/P90	2.08	3.33	2.24
P90/P50	2.91	4.42	3.42
P50/P25	3.93	11.26	4.08
Wealth Differences ($1,000)			
P99 - P90	357	1,427	1,337
P90 - P50	216	474	765
P50 - P25	84	127	239

Notes: The table reports the value of wealth, computed as the value of all assets minus all liabilities, at different percentiles of the wealth distribution. The columns named 'data' report the actual values for those years, whereas the column 'no capital gains' reports values in 2015 from the counterfactual scenario where there are no capital gains. Values are in USD, 2011 prices.

Table 3: Actual and counterfactual wealth percentiles.

has been a force toward *lower* wealth inequality in our sample period. On the other hand, if we consider differences in wealth *levels* we get nearly the opposite impression. Except at the very top of the wealth distribution, illustrated by the P99-P90 distance in Table 3, asset price growth has been a force toward *higher* wealth inequality over our sample period. While our finding that capital gains were a force towards lower relative wealth inequality may, at first, seem counterintuitive, it makes sense once we recall the importance of housing for Norwegian household balance sheets and particularly those outside the very top of the distribution (see Figure 3). Indeed, Kuhn, Schularick and Steins (2019) find that also in historical US data, and in line with our result, capital gains from housing have been a force toward lower wealth inequality.

In sum, we conclude that over our sample period, heterogeneity in net saving rates across wealth groups has been unimportant for changes in aggregate saving and the wealth distribution, whereas capital gains have been highly important.

6.2 Implications for Welfare

Our paper's main empirical finding is that net saving rates are approximately constant across the wealth distribution. At the same time, gross saving rates increase strongly with wealth because wealthier households own assets experiencing substantial capital gains and save by holding these. As we just showed, capital gains are therefore an important contributor to changes in aggregate saving and the wealth distribution. Our preferred theoretical interpretation was that the observed saving behavior simply reflects optimizing behavior in the face of rising asset prices that, however, do not come with increased cashflows but instead come with falling asset returns.

Some readers may be tempted to conclude from this discussion that capital gains induced by falling asset returns and the accompanying changes in wealth distribution should be viewed as welfare-irrelevant valuation effects, or "paper gains." This conclusion would be incorrect: while the welfare implications of such asset-price changes are subtle, they are certainly not zero. Versions of this point can be found in various contributions, going back to Paish (1940) and Whalley (1979) and more recently in Auclert (2019), Moll (2020) and Greenwald et al. (2021). For example, Moll (2020) uses a simple toy model to show that "investment plans" matter: whether investors benefits from rising asset prices depends not on the amount of assets they own, but whether they intend to buy, sell or keep their portfolios unchanged. Intuitively, rising asset prices hurt prospective buyers but benefit prospective sellers and hence redistribute from the former to the latter. We leave a more thorough exploration of these welfare effects for future work.

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 capital gains. We show that these findings are consistent with standard models of household wealth accumulation with homothetic preferences under one additional assumption: rising asset prices are accompanied by declining asset returns rather than rising dividends (cashflows).

Turning to the broader implications of our findings, we believe that both the macroeconomics and wealth inequality literatures need to consider changing asset prices with more care. Asset-price changes are not merely pesky “valuation effects” but a prevalent feature of the data: the majority of year-to-year movements in household wealth is due to asset price changes (i.e., capital gains or losses) rather than asset purchases or sales (net saving or dissaving). Fortunately, a nascent literature in macroeconomics and the study of inequality is starting to emphasize portfolio choice and asset price changes (Greenwald et al., 2021; Gomez and Gouin-Bonenfant, 2020; Feiveson and Sabelhaus, 2019; Kuhn, Schularick and Steins, 2019; Martínez-Toledano, 2019; Hubmer, Krusell and Smith, 2020; Gomez, 2018). We hope that our empirical findings and their theoretical interpretation 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

Aguiar, M. A., M. Bils, and C. Boar (2020): "Who Are the Hand-to-Mouth?," Working Paper 26643, National Bureau of Economic Research.

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. Johannesen, 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.

Auclert, A., and M. Rognlie (2016): "Inequality and Aggregate Demand," Working paper.

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 (2020): "Rich Pickings? Risk, Return, and Skill in the Portfolios of the Wealthy," *American Economic Review*, 110(9), 2703–2747.

Baker, M., S. Nagel, and J. Wurgler (2007): "The Effect of Dividends on Consumption," *Brookings Papers on Economic Activity*, 38(1), 231–292.

Bansal, R., and A. Yaron (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *The Journal of Finance*, 59(4), 1481–1509.

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.

Boerma, J., and L. Karabarbounis (2021): "Reparations and Persistent Racial Wealth Gaps," CEPR Discussion Papers 15796, C.E.P.R. Discussion Papers.

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.

Campbell, J. Y., and J. Cochrane (1999): "Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–251.

Campbell, J. Y., and R. J. Shiller (1988): "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1(3), 195–228.

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 (2021): "Stock Market Wealth and the Real Economy: A Local Labor Market Approach," *American Economic Review, forthcoming*.

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.

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.

Eika, L., M. Mogstad, and O. Vestad (2020): "What Can We Learn About Household Consumption from Information on Income and Wealth?," *Journal of Public Economics*, 189(104163).

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

Fagereng, A., C. Gottlieb, and L. Guiso (2017): "Asset Market Participation and Portfolio Choice over the Life-Cycle," *Journal of Finance*, 72(2), 705–750.

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.

——— (2020): "Heterogeneity and Persistence in Returns to Wealth," *Econometrica*, 88(1), 115–170.

Fagereng, A., M. B. Holm, and K. N. Torstensen (2020): "Housing Wealth in Norway, 1993–2015," *Journal of Economic and Social Measurement*, 45(1), 65–81.

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.

Gomez, M., and Émilien. Gouin-Bonenfant (2020): "A Q-Theory of Inequality," Working paper, Columbia University.

Greenwald, D., M. Leombroni, H. N. Lustig, and S. V. Nieuwerburgh (2021): "Financial and Total Wealth Inequality with Declining Interest Rates," Discussion paper.

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 (2021): "Housing Wealth Effects: The

Long View," *Review of Economic Studies*, 88(2), 669–707.

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 A. A. Smith (2020): "Sources of US wealth inequality: Past, present, and future," *NBER Macroeconomics Annual*, 35.

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

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.

Kaymak, B., D. Leung, and M. Poschke (2020): "Accounting for Wealth Concentration in the US," .

Kaymak, B., and M. Poschke (2016): "The evolution of wealth inequality over half a century: The role of taxes, transfers and technology," *Journal of Monetary Economics*, 77, 1–25.

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).

Kumhof, M., R. Rancière, and P. Winant (2015): "Inequality, leverage, and crises," *American Economic Review*, 105(3), 1217–45.

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.

Melcangi, D., and V. Sterk (2020): "Stock Market Participation, Inequality, and Monetary Policy," *FRB of New York Staff Report*, (932).

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

Mian, A. R., L. Straub, and A. Sufi (2021): "Indebted Demand," *Quarterly Journal of Economics, forthcoming*.

Moll, B. (2020): "Comment on "Sources of U.S. Wealth Inequality: Past, Present, and Future"," *NBER Macroeconomics Annual*, 35.

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.

Paish, F. W. (1940): "Capital Value and Income," *Economica*, 7(28), 416–418.

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

Rachel, L., and L. H. Summers (2019): "On Secular Stagnation in the Industrialized World,"

Brookings Papers on Economic Activity, 1(Spring), 1–76.
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.
SMITH, M., O. ZIDAR, AND E. ZWICK (2020): "Top Wealth in America: New Estimates and Implications for Taxing the Rich," Working Papers 264, Princeton University, Department of Economics, Center for Economic Policy Studies.
STRAUB, L. (2018): "Consumption, Savings, and the Distribution of Permanent Income," Working paper, Harvard.
VAN BINSBERGEN, J. H. (2020): "Duration-Based Stock Valuation: Reassessing Stock Market Performance and Volatility," NBER Working Papers 27367, National Bureau of Economic Research, Inc.
VON SCHANZ, G. (1896): "Der Einkommensbegriff und die Einkommensteuergesetze," *FinanzArchiv / Public Finance Analysis*, 13(1), 1–87.
WHALLEY, J. (1979): "Capital Gains Taxation And Interest Rate Changes: An Extension Of Paish's Argument," *National Tax Journal*, 32(1), 87–91.