

Asset-Price Redistribution

Andreas Fagereng

BI Norwegian Business School

Matthieu Gomez

Columbia University

Émilien Gouin-Bonenfant

Columbia University

Martin Holm

University of Oslo

Benjamin Moll

London School of Economics

Gisle Natvik

BI Norwegian Business School

Asset valuations across many asset classes have increased substantially over the past several decades. While these rising valuations had important effects on the distribution of wealth, little is known regarding their redistributive effects in terms of welfare. To make progress on this question, we develop a sufficient statistic for the money-metric welfare gain of deviations in asset valuations. This welfare gain depends on the present

Electronically published September 19, 2025

Journal of Political Economy, volume 133, number 11, November 2025.

© 2025 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/736769>

value of an individual's net asset sales rather than asset holdings: higher asset valuations benefit prospective sellers and harm prospective buyers. We estimate this quantity using panel microdata covering the universe of financial transactions in Norway from 1994 to 2019. We further demonstrate how to adapt our baseline statistic to account for important considerations, such as incomplete markets and collateral constraints. We find that the rise in asset valuations had large redistributive effects: it redistributed from the young to the old and from the poor to the wealthy.

I. Introduction

The past few decades have seen large increases in asset valuations across many asset classes.¹ These rising valuations had important effects on the distribution of wealth. This raises the questions: What are the welfare consequences of such asset price changes? Who wins and who loses from a rise in asset valuations?

One view is that any rise in asset prices represents a welfare-improving shift of resources toward the wealthy and should be taxed as such (e.g., Piketty and Zucman 2014; Saez, Yagan, and Zucman 2021).² An opposite view is that a rise in asset prices, without a corresponding rise in cash flows, simply generates “paper gains,” with no effect on actual income and therefore welfare (e.g., Cochrane 2020; Krugman 2021).³ Which (if any) of these two opposing views is correct?

To make progress on this question, we develop a sufficient statistic approach that quantifies the individual (money-metric) welfare effect of a

We wish to thank John Campbell, Jeanne Commault, Eduardo Dávila, Katya Kartashova, Camille Landais, John Leahy, Ian Martin, Clara Martínez-Toledano, Daniel Reck, Juliana Salomao, Lukas Schmid, and Stijn Van Nieuwerburgh for helpful discussions. Matthieu Gomez and Émilien Gouin-Bonenfant acknowledge support from the National Science Foundation under grant number SES-2117398. Benjamin Moll acknowledges support from the Leverhulme Trust and the European Union's Horizon 2020 research and innovation program (grant no. 865227). Andreas Fagereng, Martin Holm, and Gisle Natvik acknowledge support from the European Research Council under the European Union's Horizon 2020 research and innovation program (grant no. 851891). This work was edited by Andrew Atkeson.

¹ See, e.g., Farhi and Gourio (2018), Greenwald, Lettau, and Ludvigson (2019), or Van Binsbergen (2020) for empirical evidence.

² For example, Piketty and Zucman (2014) write: “Because wealth is always very concentrated . . . [a] high [wealth-to-income ratio] implies that the inequality of wealth, and potentially the inequality of inherited wealth, is likely to play a bigger role for the overall structure of inequality in the twenty-first century than it did in the postwar period. This evolution might reinforce the need for progressive capital taxation.”

³ Cochrane (2020) writes that “much of the increase in ‘wealth inequality’ . . . reflects higher market values of the same income flows, and indicates nothing about increases in consumption inequality.” Krugman (2021) discusses the hypothetical effect of declining interest rates on large fortunes in 19th-century England: “So since the ownership of land, in particular, was concentrated in the hands of a narrow elite, would falling interest rates and rising land prices have meant increased inequality? Clearly not. . . . The paper value of their estates would have gone up, but so what? The distribution of income wouldn't have changed at all.”

deviation in asset prices. We then operationalize this approach using Norwegian administrative panel data on asset transactions from 1994 to 2019 to quantify the redistributive effects of the rise in asset valuations over this time period.

We ask the following question: In monetary terms, how much does an individual value a deviation in the trajectory of asset prices, holding everything else (including asset cash flows) constant? The answer to this question is given by the following formula:

$$\text{Welfare Gain}_i = \sum_{t=0}^T R^{-t} \times \text{Sales}_{i,t} \times \text{Price Deviation}_{i,t}, \quad (1)$$

where i denotes the individual, T is the length of the sample period, $R > 1$ is a discount rate, $\text{Sales}_{i,t}$ are the net sales of the asset by individual i in year t , and $\text{Price Deviation}_{i,t}$ is the deviation of the price of the asset relative to a baseline scenario. In words, the welfare gain equals the net present value (NPV) of the trading profits due to the deviation in asset prices. The formula follows from applying the envelope theorem and thus holds for small price deviations, a point we discuss in more detail below. The welfare gain is in dollar terms and corresponds to the individual willingness to pay for the deviation in asset prices at time $t = 0$ (equivalent variation). The formula is for the case of one asset, but the extension to multiple assets is straightforward. Finally, this version of the formula abstracts from a number of important considerations such as incomplete markets and collateral constraints, which we take into account below.

Our formula for the welfare gains of asset-price changes ([1]) highlights that these welfare effects depend on asset transactions, not asset holdings. Intuitively, higher asset valuations are good news for prospective sellers (those with $\text{Sales}_{i,t} > 0$) and bad news for prospective buyers (those with $\text{Sales}_{i,t} < 0$). A particularly interesting case is an individual who owns assets but does not plan to buy or sell (i.e., $\text{Sales}_{i,t} = 0$). For such an individual, rising asset valuations are merely “paper gains,” with no effect on welfare.

It is useful to contrast these results with the two polar views described earlier. The first view posited that higher asset valuations redistribute toward existing asset holders. Our formula shows that it is sellers that benefit, not holders: if asset holders never sell, they do not benefit from the unrealized capital gains generated by the price deviation. The second view held that all (or at least most) of rising asset valuations are irrelevant for welfare. As our formula shows, this is only true if assets are not traded (e.g., in an economy with a representative agent). However, when heterogeneous individuals buy and sell assets as they do in the real world, fluctuations in asset prices do generate welfare gains and losses. In short, both views are incomplete.

As we show in the paper, the formula easily extends to multiple assets including bonds and long-lived assets subject to transaction costs (e.g., housing). Our key contribution is an empirical implementation of this welfare formula for the Norwegian economy. We compute welfare gains and losses due to the observed path of asset prices from 1994 to 2019 relative to a baseline where asset prices grew in tandem with dividends (i.e., relative to a balanced growth path as the baseline). Formally, we compute the relative price deviation in (1) as the relative difference between the actual price-dividend ratio PD_t and a baseline price-dividend ratio $\overline{PD}$:

$$\text{Deviation Price}_t = \frac{PD_t - \overline{PD}}{PD_t}. \quad (2)$$

For our application, we use the 1992–96 average price-dividend ratio as the baseline (i.e., a 5-year window around the beginning of the sample). Importantly, all of the variables in (1) and (2) are readily observable in our data. Price deviations in Norway have been particularly large for real estate (i.e., house prices have grown much faster than rents) and debt (i.e., real interest rates have declined sharply).

Our main findings are as follows. First, rising asset valuations have had large redistributive effects. While the average individual-level money metric welfare gain is around \$10,000, it is $-\$185,000$ at the 1st percentile and $\$273,000$ at the 99th percentile (in 2011 US dollars). As a fraction of total wealth (i.e., financial wealth plus human wealth), the average welfare gain is 0.0%, while it is -30% at the 1st percentile and 27% at the 99th percentile. Importantly, the distribution of welfare gains differs substantially from the distribution of revaluation gains (defined as the discounted sum of asset holdings times the changes in asset valuations), which are positive for almost everyone (and, in magnitude, equal to 16.4%, on average).

Second, we quantify the amount of redistribution across cohorts. Overall, we find a large amount of redistribution from young to old. For instance, the average welfare gain is approximately $-\$13,000$ for the cohorts aged 15 or younger at the end of 1993 (millennials) and around $\$22,000$ for the cohorts aged 30 and older at the end of 1993 (baby boomers). This intergenerational redistribution is primarily due to the fact that the young are net buyers of housing. Declining interest rates of mortgage debt offset the welfare losses of the young due to rising house prices but do so only partially.

Third, we quantify the amount of redistribution across the wealth distribution. We rank adults according to their total initial wealth (measured at the end of 1993) within cohorts and find that welfare gains have been concentrated at the top of the wealth distribution. The wealthiest 1% experienced, on average, a $\$73,000$ welfare gain, while the corresponding number is nearly zero at the 10th percentile, reflecting the fact that (perhaps

surprisingly) the wealthy tend to be net sellers of equity and borrowers. However, average welfare gains track total wealth almost one-for-one along most of the wealth distribution: the average welfare gain as a fraction of total wealth remains approximately constant from the 20th through the 80th percentile, at around 1.8%. This reflects that transactions are roughly proportional to wealth in that part of the wealth distribution.

Norwegian households trade not just with each other but also with the rest of the world and the government. We show that the net welfare gain of the household sector came at the expense of the Norwegian government, which, through the sovereign wealth fund, is a net saver. The government intertemporal budget constraint implies that Norwegian households will eventually have to bear the cost of this “government welfare loss” through lower future net transfers.

Our baseline welfare gain formula is derived in a deterministic model without borrowing constraints. Taking advantage of the envelope theorem’s flexibility, we consider several model extensions and explain how they affect our formula. Building on these theoretical results, we then empirically implement a version of our sufficient statistic to address what we view as the most important omissions of our baseline empirical exercise: borrowing constraints with collateral effects, incomplete markets, second-order effects from the large observed asset-price changes, and valuation changes beyond the end of our sample period. These generalizations affect our estimated welfare gains and losses quantitatively but not qualitatively. More specifically, considering incomplete markets and valuation changes beyond the end of our sample tends to dampen the welfare loss of young generations. We also discuss the interpretation of our sufficient statistic in more general environments, particularly when asset prices are determined in general equilibrium.

Literature.—Our paper contributes to several strands of literature. In recent decades, there has been a sustained rise in valuations across many asset classes (e.g., Piketty and Zucman 2014; Farhi and Gourio 2018; Greenwald, Lettau, and Ludvigson 2019). As a response to this trend, a growing literature focuses on understanding the effect of rising asset prices (and declining interest rates) on wealth inequality (e.g., Catherine, Miller, and Sarin 2020; Kuhn, Schularick, and Steins 2020; Cioffi 2021; Greenwald et al. 2021; Wolff 2022; Gomez and Gouin-Bonenfant 2024; Gomez 2025). Relative to this literature, our contribution is to study the heterogeneous effect of rising asset prices on welfare.⁴ More broadly, we contribute to a large literature that uses microdata to study the heterogeneity in saving

⁴ Our theoretical results build on Moll (2020), who studied a two-period model similar to that described in sec. II.A. Our result that the welfare of an individual who never buys or sells an asset is unaffected by a change in asset price is related to (but different from) a result by Sinai and Souleles (2005) that an individual with an infinite expected residence spell is insulated from house price risk.

and portfolio choices over the life cycle (e.g., Berger et al. 2018; Feiveson and Sabelhaus 2019; Calvet et al. 2021; Black et al. 2022) and along the wealth distribution (e.g., Bach, Calvet, and Sodini 2017, 2020; Fagereng et al. 2019; Mian, Straub, and Sufi 2020).

Dávila and Korinek (2018) study the externalities associated with asset-price fluctuations in economies with financial frictions. In this context, they obtain a similar formula as our sufficient statistic for the welfare effect of small asset-price changes (see lemma 1 in that paper). We generalize this expression along empirically relevant dimensions (e.g., more than three time periods, intergenerational linkages, the government sector, and financial transactions done via businesses), we develop a methodological framework to measure these welfare effects at the individual level, and we implement it using household-level transaction data. Dávila and Korinek (2018) show that deviations in asset prices generate two types of externalities: distributive externalities (when agents do not equate their marginal rates of substitutions across states or times) and collateral externalities (when asset prices matter for financial constraints). Building on these two insights, we stress that, while our baseline measures of welfare gains aggregate to zero in the population, they no longer do when we modify the formula to take into account incomplete markets (sec. V.A) or collateral constraints (sec. V.B).

Our formula for welfare gains is also related to Auclert (2019), who derives the welfare and consumption effects of deviations in interest rates. Relatedly, Greenwald et al. (2021) stress that the welfare effect of a permanent decline in interest rates can be measured as the duration mismatch between consumption and income, which they estimate using US data. While there is a profound connection between the two approaches, our sufficient statistic has two main advantages for our application.⁵ First, it allows us to consider the welfare effect of arbitrary valuation changes across asset classes rather than the ones induced by a uniform shift in discount rates in all asset classes. Second, it allows us to measure welfare gains using financial transactions, which we observe directly, rather than in terms of the path of consumption and income, which is typically harder to observe. Finally, our focus on the heterogeneous welfare effect of asset-price fluctuations connects this paper to Doepke and Schneider (2006), who estimate the redistributive effect of inflation episodes using data from the Survey of Consumer Finances, as well as Kiyotaki, Michaelides, and Nikolov (2011) and Glover et al. (2020), who study the redistributive effect of asset-price fluctuations using calibrated models.

More generally, our paper is related to a large asset pricing literature on the role of discount rate shocks. One key finding in the literature is that

⁵ We discuss the precise mapping between the two approaches in app. sec. E.4 (apps. A–E are available online).

discount rate shocks account for most asset-price fluctuations (Shiller 1981; Campbell and Shiller 1988). The distinction between cash flow and discount rate shocks has important implications for portfolio allocation (e.g., Merton 1973; Campbell and Viceira 2002; Campbell and Vuolteenaho 2004; Catherine et al. 2022). Relative to these papers, we examine the effect of discount rate shocks on welfare, both theoretically and empirically.

Finally, our emphasis that rising asset valuations benefit sellers and not asset holders has some historical precedent in the works of Paish (1940), Kaldor (1955), and Whalley (1979) who were, in turn, part of a debate in the public finance literature whether unrealized capital gains are a form of income and should therefore be taxed (Haig 1921; Simons 1938).⁶

Road map.—This paper is organized as follows. In section II, we present our theoretical framework to quantify the welfare effect of a deviation in asset prices. In section III, we discuss the implementation of our sufficient statistic approach using administrative data from Norway. In section IV, we report our estimates for the redistributive effects of asset-price changes. Finally, we discuss generalizations of our sufficient statistic approach in section V.

II. Theoretical Framework

This section presents our sufficient statistic approach. To focus on the intuition, we first examine the welfare effect of asset-price deviations in a two-period model with only one asset in section II.A. We then generalize the result to an infinite horizon model with multiple assets and adjustment costs in section II.B. We then discuss some extensions of our results in section II.C to more general models.

A. Intuition in a Two-Period Model

Time is discrete with two time periods $t = 0, 1$. Individual i receives labor income $Y_{i,0}$ at time 0 and $Y_{i,1}$ at time 1. There is one asset available for trading at time $t = 0$ with price $P_0 > 0$, which pays a dividend $D_1 > 0$ at time 1. Individuals have time-separable preferences with a differentiable utility function $U(\cdot)$ that is increasing and strictly concave and with a subjective discount factor $\beta < 1$.

Individual problem.—Denote by $C_{i,t}$ the consumption of individual i at time t and $N_{i,t}$ the number of shares owned at the end of period t . Given

⁶ For example, Kaldor (1955) writes: “We may now turn to the other type of capital appreciation which [comes] without a corresponding increase in the flow of real income accruing from that wealth . . . [insofar] as a capital gain is realized and spent . . . the benefit derived from the gain is equivalent to that of any other casual profit. If however it is not so realized, there is clearly only a smaller benefit.”

initial asset holdings $N_{i,-1}$, the problem of the individual is to choose consumption and asset holdings to maximize utility,

$$V_{i,0} \equiv \max_{\{C_{i,0}, C_{i,1}, N_{i,0}\}} U(C_{i,0}) + \beta U(C_{i,1}), \quad (3)$$

subject to the following budget constraints:

$$\begin{aligned} C_{i,0} + (N_{i,0} - N_{i,-1})P_0 &= Y_{i,0}, \\ C_{i,1} &= N_{i,0}D_1 + Y_{i,1}. \end{aligned} \quad (4)$$

These budget constraints say that in each period t , consumption plus net asset purchases (the left-hand side) must equal income (the right-hand side).⁷

Welfare effect.—Consider a small change in the price of the asset at time $t = 0$, holding everything else constant. We are interested in the welfare gain associated with this change in asset prices, which we define as the amount of money that would have an equivalent effect on individual welfare (“equivalent variation”). For brevity, we will simply refer to this quantity as “welfare gain” in the rest of the paper, but it is important to keep in mind that it is a money metric.⁸ For an infinitesimal price change dP_0 , the welfare gain simply corresponds to the change in welfare $dV_{i,0}$ scaled by the marginal utility of consumption $U'(C_{i,0})$. Applying the envelope theorem yields the following expression for the welfare gain:⁹

$$\frac{dV_{i,0}}{U'(C_{i,0})} = (N_{i,-1} - N_{i,0})dP_0. \quad (5)$$

The effect of a rise in P_0 is given by the extent to which it relaxes the budget constraint at $t = 0$, namely, asset sales $N_{i,-1} - N_{i,0}$. More precisely, a rise in the price of the asset benefits individuals who plan to sell the asset (i.e., $N_{i,0} < N_{i,-1}$) and hurts individuals who plan to buy the asset (i.e., $N_{i,0} > N_{i,-1}$). Importantly, a rise in the price of the asset does not affect individuals who do not plan to trade (i.e., $N_{i,0} = N_{i,-1}$): for those individuals, the rise in the price of the asset is merely a “paper gain” with no corresponding effect on welfare. Similar expressions for the welfare effect of asset prices were previously obtained by Dávila and Korinek (2018) and Moll (2020) in similar two- and three-period environments.

⁷ Recall that the environment has only two periods, with no market for transactions at time $t = 1$ (alternatively, the price of the asset is zero at $t = 1$). We consider the multiperiod case below, in which we add appropriate terminal conditions.

⁸ Consistently with standard consumer theory, we focus on a money-metric measure of welfare to respect the notion that preferences are ordinal, rather than cardinal, in nature (Mas-Colell, Whinston, and Green 1995; Baqaee and Burstein 2023). Our measure is similar to the welfare-equivalent increase in consumption defined in Lucas (2000). (For more details, see proposition A13 [propositions A1–A15 are available online].)

⁹ See app. sec. E.1 for the explicit derivation.

Welfare versus revaluation gains.—The result in equation (5) may be surprising at first. How can an asset holder not benefit from a price rise given that the market value of their initial wealth unambiguously increases? The reason is that we consider a rise in P_0 holding everything else (in particular, the dividend of the asset D_1) constant.¹⁰ While a rise in P_0 increases the return of holding the asset at time $t = 0$, it simultaneously decreases the return of holding the asset at $t = 1$. On net, only individuals whose holdings decline over time (i.e., sellers) end up benefiting from the rise in asset price.

To see this formally, denote R_t the return of the asset at time t ; that is, $R_0 = P_0/P_{-1}$ and $R_1 = D_1/P_0$. Note that a rise in P_0 increases R_0 , via a higher capital gain, but decreases R_1 , via a lower dividend yield:

$$\frac{dR_0}{dP_0} = 1/P_{-1} > 0, \quad \frac{dR_1}{dP_0} = -R_1/P_0 < 0. \quad (6)$$

The welfare gain due to this change in asset returns can then be written as¹¹

$$\begin{aligned} \frac{dV_{i,0}}{U'(C_{i,0})} &= \underbrace{N_{i,-1} P_{-1} \times dR_0}_{\text{contribution of return at } t=0} + \underbrace{R_1^{-1} N_{i,0} P_0 \times dR_1}_{\text{contribution of return at } t=1} \\ &= N_{i,-1} dP_0 - N_{i,0} dP_0, \end{aligned} \quad (7)$$

where the second line is obtained via (6). This alternative derivation highlights that the welfare effect (5) can be seen as the sum of two terms: the first term $N_{i,-1} dP_0$ accounts for the positive effect of a rise in P_0 on today's return (via a higher capital gain), while the second term, $-N_{i,0} dP_0$, accounts for the negative effect of a rise in P_0 on tomorrow's return (via a lower dividend yield). For an individual who does not trade, the two terms offset each other; as a result, a change in asset prices has no welfare effect. We will illustrate the difference between the welfare effect of a deviation in asset prices—the left-hand side of equation (7)—and its revaluation effect—the first term on the right-hand side of equation (7) in our empirical application.

As we discuss in more detail in our multiperiod model, when dividend income D_1 rises as well, it remains true that a rising asset price P_0 benefits

¹⁰ To put this more precisely, it is useful to adopt the asset pricing perspective that the asset price at $t = 0$ is the present discounted value of future cash flows: $P_0 = D_1/R_1$ where R_1 is the asset required rate of return, which we take as exogenous. An increase in the price P_0 without a change in the dividend D_1 is equivalent to a fall in the required rate of return R_1 . We develop this general point in app. sec. E.4.

¹¹ This follows from rewriting the budget constraints (4) as $C_{i,0} + A_{i,0} = R_0 A_{i,-1} + Y_{i,0}$ and $C_{i,1} = R_1 A_{i,0} + Y_{i,1}$, where $A_{i,t} \equiv N_{i,t} P_t$, and using the envelop theorem to compute the effect of deviations in R_0 and R_1 on the value function.




FIG. 1.—Welfare effect of an increase in the asset price P_0 on the welfare of a seller (A) and that of a buyer (B). The red lines represent the agent's present-value budget constraints ($4''$), which go through the endowment points $C_{i,0} = Y_{i,0}$ and $C_{i,1} = Y_{i,1} + N_{i,-1}D_1$ and have slope $-D_1/P_0$. In both panels, the solid budget constraint and indifference curve correspond to the allocation at the initial asset price, and the dotted lines are those at the new, higher price. When the asset price P_0 increases, the budget constraint slope $-D_1/P_0$ flattens, rotating through the endowment point. The seller's welfare increases (A) and the buyer's welfare decreases (B).

sellers and not holders. However, there is now an additional effect: a rise in dividend income directly benefits all asset holders. Equivalently, it offsets the decline in the dividend yield and hence the return $R_1 = D_1/P_0$.

Graphical intuition.—Building on Whalley (1979), figure 1 presents a graphical intuition for the welfare consequences of asset-price changes based on the Fisher diagram, the standard graphical apparatus for intertemporal consumption choice problems. The red line represents the present-value budget constraint of the agent's problem, with slope $-D_1/P_0$, while the black curve represents the agent's indifference curve.¹²

Consider the welfare consequences of a rise in the asset price P_0 for a hypothetical seller (panel A) and buyer (panel B). When the asset price P_0 rises, the budget constraint rotates through the endowment point and becomes flatter (the slope is $-D_1/P_0$). The figure shows that the seller ends up on a higher indifference curve (increase in welfare), whereas the buyer ends up on a lower indifference curve (decrease in welfare).¹³

¹² More precisely, the present value budget constraint for the agent problem is given by

$$C_{i,0} + \frac{P_0}{D_1} C_{i,1} = Y_0 + \frac{P_0}{D_1} Y_{i,1} + N_{i,-1}P_0. \quad (4)$$

¹³ In fact, our notion of money-metric welfare gain corresponds, at the first order, to the horizontal distance between the initial $C_{i,0}$ and the new budget line (as indicated by the solid arrows), as this distance measures the extent to which $C_{i,0}$ would need to adjust if $C_{i,1}$ was held constant: $\Delta C_{i,0} = (N_{i,-1} - N_{i,0})\Delta P_0$.

B. Baseline Model

We now extend this simple intuition to an infinite horizon deterministic economy with multiple assets and adjustment costs (hereafter, the “baseline model”), which is key to bringing the theory to the data.

Financial markets.—There is a sequence of liquid one-period bonds with a face value of one and price $Q_t > 0$ available for trading. Purchasing a one-period bond is equivalent to investing in a deposit account with an interest rate $R_{t+1} = 1/Q_t$ between time t and $t + 1$. Denote by $R_{0 \rightarrow t} = R_1 \cdot R_2 \cdots R_t$ the cumulative return of these one-period bonds between time 0 and t . There are also K long-lived assets available for trading (e.g., housing, stocks, private businesses, or long-term bonds). A share of asset $1 \leq k \leq K$ is a claim to a stream of dividends $\{D_{k,t}\}_{t=0}^{\infty}$, with price $P_{k,t}$ at the end of period t . The return of asset k between t and $t + 1$ is thus $R_{k,t+1} \equiv (D_{k,t+1} + P_{k,t+1})/P_{k,t}$.

We assume that trading these long-lived assets is subject to adjustment costs, which may be large or small depending on the asset. These adjustment costs capture that some assets, such as houses and privately traded equity, are illiquid. For other assets, such as publicly traded equity, the adjustment costs—which may be arbitrarily small—are instead a technical assumption required to have well-defined asset-demand functions in a deterministic economy.¹⁴ We assume that the adjustment costs, denoted $\chi_k(\cdot)$, are continuous functions of the number of assets purchased each period. Still they can be kinked (nondifferentiable) to capture infrequent adjustment and inaction regions (as in Bertola and Caballero [1990] or Kaplan, Moll, and Violante [2018]).

Individual problem.—Individuals have time-separable preferences with a differentiable utility function $U(\cdot)$ that is increasing and strictly concave and a subjective discount factor $\beta \in (0, 1)$. They receive labor income $Y_t > 0$ at time t , and they can trade financial assets: we denote by B_t the holdings of the one-period bond and by $N_{k,t}$ those of asset k at the end of period t . Individuals take asset prices as given and choose a path of consumption and asset holdings to maximize utility

$$V_{i,0} \equiv \max_{\{C_{i,t}, B_{i,t}, \{N_{i,k,t}\}_k\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t U(C_{i,t}), \quad (8)$$

subject to initial asset holdings $B_{i,-1}$ and $\{N_{i,k,-1}\}_k$, as well as a sequence of budget constraints

¹⁴ These adjustment costs are no longer necessary when the economy is stochastic (if assets have heterogeneous risk profiles, see app. sec. A.2) or when agents have nonmonetary benefits of owning certain assets (e.g., owning a house vs. renting it, see app. sec. A.4.1).

$$\begin{aligned}
C_{i,t} + \sum_{k=1}^K (N_{i,k,t} - N_{i,k,t-1})P_{k,t} + B_{i,t}Q_t + \sum_{k=1}^K \chi_k (N_{i,k,t} - N_{i,k,t-1}) \\
= \sum_{k=1}^K N_{i,k,t-1}D_{k,t} + B_{i,t-1} + Y_{i,t}.
\end{aligned} \tag{9}$$

The budget constraint says that consumption plus net purchases of financial assets (the left-hand side) must equal total income in each period t (the right-hand side), which is the sum of dividend, interest, and labor income.

Because of the infinite-horizon setup, we also assume the following technicality conditions: a bound on asset holdings $N_{i,k,t} \in \Theta_k$, where Θ_k is compact; a lower bound on the price of one-period bonds $\liminf_{T \rightarrow \infty} Q_T > 0$; a no-bubble condition $\lim_{T \rightarrow \infty} R_{0 \rightarrow T}^{-1} P_{k,T} = 0$; and a no-Ponzi condition $\lim_{T \rightarrow \infty} R_{0 \rightarrow T}^{-1} (B_{i,T}Q_T + \sum_{k=1}^K N_{i,k,T}P_{k,T}) \geq 0$.¹⁵ Finally, we assume that there exists a unique solution $\{C_{i,t}, B_{i,t}, \{N_{i,k,t}\}_{k=1}^K\}_{t=0}^\infty$ and that it is continuous with respect to asset prices.

Welfare effect.—We are interested in the welfare effect of a small perturbation in the path of asset prices. Formally, we consider an infinitesimal deviation of the path of asset prices, denoted by $\{dQ_t, \{dP_{k,t}\}_{k=1}^K\}_{t=0}^\infty$, holding everything else constant.¹⁶ We assume that the deviation does not explode over time, that is, that it satisfies the no-bubble condition $\lim_{T \rightarrow \infty} R_{0 \rightarrow T}^{-1} dQ_T = \lim_{T \rightarrow \infty} R_{0 \rightarrow T}^{-1} dP_{k,T} = 0$. As in the case of the two-period model above, we define the welfare gain of the deviation as the amount of money received at $t = 0$ that would generate an equivalent change in individual welfare (equivalent variation). For an infinitesimal deviation, it corresponds to the deviation in welfare $dV_{i,0}$ scaled by the initial marginal utility of consumption $U'(C_{i,0})$.¹⁷

PROPOSITION 1 (Welfare gain). The welfare gain implied by a price deviation $\{dQ_t, \{dP_{k,t}\}_{k=1}^K\}_{t=0}^\infty$ is

$$\frac{dV_{i,0}}{U'(C_{i,0})} = \sum_{t=0}^\infty R_{0 \rightarrow t}^{-1} \left(\sum_{k=1}^K (N_{i,k,t-1} - N_{i,k,t}) dP_{k,t} - B_{i,t} dQ_t \right). \tag{10}$$

The proposition, proved in appendix section A.1, says the welfare gain corresponds to the present value of the deviation in trading profits induced by the deviation in the path of asset prices. As in the two-period model, the welfare gain of a deviation in asset prices depends on financial transactions rather than holdings. Note, however, that for the liquid asset,

¹⁵ Appendix sec. E.3 discusses the implied present-value budget constraint, whereas app. sec. A.4.2 discusses the finite-horizon case.

¹⁶ This deviation can be seen as a comparative statics on the type of the economy the agent is born in, or, equivalently, as the realization of an unexpected “MIT” shock.

¹⁷ In proposition A13, we state and prove a list of alternative interpretations of our concept of welfare gains. In particular, it also corresponds to the present value of the change in individual consumption in response to the deviation in asset prices.

transactions and holdings coincide, given that the asset must be continuously rolled over. Thus, declining interest rates (i.e., $dQ > 0$) benefit individuals holding short-term debt (i.e., $B_{i,t} < 0$) because lower debt payments relax their budget constraint. Finally, note that the adjustment-cost function does not appear in the welfare formula as a consequence of the envelope theorem.¹⁸

The thought experiment of proposition 1 corresponds to a pure deviation in asset prices, that is, holding dividend and labor income fixed. In the financial literature, this is often described as a deviation in asset discount rates. We formalize the mapping between deviations in asset prices and deviations in asset discount rates in appendix section E.4. We also discuss the connection between our formula and the ones obtained by Auclert (2019) and Greenwald et al. (2021), who study the welfare effect of changes in interest rates, in appendix section E.5.

Aggregation.—One implication of proposition 1 is that welfare gains aggregate to zero in an economy composed of households trading with each other. Formally, indexing households by $i = 1, \dots, I$, $\sum_{i=1}^I (N_{i,k,t-1} - N_{i,k,t}) = 0$ for all k and $\sum_{i=1}^I B_{i,t} = 0$ implies $\sum_{i=1}^I dV_{i,0}/U'(C_{i,0}) = 0$. This property reflects that, for every seller that benefits from a rise in asset prices, there is a buyer that is equally hurt (in monetary terms), so asset-price deviations are purely redistributive. While this result is important to keep in mind, two remarks are in order. First, the fact that welfare gains aggregate to zero says nothing about the desirability of asset-price deviations from the point of view of a social planner, who may assign different weights to the value of additional dollars for various individuals. More specifically, the effect of a price deviation on social welfare can be positive or negative, depending on whether the welfare weights assigned by the planner to individuals covary positively or negatively with individual welfare gains.¹⁹ Second, this result hinges on two important facts in our baseline economy: (i) agents equalize their marginal rates of substitutions across states/times, and (ii) asset prices do not appear in the agent problems outside of their budget constraints. In section V, we will relax these assumptions by considering economies with uninsurable shocks and/or borrowing constraints with collateral effects, in which cases, welfare gains no longer aggregate to zero.²⁰

¹⁸ To apply the envelope theorem, we assumed that the solution of the optimization problem was locally continuous with respect to prices. While this does not rule out kinked adjustment costs (as in Kaplan, Moll, and Violante 2018), this does rule out adjustment-cost functions that lead to discrete adjustments in response to infinitesimal price changes. Finally, while the particular functional form for χ_t does not matter for the first-order effect of asset-price deviations on welfare, it would matter for higher-order effects, as discussed in sec. V.

¹⁹ We will emphasize this point in sec. IV by aggregating individual welfare gains with different sets of welfare weights.

²⁰ In the language of Dávila and Korinek (2018), the welfare gains of a deviation in asset prices no longer aggregate to zero in the presence of distributive externalities (when

Deviation in dividend and labor income.—We can easily extend our proposition to compute the welfare effect of a joint deviation in asset prices, dividend income, and labor income. Proposition A1 expresses the resulting welfare effect as

$$\frac{dV_{i,0}}{U'(C_{i,0})} = \sum_{t=0}^{\infty} R_{0 \rightarrow t}^{-1} \left(\underbrace{\sum_{k=1}^K (N_{i,k,t-1} - N_{i,k,t}) dP_{k,t} - B_{i,t} dQ_t}_{\text{effect of asset-price changes}} + \underbrace{\sum_{k=1}^K N_{i,k,t-1} D_{k,t} + dY_{i,t}}_{\text{effect of income changes}} \right). \quad (11)$$

Relative to our baseline formula, the expression for welfare gains is augmented with an additional term: the present value of the deviation in income. This equation emphasizes the key distinction between a deviation in asset prices and a deviation in asset income: while only asset sellers benefit from a rising asset price, all asset holders benefit from a rise in dividend income. This formula is helpful to quantify the redistributive effect of arbitrary shocks to the economy, which typically jointly affect income and asset prices in equilibrium. To give a concrete example, in appendix section A.3, we use the formula to analyze the redistributive effect of productivity shocks in a general equilibrium production economy through its impact on income and asset prices.²¹

C. Extensions

The baseline model is deliberately stylized and abstracts from several potentially important features of the real world. Before we bring our theory to the data, we consider a number of model extensions. In the rest of this section, we briefly summarize how the extension affects our welfare gain formula, (10), as well as its interpretation.

Stochastic environment.—So far, we have focused on deterministic economies. In reality, individuals do not have perfect foresight over the future. In appendix section A.2.1, we show that, in this case, the welfare gain of a deviation in asset prices (i.e., the amount of money that would generate the same increase in welfare from an ex ante perspective at $t = 0$) is modified along two dimensions. First, what matters is the expectation of future

agents do not equal their marginal rates of substitutions across dates or times) and/or collateral externalities (when asset prices matter for financial constraints).

²¹ More precisely, we focus on a tractable two-asset case with one long-lived asset in fixed supply (i.e., land) as well as physical capital with an AK technology (i.e., firms). Focusing on a two-period life-cycle to obtain closed-form solutions for prices, we then decompose the total welfare gain of the old and the young into the contribution of changes in land prices and changes in income.

financial transactions multiplied by the deviation in asset prices. Second, these trading profits need to be discounted using an individual-specific marginal rate of substitution $\beta^t U'(C_{i,t})/U'(C_{i,0})$, a random variable that no longer equals $R_{0 \rightarrow t}^{-1}$ in the presence of uninsurable idiosyncratic shocks. We will quantify the effect of these adjustments for welfare gains in section V.A, theoretically and empirically.

While the version of the welfare-gains formula in a stochastic environment differs from its deterministic counterpart as just discussed, in our baseline results, we will empirically implement the deterministic formula (10), which discounts realized transactions using a constant discount rate. One reason is simplicity. Another reason is that the statistic has a nice interpretation, even in stochastic environments: it corresponds to (minus) the amount of money received at time $t = 0$ that would have allowed agents facing the deviation in asset prices to maintain their original consumption paths.²²

Finally, our results remain valid in the case in which individuals can trade financial assets in some ex ante stage to try to insure themselves against the deviation in asset prices. Intuitively, while the ability to choose one's portfolio in anticipation of the deviation in asset prices affects individual trading patterns and the welfare effect of the deviation, it does not affect the formula for welfare gains given these trading patterns.²³

Borrowing constraints and collateral effects.—In the baseline model, individuals can take unrestricted positions in the liquid asset (i.e., long and short). In reality, there are limits on how much individuals can borrow. These borrowing constraints affect our welfare gain formula via two distinct channels.²⁴ First, agents facing a borrowing constraint are not on their Euler equations, and they tend to discount future dollars by more than the rate of return on their debt (a “discount rate” channel). Second, in models where the borrowing constraint depends on collateral values (e.g., Kiyotaki and Moore 1997; Miao and Wang 2012), higher asset prices have an additional effect on welfare by relaxing borrowing constraints (a “collateral” channel). Importantly, the strength of this collateral channel depends on asset holdings and not just asset sales. We will quantify the effect of these adjustments for welfare gains in section V.B, theoretically and empirically.

Individual preferences.—In the baseline model, we specified a utility function that depends only on consumption. In reality, individuals may also care about the quantity of assets they own. An important example is owning and

²² See proposition A13 for more results on the different interpretations of welfare gains, both in deterministic and stochastic economies.

²³ We refer the reader to app. sec. A.2 for the general analysis of welfare gains in a fully stochastic environment where labor income, dividends, asset prices, and asset-price deviations themselves are stochastic.

²⁴ For a formal statement, see proposition 3 (case of soft borrowing constraints) and proposition A11 (case of hard borrowing constraints).

living in a house that generates a direct utility flow. In appendix section A.4.1, we consider an extension of the baseline model where asset holdings enter the utility function directly. We show that as long as the utility function depends on the quantity of assets owned, this “joy of ownership” channel does not affect our welfare gain formula.²⁵ Similarly, our sufficient statistic formula is robust to preferences for leisure and endogenous labor supply.

Finite lives and bequests.—In the baseline model, we abstract from life-cycle considerations, intergenerational linkages, and bequests. In practice, bequests are an important determinant of saving decisions (De Nardi 2004). In appendix section A.4.2, we consider an extension of the baseline model where individuals have finite lives and give assets to their heirs (as well as potentially receive inheritance from their parents).

Finite lives by themselves do not change our formula for the welfare effect of asset prices.²⁶ We then study the welfare effect of asset prices when agents have altruistic preferences (i.e., agent i directly cares about some other agent j). When defining the welfare gain of individual i as the amount of money that makes them indifferent to the deviation in asset prices, assuming agent j is already compensated for it, our sufficient statistic formula remains the same. We also study the case in which agents have “warm glow” preferences instead (i.e., bequest in the utility function); we show that our formula for welfare gains remains the same as long as the bequest function depends on the quantity of assets bequeathed rather than their market prices per se.²⁷

Finally, we emphasize that, in the presence of inter vivos transfers, the welfare effect of a deviation in asset prices depends only on the number of shares sold by the individual rather than on the overall change in the number of shares (that may come from bequests or inheritance).²⁸ This distinction is relatively easy to deal with in our empirical setting since we directly observe housing transactions among individuals.

Businesses.—In the baseline model, individuals directly own and trade financial assets. In reality, individuals typically own businesses that themselves own and trade financial assets (this includes, in particular, debt issued by businesses and share repurchases).

²⁵ It is only when individuals directly care about the market price of their assets per se that the welfare gain formula gains an additional term. We do not attempt to quantify such a channel in our empirical implementation.

²⁶ This is consistent with our discussion of the two-period model.

²⁷ This result aligns with the extension for asset holdings in the utility function discussed above. As in that case, we do not take into account the additional effect of having asset prices directly in the utility function in our empirical application. We discuss further issues related to the use of “warm glow” preferences for welfare assessment in app. sec. A.4.2.

²⁸ In particular, for an individual who inherits a house and plans to live in it forever, there is no change in welfare from higher house prices. However, higher house prices do hurt individuals who do not inherit a house but are planning to buy one in the future. Thus, higher asset prices increase the relative difference between those that inherit and those that do not.

In appendix section A.4.3, we show that the sufficient statistic formula remains valid in the presence of a business sector, provided transactions conducted by businesses are attributed to their ultimate owners. Intuitively, it is irrelevant whether financial transactions are undertaken directly by individuals or indirectly through the businesses they own. Similarly, it is immaterial whether a business distributes dividends or repurchases its shares; what ultimately matters is its cash flow stream (profits net of investment). In our empirical implementation, we will account for indirect financial transactions conducted by businesses owned by each individual when implementing our sufficient statistic.

Government.—In appendix section A.4.4, we study an extension of the baseline model with a government that taxes and makes transfers and is allowed to run surpluses and deficits (subject to a no-Ponzi condition, as in the individual problem). We do not assume that the government maximizes a social welfare function. Instead, we make a weaker assumption on cost minimization (i.e., the marginal return of investing in the different assets is equalized). We obtain two main results.

First, relative to the individual welfare gain formula in the baseline model, there is an additional term that accounts for the present value of changes in net government transfers. The idea is that, in general, the government will adjust taxes and transfers in response to a change in asset prices. Second, summing over all individuals, we show that the aggregate present value of changes in net government transfers is precisely equal to the “welfare gain of the government” (i.e., eq. [10] in the baseline model). This result is intuitive and follows directly from the government budget constraint. For instance, if the government is a borrower and its cost of borrowing increases (i.e., negative government welfare gain), then there are fewer resources available for making net transfers to individuals. Finally, we also examine the role of taxes that are indexed on asset prices in appendix section A.4.4.

III. Empirical Framework

We now discuss how we implement our sufficient statistic formula to estimate the distribution of welfare gains due to the rise in asset valuations across individuals in Norway. We first define the counterfactual we use for asset prices. We then describe the combination of administrative and publicly available data from Norway to quantify our sufficient statistic formula. A more detailed description can be found in appendix B.

A. Implementation

We now discuss how we bring the theory to the data to estimate the distribution of welfare gains due to the secular rise in asset valuations in Norway.

First-order approximation.—Proposition 1 gives a formula for the infinitesimal welfare gain associated with an arbitrary infinitesimal deviation in prices $\{dQ, \{dP_{k,t}\}_k\}_{t=0}^\infty$. We use this formula to obtain a first-order approximation of the welfare effect of a noninfinitesimal deviation in the price of different assets $\{\Delta Q, \{\Delta P_{k,t}\}_k\}_{t=0}^\infty$:

$$\text{Welfare Gain}_i = \sum_{t=0}^{\infty} R_{0 \rightarrow t}^{-1} \left(\sum_{k=1}^K (N_{i,k,t-1} - N_{i,k,t}) \Delta P_{k,t} - B_{i,t} \Delta Q_t \right). \quad (12)$$

We use the term sufficient statistic as this expression depends only on the observable path of financial transactions $(N_{i,k,t-1} - N_{i,k,t})P_{k,t}$ and $B_{i,t}Q_t$ —in particular, it does not require researchers to understand what drives these financial transactions over time or how they react to deviations in asset prices.

The latter is true only because we focus on the first-order approximation of the welfare effect of a deviation in asset prices. The accuracy of this first-order approximation to measure the equivalent variation depends on the extent to which asset transactions respond to changes in asset prices. In our empirical settings, we will focus on asset-price deviations across broad asset classes, in which case we can expect these responses to be low; for instance, Gabaix and Koijen (2021) provide evidence that demand elasticities at the asset class-level (e.g., stocks vs. bonds) are much lower than the demand elasticities within asset classes (e.g., stock A vs. stock B). We will explore this topic more formally in section V.C.

Asset classes.—One can rewrite this formula for welfare gains using a deviation of asset prices in relative terms:

$$\text{Welfare Gain}_i = \sum_{t=0}^{\infty} R_{0 \rightarrow t}^{-1} \left(\sum_{k=1}^K (N_{i,k,t-1} - N_{i,k,t}) P_{k,t} \times \frac{\Delta P_{k,t}}{P_{k,t}} - B_{i,t} Q_t \times \frac{\Delta Q_t}{Q_t} \right). \quad (13)$$

The term $(N_{i,k,t-1} - N_{i,k,t})P_{k,t}$ corresponds to the financial transactions corresponding to asset k , while $\Delta P_{k,t}/P_{k,t}$ corresponds to the percentage deviation in the asset price. Similarly, the term $B_{i,t}Q_t$ corresponds to the total amount of one-period bonds, and $\Delta Q_t/Q_t$ corresponds to the percentage deviation in the price of these bonds.

We now specify our counterfactual for asset prices. First, we consider the same relative price deviations for all assets within a given asset class: equity, housing, and debt.²⁹ Our approach answers the following question: What are the welfare gains associated with an $x\%$ deviation in the price of all assets within the same asset class? Because all financial transactions within a given asset class are multiplied by the same relative price deviation, we can

²⁹ Below, we further split debt holdings into mortgages and deposits. To be clear, we allow different households to earn heterogeneous returns within a given asset class. The only key assumption is that in the counterfactual we examine, all assets within the same asset class experience the same deviation in relative prices.

aggregate financial transactions within the same asset class. Equivalently, our thought experiment allows us to reinterpret K as the number of asset classes rather than the number of assets.

Second, we take as the baseline a world in which asset prices increase at the same rate as dividends. Hence, our approach answers the following question: What are the welfare gains of the realized path of asset prices compared to a baseline scenario in which they grew proportionally to dividends? This is a natural question because, on a balanced growth path, asset prices grow at the same rate as asset dividends (i.e., price-dividend ratios are constant).³⁰ Put differently, our thought experiment can be understood as measuring the welfare effect of movements in the path of detrended asset prices (where asset prices are detrended by their cash flows).

A large literature in finance argues that fluctuations in price-dividend ratios are mostly driven by fluctuations in future asset discount rates rather than in future expected dividend growth (for a seminal paper, see Campbell and Shiller 1988; for a recent examination across asset classes and countries, see Kuvshinov 2023). Suppose one is willing to assume that all of the rise in the price-dividend ratios in our sample comes from a decline in asset discount rates rather than an increase in expected dividend growth. In that case, our approach can be interpreted as answering the following question: What are the welfare gains of the rise in asset prices due to declining discount rates?³¹

Formally, we denote by $PD_{k,t} \equiv P_{k,t}/D_{k,t}$ the aggregate price-dividend ratio for asset class k . Given a baseline value $\overline{PD}_k$, we consider the following price deviation for asset class k :

$$\Delta P_{k,t} = P_{k,t} - \overline{PD}_k \times D_{k,t} \Rightarrow \frac{\Delta P_{k,t}}{P_{k,t}} = \frac{PD_{k,t} - \overline{PD}_k}{PD_{k,t}}. \quad (14)$$

This equation is the same as equation (2) discussed in the introduction. As a motivating example, figure 2 plots the index of house prices in Norway and the index of house rents. Starting around the mid-1990s, housing prices have grown faster than rents. In this case, the price deviation corresponds to the difference between realized prices $\{P_{H,t}\}_{t=0}^T$ and the counterfactual price path associated with a constant price-to-rent ratio $\{\overline{PD}_H \times D_{H,t}\}_{t=0}^T$. For the liquid asset (i.e., the sequence of one-period bonds), we consider a deviation of the price of one-period bonds from a constant baseline value $\bar{Q}$ (i.e., $\Delta Q_t/Q_t = (Q_t - \bar{Q})/Q_t$).

³⁰ In particular, price-dividend ratios are constant in models where asset discount rates and expected dividend growth rates are constant over time (Campbell and Shiller 1988).

³¹ See app. sec. B.1 for more details. Throughout the paper, we remain silent on the fundamental driver behind this decline in discount rates, which is very much an open question. As discussed more precisely in app. sec. A.3, under the assumption that the drivers of this decline did not directly impact the dividend or labor income of Norwegian households, our sufficient statistic formula (12) entirely captures the welfare effect of these drivers. If not, it captures only the effect operating through the deviation in asset prices.




FIG. 2.—Graphical representation of the price deviation $\Delta P_{H,t}$. The figure plots the house price index in Norway from Norges Bank's project on Historical Monetary Statistics (solid line) as well as the rental price index from Statistics Norway (dashed line). Both are adjusted for inflation and normalized to one in 1980. The difference between the two can be interpreted as a deviation $\Delta P_{H,t}$ between the realized price path $P_{H,t}$ and a counterfactual price path with constant price-to-rent ratio $\bar{P}D_H \times D_{H,t}$.

Time horizon.—While formula (12) depends on all transactions done by the individual, we observe only price deviations and financial transactions over a finite sample period.

Our solution to this issue is to do only the summation from $t = 0$ to $t = T$, where T denotes the length of the sample period. In this case, the sufficient statistic should be interpreted as the welfare effect of asset-price deviations up to time T . This truncation is inconsequential if either (i) the price deviation reverts to zero after T or (ii) if there is no trade after year T . More generally, if the price deviation remains positive after T , truncation overestimates the welfare gain for individuals who plan to buy financial assets after the truncation time T , while underestimating the welfare gain for individuals who tend to sell after T . Still, note that the bias due to truncation averages to zero in the entire population since there are as many sales as there are purchases after time T .³²

³² To fix ideas on the size of the bias, it is helpful to consider the case of an individual who buys $N_{i,0}$ shares of some asset at time 0 and resells them at some time $t > T$. While the net welfare gain of these transactions is $N_{i,0}(R_{0 \rightarrow t}^{-1} dP_t - dP_0)$, a researcher observing transactions up to time T will estimate a welfare gain of $-N_{i,0} dP_0$ (i.e., a welfare loss), thereby underestimating the actual welfare gain by $N_{i,0} R_{0 \rightarrow t}^{-1} dP_t$. Note that the bias depends on three distinct forces: (i) how large the truncation time T is (ii) how large the discount rate is relative to the baseline growth of house prices (i.e., how quickly $R_{0 \rightarrow t}^{-1} dP_t$ decays to zero as $t \rightarrow \infty$), and (iii) how persistent are house price deviations after T (i.e., how large dP_t/P_t is $t > T$).

As an alternative to truncating the infinite sum (12), we also construct hypothetical price deviations and financial transactions after year T in section V. We show that these alternative measures give similar results to our truncated measure under a wide range of scenarios about the path of future asset prices, given that we observe a relatively long time sample ($T = 25$ years).

Sufficient statistic.—Combining the first-order approximation of welfare gains (12) with the empirical price deviations (14) and truncating the formula at time horizon T , we obtain a sufficient statistic for the welfare gain of individual i due to the realized deviation of asset prices from balanced growth:

$$\text{Welfare Gain}_i = \sum_{t=0}^T R_0^{-1} \left(\sum_{k=1}^K (N_{i,k,t-1} - N_{i,k,t}) P_{k,t} \times \frac{PD_{k,t} - \overline{PD}_k}{PD_{k,t}} - B_{i,t} Q_t \times \frac{Q_t - \bar{Q}}{Q_t} \right). \quad (15)$$

This formula forms the core of our empirical implementation using administrative data.³³

We estimate equation (15) using data covering the 1994–2019 period. The reference year (i.e., $t = 0$) is 1994, and the sample length (T) is, therefore, 25 years. Our data cover the universe of individuals in Norway who were at least 18 years old for at least 1 year in the 1994–2019 period. We consider four asset classes: housing, debt, deposits, and equity, corresponding to the four main asset classes traded by Norwegian individuals. Note that we do not need to account for fully illiquid forms of wealth such as human wealth and defined-benefit pensions since they are not traded (i.e., they have no market price).

Given this, we estimate our sufficient statistic as follows:

$$\begin{aligned} \text{Welfare Gain}_i &= \sum_{k \in \{\text{housing, debt, deposit, equity}\}} \text{Welfare Gain}_{i,k}, \\ \text{Welfare Gain}_{i,\text{housing}} &= \sum_{t=0}^{25} R^{-t} (N_{i,H,t-1} - N_{i,H,t}) P_{H,t} \times \frac{PD_{H,t} - \overline{PD}_H}{PD_{H,t}}, \\ \text{Welfare Gain}_{i,\text{debt}} &= \sum_{t=0}^{25} R^{-t} (-B_{i,M,t} Q_{M,t}) \times \frac{Q_{M,t} - \bar{Q}_M}{Q_{M,t}}, \\ \text{Welfare Gain}_{i,\text{deposit}} &= \sum_{t=0}^{25} R^{-t} (-B_{i,D,t} Q_{D,t}) \times \frac{Q_{D,t} - \bar{Q}_D}{Q_{D,t}}, \\ \text{Welfare Gain}_{i,\text{equity}} &= \sum_{t=0}^{25} R^{-t} (N_{i,E,t-1} - N_{i,E,t}) P_{E,t} \times \frac{PD_{E,t} - \overline{PD}_E}{PD_{E,t}}, \end{aligned} \quad (16)$$

³³ This corresponds to the combination of formulas (1) and (2) in the introduction, generalized to multiple assets.

where $\overline{PD}_H$, $\bar{Q}_M$, $\bar{Q}_D$, and $\overline{PD}_E$ represent the average valuation of housing, debt, deposits, and equity (respectively) over 1992–96.³⁴

Our empirical implementation (16) also assumes that the discount rate in equation (15) is constant, $R_t = R$, and hence $R_{0 \rightarrow t}^{-1} = R^{-t}$. We set the discount rate to 5% (i.e., $R = 1.05$), which roughly corresponds to the average of the deposit and mortgage rates in a 5-year window around the start of our sample.³⁵

Computing these welfare gains requires data on valuation ratios for each asset class (to compare actual valuations to a baseline) and on the market value of financial transactions at the individual level. We now discuss each component separately.

B. Aggregate Data on Valuations

We rely on publicly available data sources for asset prices. For interest rates on debt and deposits (i.e., the inverse of the price of one-period bonds Q in the theory), we use Statistics Norway's database on interest rates on loans and deposits offered by banks and mortgage companies.³⁶ More than 90% of Norwegian mortgage debt in our sample has adjustable interest rates so that year-to-year variation in bank-level interest rates immediately affects individuals' interest costs.³⁷ Put differently, given that mortgage debt is mostly floating rate, we interpret the outstanding balance of the mortgage as a negative position in 1-year bonds.

For the price-to-rent ratio in the Norwegian housing market (i.e., the price-dividend ratio $PD_{H,t} = P_{H,t}/D_{H,t}$ in the theory), we combine data from different sources. We combine two indexes, one for house prices

³⁴ Relative to formula (15), we split the total amount of one-period bonds into two terms: $B_{D,t}Q_{D,t}$ the amount held in deposits, and $B_{M,t}Q_{M,t}$ the amounts held in debt, which is negative if individuals are net borrowers.

Note we use the same price deviation $(PD_{k,t} - \overline{PD}_k)/PD_{k,t}$ for all assets within an asset class. Hence, the welfare gain for asset class k should be interpreted as the welfare gain due to a common deviation in the relative price of all assets within this asset class (the one given by the deviation in the aggregate price-dividend ratio of the asset class).

³⁵ We pick a discount rate equal to the interest rate at the start of our sample $R = 5\%$, as a compromise between two opposite forces. On the one hand, to account for the effect of market incompleteness and borrowing constraints, secs. V.A and V.B suggest that we use a discount rate that is higher than the rate of return on the liquid asset. On the other hand, to obtain an approximation of welfare gains that is valid at the second order, sec. V.C suggests that we use a discount rate equal to the average rate of return between the baseline and counterfactual economy, which would give a lower value $R \approx 2.5\%$. We explore the robustness of our results to these extensions in sec. V.

³⁶ These data are available on Statistics Norway's website, <https://www.ssb.no/en/statbank/table/08175/>.

³⁷ Mortgage contracts in Norway typically are annuity loans with 25-year repayment schedules. When interest rates change, the payment schedule adjusts so that the sum of monthly debt repayment and interest costs remains constant at a new level throughout the remaining period of the contract. Such adjustments happen frequently, normally whenever the Central Bank policy rate changes.

and one for housing rents, to obtain our price-to-rent series. The rental index comes from Statistics Norway and is part of the official Consumer Price Index (CPI). The house price series comes from Norges Bank's project on Historical Monetary Statistics (Eitrheim and Erlandsen 2005).³⁸ As these two series are indexes, we scale their ratio so that it equals the price-to-rent ratio for Norwegian residential real estate for 2013 reported in MSCI (2016). We explore alternative constructions for the price-to-rent ratio in appendix section B.4.

We now turn to equity valuation (i.e., the price-dividend ratio for equity $PD_{E,t} = P_{E,t}/D_{E,t}$ in the theory). As explained in appendix section A.4.3, we focus on a valuation ratio for the overall corporate sector (i.e., unlevered equity). We measure it as the ratio between an aggregate measure of enterprise value (i.e., market value of equity plus debt) and the total cash flows distributed to equity and debt holders among publicly listed nonfinancial Norwegian firms using data from WorldScope.³⁹ Note that, unlike the price-dividend ratio, our equity-valuation ratio is unaffected by the relative importance of dividend payouts versus share repurchases as well as firms' capital structure (i.e., debt vs. equity financing). We account for the fact that firms have financial liabilities besides equity (e.g., debt for most firms and deposits for private banks) by allocating these indirectly held assets to the equity holders (for more details on the theoretical motivation, see app. sec. A.4.3; for more details on our implementation, see app. sec. B.2.2).

Figure 3 plots the yield of each asset class over time (i.e., $1/Q$ for debt and deposits and $D_{k,t}/P_{k,t}$ for long-lived assets $k \in \{H, E\}$), which are the inverse of the valuation ratios in equation (16). The notches on the vertical line marking the year 1993 correspond to our baseline values for each asset class. All yields decline substantially over time (i.e., valuations increase). On average, over our time sample, the housing yield fell by 5.7 percentage points, mortgage interest rates by 2.5 percentage points, deposit interest rates by 1.3 percentage points, and the equity yield by 0.7 percentage points. In particular, note that the equity yield has decreased less in Norway relative to the United States.

³⁸ This house price index is derived from data compiled by the Norwegian Real Estate Broker's Association, the private consulting firm Econ Poyry, and listings at the leading platform for house transactions Finn.no. Norges Bank updates these data regularly and provides them online, currently at <https://www.norges-bank.no/en/topics/Statistics/Historical-monetary-statistics/>.

³⁹ We use a valuation ratio for Norwegian firms, as opposed to foreign firms, as Norwegians mostly own and sell domestic equity (more precisely, Norwegians' holdings of domestic equity account for 100% of their private equity holdings and 72% of their public equity holdings). This contrasts with the Norwegian government, which mainly owns and buys foreign equity. Appendix C will discuss how using separate price indexes for domestic and foreign equity changes our estimates of welfare gains at the sectoral level.




FIG. 3.—Evolution of yields in Norway. The figure plots the yield of each asset class over time, that is, the inverse of the valuation ratios in equation (16). For debt and deposit, the yield corresponds to the average real interest rate on mortgages and debt, respectively. Nominal yields from Statistics Norway are adjusted for expected inflation using the average rate in the preceding 4 years. The housing yield corresponds to the rent-to-price ratio (see text for details). The equity yield corresponds to the aggregate ratio of cash flows to enterprise value among publicly listed Norwegian firms from WorldScope.

To compute the welfare gains of asset-price deviations, equation (16) requires a measure of the relative difference between valuations at time t and their average baseline value (i.e., their averages over the 1992–96 period). Figure A2 (figs. A1–A13 are available online) visualizes these price deviations.

C. Microdata on Transactions

We combine data from various Norwegian administrative registries covering the universe of Norwegians from the end of 1993 to the end of 2019. These data come with identifiers at the individual, household, and firm levels, as well as information linking parents and children. In particular, we use registries for individual tax payments, holdings of equity shares (listed and unlisted corporations), private business balance sheets, and housing transactions. Flow variables are measured annually, whereas assets and liabilities are valued at the end of the year. The data are uncensored (i.e., no top coding), and the only sources of attrition are mortality and emigration. The income and wealth data are largely third-party reported (i.e., employers and financial intermediaries) and scrutinized by the tax authority, as they are used for income and wealth tax purposes.

Data on holdings.—On individual balance sheets, we observe bank deposits, bond holdings (corporate, sovereign, mutual, and money market funds), debt, vehicles (cars and boats), stock mutual funds, publicly listed and private businesses, housing, and other forms of estate holdings. The values of these asset classes' holdings are available from the end of 1993.

In principle, we observe each individual's holdings. However, while financial holdings are registered at the individual level, they are taxed at the household level. The reported allocation of assets between individuals within the household is, therefore, somewhat arbitrary and can vary substantially from year to year. To compute a consistent measure of individual holdings across time, we therefore aggregate holdings at the household level and distribute it equally across adult household members.⁴⁰

We construct five main variables that cover most of individuals' financial wealth: "debt" (mortgages, student loans, and unsecured credit); "deposits" (bank deposits and bonds); "housing" (principal residence, secondary homes, and recreational estates); "private business equity" (equity in private businesses); and "public business equity" (listed stocks and stock funds). All of these variables are recorded at market value at the end of the year, except for private business equity, which is a tax-assessed value (i.e., the value reported to the tax authority, which is typically higher than the book value of equity; see app. sec. B.2.2). For housing, we use a valuation approach that combines transaction data and registered housing characteristics to estimate a value for each house in every year (for details on the valuation methodology, see Fagereng, Holm, Torstensen 2020). Note that this will only matter when reporting our welfare gains relative to total wealth.

Some individuals own private businesses. These firms directly hold financial assets and liabilities but often also own shares in other firms. To properly account for individuals' ownership, we must include their indirect asset positions held through private businesses. Our procedure is as follows. First, we compute each individual's direct and indirect ownership of private businesses. For instance, if an individual owns 80% of firm A, which in turn holds 50% of firm B, the individual effectively owns 80% of firm A and 40% of firm B. If firm B owns 25% of firm C, the individual then indirectly owns 10% of firm C as well. We calculate indirect ownership by going through 10 such layers of firm holdings. Equipped with these ownership shares, private firms' balance sheets, and publicly available data on public firms' balance sheets, we then allocate holdings and transactions conducted by firms to their ultimate owners (for details, see app. sec. B.2.2).

⁴⁰ Our definition of a *household* is either a single individual or a married or cohabitant (with children) couple. Offspring older than 18 living with their parents are considered to be in a separate household.

This approach enables us to treat financial transactions conducted directly and indirectly (via owned firms) in a consistent manner.⁴¹

Our notion of welfare gain can be interpreted as the present value of the deviation in consumption due to the deviation in asset prices (see proposition A13). Therefore, it is natural to express it as a share of the present value of consumption. However, we do not observe consumption directly in our sample. Instead, in some exercises, we will scale the welfare gain by “total wealth,” which is defined as the sum of financial wealth (i.e., debt, deposits, housing, and equity) and human wealth (i.e., the present value of earned income, defined as future labor income plus net government transfers received between 1994 and 2019, discounted at 5% annually). We also set the minimum value of earned income to twice the base amount in the social security system.⁴²

Table A2 (tables A1–A11 are available online) summarizes the data. Throughout the paper, we express all values in real terms (2011 Norwegian krone using the CPI) and then convert them to US dollars using a fixed exchange rate of 5.607. In appendix section B.2.1, we show that our aggregated microdata closely aligns with publicly available data on households’ asset holdings from the national accounts.

Data on transactions.—Equation (16) highlights the fact that we need data on holdings for debt and deposits and net transactions for housing and equity.

For housing, we observe the annual value of market transactions in the housing market at the individual level. Thus, net transactions in housing are directly observed. For public equities, we observe holdings at the beginning and end of the year and a price index. We then compute a measure of unrealized capital gains by assuming that all transactions are in the same direction and uniformly distributed within a year. Net transactions are thus constructed as the change in market value minus imputed capital gains. The price index used for imputation differs between assets. For listed stocks, the method varies depending on the available information. Starting in 2005, we have information on individual stock ownership and use market prices on individual stocks to impute capital gains. Before 2005, we lack information on individual stock ownership and use capital gains from the financial accounts to impute capital gains on listed stocks at the individual level. We also use capital gains from the financial accounts to impute individual capital gains for mutual funds.

For equity in private businesses, we impute the value of transactions using the data on ownership shares described earlier. In particular, if we see

⁴¹ We outline how theory motivates our consolidation of firms’ financial transactions in app. sec. A.4.3.

⁴² As with financial holdings, an individual’s human wealth is computed based on their household’s human wealth.

that an individual owns 50% of a private business in a given year and 25% the following year, this implies that the individual sold a 25% stake of the business.⁴³ In appendix section B.2.2, we describe this methodology in detail. Private business equity transactions are infrequent and not quantitatively important. As a result, private business owners are not meaningfully exposed to private equity-valuation changes. It is worth stressing that, even in a world in which business owners never sell their stakes in their businesses, they are still exposed to asset-price changes via the financial transactions made by the firms they own. For instance, if the interest rate on debt declines, the owner of a levered business will incur a positive welfare gain. This phenomenon is particularly important for individuals at the top of the wealth distribution, as they hold a lot of assets through their private firms.

Bequest events pose two challenges when computing net transactions. First, housing transactions may be problematic at the time of death. In most cases, when an individual dies, the estate is transferred to the heirs. In this case, the heirs sell the property, and net transactions are computed correctly. But in a few cases, parts of the estate are sold after death but before it is transferred to the heirs. In this case, we allocate the transaction to the living children of the deceased, in accordance with the Norwegian inheritance law.⁴⁴

Second, because our imputation of net transactions in equity is based on changes in holdings net of imputed capital gains, a bequest event may be problematic because wealth transfers may be counted as transactions. For example, if one individual gives 100 equity shares to another individual, this should be reported neither as a purchase by the recipient nor a sale by the giver. To address this issue, we allocate all imputed equity transactions of givers to recipients when there is a bequest event. A bequest event is defined as any transfer reported in the inheritance tax registry (both *inter vivos* and at death).⁴⁵

IV. Asset-Price Redistribution

We now estimate our sufficient statistic (16) for all Norwegians who were at least 18 years old at some point between 1994 and 2019. More precisely, we describe the heterogeneity in welfare gains across individuals in section IV.A,

⁴³ Alternatively, the business might have issued new equity, leading to a dilution of existing owners. In terms of welfare exposure to equity prices, those two scenarios are equivalent (see app. sec. A.4.3).

⁴⁴ By law, inheritance is split equally between all direct descendants unless explicitly specified otherwise in a will.

⁴⁵ Before 2014, there was an inheritance tax in Norway, and the tax authority collected information on sender, receiver, and the amount transacted. However, this register does not contain information on the types of assets transferred.

TABLE 1
SUMMARY STATISTICS ON TRANSACTIONS (NET PURCHASES)

Asset Type	Average	SD	p1	p10	p25	p50	p75	p90	p99
Housing	.9	116.3	-190.1	.0	.0	.0	.1	.9	220.4
Debt	-73.6	3,043.1	-602.6	-222.0	-127.5	-36.1	.0	9.5	348.2
Deposits	19.6	1,848.8	-157.0	-1.8	1.0	7.6	28.7	76.7	339.6
Equity	-.4	414.0	-27.1	-.6	.0	.0	.0	1.4	30.3

NOTE.—All values are in thousands of 2011 US dollars.

across cohorts in section IV.B, across the wealth distribution in section IV.C, and across sectors (i.e., households, government, and foreigners) in section IV.D.

A. Redistribution across Individuals

Transactions.—We start by documenting the heterogeneity in financial transactions. Table 1 reports summary statistics for transactions across the population, computing them every year and averaging them across all years in our sample. Compared with table A2, we also include indirect transactions via firms owned by individuals.

Housing transactions are very lumpy, and most people hold debt and deposits. The magnitude of equity transactions is much smaller than housing transactions, reflecting that housing holdings dominate equity holdings for Norwegian individuals (see table A2). Also, deposits are negative (and debt is positive) for a substantial fraction of the population, which comes from the fact that we report consolidated holdings and transactions: individuals who own equity in financial firms (e.g., banks) indirectly hold long positions in debt and short positions in deposits. Finally, financial transactions do not exactly average to zero: as we will discuss below, this reflects the fact that individuals in our sample also trade with the Norwegian government and the rest of the world.

Welfare gains.—Figure 4 presents the histogram of total welfare gains. Note that the average welfare gain is close to zero, reflecting that for every seller benefiting from higher asset prices, there is a seller equally harmed in monetary terms.⁴⁶ However, there is substantial heterogeneity: the welfare gain is -\$185,000 at the 1st percentile and \$273,000 at the 99th percentile, with an interquartile range of \$31,000. There is a large mass around zero, reflecting that consumption is close to income for a large fraction of individuals. As mentioned, financial transactions within the household sector do not average to zero in our sample. As a result, welfare gains do not

⁴⁶ The reason our baseline statistic does not average to zero across individuals is that Norwegian households do not trade exclusively with one another; they also trade with the government and foreign entities (see sec. IV.D for more details).




FIG. 4.—Distribution of welfare gains. This figure plots the density of individual welfare gains, as defined in (16), across individuals in Norway. More precisely, the figure plots the relative mass of individuals within equally spaced bins of welfare gains (width of \$1,000). Panel A plots welfare gains in levels (in 2011 US dollars), while panel B plots welfare gains as a percent of total wealth, which is defined as the sum of financial wealth and human capital at the end of 1993 (i.e., the present value of labor income and government benefits received between 1994 and 2019).

average to zero either: they average to \$10,000, which is slightly positive. In appendix C, we will show that this positive welfare gain corresponds to a welfare loss for the Norwegian government and for foreigners. The Kelly skewness of the distribution is fairly small, at 0.08, reflecting the fact that the distribution of welfare gains is fairly symmetrical around its mean.⁴⁷

To understand which asset class contributes the most to redistribution, table 2 decomposes the average welfare gain into different percentile groups of the welfare-gain distribution. More precisely, for each percentile group, the table reports the average welfare gain, as well as the average welfare gain due to each asset class. Housing is, by far, the asset class that generates the most redistribution. This comes from the fact that, even though housing transactions tend to be smaller than debt or deposit holdings (table 1), the price deviations associated with housing are much larger than the price deviations associated with debt and deposits (fig. A1). Nevertheless, debt is also an important (and almost always positive) contributor, with a relatively large magnitude both at the top and bottom of the welfare-gain distribution. Similarly, deposits make a very small and almost always negative contribution. Welfare gains due to equity are small, reflecting the fact that there are fewer equity transactions in our sample (table 1) and that the run-up in equity prices was smaller than the run-up in house prices (fig. A1).

Welfare gains as a percent of total wealth.—We now evaluate the dispersion of welfare gains relative to total wealth, defined as the sum of financial and

⁴⁷ Kelly skewness is defined as $(p_{90} + p_{10} - 2 \times p_{50}) / (p_{90} - p_{10})$ where p_{10} , p_{50} , and p_{90} are the 10th, 50th and 90th percentiles of the distribution under consideration.

TABLE 2
DECOMPOSITION OF WELFARE GAINS BY PERCENTILE GROUPS

Asset	Average	Average by Percentile Groups of Welfare Gains					
		p0 – p1	p1 – p10	p10 – p50	p50 – p90	p90 – p99	p99 – p100
Housing	–4.7	–294.6	–99.6	–18.2	2.3	59.4	414.4
Debt	16.9	–218.0	18.0	8.8	16.2	39.5	202.6
Deposits	–2.4	67.1	–4.4	–2.9	–2.1	–4.3	–15.8
Equity	.2	–61.6	–2.4	–.4	–.1	.5	61.2
Total	10.0	–507.2	–88.3	–12.8	16.3	95.1	662.4

NOTE.—For each percentile group of welfare gains, the table reports the average welfare gain, and the average welfare gain due to each asset class, as defined in (16). All numbers are in thousands of 2011 US dollars.

human wealth (see sec. III.C). As discussed in appendix section E.2, welfare gains can be interpreted as the present value of the change in consumption. Consequently, this normalized version of welfare gains can be interpreted as the relative change in consumption due to asset-price deviations.⁴⁸ In this exercise and the ones below, we winsorize total wealth at the bottom 1% within each cohort to limit the influence of observations with very small total wealth.

Figure 4 shows significant heterogeneity in welfare gains, even after normalizing by initial wealth. The normalized welfare gain is –30% at the 1st percentile and 27% at the 99th percentile, with an interquartile range of 5.0%. While the Kelly skewness of the distribution is close to zero (–0.06), reflecting a symmetric distribution, the kurtosis of the distribution is 11, reflecting a larger mass in the tails relative to the normal distribution.

Aggregation.—Our notion of individual welfare gain represents the amount of cash, received in the baseline economy at time $t = 0$, that would make the individual indifferent between the baseline and perturbed paths of asset prices. As discussed by Saez and Stantcheva (2016), one can aggregate these individual welfare gains, together with a set of social marginal welfare weights, to compute the “social” welfare gain associated to the deviation in asset prices.⁴⁹ The key point is that, given a specific set of social marginal welfare weights that represent how much society values different individuals’ marginal consumption, our measures of individual welfare gains are the only inputs needed to compute the associated social welfare gain.

⁴⁸ Another way to interpret this number is that it corresponds to the relative increase in consumption every period that would be welfare equivalent to the change in asset prices (see proposition A13 for details).

⁴⁹ More precisely, Saez and Stantcheva (2016) define the “social” welfare gain as $\sum_{i=1}^I g_i dV_{i,0}/U'(C_{i,0})$, where g_i corresponds to the social marginal welfare weight on individual i and $dV_{i,0}/U'(C_{i,0})$ corresponds to our (money-metric) notion of welfare gain of individual i . For the special case of a utilitarian social planner, g_i corresponds to the Pareto weight for individual i times the marginal utility of consumption.




FIG. 5.—Social welfare gain as a function of inequality aversion. The figure plots the social welfare gain $\sum_{i=1}^I g_i \times \text{Welfare Gain}_i$, where $g_i \equiv \text{Total Wealth}_i^{-\sigma} / \sum_j \text{Total Wealth}_j^{-\sigma}$ denotes the social marginal welfare weight associated to individual i . Because our marginal social welfare weights sum up to one, we can interpret the result in dollar terms: from the social planner's point of view, a social welfare gain of $\$X$ is equivalent to giving $\$X$ to each individual.

As an example, we plot in figure 5 the social welfare gain obtained by social marginal welfare weights equal to individual total wealth at the power $-\sigma$, where σ can be interpreted as an index of social aversion for inequality.⁵⁰ When $\sigma = 0$, the social welfare gain is the average welfare gain in the population, which is roughly \$10,000. As σ increases, the social welfare gain decreases and ultimately becomes negative, reflecting the fact that the statistic weighs more and more the welfare gains of poorer individuals relative to more affluent individuals (and that, as we will see shortly, the rise in asset prices redistributed from the poor toward the wealthy).

Revaluation gains.—We now compare welfare gains with revaluation gains, defined as the (present value of the) effect of the deviation in asset prices on wealth:

$$\text{Revaluation Gain} = \sum_{t=0}^T R_{0 \rightarrow t}^{-1} \sum_{k=1}^K N_{k,t-1} P_{k,t-1} \Delta \left(\frac{P_{k,t}}{P_{k,t-1}} \right), \quad (17)$$

where we define $\Delta(P_{k,t}/P_{k,t-1}) \equiv (P_{k,t}/P_{k,t-1})(\Delta P_{k,t}/P_{k,t} - \Delta P_{k,t-1}/P_{k,t-1})$ as the deviation in the capital gains component $P_{k,t}/P_{k,t-1}$ of asset returns caused by the price deviation $\{\Delta P_{k,t}\}_{t \geq 0}$.

⁵⁰ Alternatively, this can be interpreted as the welfare change of a utilitarian social planner that aggregates equally the utility of individuals who have homothetic utility functions with parameter σ (since, in this case, consumption is proportional to total wealth).

Welfare gains are different from revaluation gains. This is because revaluation gains capture only the positive effect of rising valuations on returns through higher capital gains, while welfare gains also take into account the negative effects of higher valuations on returns through lower dividend yields. In particular, revaluation gains systematically overestimate welfare gains in a time of inflated asset prices. We derive a formal expression for the difference between welfare and revaluation gains in appendix section E.6.

Figure 6A compares the density of welfare and revaluation gains, both as a percent of initial (total) wealth. As discussed above, welfare gains are centered around zero (0.0% on average). In contrast, revaluation gains are centered around a large positive value (16.4%, on average). This reflects the fact that revaluation gains are positive for all asset holders, while welfare gains are only positive for asset sellers.

Do individuals with higher revaluation gains also tend to have higher welfare gains? To answer this question, we now focus on the ordinal relationship between the two variables. Figure 6B plots a heatmap for the joint density of ranks of welfare gains and ranks of revaluation gains. Overall, we find that the Spearman rank correlation between welfare gains and revaluation gains is 0.19, which shows that there is a substantial difference between those who get richer from the rise in asset prices and those who truly benefit from it. Some individuals with large asset positions buy and hence lose in welfare terms; conversely, others with small positions sell and thus win.


B. Redistribution across Cohorts

In the previous section, we documented a large amount of heterogeneity in welfare gains across individuals. We now focus on describing the heterogeneity in welfare gains across one observable characteristic: the age of each individual at the end of 1993 (or, alternatively, the cohort they belong to). Indeed, the existing literature on household finance has documented large differences in portfolio holdings over the life cycle (e.g., Cocco, Gomes, and Maenhout 2005; Flavin and Yamashita 2011). This heterogeneity may naturally generate heterogeneity in financial transactions and, therefore, in welfare gains.

Transactions.—Figure 7A plots the average (consolidated) financial transactions in equity and housing by age. Importantly (though unsurprisingly), younger individuals tend to be net buyers of housing and equity, whereas older individuals tend to be net sellers. Figure 7B plots the average holdings of debt and deposits by age, as they also enter the sufficient statistic (16). Younger individuals hold a large amount of debt, primarily mortgage debt.

Welfare gains.—Figure 7C plots the average welfare gain for different cohorts, indexed by individuals' age at the end of 1993. The main pattern is that welfare gains are negative for the young and positive for the old, meaning that rising asset prices redistributed from the young toward the

A Marginal distributions



B Rank correlation




FIG. 6.—Welfare gains versus revaluation gains as percent of total wealth. Panel A plots the marginal distributions of welfare gains defined in (16), in black lines, and of revaluation gains defined in (17), in gray shading, across individuals in Norway. Panel B plots the joint density of the rank of welfare and revaluation gains; that is, the fraction of individuals within each quintile of welfare and revaluation gains. By definition of quintiles, numbers within each row (or column) aggregate to $1/5$. Both welfare and revaluation gains are expressed as a percentage of initial total wealth, defined as the sum of financial wealth and human capital at the end of 1993 (i.e., the present value of labor income and government benefits received between 1994 and 2019).




FIG. 7.—Financial transactions and welfare gains by age group. Panels A and B plot (consolidated) financial transactions (net purchases) per capita by age, averaged across all years in our sample. Specifically, for each asset class and year in our sample, we calculate the average transaction value within groups of individuals belonging to the same 3-year age range as of year-end. We then average this quantity across all years in our sample. Panel C plots the average welfare gain (16) for individuals in each cohort (individuals belonging to the same 3-year age range at the end of 1993). All numbers are in 2011 US dollars.

old. This is consistent with standard life cycle models of savings: the young save for retirement by purchasing financial assets, while the old sell their financial assets to consume.

Quantitatively, the average welfare gain is $-\$13,000$ for individuals below 15 years old in 1993 (millennials) and around $\$22,000$ for individuals above 50 years old in 1993 (baby boomers). The figure also decomposes welfare gains into each asset class's contribution, revealing interesting patterns. On the one hand, higher house prices redistribute from young to old, as the young tend to buy houses from the old. On the other hand, lower mortgage rates redistribute from old to young, as the young tend to borrow from the old.⁵¹ Overall, the effect of higher house prices dominates the

⁵¹ As we discuss in app. C, the household sector, as a whole, is a net debtor. Therefore, the young borrow not only from the old but also from foreigners and, indirectly, from the government. Also note that, while life-cycle mortgage balances peak around age 30 (fig. 7B), the welfare effect of lower mortgage rates is highest for individuals who are 20 years old in 1993

effect of lower mortgage rates for two reasons. First, and most importantly, the housing yield decreased more than the interest rate on debt (see fig. 3). Second, as young people build equity in their houses, they decrease their mortgage balances over time, which means they benefit relatively less from the decline in mortgage rates as they age.

C. *Redistribution across Wealth Percentiles*

A growing literature has emphasized that rising asset valuations affect the distribution of wealth (e.g., Kuhn, Schularick, and Steins 2020; Greenwald et al. 2021; Gomez 2025). A natural question is, Are these revaluation gains actually welfare gains? To answer this question, we compare revaluation and welfare gains across percentiles of the initial wealth distribution at the end of 1993. More precisely, we rank individuals according to their total initial wealth within their cohort. We then compare average revaluation and welfare gains at these different percentiles.

Transactions.—Figure 8A plots the average consolidated equity and housing transactions across different percentiles of the wealth distribution. To make it more easily comparable across different percentiles, we normalize average transactions by the total wealth at the end of 1993 at each percentile. The key observation is that richer individuals are, on average, net sellers of equity, while poorer individuals are, on average, net buyers. In contrast, housing net purchases are mildly positive across most of the wealth distribution (consistent with the mildly positive aggregate housing net purchases by households; see table 2).

Figure 8B plots the consolidated holdings of debt and deposits across the wealth distribution. As a proportion of financial wealth, the level of debt decreases (in absolute value) with the level of wealth, while the level of deposits increases. The negative value of deposits at the top 1% reflects that richer individuals tend to hold more equity, and as a result, they indirectly hold negative positions in deposits through their ownership of Norwegian banks. Finally, the top 1% holds little debt on a consolidated basis.⁵²

Welfare gains.—Figure 8C plots the average welfare gains at different wealth percentiles. Welfare gains increase with total wealth: the top 1% experienced, on average, a \$73,000 welfare gain, while the corresponding number is \$8,000 at the bottom 1%. Figure 8D plots welfare gains as a percent of the average total wealth in each percentile. The main pattern is

(fig. 7C). This phenomenon is due to two forces: (i) mortgage rates are mostly flat at the beginning of our sample and only start declining in 2001 (fig. 3), and (ii) this cohort spends a more extended amount of time with mortgage debt than the older cohorts aged around 30 in 1993 (fig. 7B).

⁵² While richer individuals issue debt through their ownership in nonfinancial businesses, they also buy this debt through their ownership in financial businesses.




FIG. 8.—Financial transactions and welfare gains by wealth percentile. Panels A and B plot net transactions per capita, averaged across years and 3 percentile groups of total wealth at the end of 1993, and divided by average total wealth measured at the end of 1993. Panel C plots the average welfare gain, as defined in (16), in 2011 US dollars. Panel D plots welfare gains divided by total wealth at the end of 1993, both averaged across 3-percentile groups of total wealth at the end of 1993 (except for the top 1%, plotted separately). Wealth percentiles are constructed by ranking individuals within each cohort based on total wealth at the end of 1993, defined as the sum of financial wealth and human capital (i.e., the present value of labor income and government benefits received from 1994 to 2019). When generating these figures, we exclude eight individuals with low initial measured total wealth but extremely high subsequent wealth, likely due to inheritance, as they generate discrete spikes in transactions and welfare gains divided by initial total wealth.

these “normalized” welfare gains tend to be stable across the wealth distribution, except for the top 1%. Individuals in the top 1% of their cohort experience a welfare gain of roughly 3.1% (as a percent of total wealth), which is higher than the population average of 1.5%. Moreover, most of the relatively higher welfare gains for the top 1% come from equity, reflecting that they tend to be net sellers in this asset class.

Revaluation gains.—Finally, figure 9 contrasts revaluation and welfare gains. Similarly to welfare gains, revaluation gains increase with top percentiles, which reflects the importance of revaluations for the rise in wealth inequality. However, the figures show that the magnitude of revaluation gains (44.5% of total wealth for the top 1%) is much bigger than the magnitude of welfare gains (3.1% of total wealth for the top 1%). Put differently, only a small part of these revaluation gains are welfare relevant.




FIG. 9.—Welfare and revaluation gains across wealth percentiles. This figure plots the average welfare and revaluation gains, as defined in (16), across 3-percentile groups of total wealth at the end of 1993. Panel A reports the two quantities in level (dollar terms), while panel B reports the two quantities as a percent of total wealth, as measured at the end of 1993. Wealth percentiles are constructed by ranking individuals within each cohort based on total wealth at the end of 1993, defined as the sum of financial wealth and human capital (i.e., the present value of labor income and government benefits received from 1994 to 2019).

D. Redistribution across Sectors

As discussed in the previous sections, our baseline measure of welfare gains does not aggregate to zero across households. This is because Norwegian households do not trade exclusively with one another; they also trade with the government and foreign entities.

In appendix C, we use data on sectoral financial transactions from Norwegian national accounts to analyze the redistributive effects of asset prices across sectors. We show that the positive average welfare gain of Norwegian households is counterbalanced by a negative welfare gain of the Norwegian government. Indeed, while Norwegian households are net debtors on average, the consolidated government (through Norway's sovereign wealth fund) is a net saver.

As discussed in the final subsection of section II.C ("Government"), a welfare loss for the government represents a loss of real resources available for net transfers to the household sector. While it is beyond the scope of this paper to quantify how the Norwegian government has adjusted (and will adjust) net transfers in response to persistently lower interest rates and higher asset prices, it is possible that the very individuals who experienced welfare losses in our exercise (i.e., the young) will also be the ones to bear the brunt of future reductions in government transfers such as pension benefits.

IV. Generalizations of the Baseline Sufficient Statistic Approach

We now implement several extensions and generalizations of our baseline sufficient statistic approach. In particular, we modify our sufficient statistic

approach to take into account (i) uninsurable income risk, (ii) borrowing constraints with collateral effects, (iii) second-order effects, and (iv) extrapolation beyond the end of the sample. In each case, we discuss the theoretical difference relative to our baseline formula, our methodology to implement the correction, and its quantitative effect.

For the sake of transparency, we analyze each extension separately. As a preview of our results, table 3 reports the effect of each generalization for the distribution of welfare gains across cohorts. Overall, we do find that each of these effects matters quantitatively. The last line of the table reports the effect of combining all these extensions. We find that average welfare gains increase across the wealth distribution, with a more significant increase for the 20–40 cohorts.

A. *Uninsurable Income Risk*

We have derived our sufficient statistic formula in a deterministic model. In reality, agents are exposed to both individual-specific and economy-wide shocks. In this section, we study theoretically and empirically the effect of uninsurable labor income risk on our sufficient statistic formula. We refer the reader to appendix section A.2 for a more general analysis of welfare gains in a fully stochastic environment where not just labor income but also dividends, asset prices, and asset-price deviations themselves are stochastic.

Theory.—The environment is the same as in the baseline model except that individual labor income $Y_{i,t}$ is now subject to idiosyncratic shocks. The individual chooses a stochastic path of consumption and asset holdings to maximize the expected utility of consumption:

TABLE 3
WELFARE GAINS ACROSS COHORTS: GENERALIZATIONS OF OUR BASELINE APPROACH

	Welfare Gains by Age Group				
	Mean	0–20	20–40	40–60	60–80
Baseline	11.9	–12.6	25.4	25.3	10.6
Additional effect of:					
Uninsurable income risk	+2.3	+4.2	+1.2	+2	–.1
Borrowing constraints and collateral effects	–.1	+1.9	–1.4	–1.1	+1
Second-order effects	+3.5	–9.5	+8.8	+12.0	+4.8
Extrapolation	+4.4	+6.6	+6.7	+1.2	+2
Combining all extensions	23.7	–4.3	42.8	39.3	10.9

NOTE.—The age group refers to the age of the cohort at the end of 1993. (See the replication data for further details of the breakdown.) “Uninsurable income risk” reports the results obtained in sec. V.A with $\gamma = 1$. “Borrowing constraints and collateral effects” reports the results obtained in sec. V.B with $\xi = 0.01$. “Second-order effects” reports the results obtained in sec. V.C. “Extrapolation” reports the results obtained in sec. V.D. with $\phi = 0.9$. “Combining all extensions” reports the results obtained in app. sec. D.5. All numbers are in thousands of 2011 US dollars.

$$V_{i,0} = \max_{\{C_{i,t}, B_{i,t}, \{N_{i,k,t}\}_k\}_{t=0}^{\infty}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U(C_{i,t}) \right],$$

subject to initial asset holdings $B_{i,-1}$ and $\{N_{i,k,-1}\}$ and the usual sequence of budget constraints:

$$\begin{aligned} C_{i,t} + \sum_{k=1}^K (N_{i,k,t} - N_{i,k,t-1})P_{k,t} + B_{i,t}Q_t + \sum_{k=1}^K \chi_k (N_{i,k,t} - N_{i,k,t-1}) \\ = \sum_{k=1}^K N_{i,k,t-1}D_{k,t} + B_{i,t-1} + Y_{i,t}. \end{aligned}$$

The next proposition characterizes the welfare gain of a deviation in asset prices in this stochastic environment, defined as the individual's willingness to pay for the deviation in asset prices at $t = 0$.

PROPOSITION 2. In the presence of uninsurable income risk, the welfare gains from a deviation in asset prices for individual i is

$$\frac{dV_{i,0}}{U'(C_{i,0})} = \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \frac{\beta^t U'(C_{i,t})}{U'(C_{i,0})} \left(\sum_{k=1}^K (N_{i,k,t-1} - N_{i,k,t}) dP_{k,t} - B_{i,t} dQ_t \right) \right]. \quad (18)$$

There are two differences with the baseline welfare gain formula. The first is that, in a stochastic environment, what matters for the ex ante welfare effect of an asset-price deviation is the expected path of net asset sales. The second is that this expectation is under the individual's risk-neutral measure, which tilts the objective measure by the growth of the individual's marginal utility of consumption $\beta^t U'(C_{i,t})/U'(C_{i,0})$ (i.e., the individual marginal rate of substitution). Because of uninsurable idiosyncratic shocks, this adjustment is individual specific. To emphasize the role of this adjustment, we use the Euler equation $\mathbb{E}_0[\beta^t U'(C_{i,t})/U'(C_{i,0})] = R_{0 \rightarrow t}^{-1}$ to rewrite the welfare-gains formula as a sum of two terms:

$$\begin{aligned} \frac{dV_{i,0}}{U'(C_{i,0})} = & \underbrace{\sum_{t=0}^{\infty} R_{0 \rightarrow t}^{-1} \mathbb{E}_0 \left[\sum_{k=1}^K (N_{i,k,t-1} - N_{i,k,t}) dP_{k,t} - B_{i,t} dQ_t \right]}_{\text{baseline}} \\ & + \underbrace{\sum_{t=0}^{\infty} \text{cov}_0 \left(\frac{\beta^t U'(C_{i,t})}{U'(C_{i,0})}, \sum_{k=1}^K (N_{i,k,t-1} - N_{i,k,t}) dP_{k,t} - B_{i,t} dQ_t \right)}_{\text{covariance term}}. \end{aligned} \quad (19)$$

The first term captures the welfare gain due to the expected path of asset transactions (in the objective measure). The second term captures the welfare gain due to the covariance between the growth rate of marginal utility and net asset sales. In our context, we can expect this covariance term to be positive, as labor income shocks generate a positive comovement

between the marginal utility of consumption and asset sales (e.g., individuals jointly reduce consumption and savings after a negative income shock). One implication of this covariance term is that welfare gains no longer aggregate to zero in the population. While it is still the case that higher asset prices are purely redistributive from ex post buyers to ex post sellers (i.e., transactions sum up to zero in every state of the world), agents disproportionately weight the states in which they are sellers from an ex ante perspective, meaning that welfare gains aggregate to a net positive sum across the population.

Finally, note that, even in the presence of uninsurable income risk, our baseline sufficient statistic (10), which discounts realized transactions using a constant discount rate, still has a valid interpretation as (minus) the amount of money, received at time $t = 0$, that would have allowed the individual facing the deviation in asset prices to maintain their original path of consumption.⁵³

Implementation.—We now adjust our sufficient statistic approach to quantify the contribution of uninsurable labor income for ex ante welfare. As seen in (19), the key empirical object that governs the effect of market incompleteness is the covariance between the growth of marginal utility of consumption and future asset sales at each horizon $t \geq 1$. To estimate this incomplete market adjustment term in the data, we assume that individuals have CRRA (constant relative risk aversion) utility with a coefficient of relative risk aversion γ . The covariance term for asset k in (19) can be approximated as

$$\begin{aligned} \text{cov}_0 \left(\frac{\beta^t U'(C_{i,t})}{U'(C_{i,0})}, (N_{i,k,t-1} - N_{k,t}) P_{k,t} \right) \\ \approx R_{0 \rightarrow t}^{-1} \times \gamma \times \text{cov}_0 \left(\log \left(\frac{C_{i,t}}{C_{i,0}} \right), (N_{i,k,t} - N_{i,k,t-1}) P_{k,t} \right) \end{aligned} \quad (20)$$

using a log-linear approximation in consumption growth.⁵⁴ While this approximation is not strictly necessary, it makes the statistic more robust in the data, as consumption growth can have extreme outliers at the individual level due to fat-tailed events or measurement errors (Toda and Walsh 2015).

⁵³ See proposition A13.

⁵⁴ More precisely, we have

$$\begin{aligned} \text{cov}_0 \left(\frac{\beta^t U'(C_{i,t})}{U'(C_{i,0})}, (N_{i,k,t-1} - N_{k,t}) P_{k,t} \right) &= \text{cov}_0 \left(\beta^t \left(\frac{C_{i,t}}{C_{i,0}} \right)^{-\gamma}, (N_{i,k,t-1} - N_{k,t}) P_{k,t} \right) \\ &= R_{0 \rightarrow t}^{-1} \text{cov}_0 \left(\left(\frac{C_{i,t}}{C_{i,0}} \right)^{-\gamma}, (N_{i,k,t} - N_{i,k,t-1}) P_{k,t} \right). \end{aligned}$$

Approximating at the first order in $c_{i,t} \equiv \log(C_{i,t}/C_{i,0})$ around $c^* \equiv (-1/\gamma) \log \mathbb{E}_0[(C_{i,t}/C_0)^{-\gamma}]$ gives the result.

In our particular settings, we can construct a measure of individual consumption as a residual from the budget constraint (9), that is, total net income minus net asset purchases. However, this measure has two limitations. First, measurement error in either income or asset purchases generates a mechanical negative correlation between consumption growth and asset purchases, leading us to underestimate the effect of incomplete markets (since we expect the covariance term to be positive). Second, our measure captures total spending rather than nondurable consumption, which would be the appropriate quantity in this context (see, e.g., Vissing-Jørgensen 2002). To partially address these two issues, we substitute our measure of consumption growth with its projection on log income growth. See appendix section A.4.3 for more details on our implementation.

Results.—Figure 10A reports the sum (over time) of the covariances (20) for each cohort and for each asset class. The positive “total” (i.e., the sum of the asset-specific covariances) reflects that households with a high consumption level—relative to others with the same observables in 1994—tend to purchase more housing and hold more debt. For any positive level of relative risk aversion γ , this is a force that will dampen the (ex ante) welfare loss associated with rising house prices (and declining interest rates), given that housing purchases (and borrowing) disproportionately occur in idiosyncratic states in which individuals have high income and low marginal utility. Note that the covariances tend to decay with cohort age, reflecting the fact that the retirement income of Norwegians is pretty stable over time.⁵⁵ Figure 10B reports average welfare gains across cohorts, including the incomplete-market adjustment term (19) for different values of the risk aversion parameter γ . We find that the effect of uninsurable labor income risk is particularly important for younger cohorts, who face more uncertainty over their lifetimes. In particular, we find that the incomplete-market adjustment term offsets some of the welfare loss for the young: the average welfare gain for the cohort of individuals who are 10 years old in 1994 increases from $-\$17,000$ when $\gamma = 0$ (baseline) to $-\$13,000$ when $\gamma = 1$, up to $-\$6,000$ when $\gamma = 3$.

Calibration approach.—Overall, our results suggest that uninsurable labor income shocks only moderately affect our welfare gain formula in Norway. How general is this result? To answer this question, we take a more standard model-based approach and, in appendix section D.1.2, we study the welfare effect of asset-price deviations in a Bewley-type model in which agents face a realistic labor income process with both transitory and permanent labor income shocks. We show that market incompleteness generates a relatively minor correction to the baseline sufficient statistic across a wide range of calibrations.

⁵⁵ Put differently, while saving decisions still react to income changes, there is little variability in income after retirement.




FIG. 10.—Accounting for uninsurable labor income risk. Panel A plots the covariances (20) aggregated over the horizons for each cohort and asset class. Panel B plots the average welfare gain by cohort, including the incomplete-market adjustment term (19). Note that the average welfare gain for $\gamma = 0$ (risk neutrality) is the same as the baseline one plotted in figure 7. Units are 2011 US dollars.

In appendix section D.1.2, we also consider a model in which individuals face idiosyncratic risk in portfolio returns, which is the dominant source of risk at the top of the wealth distribution (e.g., Fagereng et al. 2020; Gomez 2023). In this particular case, we can obtain simple closed-form formulas for the effect of market incompleteness on welfare gains: return risk effectively increases individual discount rates by the product of their relative risk aversion and the variance of return shocks. This effect is small for realistic calibrations: for instance, with $\gamma = 1$ and $\sigma = 10\%$, the effective increase in the discount rate is 1 percentage point ($\approx 1 \times 0.1^2$). Moreover, this formula makes it easy to adjust our baseline sufficient statistic to take this effect into account since it simply requires adjusting individual discount rates upward.⁵⁶

B. Borrowing Constraints and Collateral Effects

In the baseline model, individuals can take unrestricted positions in the liquid asset. In reality, individuals often face constraints on how much debt they can incur. More generally, the interest rate charged to an individual may increase with the debt level or decrease with the value of its assets. We now examine the effect of these borrowing constraints on our formula for welfare gains.

Theory.—For simplicity, we consider a two-asset version of the baseline model. The agent maximizes

$$V_{i,0} = \max_{\{C_{i,t}, N_{i,t}, B_{i,t}\}} \sum_{t=0}^{\infty} \beta^t U(C_{i,t}), \quad (21)$$

⁵⁶ Consistently with this idea, in our baseline approach, we use a relatively high discount rate of 5%, which is larger than the average rate of return on the deposits and debt over our time sample (see also the discussion in n. 36).

subject to budget constraints at each period $t \geq 0$,

$$C_{i,t} + (N_{i,t} - N_{i,t-1})P_t + B_{i,t}Q_{i,t} + \chi(N_{i,t} - N_{i,t-1}) = Y_{i,t} + B_{i,t-1} + N_{i,t-1}D_t. \quad (22)$$

The key difference, relative to the baseline model, is that we allow the price of the liquid asset $Q_{i,t}$ to be individual specific. More precisely, we assume that individuals face an interest-rate schedule (or “credit surface,” in the language of Geanakoplos [2016]):

$$Q_{i,t} = F(Q_t, B_{i,t}, N_{i,t}P_t), \quad (23)$$

where F is a smooth function of economy-wide reference price Q_t (e.g., “prime rate”), individual bond holdings $B_{i,t}$, and the market value of asset holdings $N_{i,t}P_t$. The dependence of $Q_{i,t}$ on bond holdings, $\partial Q_{i,t}/\partial B_{i,t}$, captures the idea that the interest rate faced by individuals may increase with individual debt balances. In contrast, the dependence of the interest rate on asset values, $\partial Q_{i,t}/\partial(N_{i,t}P_t)$, captures “collateral effects.” In particular, when $\partial Q_{i,t}/\partial(N_{i,t}P_t) > 0$, a higher value of asset holdings allows the individual to issue bonds at a higher price $Q_{i,t}$, that is, to borrow at a lower interest rate $Q_{i,t}^{-1}$, thereby capturing the key idea in collateral-constraint models that higher asset prices relax financial frictions. While we focus on smooth interest-rate schedules in the main text, we also study the case where individuals face “hard” borrowing constraints in appendix section D.2.3, which can be seen as a limiting case. The next proposition expresses the effect of borrowing constraints on the welfare gains of a deviation in asset prices.

PROPOSITION 3. In the presence of the interest-rate schedule (23), the welfare gain of individual i is

$$\frac{dV_{i,0}}{U'(C_{i,0})} = \sum_{t=0}^{\infty} \tilde{R}_{i,0 \rightarrow t}^{-1} \left((N_{i,t-1} - N_{i,t})dP_t - B_{i,t} \left(\frac{\partial Q_{i,t}}{\partial Q_t} dQ_t + \frac{\partial Q_{i,t}}{\partial(N_{i,t}P_t)} N_{i,t}dP_t \right) \right), \quad (24)$$

where

$$\tilde{R}_{i,0 \rightarrow t}^{-1} \equiv \prod_{s=0}^{t-1} \left(Q_{i,s} \left(1 + \frac{B_{i,s}}{Q_{i,s}} \frac{\partial Q_{i,s}}{\partial B_{i,s}} \right) \right).$$

This proposition shows that borrowing constraints affect our welfare gain formula in two ways. First, when $\partial Q_{i,t}/\partial B_{i,t} > 0$, meaning that the interest rate increases with the amount of debt ($Q_{i,t}^{-1}$ increases as $B_{i,t}$ becomes more negative), any increase in individuals’ debt level increases the interest rate on their entire debt balance so that they effectively face a higher marginal interest rate.⁵⁷ As a result, individuals discount more

⁵⁷ In fact, the schedule for the average interest payment $Q_{i,t}$ (23) implies the following schedule for the marginal interest payment

$$\partial_B(Q_{i,t}B_{i,t}) = Q_{i,t} + B_{i,t}\partial_B Q_{i,t} = Q_{i,t} \left(1 + \frac{B_{i,t}}{Q_{i,t}} \frac{\partial Q_{i,t}}{\partial B_{i,t}} \right).$$

heavily the future, which dampens the welfare effect of future deviations in asset prices: we call this the “discount rate channel.” Second, when $\partial Q_{i,t}/\partial(N_{i,t}P_t) > 0$, agents who hold levered positions in the asset benefit from a rise in asset prices through lower debt payments; this is what we call the “collateral channel.”⁵⁸ Importantly, in the presence of the collateral channel, asset holdings $N_{i,t}$ matter for the welfare effects of asset-price changes (in contrast to our baseline results, in which only asset sales $N_{i,t-1} - N_{i,t}$ mattered).

To formalize these two channels, we can rewrite the expression for welfare gains in the presence of borrowing constraints, as given in proposition 3, as a sum of three terms that capture, respectively, the welfare gains in the baseline model, the effect of the discount-rate channel, and the effect of the collateral channel:

$$\begin{aligned} \frac{dV_{i,0}}{U'(C_{i,0})} = & \underbrace{\sum_{t=0}^{\infty} R_{0 \rightarrow t}^{-1} \left((N_{i,t-1} - N_{i,t}) dP_t - B_{i,t} \frac{\partial Q_{i,t}}{\partial Q_t} dQ_t \right)}_{\text{baseline}} \\ & + \underbrace{\sum_{t=0}^{\infty} (\tilde{R}_{i,0 \rightarrow t}^{-1} - R_{0 \rightarrow t}^{-1}) \left((N_{i,t-1} - N_{i,t}) dP_t - B_{i,t} \frac{\partial Q_{i,t}}{\partial Q_t} dQ_t \right)}_{\text{discount-rate channel}} \\ & + \underbrace{\sum_{t=0}^{\infty} \tilde{R}_{i,0 \rightarrow t}^{-1} \left(-B_{i,t} \frac{\partial Q_{i,t}}{\partial(N_{i,t}P_t)} N_{i,t} dP_t \right)}_{\text{collateral channel}}. \end{aligned} \quad (25)$$

Implementation.—We now assume a specific parametric form for the interest-rate schedule (23). More precisely, we assume that the individual-specific (log) interest rate increases linearly with the loan-to-value ratio

$$Q_{i,t} = Q e^{-\xi \times \text{LTV}_{i,t}}, \quad (26)$$

where $\text{LTV}_{i,t} \equiv -B_{i,t}/(N_{i,t}P_{i,t})$. The parameter ξ governs the sensitivity of the interest rate to the loan-to-value ratio (and so the importance of borrowing constraints). The case $\xi = 0$ corresponds to the baseline model without borrowing constraints. Plugging the parametric form (26) into proposition 3 gives the following simplified formula for welfare gains.

COROLLARY 4. In the presence of a loan-to-value constraint represented by the interest-rate schedule (26), the welfare gain of individual i is

⁵⁸ While we do not discuss aggregation in the context of the collateral effects extension, banks charging lower mortgage interest rates in response to higher home values may also generate some losers, in particular, bank shareholders who indirectly hold mortgage debt as an asset. An offsetting effect is that lower loan-to-value ratios may lower bank monitoring costs, so bank shareholders may not be impacted much overall.

$$\frac{dV_{i,0}}{U'(C_{i,0})} = \sum_{t=0}^{\infty} \tilde{R}_{i,0 \rightarrow t}^{-1} \left((N_{i,t-1} - N_{i,t}) dP_t - B_{i,t} Q_{i,t} \left(\frac{dQ_t}{Q_t} + \xi \times \text{LTV}_{i,t} \times \frac{dP_t}{P_t} \right) \right), \quad (27)$$

where $\tilde{R}_{i,0 \rightarrow t}^{-1} \equiv \prod_{s=0}^{t-1} (Q_{i,s} (1 - \xi \times \text{LTV}_{i,s}))$.

This equation gives simple closed-form expressions for the effect of borrowing constraints on the welfare gains of a deviation in asset prices with two key modifications relative to the baseline: first, borrowing constraints increase the effective discount rate of agent i by $2\xi \times \text{LTV}_{i,t}$ (discount-rate channel); second, they increase the welfare exposure of asset-holders to rising asset prices by an amount equivalent to an increase in their annual rate of asset sales of $\xi \times \text{LTV}_{i,t}^2$ (collateral channel).⁵⁹

Results.—We estimate the parameter ξ by examining the relationship between individual mortgage interest rates and the ratio of mortgage debt to house value. Figure 11A presents a binned scatterplot of these two variables in the data. A clear positive relationship is visible: as loan-to-value ratios increase from 0 to 100%, mortgage interest rates increase by around 0.2 percentage points, from around 5% to 5.20%. Consistently with our parametric assumption (26), the relationship is approximately linear. In appendix section D.2.2, we estimate this relationship more formally using panel regressions and obtain values for ξ between 0.0025 and 0.004, depending on the controls included. The interpretation is that an increase of the loan-to-value ratio from zero to one is associated with a 0.25 percentage point to 0.4 percentage point (25–40 basis points) higher mortgage interest rate.

One potential concern, however, is that measurement error in the loan-to-value ratio or omitted variables may bias this coefficient downward. To deal with measurement errors, we also collect direct evidence of the interest-rate schedule posted by one of the Norwegian banks (Bulder Bank), which indicates a higher value of $\xi \approx 0.01$ (i.e., an increase in the loan-to-value ratio from zero to one implies a 1 percentage point rise in the interest rate).

We then implement the expression for welfare gains (27) in the data.⁶⁰ Figure 11B reports the average welfare gains in each cohort. Given the uncertainty regarding the value of ξ , we report results for a range of values $\xi \in \{0, 0.005, 0.01\}$, where the case $\xi = 0$ corresponds to the baseline welfare gain formula (i.e., same welfare gains as in fig. 7). We find that the effect of borrowing constraints is small. Figure A9 plots separately

⁵⁹ The first statement comes from the fact that the effective discount rate of agent i between t and $t + 1$ is $Q_{i,t}(1 - \xi \times \text{LTV}_{i,t}) = Q_t e^{-\xi \text{LTV}_{i,t}} (1 - \xi \times \text{LTV}_{i,t}) \approx Q_t e^{1 - 2\xi \times \text{LTV}_{i,t}}$. The second statement comes from the fact that the effect of collateral constraints at time t in (27) is $(-B_{i,t} Q_{i,t}) \times \xi \times \text{LTV}_{i,t} \times (dP_t/P_t) \approx \xi \times \text{LTV}_{i,t}^2 \times N_{i,t} dP_t$. In words, collateral constraints mean that every asset holder gains an additional exposure to a rise in asset prices that is equivalent to a “shadow” increase in their annual selling rate by $\xi \times \text{LTV}_{i,t}^2$.

⁶⁰ See app. sec. D.2.2 for more details.




FIG. 11.—Accounting for borrowing constraints and collateral effects. Panel A contains a binned scatterplot of the interest rate on mortgages and the ratio of mortgage debt to house value across individuals over the years 1994–2019. To do this figure, we focus on the sample of individuals with an interest rate in the 5%–95% range every year and a mortgage debt to housing value ratio between 0.2 and 0.99. We then demean the individual-specific interest rate and loan-to-value ratios by their average values each year, adding back the average interest rate over 1994–2019. Each dot represents a percentile of the sample, ranked according to their loan-to-value ratio. Panel B plots the average welfare gain by cohort, including the collateral effect adjustment term (27). The welfare gain with $\xi = 0$ is the same as in figure 7. Units are 2011 US dollars.

the impact of the discount-rate channel and the collateral channel; the two effects tend to have opposite signs (the discount-rate channel tends to be negative, while the collateral channel is always positive), and so the two effects tend to cancel out.

In appendix section D.2.2, we also estimate the effect of borrowing constraints at the individual level. We document a sizable dispersion across individuals (in particular, within cohorts). While the average effect of borrowing constraints is close to zero, they can increase welfare by up to \$98,000 for the top 1% of individuals most impacted via these constraints. Still, despite this sizable dispersion, the correction due to borrowing constraints remains small relative to the dispersion in baseline welfare gains (table 2).

Calibration approach.—Overall, our results suggest that borrowing constraints only moderately affect our welfare gain formula in Norway. How general is this result? To answer this question, we use the closed-form formula (27) to evaluate the effect of changing ξ (the elasticity of interest rates to the loan-to-value ratio) for welfare. Consider, for example, an economy with 10 times our baseline value for ξ ; that is, $\xi = 0.1$ (an increase in the loan-to-value ratio by one increases the interest rate by 10%). In this case, borrowing constraints would effectively increase the discount rate of a borrower with $LTV_{i,t} = 0.26$ (the average loan-to-value ratio in our sample) by $2\xi \times LTV_{i,t} = 5$ percentage points; that is, from 5% to 10%. Moreover, due to the collateral channel, the same borrower would experience an additional welfare gain equivalent to a net increase in house

sales by $\xi \times \text{LTV}_{i,t}^2 \approx 0.6$ percentage points, which remains modest.⁶¹ Overall, this type of computation makes it possible to assess the quantitative importance of borrowing constraints in different economic models or empirical settings.

How robust are these results to the assumption of a smooth interest-rate schedule? To answer this question, we study the case where individuals face “hard” borrowing constraints in appendix section D.2.3. As in the case of our interest-rate schedule (26), we derive simple closed-form formulas for the effect of these borrowing constraints on welfare gains in terms of three key parameters: (i) the proportion of individuals at the constraint, (ii) the wedge between the marginal rate of substitution across times and the interest rate for individuals at the constraint, and (iii) the loan-to-value ratio at the constraint. Overall, hard borrowing constraints lead to similar adjustments, both in spirit and in magnitude, to those derived in the baseline case where agents face a smooth interest-rate schedule.⁶²

C. Second-Order Effects

Proposition 1 characterizes the welfare gains of an infinitesimal deviation in asset prices. Hence, our sufficient statistic captures only the first-order effect of a noninfinitesimal deviation in asset prices. We now discuss theoretically and empirically the importance of higher-order effects.

Theory.—We first derive a formula for the welfare gain corresponding to a noninfinitesimal deviation in asset prices. As in the proof of proposition 1, we consider a noninfinitesimal deviation in the path of asset prices $\{\Delta Q_t, \{\Delta P_{k,t}\}_k\}_{t=0}^\infty$ and consider a continuum of intermediate economies where the deviation in asset prices is scaled by θ : $Q_t(\theta) = Q_t + \theta \Delta Q_t$ and $P_{k,t}(\theta) = P_{k,t} + \theta \Delta P_{k,t}$.

The money-metric welfare gain (i.e., the equivalent variation) of the deviation in asset prices indexed by θ is then the integral of infinitesimal welfare gains between 0 to θ :

$$\text{EV}_i(\theta) = \int_0^\theta \sum_{t=0}^\infty R_{0 \rightarrow t}^{-1}(u) \left(\sum_{k=1}^K (N_{i,k,t-1}(u) - N_{i,k,t}(u)) \Delta P_{k,t} - B_{i,t}(u) \Delta Q_t \right) du, \quad (28)$$

where $\{B_{i,t}(u), \{N_{i,k,t}(u)\}_k\}_{t=0}^\infty$ denote the path of asset holdings in the economy indexed by u after adjusting individual wealth at $t = 0$ to keep

⁶¹ Figure 7A and table A2 imply that the typical annual rate of home sales in a cohort is $\pm 4\%$.

⁶² The key difference between the two types of constraints is that soft borrowing constraints (interest-rate schedule) affect the formula for welfare gains for everyone. In contrast, hard borrowing constraints (strict borrowing limit) affects the formula for welfare gains for a limited number of agents—those hitting the limit. This stark distinction vanishes in a model with idiosyncratic risk, where all agents have some probability of hitting the constraint.

individual welfare fixed (i.e., Hicksian demands). When $u = 0$, this corresponds to the path of asset holdings in the baseline economy.

Using a trapezoidal approximation, we can then obtain a second-order approximation of welfare gains:⁶³

$$\begin{aligned} EV_i(\theta) = \sum_{t=0}^{\infty} R_{0 \rightarrow t}^{-1} \left(\frac{\theta}{2} \right) & \left\{ \sum_{k=1}^K \left(\frac{N_{i,k,t-1}(0) - N_{i,k,t}(0)}{2} + \frac{N_{i,k,t-1}(\theta) - N_{i,k,t}(\theta)}{2} \right) \Delta P_{k,t}(\theta) \right. \\ & \left. - \frac{B_{i,t}(0) + B_{i,t}(\theta)}{2} \Delta Q_t(\theta) \right\} + o(\theta^2). \end{aligned} \quad (29)$$

Compared with the first-order approximation in (12), this second-order approximation requires knowing how asset transactions respond to changes in asset prices (e.g., portfolio reshuffling). In particular, second-order effects are positive for individuals who respond to higher asset prices by selling more assets.⁶⁴ One implication is that the accuracy of our baseline first-order approximation depends on the extent to which the financial transactions of individuals react to deviations in asset prices. Another difference with the baseline first-order approximation is that one should use the average interest rate between the baseline and counterfactual economy, $R_t(\theta/2)$, to discount future transactions.

Implementation.—The empirical implementation of this second-order approximation requires additional assumptions: in contrast to the first-order approximation (12), we now need to specify what financial transactions would be if asset valuations had remained at their 1994 level. One way to do so would be to specify parametric forms for the utility function, the adjustment-cost functions, and individuals' beliefs about future asset prices.

Instead, we simply assume that, had valuations remained at their 1994 level, the quantity of transactions of a 30-year-old in each year would be the same as the transactions of a 30-year-old in 1994. Formally, we assume that the counterfactual transactions of individuals of age a are given by the following:⁶⁵

$$\begin{aligned} \bar{N}_{a,k,t}(\theta) - \bar{N}_{a,k,t-1}(\theta) &= \bar{N}_{a,k,0}(0) - \bar{N}_{a,k,-1}(0) \\ \bar{B}_{a,t}(\theta) &= G^t \bar{B}_{a,0}(0), \end{aligned} \quad (30)$$

⁶³ We use the notation $o(\theta^2)$ to denote a term that converges to zero faster than θ^2 as $\theta \rightarrow 0$. Note that

$$\int_0^\theta f(u)g(u)du = f(\theta/2) \frac{g(0) + g(\theta)}{2} \theta + o(\theta^2)$$

for any functions $f(\cdot)$, $g(\cdot)$ (this can be proven formally by showing that both sides of the formula have the same first and second derivatives with respect to θ at zero). Setting $f(u) = R_{0 \rightarrow t}^{-1}(u)$, and $g(u) = N_{i,k,t-1}(u) - N_{i,k,t}(u)$, gives (29).

⁶⁴ Martínez-Toledano (2022) empirically studies the effect of this type of market timing on wealth inequality.

⁶⁵ See app. sec. D.3 for more detail on the implementation.

where $\{\bar{B}_{a,t}(\theta), \{\bar{N}_{a,k,t}(\theta)\}_{k=0}^{\infty}\}$ denotes the average asset holdings of individuals of age a in year $1994 + t$ in the economy indexed by θ and $G = 1.01$ denotes the real per-capita growth rate of the economy in our sample period.

In alignment with (29), we discount future transactions using the discount rate $R(\theta/2) = 1.025$, which represents a midpoint between the net debt and deposit rates at the start of our sample (5%) and the rates at the end of our sample (0%; fig. 3). This adjustment effectively magnifies welfare gains relative to our baseline $R = 1.05$, as it implies that individuals discount less the profits or losses associated with future transactions.

Results.—We now examine how these counterfactual transactions differ from actual transactions. Figure 12A compares the actual and counterfactual housing and equity transactions for different age groups. The two quantities are very close, reflecting that real net housing and equity purchases have remained roughly constant over time. Figure 12B compares




FIG. 12.—Accounting for second-order effects. Panels A and B compare actual versus counterfactual transactions (if valuations had remained at their baseline level). More precisely, panel A plots $(\bar{N}_{a,k,t} - \bar{N}_{a,k,t-1})P_{k,t}$ and $(\bar{N}_{a,k,t}(\theta) - \bar{N}_{a,k,t-1}(\theta))P_{k,t}(\theta)$, averaged across years, for housing and equity. Panel B plots $\bar{B}_{a,t}Q_t$ and $\bar{B}_{a,t}(\theta)Q_t(\theta)$, averaged across years. Counterfactual asset transactions and bond holdings are estimated using (30). Panel C plots the average welfare gain at the first order and at the second order for individuals in each cohort (indexed by their age at the end of 1994), while panel D plots it separately for each asset class. Units are 2011 US dollars.

the actual and counterfactual debt balances. Net debt (debt minus deposits) has increased much more rapidly than one could expect from economic growth. Intuitively, the young must now borrow more to finance the purchase of houses whose values have grown faster than the economy. Still, overall, we find that counterfactual transactions are relatively similar to actual transactions, which suggests that second-order effects are likely to be moderate in our settings.

Figure 12C plots the total second-order welfare gain computed using (29), while figure 12D plots the second-order welfare gain class by asset class. The figure confirms this intuition: the overall effect of the second-order adjustment is small, and the results are quantitatively similar to those using our first-order approximation. Most of the effect is driven by the fact that we now use a much lower discount rate to discount the future (2.5% instead of 5%), which means that our second-order approximation tends to magnify the present value of gains and losses obtained with our baseline (first-order) approach. One additional negative effect for the young is driven by the cross elasticity of mortgage balance to house prices. As we have discussed, low mortgage rates have an important offsetting effect on home buyers who are hurt by rising house prices. If house prices had remained at their initial values, the young would have had lower mortgage balances, and as a result, they would have benefited less from the decrease in mortgage rates (see fig. 12D for a plot of the second-order correction by asset class).

D. *Extrapolation*

Our measure of welfare gains in proposition 1 expresses the welfare gains as the present value of all future transactions, multiplied by the path of future price deviations. However, as discussed in section III, we apply our formula only to a finite sample that ends in the year 2019 ($T = 25$). Therefore, our formula should be interpreted as the welfare gain associated with price deviations equal to zero after 2019 (i.e., assuming that valuations revert to the baseline in which asset prices grow at the same rate as dividends after 2019).

Implementation.—How important is this truncation for our results? To examine this question, we recompute our welfare gains with different assumptions about the behavior of asset prices after 2019. More precisely, we assume that, after the end of the sample, valuations revert to their baseline level according to a mean reversion parameter $\phi \in [0, 1]$. Formally, we assume that the valuation of asset class k at $t > T$ is given by⁶⁶

⁶⁶ See Campbell (2018) for an example of such an AR(1) specification for the logarithmic price-dividend ratio.

$$\log\left(\frac{PD_{k,t}}{\overline{PD}_k}\right) = \phi^{t-T} \log\left(\frac{PD_{k,T}}{\overline{PD}_k}\right), \quad \log\left(\frac{Q_t}{\overline{Q}}\right) = \phi^{t-T} \log\left(\frac{Q_T}{\overline{Q}}\right), \quad (31)$$

where $PD_{k,T}$ denotes the asset valuation in year 2019 and $\overline{PD}_k$ denotes the baseline level of the asset valuation defined in section III. Our baseline summary statistic, which considers asset-price deviations that stop after T , can be seen as the limit case $\phi = 0$. Figure 13A plots house prices obtained using this methodology up to 2060, for values of ϕ between 0 and 1. Note that, in all scenarios, we assume that housing valuations ultimately revert to their initial value ($\phi < 1$), consistent with the fact that asset valuations are stationary processes (Campbell and Shiller 1988).

To implement the sufficient statistic formula, we must also predict individuals' future transactions. To do so, we assume that the quantity of assets sold by a given cohort in any year after 2019 equals the average quantity sold by individuals of the same age in our sample, adjusted for economic growth (see app. sec. D.4 for details). This assumption is motivated by the fact that the quantity of transactions by age group has remained remarkably stable over our sample period, as discussed above (sec. V.C).

Results.—Figure 13 plots our estimated values for the average welfare gain in each cohort for different values of ϕ . As ϕ increases, two things happen. First, the graph of welfare gains is translated to the left. Intuitively, a high ϕ means that aging individuals sell more assets at elevated prices beyond our sample period, thereby increasing their welfare gains. However, this comes at the expense of young generations, unborn in 1994, who will ultimately purchase these assets. Second, the graph of welfare gains shifts up. This is because, as we show in the sectoral analysis in appendix C, individuals in Norway benefit on net from the rise in asset




FIG. 13.—Accounting for extrapolated changes in asset prices beyond 2019. Panel A plots the path of future house prices for different values of ϕ , constructed using (31). All paths are adjusted for inflation and normalized to one in 1980. Panel B plots the average welfare gain in each cohort with different assumptions about the future path of asset prices. Units are 2011 US dollars.

prices because they hold a positive amount of debt in the aggregate. As ϕ increases, higher valuations last longer, which means that, on average, welfare gains increase. However, doing the same exercise for sectoral welfare gains would reveal that this comes at the cost of a decrease in the total welfare gains for the government. Figure A12 decomposes the welfare gains by asset class. The decomposition shows that, as ϕ increases, most of the higher welfare gains in the population come from lower interest rates on debt.

VI. Conclusion

The main contribution of our paper is to provide a simple framework to quantify the welfare effects of fluctuations in asset prices. The core economic idea is that the welfare effect of changes in asset prices can be measured from the path of realized financial transactions: rising asset valuations benefit sellers and harm buyers. We implement our sufficient statistic formula using administrative data on financial transactions to quantify welfare gains and losses in Norway from 1994 to 2019.

Our empirical implementation generates four main findings. First, the rise in asset valuations had large redistributive effects; that is, they resulted in significant welfare gains and losses. At the same time, welfare gains differed substantially from naively calculated revaluation gains; in particular, individuals with the highest revaluation gains were not necessarily the ones with the highest welfare gains. Second, rising asset prices redistributed across cohorts, with the old benefiting at the expense of the young. Third, they redistributed across the wealth distribution, from the poor to the wealthy. Fourth, they also redistributed across sectors: declining interest rates benefited Norwegian households at the expense of the Norwegian government.

While our sufficient statistic approach is general, our empirical results are country specific. Differences in institutions, regulations, and norms shape the exposure of household welfare to asset-price changes. For instance, in Norway, public equities represent merely 3% of household wealth (see table A2), mortgages essentially all have floating interest rates, and the government is a net saver (through the sovereign wealth fund; see app. C). One can expect the welfare effect of deviations in asset prices to be different in countries such as the United States, where public equities represent roughly 20% of household wealth, mortgages tend to have fixed interest rates, and the government is a net debtor (see Greenwald et al. 2021).

Recent work building on our methods suggests that our sufficient statistic approach may also be helpful in other contexts. Del Canto et al. (2023) and Pallotti et al. (2024) study the money-metric welfare gains and losses from inflationary shocks of US and Euro-area households and

implement the corresponding welfare formulas using microdata. Similarly, Crawley and Gamber (2023) study the welfare consequences of the large asset-price and interest-rate changes on US households over the 2021–23 period rather than the longer-run trends considered here. Another valuable exercise would be to systematically quantify the welfare consequences of higher-frequency asset-price booms and busts that the literature has emphasized as essential drivers of wealth inequality dynamics (Kuhn, Schularick, and Steins 2020; Cioffi 2021; Martínez-Toledano 2022; Gomez 2025).

Finally, our results on the redistributive effect of asset prices raise important questions for optimal capital gains and wealth taxation. Answering such questions requires studying environments with changing asset prices using the tools from public finance. Aguiar, Moll, and Scheuer (2024) take some steps in this direction.

Data availability

Code replicating the tables and figures in this article can be found in Fagereng et al. (2024) in the Harvard Dataverse: <https://doi.org/10.7910/DVN/TJD0VI>.

References

- Aguiar, Mark, Benjamin Moll, and Florian Scheuer. 2024. “Putting the ‘Finance’ into ‘Public Finance’: A Theory of Capital Gains Taxation.” Working paper, London School Econ.
- Auclert, Adrien. 2019. “Monetary Policy and the Redistribution Channel.” *A.E.R.* 109 (6): 2333–67.
- Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2017. “From Saving Comes Having? Disentangling the Impact of Saving on Wealth Inequality.” Research Paper no. 18-8, Swedish House Finance, Stockholm.
- . 2020. “Rich Pickings? Risk, Return, and Skill in Household Wealth.” *A.E.R.* 110 (9): 2703–47.
- Baqae, David R., and Ariel Burstein. 2023. “Welfare and Output with Income Effects and Taste Shocks.” *Q.J.E.* 138 (2): 769–834.
- Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2018. “House Prices and Consumer Spending.” *Rev. Econ. Studies* 85 (3): 1502–42.
- Bertola, Giuseppe, and Ricardo J. Caballero. 1990. “Kinked Adjustment Costs and Aggregate Dynamics.” *NBER Macroeconomics Ann.* 5:237–88.
- Binsbergen, Jules H. Van. 2020. “Duration-Based Stock Valuation: Reassessing Stock Market Performance and Volatility.” Working Paper no. 27367, NBER, Cambridge, MA.
- Black, Sandra E., Paul J. Devereux, Fanny Landaud, and Kjell G. Salvanes. 2022. “The (Un)Importance of Inheritance.” Working Paper no. 29693, NBER, Cambridge, MA.
- Calvet, Laurent E., John Y. Campbell, Francisco Gomes, and Paolo Sodini. 2021. “The Cross-Section of Household Preferences.” Working Paper no. 28788, NBER, Cambridge, MA.

- Campbell, John. 2018. *Financial Decisions and Markets*. Princeton, NJ: Princeton Univ. Press.
- Campbell, John Y., and Robert J. Shiller. 1988. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *Rev. Financial Studies* 1 (3): 195–228.
- Campbell, John Y., and Luis M. Viceira. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford: Oxford Univ. Press.
- Campbell, John Y., and Tuomo Vuolteenaho. 2004. "Bad Beta, Good Beta." *A.E.R.* 94 (5): 1249–75.
- Canto, Felipe N. Del, John R. Grigsby, Eric Qian, and Conor Walsh. 2023. "Are Inflationary Shocks Regressive? A Feasible Set Approach." Working Paper no. 31124, NBER, Cambridge, MA.
- Catherine, Sylvain, Max Miller, James D. Paron, and Natasha Sarin. 2022. "Who Hedges Interest-Rate Risk? Implications for Wealth Inequality." Working paper.
- Catherine, Sylvain, Max Miller, and Natasha Sarin. 2020. "Social Security and Trends in Inequality." Working paper.
- Cioffi, Riccardo. 2021. "Heterogeneous Risk Exposure and the Dynamics of Wealth Inequality." Working paper.
- Cocco, Joao, Francisco Gomes, and Pascal Maenhout. 2005. "Consumption and Portfolio Choice over the Life Cycle." *Rev. Financial Studies* 18 (2): 491–533.
- Cochrane, John. 2020. "Wealth and Taxes, Part II." Blog post.
- Crawley, Edmund, and William L. Gamber. 2023. "Winners and Losers from Recent Asset Price Changes." FEDS Notes 2023-05-12, Board Governors Federal Res. System (US) May 2023.
- Dávila, Eduardo, and Anton Korinek. 2018. "Pecuniary Externalities in Economies with Financial Frictions." *Rev. Econ. Studies* 85 (1): 352–95.
- De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *Rev. Econ. Studies* 71 (3): 743–68.
- Doepke, Matthias, and Martin Schneider. 2006. "Inflation and the Redistribution of Nominal Wealth." *J.P.E.* 114 (6): 1069–97.
- Eitheim, Øyvind, and Solveig Erlandsen. 2005. "House Price Indices for Norway 1819–1989." *Scandinavian Econ. Hist. Rev.* 53:7–33.
- Fagereng, Andreas, Matthieu Gomez, Emilien Gouin-Bonenfant, Martin Holm, Benjamin Moll, and Gisle Natvik. 2024. "Replication Data for: 'Asset-Price Redistribution.'" Harvard Dataverse, <https://doi.org/10.7910/DVN/TJD0VI>.
- Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. "Heterogeneity and Persistence in Returns to Wealth." *Econometrica* 88 (1): 115–70.
- Fagereng, Andreas, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik. 2019. "Saving Behavior across the Wealth Distribution: The Importance of Capital Gains." Technical report, NBER, Cambridge, MA.
- Fagereng, Andreas, Martin Blomhoff Holm, and Kjersti N. Torstensen. 2020. "Housing Wealth in Norway, 1993–2015." *J. Econ. and Soc. Measurement* 45 (1): 65–81.
- Farhi, Emmanuel, and François Gourio. 2018. "Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia." Technical report, NBER, Cambridge, MA.
- Feiveson, Laura, and John Sabelhaus. 2019. "Lifecycle Patterns of Saving and Wealth Accumulation." Finance and Economics Discussion Series 2019-010. Washington, DC: Board Governors Fed. Reserve System, <https://doi.org/10.17016/FEDS.2019.010r1>.

- Flavin, Marjorie, and Takashi Yamashita. 2011. "Owner-Occupied Housing: Life-Cycle Implications for the Household Portfolio." *A.E.R.* 101 (3): 609–14.
- Gabaix, Xavier, and Ralph S. J. Koijen. 2021. "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis." Technical report, NBER, Cambridge, MA.
- Geanakoplos, John. 2016. "The Credit Surface and Monetary Policy." In *Progress and Confusion: The State of Macroeconomic Policy*, 143–53. Cambridge, MA: MIT Press.
- Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull. 2020. "Intergenerational Redistribution in the Great Recession." *J.P.E.* 128 (10): 3730–78.
- Gomez, Matthieu. 2023. "Decomposing the Growth of Top Wealth Shares." *Econometrica* 91 (3): 979–1024.
- . 2025. "Wealth Inequality and Asset Prices." *Rev. Econ. Studies*, <https://doi.org/10.1093/restud/rdaf008>.
- Gomez, Matthieu, and Émilien Gouin-Bonenfant. 2024. "Wealth Inequality in a Low Rate Environment." *Econometrica* 92 (1): 201–46.
- Greenwald, Daniel L., Matteo Leombroni, Hanno Lustig, and Stijn Van Nieuwerburgh. 2021. "Financial and Total Wealth Inequality with Declining Interest Rates." Technical report, NBER, Cambridge, MA.
- Greenwald, Daniel L., Martin Lettau, and Sydney C. Ludvigson. 2019. "How the Wealth Was Won: Factors Shares as Market Fundamentals." Technical report, NBER, Cambridge, MA.
- Haig, Robert M. 1921. "The Concept of Income—Economic and Legal Aspects." In *The Federal Income Tax*, 1–21. New York: Columbia Univ. Press.
- Kaldor, Nicholas. 1955. *An Expenditure Tax*. London: Routledge.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy according to HANK." *A.E.R.* 108 (3): 697–743.
- Kiyotaki, Nobuhiro, Alexander Michaelides, and Kalin Nikolov. 2011. "Winners and Losers in Housing Markets." *J. Money, Credit and Banking* 43 (2–3): 255–96.
- Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit Cycles." *J.P.E.* 105 (2): 211–48.
- Krugman, Paul. 2021. "Pride and Prejudice and Asset Prices." Blog post.
- Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins. 2020. "Income and Wealth Inequality in America, 1949–2016." *J.P.E.* 128 (9): 3469–519.
- Kuvshinov, Dmitry. 2023. "Asset-Class-Specific Discount Rates." Working paper.
- Lucas, Robert E. 2000. "Inflation and Welfare." *Econometrica* 68 (2): 247–74.
- Martínez-Toledano, Clara. 2022. "House Price Cycles, Wealth Inequality and Portfolio Reshuffling." Working paper.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. New York: Oxford Univ. Press.
- Merton, Robert C. 1973. "An Intertemporal Capital Asset Pricing Model." *Econometrica* 41 (5): 867–87.
- Mian, Atif R., Ludwig Straub, and Amir Sufi. 2020. "The Saving Glut of the Rich." Technical report, NBER, Cambridge, MA.
- Miao, Jianjun, and Pengfei Wang. 2012. "Bubbles and Total Factor Productivity." *A.E.R.* 102 (3): 82–87.
- Moll, Benjamin. 2020. "Comment on 'Sources of US Wealth Inequality: Past, Present, and Future.'" *NBER Macroeconomics Ann.* 35.
- MSCI. 2016. "Real Estate Analytics Portal." <https://www.msci.com/indexes/private-asset-indexes>.
- Paish, F. W. 1940. "Capital Value and Income." *Economica* 7 (28): 416–18.

- Pallotti, Filippo, Gonzalo Paz-Pardo, Jiri Slacalek, Oreste Tristani, and Giovanni L. Violante. 2024. "Who Bears the Costs of Inflation? Euro Area Households and the 2021–2023 Shock." *J. Monetary Econ.* 148:103671.
- Piketty, Thomas, and Gabriel Zucman. 2014. "Capital Is Back: Wealth-Income Ratios in Rich Countries 1700–2010." *Q.J.E.* 129 (3): 1255–310.
- Saez, Emmanuel, and Stefanie Stantcheva. 2016. "Generalized Social Marginal Welfare Weights for Optimal Tax Theory." *A.E.R.* 106 (1): 24–45.
- Saez, Emmanuel, Danny Yagan, and Gabriel Zucman. 2021. "Capital Gains Withholding." Technical report, Univ. California, Berkeley.
- Shiller, Robert. 1981. "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *A.E.R.* 71 (3): 421–36.
- Simons, Henry C. 1938. *Personal Income Taxation: the Definition of Income as a Problem of Fiscal Policy*. Chicago: Univ. Chicago Press.
- Sinai, Todd, and Nicholas S. Souleles. 2005. "Owner-Occupied Housing as a Hedge Against Rent Risk." *Q.J.E.* 120 (2): 763–89.
- Toda, Alexis Akira, and Kieran Walsh. 2015. "The Double Power Law in Consumption and Implications for Testing Euler Equations." *J.P.E.* 123 (5): 1177–200.
- Vissing-Jørgensen, Annette. 2002. "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution." *J.P.E.* 110 (4): 825–53.
- Whalley, John. 1979. "Capital Gains Taxation and Interest Rate Changes: An Extension of Paish's Argument." *National Tax J.* 32 (1): 87–91.
- Wolff, Edward N. 2022. "The Stock Market and the Evolution of Top Wealth Shares in the United States." *J. Econ. Inequality* 20:53–66.