IARIW 2022

Monday 22 - Friday 26 August

The Business Cycle Dynamics of the Wealth Distribution

Jesse Bricker
(Federal Reserve Board, United States)

Kamila Sommer
(Federal Reserve Board, United States)

Joseph Briggs
(Goldman Sachs, United States)

Sarah Friedman
(University of Chicago, United States)

Paper prepared for the 37th IARIW General Conference

August 22-26, 2022

Session 6A-1, Dynamics of the Wealth Distribution around the World I

Time: Friday, August 26, 2022 [9:00-10:30 CEST]

The Business Cycle Dynamics of the Wealth Distribution*

Jesse Bricker[†] Joseph Briggs[‡] Sarah Friedman[§] Kamila Sommer[¶]

July 21, 2022

Abstract

We develop new household balance sheet measures to quantify the business cycle dynamics of the wealth distribution in the US since 1989. After introducing these data and establishing their credibility, we show that heterogeneous exposures to aggregate price risk accounts for most cyclical variation in the wealth distribution, but that group-specific factors also have a non-negligible effect. We also show that increases in the output gap and unemployment rate increase wealth inequality, as do accommodative monetary policy shocks. Finally—unlike after the Great Recession—our data suggest that household balance sheets should provide a tailwind to the recovery from the pandemic.

1 Introduction

Recent macroeconomic research suggests heterogeneity in household income and wealth is important in propagating aggregate shocks and determining macroeconomic outcomes (see, e.g., Krueger et al. (2016), Kaplan et al. (2018)). Despite recent advances in measuring long-term trends in the wealth distribution (Saez and Zucman (2016), Kuhn et al. (2018), Smith et al. (2019)), the income

*We are grateful for comments at the 2021 Society for Economic Dynamics (SED) conference and seminars at the Dutch Central Bank and European Central Bank. We thank Ella Deeken and Jackie Blair for excellent research assistance. The analysis and conclusions set forth here are those of the authors and do not indicate concurrence by other members of the research staff, the Board of Governors, or the Federal Reserve System.

[†]Federal Reserve Board, jesse.bricker@frb.gov

[‡]Goldman Sachs, jsbriggs87@gmail.com

[§]University of Chicago, sarah.evelyn.f@gmail.com

[¶]Federal Reserve Board, kamila.sommer@frb.gov

distribution (Piketty et al. (2018)), as well as business cycle dynamics of the income distribution (Guvenen et al. (2012), Guvenen et al. (2019)), data on the business cycle dynamics of household portfolios has until now been limited. This scarcity of such data has limited the understanding of the evolution of household balance sheets during recessions, as well as the ability to test and discipline theoretical advances in heterogeneous agent models featuring aggregate risk (Ahn et al. (2017)).

This paper bridges this empirical gap by providing the first analysis of the cyclical dynamics of the wealth distribution. Our analysis relies on the Distributional Financial Accounts (DFAs), a new data set now published quarterly by the Federal Reserve Board. The DFAs contain quarterly time series of household balance sheets (including disaggregated assets and liabilities) for different segments of the wealth distribution from 1989 to the present. The data's quarterly measures permit higher frequency observation of household portfolios than are available in other data sets, and therefore allow a unique look at the wealth distribution's evolution through recent business cycles.

Our paper proceeds in several stages. Since the credibility of our analysis hinges critically on the accuracy of the DFA data, we first describe the methodology underlying our data. The DFAs infer quarterly changes across the wealth distribution by applying established "temporal disaggregation" methods (see, e.g., Chow and Lin (1971), Fernandez (1981), Litterman (1983), Mönch and Uhlig (2005)) to the Survey of Consumer Finances (a triennial measure of the household wealth distribution) and the Financial Accounts of the United States (a quarterly measure of aggregate household wealth). To demonstrate that our data are credible, we show that our methodology recovers out-of-sample SCF balance sheets with reasonable accuracy. We also provide standard errors to show that our estimates are well identified, show that our methodology is robust to alternative error assumptions and estimation methods, and confirm that trends in inequality established in other studies are also well captured by our data.[1] Taken altogether, this analysis strongly suggests that our data provides credible insights into higher-frequency dynamics of the wealth distribution.

We then provide a first look at the cyclical dynamics of the wealth distribution captured by our data. We find that wealth inequality is strongly pro-cyclical, since wealth gains and losses incur disproportionately to the wealthiest of households during the economic expansions and downturns.

[1]See, e.g., Wolff et al. (2012), Piketty (2013), Bricker et al. (2016), Saez and Zucman (2016), Ríos-Rull and Kuhn (2016), Kuhn et al. (2018)), Smith et al. (2019)

These cyclical dynamics can be explained in large part by systematic differences in portfolio composition and exposure to aggregate price changes across the wealth distribution.[2] Intuitively, the wealthiest 1% of households hold riskier assets—mainly in the form of business equity—which have both higher expected but also more cyclical aggregate price returns, while households outside the top 1% hold a mixture of housing and equity wealth with less exposure to cyclical risk (consistent with Kuhn et al. (2018)). However, we also show that group-specific factors affect wealth accumulation, particularly for wealthy households.

Next, we exploit our quarterly time series to estimate the wealth distribution's response to shocks to key economic variables. An active structural literature considers the redistributional aspect of aggregate shocks,[3] but empirical tests of these predictions has thus far been limited, primarily due to data limitations. Applying standard local projection method (see e.g., Jordá (2005)) to our data, we show that increases in the output gap and unemployment rate increase wealth inequality, as do accommodative monetary policy shocks. These patterns are qualitatively similar to predictions from the quantitative macro literature and provide a new set of numerical targets for models to match.

Finally—and motivated by empirical work (see, e.g., Mian et al. (2013), Mian and Sufi (2014)) that has emphasized the importance of the distribution of household balance sheet in determining aggregate economic outcomes—we consider the evolution of household balance sheets through recent recessions. We document very different trends in the Great Recession and COVID-19 pandemic, with an unprecedented deterioration in household balance sheets during the Great Recession due to a slower recovery in employment and house prices, and an unprecedented strengthening during the COVID-19 recession due to significant fiscal support. These patterns suggest that household balance sheets should provide a tailwind to the recovery from the pandemic, in notable contrast to the recovery following the Great Recession.

Taken altogether, our findings support a key message: different segments of the wealth distribution hold different portfolios and are exposed to different shocks. Accurately modeling wealth distribution dynamics, then, requires a sufficiently rich portfolio choice problem to capture the dif-

[2]See Bach et al. (2020) and Fagereng et al. (2020) for evidence of differential risk exposures in other countries.

[3]See, e.g., Doepke and Schneider (2006), Rubio (2011), Calza et al. (2013), Auclert (2019), Kaplan et al. (2018), Luetticke (2018), and Kekre and Lenel (2019), Garriga et al. (2015).

ferential exposure to shocks associated with different asset classes. The data and analysis in this paper highlight the need for a workhorse three-asset macro model with riskless bonds, illiquid real estate, and risky capital to capture both the long-run and short-run dynamics of the wealth distribution and its effects on macroeconomic aggregates (see, e.g., Kaplan and Violante (2014)) and risk premia (Kekre and Lenel (2019)) in future generations of macroeconomic modeling.

Our paper relates to several strands of economic research. First, several recent studies, including Saez and Zucman (2016) and Smith et al. (2019), have attempted to improve measurement of wealth inequality by capitalizing income tax returns, while Kopczuk and Saez (2004) do the same using estate tax filings. In Section 3.3 we benchmark our data against Saez and Zucman (2016) and Smith et al. (2019) and find that patterns are qualitatively consistent with both, but more similar quantitatively to Smith et al. (2019). More importantly, our work focuses on higher-frequency changes in wealth that are not well-measured by these prior studies that only provide annual measures and rely on taxable income that may be manipulated over the business cycle to minimize tax burdens (Dowd et al. (2019)).

A second set of empirical papers rely on the SCF to measure inequality (see, e.g., Wolff et al. (2012), Bricker et al. (2016), Ríos-Rull and Kuhn (2016)). In addition to the providing higher frequency distributional measures, we build upon this work by fully reconciling the SCF with official aggregate household balance sheet measures.[4] More recent work by Kuhn et al. (2018) combine triennial SCF data post-1983 with annual survey data from 1952-1973 to provide insight into long-term dynamics of the distribution of wealth. Our findings that house price changes affect the middle class's wealth and that changes in equity prices affect the top of the wealth distribution support these findings, but the annual data used by these authors to study higher-frequency changes in wealth are only available pre-1973. Our data therefore complements this work by examining changes in the wealth distribution at a higher frequency and during the more recent period when the SCF sample provides more insight into the balance sheets of the wealthiest families.

[4]By merging survey data with national accounting data, the DFAs provide a comprehensive new measure of the distribution of aggregate household wealth. Thus, the DFAs help overcome some of the challenges that have impeded past efforts to integrate microeconomic data with macroeconomic analysis (see Carroll (2014) for a rich discussion of this issue). Earlier attempts to reconcile the SCF and Financial Accounts include Avery et al. (1987), Eller (1994), Antoniewicz (1996), Maki and Palumbo (2001), Henriques and Hsu (2014), and Dettling et al. (2015). Our reconciliation differs from these earlier attempts because our goal is to align the SCF with the Financial Accounts as opposed to either aligning the Financial Accounts with the SCF or restricting household balance sheet to lines that can be readily compared.

Third, our work relates to a number of papers that model the dynamics of the wealth distribution, including Chatterjee (1994), Ferreira (1995), Piketty (1997), and Álvarez-Peláez and Díaz (2005), all of which focus on transition dynamics of the wealth distribution. Similarly, Benhabib et al. (2011), Benhabib et al. (2016), and Benhabib et al. (2017) characterize the long-term dynamics of the wealth distribution in models with intergenerational wealth transfers, idiosynchratic earnings risk, and persistent, and stochastic returns to capital. We complement these and similar studies by presenting model-free estimates of wealth distribution dynamics at higher frequencies.

Fourth, heterogeneous agent macro-models featuring aggregate risk are not new (see, e.g., Krusell and Smith (1998)) but gained significant traction in macroeconomic research following the Great Recession (see Kaplan and Violante (2018) for a review). In addition, recent work examines the interaction between inequality and aggregate economic measures like inflation (see, e.g., Doepke and Schneider (2006) and Meh et al. (2010), aggregate demand (see, e.g., Guerrieri and Lorenzoni (2017) and Auclert and Rognlie (2018)), and fiscal policy (see, e.g., Kaplan and Violante (2014), McKay and Reis (2016), Hagedorn et al. (2019), Bhandari et al. (2018)), as well as monetary policy transmission (see, e.g., Werning (2015), Mckay et al. (2016), Kaplan et al. (2018), Gornemann et al. (2016), Luetticke (2018) and Auclert (2019), Kekre and Lenel (2019)). Our paper contributes to this increasingly important literature by providing the data needed to test business cycle dynamics of macroeconomic models that incorporate balance sheet heterogeneity.

Our paper proceeds as follows. Section 2 introduces the methodology used in producing the DFAs. Section 3 shows that the data credibly captures changes in household balance sheets across the wealth distribution and provides an overview of patterns in the DFAs. Section 4 provides a detailed look at the dynamics of the wealth distribution, Section 5 uses our data to estimate impulse responses to various economic shocks, and Section 6 applies our data to examine the dynamics of household portfolios during the Great Recession and Covid-19 period. Finally, Section 7 concludes.

2 Methodology

The data used in this paper build on and extend the Distributional Financial Accounts (DFAs)—the quarterly distributional data published by the Federal Reserve Board. The DFAs are a new data set which track the distribution of household wealth for four broad wealth groups—which we expand upon in this paper—as well as several other socio-economic groups.[5] The data complement and expand on other existing measures of household wealth inequality (Saez and Zucman (2016), Bricker et al. (2016), Smith et al. (2019)) by providing high frequency measures of changes in the wealth distribution that are consistent with the US national economic accounts—a feature that is lacking in most studies of wealth inequality.[6]

The DFA data are constructed by integrating two data sets produced by the Federal Reserve Board: the Financial Accounts of the United States and the Survey of Consumer Finances (SCF). The Financial Accounts are U.S. national accounts that provide regular, timely, and comprehensive measures of the aggregate household balance sheet, with quarterly data releases within 10 weeks of the end of the quarter. These aggregate data, though, do not provide any information on the distribution of the balance sheet by household wealth, income, or any other characteristics. Complementing these aggregate series, the SCF collects detailed household balance sheet information for a representative cross-section of U.S. households (including of very wealthy households). However, the SCF is fielded triennially, which limits the ability of researchers to use these data to study changes in the distribution of household wealth during business cycles. The crux of the construction of the quarterly distributional data is to combine the SCF's rich distributional information with the Financial Accounts' quarterly aggregates in a manner that is conceptually consistent with both data sets.

The construction of the DFA data proceed in three key steps, the first of which is to conceptually align the aggregate balance sheet in the SCF to the Financial Accounts' household net worth Table B.101.h—a recent data product that breaks out households from the long-standing table B.101 that also includes nonprofit organizations. Second, with the reconciled triennial SCF balance sheet in hand, we calculate the shares of the total SCF wealth held by each wealth group, and then estimate the SCF balances for each asset and liability category for quarters where the SCF is not observed using temporal disaggregation methods (see Section 2.2). The estimation allows us to populate

[5]The data were first published in March 2019 and can be found at: https://www.federalreserve.gov/releases/efa/efa-distributional-financial-accounts.htm.

[6]See Deaton (2005), Carroll (2014), and Ahn et al. (2017) for discussion on the importance of distributional measures consistent with national statistics.

the missing SCF quarters with estimates for each distributional stratum, all the way to the present, as though the SCF were observed quarterly. Third, from the reconciled, quarterly SCF data, we construct the share of wealth held by each wealth group and apply this share to the aggregate Financial Accounts' balance to evaluate the amount of the national wealth held by this group.

2.1 Reconciling the Financial Accounts and the SCF

The first step in constructing the DFAs is reconciling the measurement concepts used in the SCF with those used in the Financial Accounts. Aggregate household wealth data in the Financial Accounts are found in Table B.101.h, which reports total U.S. household wealth and its 19 main components. The detailed nature of the SCF questionnaire allows us to create either direct or close analogues of most Financial Accounts balance sheet concepts. These categories include some of the largest balance sheets items (e.g., real estate, equities), but also a number of smaller asset and liability categories. Batty et al. (2019) provides full details of the reconciliation exercise. We summarize these details in Appendix A.1.

The results of the SCF-B.101.h reconciliation are summarized in Table 1 by showing the ratio of the two measures for each line of Table B.101.h, for each wave of the SCF since 1989. A ratio of 100% would indicate that the two series match exactly, while lower (higher) percentages indicate that the reconciled SCF is understated (overstated) relative to the B.101.h total.[7] For reference, the figure also shows the level of the B.101.h and SCF series in 2019 in billions of dollars.

Overall, the topline numbers (assets, liabilities, and net worth) from our reconciled SCF balance sheet are quite similar to those from B.101.h.[8] For example, in 2019, reconciled SCF assets aggregate to \$125 trillion, compared with \$123 trillion on B.101.h. Reconciled SCF liabilities aggregate to \$14 trillion, versus \$15 trillion on B.101.h. Averaging across SCF waves, aggregate SCF net worth is very close (at 104%) to B.101.h net worth.[9] Additionally, the two data sets also align reasonably well for most of the underlying asset and liability categories. Most importantly, for large

[7] An exact ratio of 100% (and double asterisks) implies that the B.101.h total is distributed to SCF respondents using an asset- or liability-specific imputation rule and the B.101.h and reconciled SCF lines match by construction. See Batty et al. (2019) for details.

[8] While the reconciliation done here is the most comprehensive to date, for reasons discussed below, even with the best effort at reconciling, one would not expect to get a perfect alignment between the two data sources, as various types of measurement error are likely to affect the alignment of the Z.1 and the reconciled SCF totals.

[9] While the match is reasonable in all years, the alignment further improves in recent years. For example, in 2019 the ratio of SCF to B.101.h assets, liabilities, and net worth are 101%, 92% and 103%.

asset categories that disproportionately affect the distribution of wealth, differences between the reconciled SCF and B.101.h balance sheets are quite small, despite the very different approaches in constructing the two data sets.

Still, there are several smaller asset and liability categories (e.g., consumer durable goods, time deposits, or debt securities, consumer credit), where the match is imperfect. We are naturally less concerned regarding differences in smaller balance sheet line that have a smaller impact on the distribution of household wealth, but nevertheless rescale SCF sample weights proportionally via a scaling factor that is independent of observables–consistent with prior evidence that mismeasurement in the SCF is not driven by any one group (Bricker et al., 2016)–to ensure that SCF totals match the Financial Accounts aggregate.[10] Batty et al. (2019) provides a battery of tests showing that the DFAs' distributional data are robust to a number of alternative reconciliation assumptions.

2.2 Estimating SCF Balance Sheets in Unobserved Quarters in the DFAs

With the reconciled SCF-B.101h in hand, the next challenge is estimating the reconciled SCF balance sheets for quarters where SCF measures are not available. This "temporal disaggregation" problem of imputing higher-frequency data from lower-frequency observations has been well-studied, beginning with Chow and Lin (1971) and extended to allow for richer error processes by Fernandez (1981), and Litterman (1983). We follow standard methods and use the empirical relationship between the SCF, the Financial Accounts, and other economic data to estimate latent reconciled SCF balance sheets in quarters when only the Financial Accounts and macroeconomic data are available.

The baseline approach adopted in the DFAs is based on methodology proposed in Chow and Lin (1971). The Chow-Lin method assumes that the target series Y (in our case, the level of each reconciled SCF balance sheet line) that requires imputation/forecasting comes from a higher-frequency underlying series X. Let B be the matrix which selects the observed elements Y from the underlying series X:

$$Y = B'X. \tag{1}$$

[10]The SCF is precluded from sampling the Forbes 400, though the SCF does cover the lower part of that set of families (due to incomplete coverage of the wealthiest in the Forbes list (Vermuelen (2018) and Bricker Hansen and Volz (2019)). However, the DFAs use weights that have been adjusted to incorporate supplemental wealth data from Forbes 400 households.

Table 1: The Ratio of the Reconciled SCF Household Balance Sheet to B.101.h

	Ratios in SCF Years												Recent Levels ($ billion)	
	1989	1992	1995	1998	2001	2004	2007	2010	2013	2016	2019	Average	FA 2019Q3	SCF 2019
Total Assets	97	91	91	97	107	104	102	106	102	108	101	101	123197	124951
Nonfinancial assets	97	93	92	99	97	106	114	117	115	110	106	104	35324	37405
Real estate (1)	108	105	102	110	105	114	123	131	127	119	114	114	29613	33715
Consumer durable goods (2)	60	48	58	58	64	67	63	62	63	66	65	61	5711	3690
Financial assets	97	91	90	96	113	103	96	101	97	107	100	99	87873	87546
Checkable deposits and currency	65	45	50	85	136	194	809	240	130	145	189	190	817	1545
Time deposits and short-term investments	60	63	59	65	58	63	51	54	42	47	45	55	9798	4437
Money market fund shares	83	80	76	59	73	102	71	93	133	128	102	91	1968	2006
U.S. government and municipal securities	70	53	54	54	106	95	94	71	81	100	77	78	4390	3359
Corporate and foreign bonds	88	51	27	31	67	60	59	45	62	105	90	62	888	799
Other loans and advances	333	123	186	63	62	43	34	52	71	54	62	98	788	485
Mortgages	110	94	91	84	97	97	102	96	177	275	155	125	81	126
Corporate equities and mutual fund shares	143	119	119	131	186	141	111	128	111	129	112	130	26685	29789
Life insurance reserves**	100	100	100	100	100	100	100	100	100	100	100	100	1719	1719
Pension entitlements (3)	101	100	100	100	100	100	100	100	100	100	100	100	27166	27145
Equity in noncorporate business	103	91	79	85	97	94	103	136	115	134	121	105	12316	14872 (5)
Miscellaneous assets**	101	101	100	100	101	100	100	100	101	100	100	100	1257	1263
Total Liabilities	79	81	79	86	82	89	84	88	87	88	92	85	15300	14059
Home mortgages (5)	81	85	85	93	89	95	87	92	95	95	103	91	10414	10744
Consumer credit	59	57	55	60	52	59	65	69	59	68	66	61	4114	2727
Depository institution loans n.e.c.	1897	3133	278	210	470	-3152	216	89	95	36	39	301	256	99
Other loans and advances	99	99	99	97	99	99	98	90	98	99	94	97	480	452
Deferred and unpaid life insurance premiums	102	102	99	100	99	99	99	98	98	99	98	99	37	36
Net worth	100	93	93	99	112	107	106	110	105	111	103	104	107897	110892

Notes:

(1) All types of owner-occupied housing including farm houses and mobile homes, as well as second homes that are not rented, vacant homes for sale, and vacant land. At market value.

(2) At replacement (current) cost.

(3) Includes public and private defined benefit and defined contribution pension plans and annuities, including those in IRAs and at life insurance companies. Excludes social security.

(4) Net worth of nonfinancial noncorporate business and owners' equity in unincorporated security brokers and dealers.

(5) Includes loans made under home equity lines of credit and home equity loans secured by junior liens.

In our application, Y is observed every 3 years, while X is quarterly.[11]

The Chow-Lin method uses higher frequency indicator series, denoted here by Z, to impute/forecast the underlying series X. It does this by supposing that X and Z have a linear relationship[12]

$$X = \beta' Z + u, \tag{2}$$

where the residual vector u is mean zero with covariance matrix $V = \mathbb{E}[uu']$. Linearity combined with Equation 1 implies that

$$Y = B' Z' \beta + B' u. \tag{3}$$

The Chow-Lin method solves the multiple regression model specified by Equations 1 and 3 to obtain an estimate of $\hat{X}$ given observations Y and Z and covariance matrix V. Appendix A.2 describes the solutions proposed by Chow and Lin (1971), Fernandez (1981), and Litterman (1983).

Mönch and Uhlig (2005) (among others) show that the Chow-Lin method can be recast in a state-space format that permits consistent estimates via maximum likelihood estimation of the resulting Kalman smoother problem. Given that this framework is more commonly employed in modern macroeconomic research, we next provide the state-space representation of the problem (assuming u is ar(1) with innovation η_t).

Let ζ_t be the state vector, which includes the most recent observations of y_t and error term u_t.[13] Then a generalized version of our state equation can be expressed as

[11]Formally, we suppose that $Y = [y_1, y_2, , y_m]'$ is observed m times, with $k-1$ unobserved periods between observations and e periods to extrapolate after the last observation of Y so that $X = [x_1, x_2, , x_n]'$ with observation y_m of Y corresponding to observation $x_{(m-1)k+1}$ of X. The $n \times m$ matrix B can thus be written as

$$B = \begin{bmatrix} \iota & \dots & 0_{(m-1)k} \\ 0_{(m-1)k} & \dots & \iota \\ 0_e & \dots & 0_e \end{bmatrix}$$

where ι represents a k-dimensional column vector with one as the first element and zero elsewhere, and where 0_j denotes a j-dimensional column vector of zeros.

[12]Z can be expressed as an $n \times q$ matrix $Z = [Z_1, Z_2, , Z_q]$, where each Z_i denotes a separate column vector $Z_i = [z_{i,1}, z_{i,2}, \dots, z_{i,n}]'$ corresponding to the i^{th} indicator series.

[13]This is defined equivalently to Equation 18 in Appendix A.2.

$$\zeta_t = \begin{pmatrix} y_t \\ u_t \end{pmatrix} \begin{pmatrix} a & \rho \\ 0 & \rho \end{pmatrix} \begin{pmatrix} y_{t-1} \\ u_{t-1} \end{pmatrix} + \begin{pmatrix} x_t'\beta \\ 0 \end{pmatrix} + \begin{pmatrix} 1 \\ 1 \end{pmatrix} \eta_t. \tag{4}$$

The observation equation is then expressed as

$$Y_t = B_t'\zeta_t \tag{5}$$

$$B_t = \begin{cases} [1,0] \;\; for\; t = 0, 12, 24... \\ [0,0] \;\; otherwise, \end{cases} , \tag{6}$$

yielding a linear system equivalent to that defined by Equations 1 and 3. Further assuming normal errors yields a likelihood function that can be estimated using a Kalman smoother. Asymptotically, both the Chow Lin and the state-space formulation provide consistent estimates but might differ in small samples and due to the assumed normal errors in our likelihood function.

Overall, the results from estimating the state-space model align closely with results from our baseline estimates; these comparisons (augmented with standard errors) are presented in Figure 4 in Section 3.1. The similarity in these two estimation methods indicates the robustness of our results to alternative, and likely more familiar, approaches to estimating latent household balance sheets for our wealth groups.

As a final step in producing the DFAs, we project the Financial Accounts data onto the reconciled quarterly SCF asset and liability share estimates. To do so, we define $\gamma_t^{j,p}$ as the level of the asset or liability indexed by balance sheet line j, for wealth quantile group p, in quarter t, and let Γ_t^j denote the corresponding line from the B.101.h balance sheet. Defining group p's asset or liability share of balance sheet line j in quarter t as its share of the total reconciled SCF balance sheet line

$$\omega_t^{j,p} = \frac{\gamma_t^{j,p}}{\sum_k \gamma_t^{j,k}}$$

and multiplying these balance sheet shares by the total B.101.h balance sheet line

$$\bar{\gamma}_t^{j,p} = \Gamma_t^j \omega_t^{j,p}$$

yields estimates of assets and liability levels for each quantile that aggregate to the Financial Accounts household balance sheet table.[14]

2.3 The Data Overview

Our core data set is presented in Figure 1, which shows the real wealth levels (Panel (a)) and shares (Panel (b)) of net wealth for households across the wealth distribution. The groups used in this paper are more granular than those used in the quarterly DFA data release. Rather than the standard bottom 50%, 50 - 90%, 90 - 99%, and top 1%, we use six groups: bottom 50%, 50 - 70%, 70 - 90%, 90 - 99%, 99 - 99.9%, and top .1%. Further dividing the middle class as well as the top 1% allows us to provide additional insight into the credibility of the data and business cycle dynamics.

The figures confirm the findings of a number of recent studies (see, e.g., Wolff et al. (2012), Piketty (2013), Bricker et al. (2016), Ríos-Rull and Kuhn (2016), Kuhn et al. (2018), and Smith et al. (2019)) that wealth inequality has increased substantially since 1989. During this time period wealth share of the top 10% of U.S. households (the teal, green, and purple portions of graph) increased from 61 to 70 percent, with this trend almost entirely driven by an increase in the top 1% wealth share (purple and green) from 23 percent to 31 percent. In contrast, the bottom 70% (yellow and orange) held very little wealth and experienced a notable decrease in wealth share over this period.

The rise in wealth concentration stems primarily from increased concentration of assets (Figure 2) rather than a decreased concentration in liabilities (Figure 3), with trends for assets largely mimicking those for overall wealth. The share of assets held by the top 10% of the wealth distribution rose from 55 percent to 64 percent since 1989, with asset shares increasing the most for the top 1% of households. These increases were mirrored by decreases for households in the 50-90th per-

[14] Because aggregated balance sheet items are constructed from B.101.h balance sheet lines and not reconciled SCF balance sheet lines, the shares of aggregated balance sheet lines for each wealth quantile do not necessarily align with the shares from the SCF.

Figure 1: Wealth Levels and Shares in the DFAs

(a) Levels **(b)** Shares

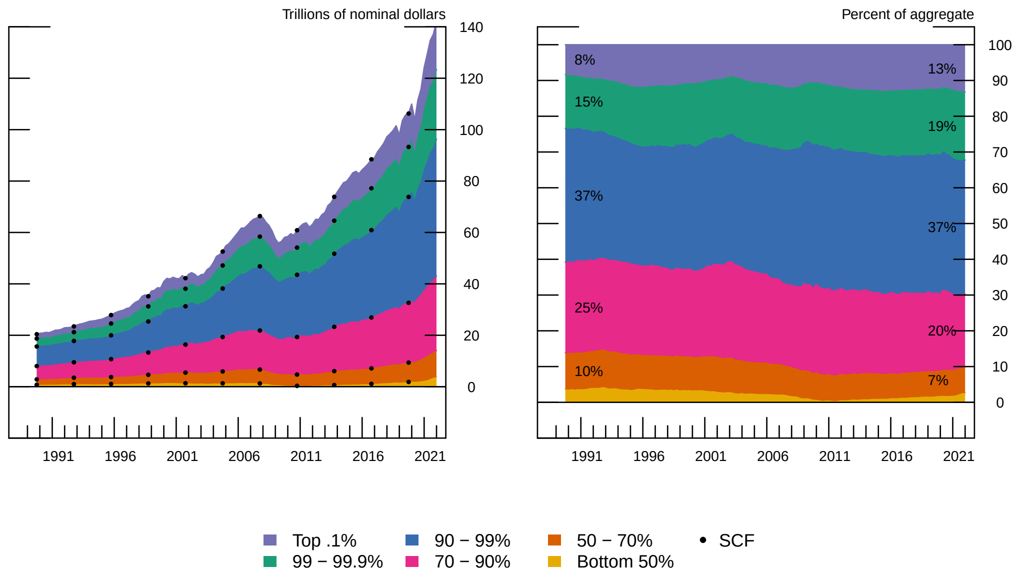

centiles of the wealth distribution.[15] Figure 3(b) shows that, in contrast to assets, the distribution of liabilities is both much more equitable than the distribution of assets and has changed relatively little since 1989.

Panel (a) in Figure 1 also shows that household wealth varies substantially between the SCF observations indicated by black dots. In particular, the figure shows the evolution of wealth for our six wealth groups during the dot-com boom and bust between 1998 and 2001, the early-2000s recession, the start of the housing boom between 2001 and 2004, the Great Recession between 2007-2010, and the ongoing Covid-19 pandemic.

The dynamics align intuitively with the economic events that drive them but are generally not captured by the lower frequency SCF. Panel (b) in Figure 1 shows that higher frequency changes in the wealth distribution generally occur during recessions and are relatively short-lived, and thus are often unobserved in lower-frequency data. For example, the top 1%'s share fell during both the 2002 and Great recessions but had mostly recovered by the 2004 and 2010 waves of the SCF. Over-

[15] Appendix A.4 shows that the inequality of holdings within asset class generally increased as well.

Figure 2: Total Assets by Wealth Percentile Group

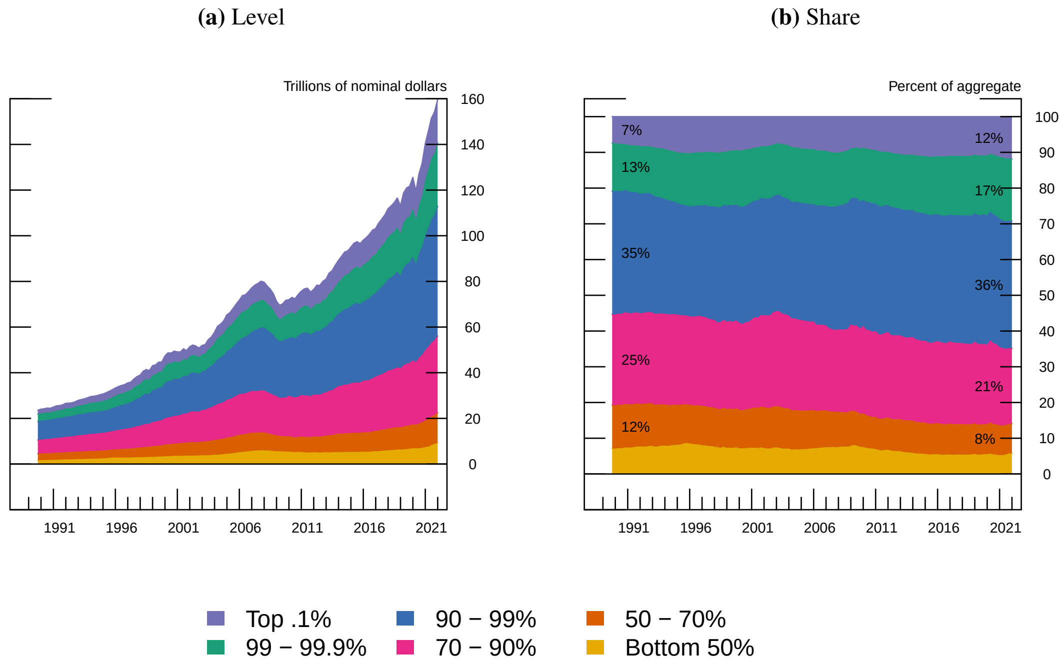

all, Figure 1 suggests our data captures meaningful variation in wealth levels and shares between SCF waves, which we exploit more fully starting in Section 4 to formally document the dynamics of the wealth distribution.

3 Credibility

The data shown in the previous section provide novel quantitative measures that, reassuringly, are qualitatively consistent with those inferred from lower frequency data and hypothesized in prior studies. However, the credibility of our quantitative findings is paramount for our analysis in starting in Section 4. Therefore, in this section, we establish that our estimates are well identified and robust to alternative modeling assumptions, that the estimation procedure works well out of sample, and that estimates align well with alternative, leading measures of wealth inequality at lower frequencies at which these measures are available.

Figure 3: Total Liabilities by Wealth Percentile Group

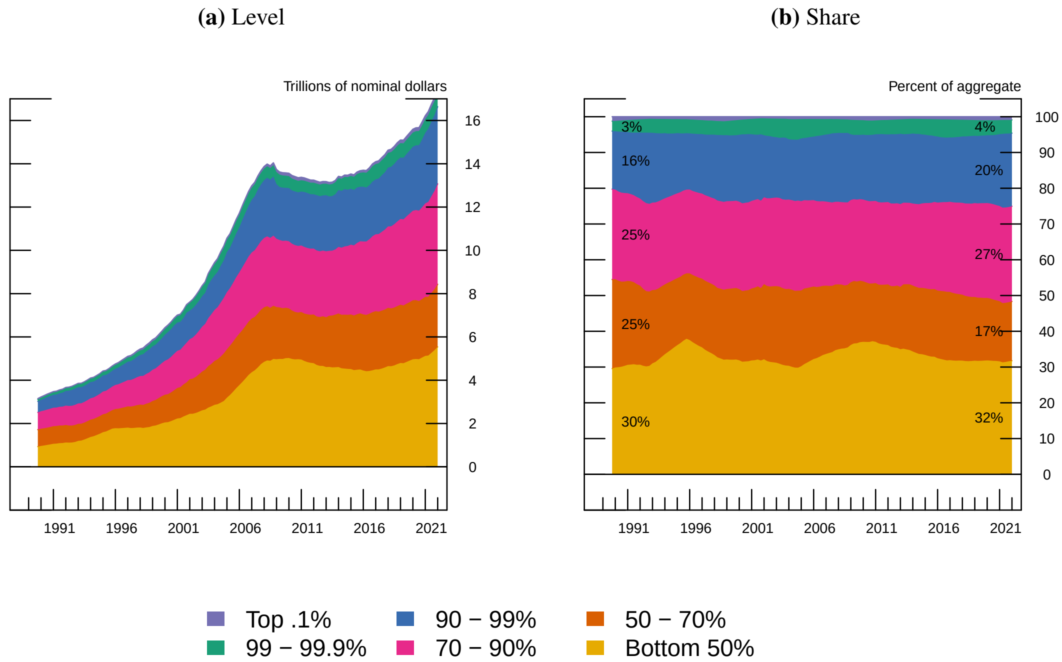

3.1 Estimate Uncertainty and Standard Errors

Our data is designed to be consistent with the Financial Accounts and SCF, so measurement error in either data source (to the extent it exists) which is inherited by the DFAs would not lower the value of our data. In contrast, if our estimates of balance sheets in quarters when the SCF is unobserved were noisy, our higher frequency measures would be less valuable. In this section we therefore calculate standard errors for our balance sheet estimates for each wealth group to show that they are reasonably well identified.[16]

The error bands for the 95% confidence interval for each group's wealth share are presented in Figure 4. During SCF years there is no uncertainty and so the standard errors are therefore zero, but estimates increase in quarters that are further away from SCF observations. And although uncertainty around estimates is sufficiently large that we generally cannot exclude the possibility that wealth adjusts linearly between SCF waves, the standard errors are not overly large and wealth

[16]See Proietti et al. (2017) for formulas for standard errors of asset and liability levels. We convert these standard errors for our level estimates to standard errors for share estimates via the delta method.

shares are generally identified within 1pp.

The relatively small standard errors from our estimation procedure are fairly intuitive: changes in net worth are primarily driven by capital gains and losses. The aggregate Financial Accounts series included as indicator series for each balance sheet line for each wealth group serve as a proxy for these capital gains and losses, thus limiting uncertainty in movements in households balance sheets. Furthermore, saving rates are generally fairly stable over time and well proxied by other indicator series. While there is some uncertainty in wealth changes due to other factors (e.g., group-specific returns), these changes are generally secondary and do not generate significant shifts in wealth shares. In short, the indicator series we have selected do a good job in capturing changes in household wealth, resulting in reasonably well identified quarterly estimates of household balance sheets and net worth.

3.2 Consistency With Untargeted Moments

Although our estimates are reasonably well identified, the DFAs success in matching actual measures of the wealth distribution not used in their construction is a better test of data quality. Opportunities for such external validity are limited, but a special SCF panel was collected in 2009Q1 as a one-time follow-up to the regular 2007 SCF and was not used in constructing our data. This sample therefore offers an opportunity to cross-check the DFA estimates against actual observations, and poses a particularly rigorous validation test given the significant economic and asset price volatility during this time period.[17]

Figure 5, panels (a)-(f), summarizes the DFA and 2009 SCF data for each of our six wealth percentile groups. In each panel, the first bar illustrates the DFA household balance sheets in the 3rd quarter of 2007 (aligned with the 2007 SCF data), the second bar presents the interpolated DFA data for 2009Q1, and the third bar presents the 2009 SCF panel household balance sheet data, adjusted so that asset class definitions align with the DFA/FOF wealth concepts. The regions of each bar above the x-axis indicate the aggregate level of assets and a general composition (real

[17] To reduce respondent burden, the 2009 panel survey questionnaire did not repeat the detail found in the 2007 cross section. However, there is enough information in the 2009 SCF panel to generally align the asset and debt categories in those data to the B.101.h (as described in section 2.1); because some of the finer detailed categories cannot be mapped, we will focus in this section on broader asset and debt categories. The 2009 SCF panel represent the set families eligible for the 2007 SCF, and not necessarily the full cross-section of 2009 families.

Figure 4: Standard Errors of Estimates

(a) Fernandez (baseline) shares and standard errors

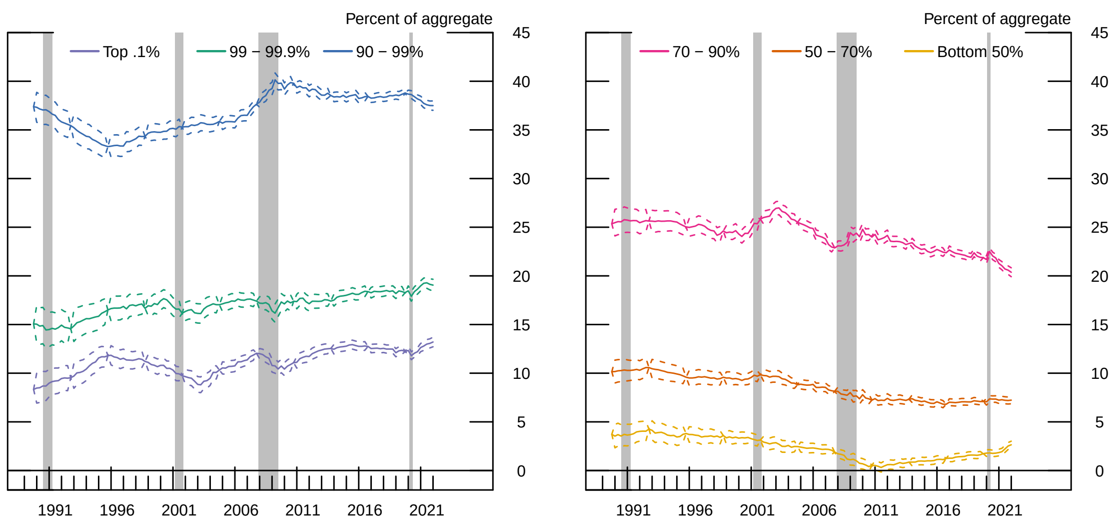

(b) Kalman shares and standard errors

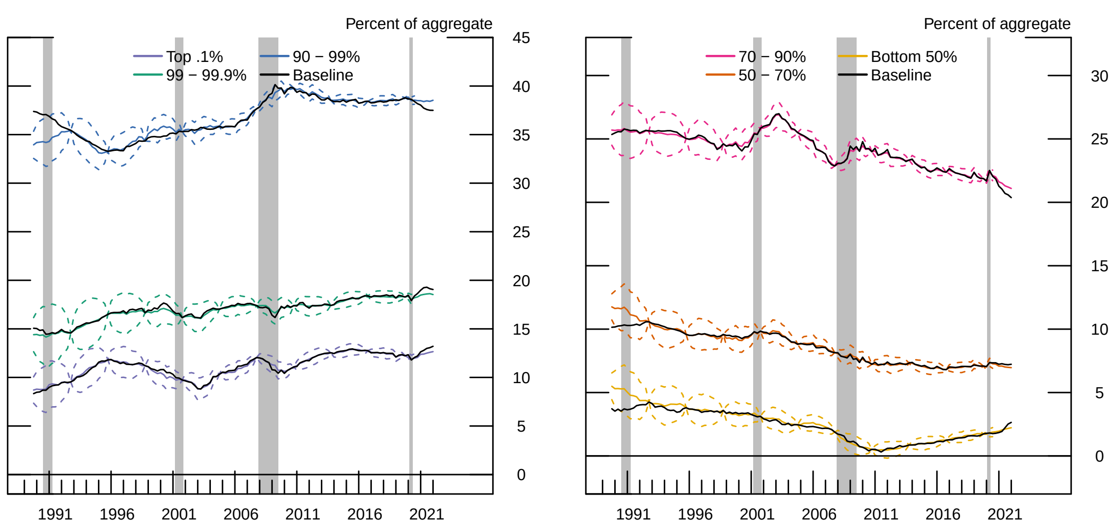

Figure 5: Comparison to 2007-2009 SCF Panel

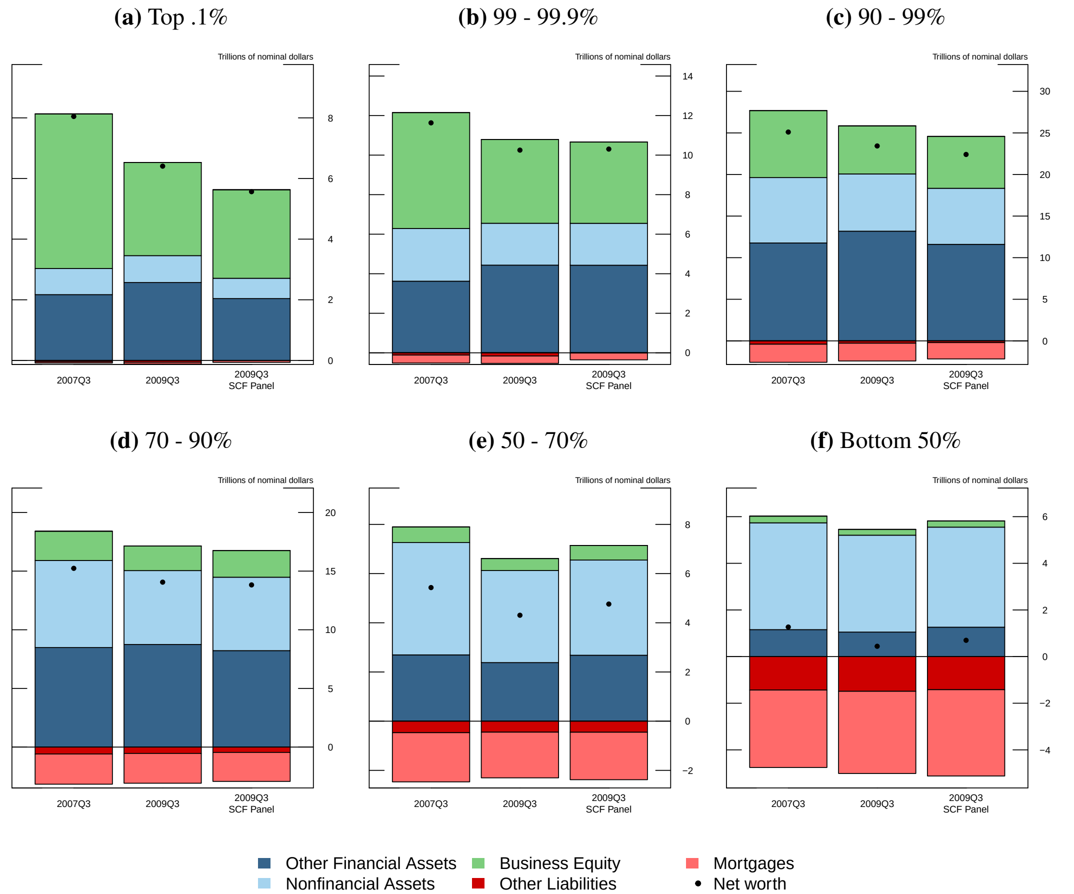

estate, other non-financial assets, and financial assets), and the regions below the x-axis indicate aggregate liabilities (mortgage and non-mortgage).

Panel (a) illustrates the two data sets for the wealthiest 0.1%. Entering the recession in the 3rd quarter of 2007, aggregate net worth for this top wealth group was about $8.2 trillion (first column), but by the third quarter of 2009 aggregate assets of this group had fallen to less than $7 trillion according to our DFA estimates (second column). The 2009 SCF panel (third column) shows a similar decline in overall net worth, with a drop in business equity (green region) driving most of the decline in both the DFA and SCF data. In addition to replicating balance sheet dynamics of the top 0.1%, the DFAs are even more accurate for the remaining wealth groups, with Panels (b)-(f) showing that balance sheet measures in the 2009 DFA and SCF data across the wealth distribution are nearly identical.[18] Overall, this exercise strongly suggests that the DFAs provide credible measures of household balance sheets and net worth across the wealth distribution, even in quarters when the SCF is not collected.

3.3 Comparison to Alternative Lower-Frequency Wealth Distribution Measures

Although the high-frequency dynamics in our data are the focus of our study, our data also capture lower frequency dynamics that have been the focus of a substantial body of research, as shown in this section.

The most directly comparable measures are from Saez and Zucman (2016; updated in 2020)—hereafter "SZ20"—and from Smith, Zidar, and Zwick (2020)—hereafter "SZZ"—both of which also distribute aggregate balance sheet measures from the Financial Accounts across households.[19] Both SZ20 and SZZ use a slightly different aggregate wealth measure (neither include consumer durables, SZ20 do not include unfunded DB pension wealth, and SZZ use alternate data for noncorporate business equity) and both distribute aggregate wealth mainly using rates of return inferred from annual income tax data. Although both SZ20 and SZZ rely on the same wealth and annual income data, they differ in the assumptions they make when imputing household balance sheets. In particular, the two studies make notably different assumptions regarding the rates of return used when estimating the stock of interest-bearing assets and in allocating noncorporate business wealth. To illustrate the similarities and differences across these data, Figure 6 compares the top 1% wealth shares for all three data sets. Reassuringly, all data sets imply a clear upward trend in wealth in-

[18]While the DFA data between SCF survey years are interpolated, the data since the most recent SCF survey are forecasted. An alternate exercise would be to test these forecasted DFA data by constructing a pseudo-forecast of the 2009 DFA data—as if they were being constructed in 2009—and compare to the 2009 SCF panel. These results (not shown, for brevity) are also similar to those shown in figure 5 and further demonstrate the ability to provide good out-of-sample predictions.

[19]The Saez and Zucman (2016) distributional estimates are also found in World Inequality Database (WID), a statistical database focused on measures of inequality in the U.S. and many other countries. For more information, see https://wid.world/ and Alvaredo et al. (2016). The most recent estimates for these models rely on administrative tax data through 2016, though the SZ20 estimates include projections through 2019. This section will only draw a comparison to the recent innovations in inferring wealth from income tax data, though this practice has a long history (see Greenwood, 1983). Wealth has also been measured by data from estate tax filings (as in Kopczuk and Saez (2004)) and directly from household surveys (such as the SCF as in Bricker et al. (2016)). As estate tax filing thresholds have increased, the recent estate tax data have relatively small sample sizes and are less representative than are the past data. A comparison of the WID data to the SCF can be found in Bricker et al. (2016).

Figure 6: Wealth Shares from the DFA, WID, Smith et al.

(a) Top 1% shares comparison across measures

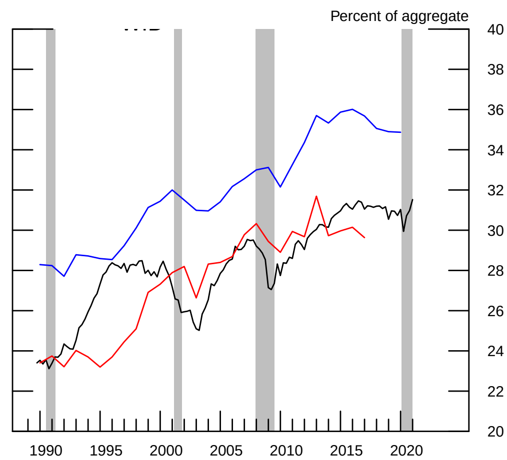

equality since 1989, although both the DFAs and SZZ imply a slower, more gradual increase than SZ20.

Additionally, the high-level drivers of the wealth distribution in the DFAs also align with patterns shown in recent work in Kuhn et al. (2018). These "historical SCF" data combine triennial SCF data post-1983 with annual survey data from 1952-1973 to show that differential changes in equity (which are predominantly held by wealthy households) and housing (which dominate portfolios of the middle class) returns have shaped postwar wealth trends. The DFAs show similar patterns. Figure 7 shows that despite clear cyclical booms and busts, business equity has increased as a share of total wealth (from 25% to 36%) while the top 1%'s share of business equity has risen (from 41% to 51%) since 1989. Thus, much of the increase in top wealth shares is explained by increased equity concentration. Figure 8 shows a similar set of charts for real estate equity and middle class families (as proxied by the 50-90th percentile). Despite a notable boom-bust cycle, housing as a share of total wealth (panel (a)) has become a smaller since 1989 at the same time that the share of housing assets held by middle class families has also fallen (panel (b)). These two patterns are major drivers of the overall decline in middle class wealth shown in Figure 1.

Figure 7: Business Equity Holdings and the Top 1%

(a) Total Business Equity as a Share of Total Wealth **(b)** Top 1%'s Share of Total Business Equity

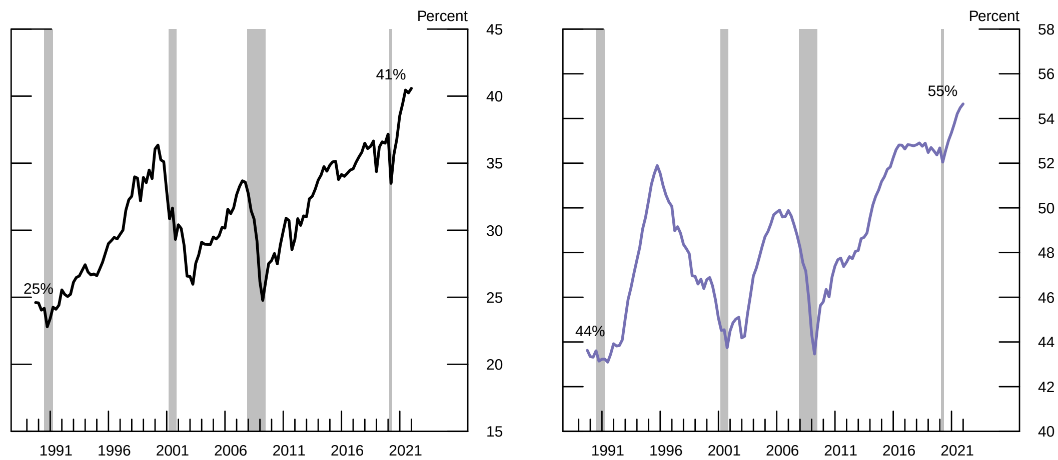

Note: Business equity is the sum of corporate equities and mutual fund shares plus equity in noncorporate business.

Figure 8: Real Estate Equity Holdings and the Middle Class

(a) Real Estate Equity as a Share of Total Wealth **(b)** 50-90th Percentiles Share of Total Real Estate Equity

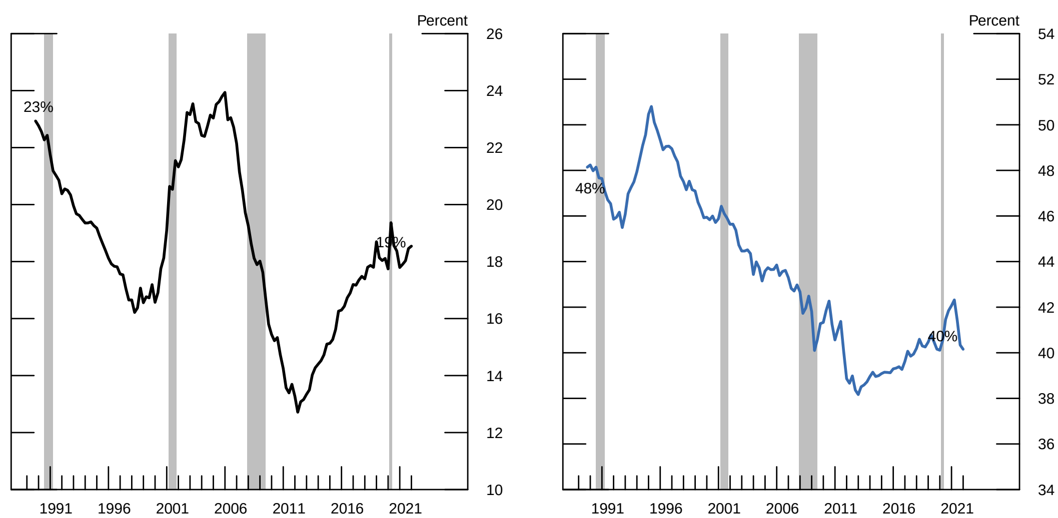

Note: Real estate equity is defined as real estate less mortgages..

Figure 9: Bottom 50% Balance Sheet

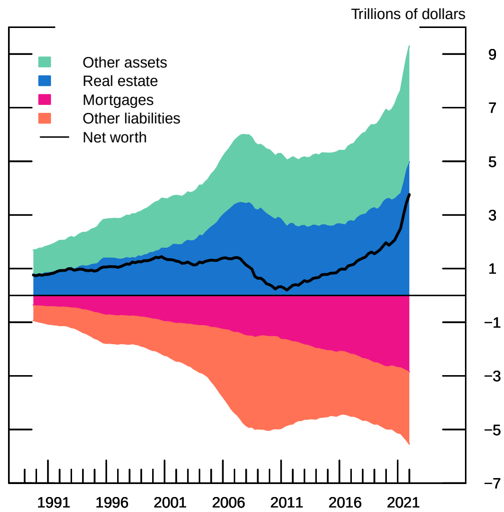

Finally, several studies (e.g., Krusell and Smith (1998), Greenwald (2018), Jones et al. (2018), Kaplan et al. (2019), Garriga and Hedlund (2019)) have shown that household debt levels–which Kuhn et al. (2018) show have increased particularly for low wealth households–have become increasingly important drivers of macroeconomic dynamics. In Figure 9, the DFAs similarly show that debt growth drove overall wealth dynamics for the bottom half of the wealth distribution over the last 30 years. For example, a rapid increase in leverage during the housing boom pulled the bottom 50%'s net worth negative for several quarters after the housing bubble burst, and elevated debt holdings have kept the bottom 50%'s nominal wealth below its level in the mid-1990s despite substantial deleveraging after the Great Recession.

Taken altogether, the DFA data are generally consistent with the secular trends documented in other studies, and therefore appear credible. However, the DFAs also offer notable advantages relative to existing data products. First, our data cover the complete wealth distribution, while estimates from income tax returns (e.g., Smith et al. (2019) and Saez and Zucman (2016)) are less suited to measure balance sheets of lower-wealth households that largely hold assets without income flows that are used to approximate asset holdings. Additionally, measures of the wealth distribution imputed from tax returns rely on *realized* capital gains and dividend payouts, which may differ from unrealized capital gains that likely affect household decisions. This is potentially problematic, as household's decisions to realize capital income and gains itself likely depends on on cyclical factors. For example, Dowd et al. (2019) estimate elasticities of realized gains to tax rates spiked during the Financial Crisis, and wealthier households that face larger tax bills have stronger incentives to shift capital income and gains across time. Such concerns are maybe not significant when studying the longer run evolution of wealth since capital income and gains must eventually be realized. However, without carefully modeling the tax incentives of households, it is unclear how these factors affect the cyclical properties of estimates imputed from tax data. Finally and most importantly, the DFAs are the only data available at a quarterly frequency, and therefore are uniquely positioned to analyze business cycle dynamics of the wealth distribution.

4 Business Cycle Dynamics of Wealth Distribution

Figure 1 showed that household wealth varies substantially between the SCF observations indicated by black dots. In this section, we will describe the dynamics of the wealth distribution at a quarterly frequency more fully.

Changes in the wealth distribution are driven by relative differences in wealth accumulation, since in a given period the wealth distribution will shift towards households that accumulate more wealth than the average household. There are two high-level reasons why wealth accumulation will vary across wealth groups: differential returns to wealth due to differences in exposure to aggregate asset returns and differential wealth accumulation due to group-specific factors, reflecting group-specific differences in returns and savings.[20] We will therefore use this framework–which is consistent with a growing body of influential research (see, e.g., Bach et al. (2020), Fagereng et al. (2020))—to frame our analysis. We first explore how exposure to different assets drive returns to wealth across wealth groups in Section 4.1 and then decompose returns to wealth and change in wealth shares into the two factors—aggregate price changes and group-specific changes—in sections 4.2 and 4.3, respectively.

[20]Our data are not well-suited to distinguish between the latter two explanations (i.e., group-specific returns vs. differences in savings).

Figure 10: Portfolio Shares

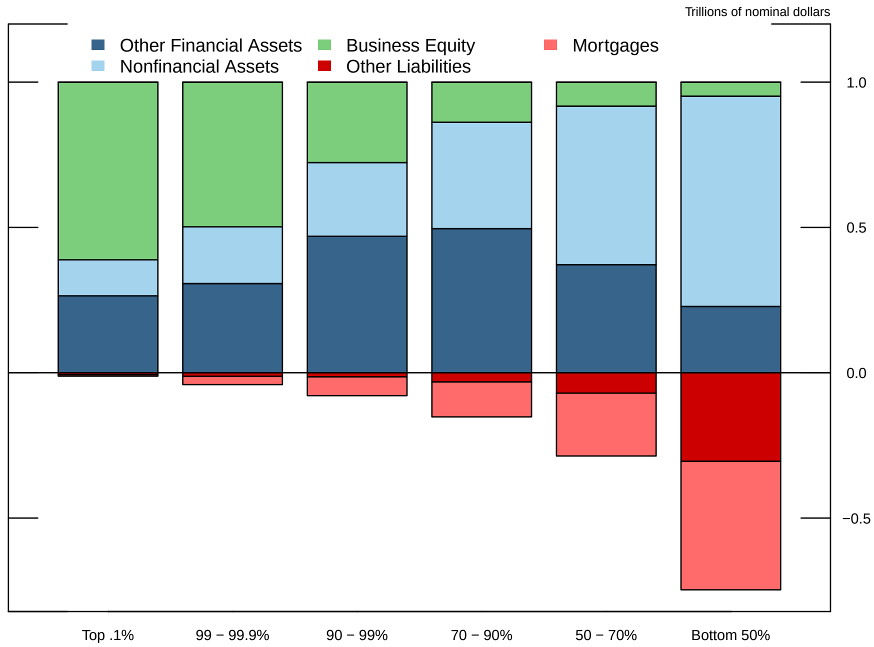

Note: The figure shows the portfolio compositions for (rescaled so that total assets equal to one) for our six wealth groups averaged across all quarters

4.1 The Determinants of Household Portfolio Returns

Intuitively, one reason why wealth accumulation differs across wealth groups is differences in portfolios and exposure to aggregate asset price risks. To demonstrate the difference in exposures simply, Figure 10 shows the portfolio compositions for our six wealth groups—averaged across all quarters and re-scaled so that total assets equal to one. The two wealthiest groups are heavily concentrated in business equity, and their wealth shares will change predominantly in response to changes in business values. Meanwhile, the middle wealth groups are better diversified across asset classes, while poorer households are highly levered and invested in non-financial assets.

To provide an initial quantitative assessment of the drivers of household returns to wealth and compare drivers of wealth accumulation across the wealth distribution, we next estimate a general factor pricing model:

$$\Delta Y_t^p = \beta_0^p + \sum_{k=1}^{K} \beta_{k,t}^p X_{k,t} + \epsilon_t^p \tag{7}$$

separately for each wealth group p. In our baseline analysis we focus on the pre-pandemic data through 2019Q4, but include results for our full sample in Appendix Table 4.

In specifying factors X_k, we choose variables that are likely to capture different sources of wealth variation. For example, we include excess equity returns (measured by the difference between realized S&P 500 returns and the one-year treasury rate) to capture exposure to cyclical assets, and excess housing returns (measured by the difference between realized corelogic housing index returns and the one-year treasury rate) to capture exposure to housing returns. We also include the one-quarter change in the Fed Funds rate to capture exposure to short term changes in bond prices and cost of credit. Finally, we also include the unemployment rate to capture changes in household savings rates due to changes in income and a time-varying precautionary saving motive.

The results of these regressions are presented in Table 2. The first key result is that the coefficient on the S&P 500 returns is positive and significant for all groups, although more so for high-wealth households. For example, a 1% excess return on the S&P 500 is associated with almost a .3% return for households in the top 1%, but the effect is less half as large for households in the bottom 70% of the wealth distribution. These patterns are consistent with the pattern shown in Figure 10 that exposure to equity returns increases significantly in wealth, as well as findings in recent academic work studying portfolio and return heterogeneity (e.g., Bach et al. (2020)).

The second key result is that the coefficient on house prices is also positive for all groups, but more so for lower-wealth households. This pattern reflects—again consistent with Figure 10 and recent research (e.g., Kuhn et al. (2018))—that real estate is the key asset in most lower-wealth households' portfolio and that exposure to house price risk decreases with wealth. For example, a 1% excess return on housing boosts returns for households in the 50-70th percentiles of the wealth distribution by .6%—over twice as much as for households in the top 1%— and is associated with a more than 2% increase in returns for households in the bottom half of the wealth distribution, with the large point estimate reflecting that real estate holdings of low-wealth households are highly

Table 2: Estimating Drivers of Wealth Gains, by Household. Sample 1989-2019.

	Dependent variable:					
	Top .1%	99-99.9%	90-99%	70-90%	50-70%	50%
	(1)	(2)	(3)	(4)	(5)	(6)
S&P 500	0.281***	0.268***	0.213***	0.149***	0.103***	0.158***
	(0.020)	(0.018)	(0.014)	(0.010)	(0.009)	(0.045)
Core Logic House Price Index	0.295***	0.300***	0.315***	0.404***	0.590***	2.175***
	(0.089)	(0.083)	(0.061)	(0.045)	(0.040)	(0.200)
FFR, 1 Quarter Change	−0.596	−0.586*	−0.502*	−0.398**	−0.368**	−0.913
	(0.377)	(0.351)	(0.259)	(0.210)	(0.170)	(0.845)
5 Year Forward Rates	−0.489	−0.569	−0.2105	0.088	0.388	2.666*
	(0.711)	(0.661)	(0.489)	(0.359)	(0.320)	(1.594)
Unemployment Rate	−0.896	−0.881	−0.551	−0.321	−0.384	−5.127***
	(0.588)	(0.547)	(0.404)	(0.2197)	(0.265)	(1.319)
Constant	0.095	0.028	−0.145	−0.168	−0.010	0.303
	(0.216)	(0.201)	(0.149)	(0.109)	(0.097)	(0.485)
Observations	103	103	103	103	103	103
R^2	0.734	0.749	0.783	0.812	0.844	0.722
Adjusted R^2	0.720	0.736	0.772	0.803	0.836	0.708
Residual Std. Error (df = 97)	1.410	1.312	0.970	0.712	0.635	3.162
F Statistic (df = 5; 97)	53.587***	57.742***	70.27***	83.945***	105.316***	50.469***

Note: *p<0.1; **p<0.05; ***p<0.01

levered.

In addition to these obvious takeaways, these regressions also highlight more subtle patterns in returns across the wealth distribution. Focusing on interest rates, we find that a increase in the effective federal funds rate (FFR) lowers changes in wealth for all households, with the effect decreasing from the top 1% (1pp change in FFR is associated with a -.6% change in returns for the top 1%) to the 50-70th percentiles (-.4%), before rising sharply for the bottom 50% (-.9%, although the effect is not significant). For households in the top half of the wealth distribution, this pattern likely reflects that the value of bond holdings generally falls as the FFR increases. For low-wealth households that don't hold many fixed-income assets, the decline likely reflects a separate, consumer credit channel. Low wealth households are more likely to have floating rate and revolving debt, so an increase in interest rates will increase their debt servicing cost and lower wealth accumulation.

Finally, we also find that a higher unemployment rate is associated with lower returns on wealth for all households, but especially for low wealth households, for whom a 1pp increase in the unemployment rent reduces excess wealth returns by 5.27%. This large, negative effect reflects that low wealth households are most exposed to cyclical layoffs, and so their income tends to be much more procyclical. As a result, low-wealth households savings drops as unemployment increases, thereby reducing wealth accumulation.

The share of wealth variation in wealth accumulation explained by our model is also of interest. We find that our four factor model can explain over 70% of the percent changes in wealth accumulation for all wealth quintiles, with greater explanatory power for houesholds in the middle of the wealth distribution. This hump-shaped pattern reflects two considerations. First, households at the top of the distribution hold riskier assets that are not as well priced by a broad equity index, including non-public assets that are less correlated with indices of public equities, and consistent with Fagereng et al. (2020). Second, for households at the bottom, idiosyncratic fluctuations in saving and policy changes affect their wealth more, which are hard to capture with aggregate risk factors.[21]

[21] In Appendix Table 4 we repeat these regressions including observations for 2020. Given that variation in explanatory variables were unusually large during the pandemic we prefer to omit them from our baseline analyses. However, patterns from these regressions are broadly the same, with the exception that the unemployment rate is no longer sig-

4.2 Decomposing Returns to Wealth into Aggregate and Group-Specific Changes, by Wealth Group

As we showed above, our factor model can explain a significant amount of wealth accumulation for all wealth groups. However, it cannot explain all changes and may overstate the importance of aggregate price changes that share a common component with group-specific factors.

For more complete treatment, we build on prior literature (Bach et al. (2020), Fagereng et al. (2020), Kuhn et al. (2018)) and introduce the intertemporal law of motion for each asset or liability j on group p's balance sheet is

$$Y_t^{j,p} = (1 + r_t^j + \epsilon_t^{j,p})Y_{t-1}^{j,p} + \theta^{j,p} + S_t^{j,p}, \tag{8}$$

where r_t^j is the aggregate return to asset j in period t common to all groups, $\epsilon_t^{j,p}$ is group's p scale-dependent idiosyncratic return for asset j in period t, $\theta^{j,p}$ is the sum of type-dependent returns for members of the wealth group p for asset class j, and $S_t^{j,p}$ is group p's saving into asset class j.[22]

This expression can be decomposed into the contribution to next period's wealth due from capital gains from aggregate price changes ($E_t^{j,p}$) and the contribution from group specific returns and savings ($U_t^{j,p}$).

$$Y_t^{j,p} = Y_{t-1}^{j,p}\left(1 + \underbrace{r_t^j}_{E_t^{j,p}} + \underbrace{\epsilon_t^{j,p} + \theta^{j,p} + S_t^{j,p}/Y_{t-1}^{j,p}}_{U_t^{j,p}}\right). \tag{9}$$

Aggregating over all assets, group p's net worth can be expressed in aggregate price changes and

nificant for low wealth households. This reflects the well-documented pattern that wealth increased sharply, particularly for low wealth households, in the COVID-19 recession due to very high fiscal transfers. Additionally, in unshown analyses we find that other proxies—for example, wage growth—were associated with higher wealth accumulation for low-wealth households although the effects were not significant and so we omitted them from our baseline analyses (not shown).

[22]Footnote on definition of type and scale dependence of returns here.

group-specific factors:

$$Y_t^{NW,p} = \sum_j Y_{t-1}^{j,p} \times (1 + E_t^{j,p} + U_t^{j,p}) \tag{10}$$

$$= Y_{t-1}^{NW,p} \sum_j \left[(1 + E_t^{j,p} + U_t^{j,p}) \times \alpha_{t-1}^{j,p} \right] \tag{11}$$

$$= Y_{t-1}^{NW,p}(1 + E_t^p + U_t^p), \tag{12}$$

where $\alpha_{t-1}^{j,p}$ is the portfolio share group p's wealth invested in asset j and E_t^p and U_t^p denote the growth in wealth due to aggregate capital gains and group-specific contributions, respectively.

Equation 12 can be transformed into a simple expression for the quarterly growth rate of group p's net worth:

$$\Delta Y_t^{NW,p} = \frac{Y_t^{NW,p}}{Y_{t-1}^{NW,p}} - 1 = E_t^p + U_t^p, \tag{13}$$

which we will use below to decompose the contributions of E_t^p and U_t^p to each group p's quarterly wealth growth. In what follows, we use Equation 13 to separate the group-specific returns to wealth ($\Delta Y_t^{NW,p}$) that are driven by aggregate prices changes (E_t^p) and group-specific factors (U_t^p). Aggregate asset returns E_j^p are taken from the re-valuation adjustment from the Federal Reserve's Z.1 Financial Accounts, while changes from group-specific factors are calculated as the residual that cannot cannot be explained by aggregate price fluctuations.

Figure 11 shows the resulting return decomposition. For each wealth group, aggregate asset price fluctuations (E_t^p, the black lines) drive most of the changes in returns to net worth ($\Delta Y_t^{NW,p}$, the blue lines). Based on its construction, differences in the portfolio composition across wealth distribution documented in Section 4.1 drive the differences in the dynamics of E_t^p. Thus, our decomposition suggests that most variation in wealth over time is explained by aggregate price movements, consistent with our finding in Table 2 that aggregate variables can explain a large share of temporary variation in household returns.

For most wealth groups, the group-specific return U_t^p (red line) is largely acyclical, consistent with the idea that a meaningful fraction of total return to wealth is explained by type-dependent returns that are time-invariant. However, U_t^p is fairly procyclical for the wealthiest 0.1% and 99-

Figure 11: Decomposition of $\Delta Y_t^{NW,p}$ into aggregate price changes, group-specific changes

(a) Top .1%

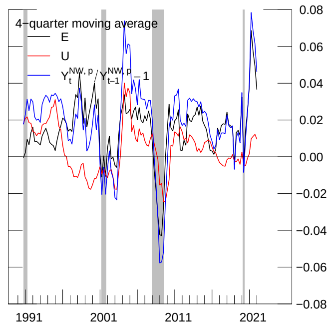

(b) 99 - 99.9%

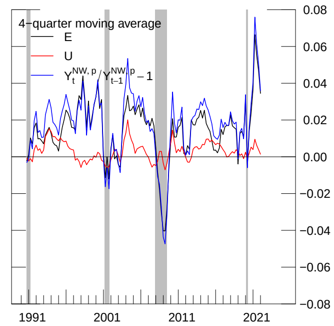

(c) 90 - 99%

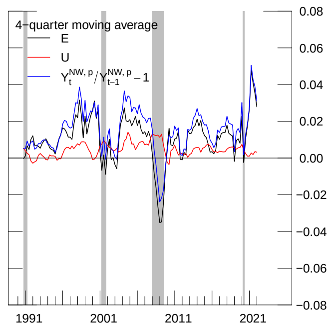

(d) 70 - 90%

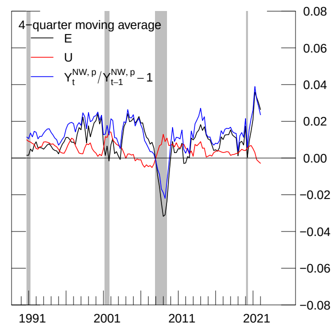

(e) 50 - 70%

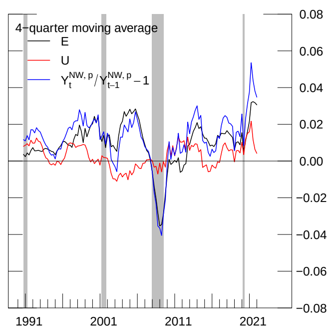

(f) Bottom 50%

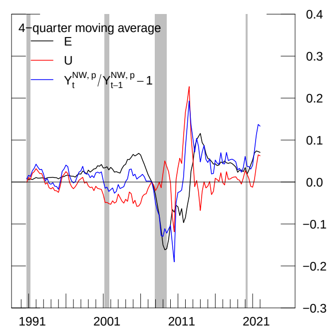

99.9% groups, likely reflecting riskier asset holdings within a given asset class, as captured by by $\epsilon_t^{j,p}$ in Equation 8. For wealthier households, both type dependence and riskier holsings within a given asset class are consistent with the findings in Fagereng et al. (2020) and Bach et al. (2020).

Figure 11 also reveals a near-monotonic spread in returns to net worth ($\Delta Y_t^{NW,p}$) across the wealth groups, meaning that the level of returns is generally higher in expansions and lower in recessions for wealthier households. For example, in the run-up to the Financial Crisis, the quarterly returns for the wealthiest 1% of households (panels (a) and (b)) reach a peak of 6 to 8 percent, roughly double of those realized by the households in the 90-99% and 70-90% of the wealth distribution (panels (c) and (d)), thereby indicating an unequal pace of wealth accumulation. While opposite patterns are generally observed in downturns, the spread across households is generally more muted, and the returns of wealthy households tend to rebound quickly early in economic recoveries.

Finally, the common return E_t^p exhibits different behavior for households in the top 50% and bottom 50%. During economic expansions, the E_t^p return increases monotonically with wealth for the top 50%, and is therefore a major driver of the increasing returns to wealth mentioned above. For the bottom 50% of households, E_t^p is highly volatile and exhibits large positive and negative swings. Although these households wealth is primarily concentrated in assets with less cyclical risk (i.e., real estate) than households at the top of the wealth distribution (i.e., business equity), these households are also highly levered due to high levels of mortgage debt and consumer credit. This leverage amplifies smaller price fluctuations, particularly following run up in household debt in the mid-2000's. Similarly, group-specific return U_t^p to wealth for the bottom 50% of households is highly volatile (again due to their substantial leverage) but exhibit little cyclical variation, which is similar to results from Swedish (Bach et al. (2020)) and Norwegian (Fagereng et al. (2020)) data.

4.3 Decomposing Changes in Wealth Shares into Aggregate and Group-Specific Price Changes

In the previous subsection, we showed that wealthier households hold higher shares of more cyclical assets like corporate and noncorporate business equity, and therefore their returns to wealth are more sensitive to changes in the values of these assets. Additionally, group-specific return components also tend to be more procyclically for wealthy households, suggesting that their overall wealth shares should also vary procyclically. In this subsection we leverage the unique quarterly frequency of our data to confirm this patterns and provide novel quantitative insights into the cycli-

Figure 12: Cyclical Variation in Wealth Shares

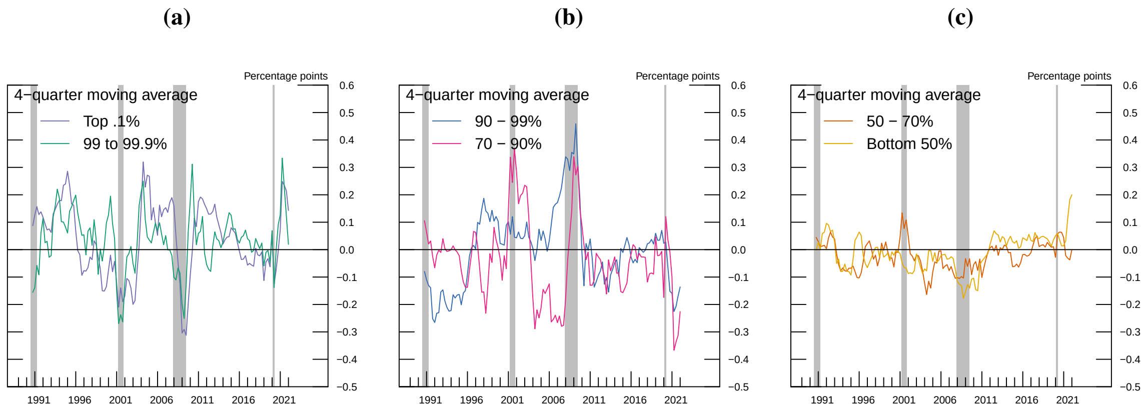

cal variation in wealth shares.

In Figure 12, we start by plotting the quarterly changes in wealth shares for each of the six considered wealth groups, with gray bars indicating NBER recessions. The figure reveals differences in the cyclical behavior of the wealth shares that vary systematically across the wealth distribution. Panel (a) in Figure 12 makes clear that the top 1% and the 0.1% shares vary procyclically, with notable drops during each of the three recessions captured in our data. In contrast, in Panel (b), the 90-99th and 70-90th percentiles of the distribution move countercyclically. Finally, Panel (c) shows that wealth shares of households in the Bottom 50% and in the 50-70th percentiles of the wealth distribution are roughly acyclical. Hence, cyclical variation in the wealth distribution is largely confined to the wealthiest 30% of households, with wealth shares shifting from the Top 1% to 70-90th and 90-99th percentile households during recessions and expansions.

Next, we decompose changes in wealth shares from Figure 12 into aggregate and group-specific price changes. Let ω_t^p represent the wealth share of group p at time t. Further transforming Equation 13, group p's quarterly *change* in wealth share can be conveniently obtained by dividing through by the total wealth across all groups. Rearranging terms yields

$$\omega_t^p - \omega_{t-1}^p = \frac{Y_{t-1}^{NW,p}}{\sum_{p'} Y_t^{NW,p'}} \left[(E_t^p + U_t^p) - \left(\frac{\sum_{p'} Y_t^{NW,p'}}{\sum_{p'} Y_{t-1}^{NW,p'}} - 1 \right) \right]. \tag{14}$$

The term inside the brackets in Equation 14 thus shows that—as long as $\frac{Y_{t-1}^{NW,p}}{\sum_{p'} Y_t^{NW,p'}} > 0$—each group p's wealth share (depicted in Figure 12) will increase between two quarters if the total net return on wealth for group p ($E_t^p + U_t^p$) is larger than the net return on aggregate wealth ($\sum_{p'} Y_t^{NW,p'} / \sum_{p'} Y_{t-1}^{NW,p'} - 1$).[23] In other words, the equation intuitively shows that wealth distribution dynamics are determined by differences in relative returns to wealth (inclusive of savings) across groups. We therefore first examine the determinants of each group's returns, and then document how they drive cyclical changes in the distribution of wealth.

Through a series of counterfactual exercises, we quantify how much the group specific and aggregate returns shown previously in Figure 11 contribute to the overall cyclical changes in wealth shares (ω_t^p) shown in Figure 12. In the first counterfactual, we assume that returns from aggregate price changes ($E_t^{j,p}$) align with each group's realized returns but that group specific returns are held fixed at their group specific average ($\bar{U}_{j,p}$) each period. In the second counterfactual, we assume that returns from aggregate price changes are constant at their group specific sample average ($\bar{E}_t^{j,p}$) each period and group specific returns align with each group's realized returns ($U_t^{j,p}$). The first counterfactual therefore captures cyclical changes in the wealth distribution only from aggregate price fluctuations (but preserving long-run changes in wealth shares), while the second captures changes only from specific factors.

The resulting quarterly changes in the counterfactual wealth shares are shown in Figure 13. For most households, changes in wealth shares would have been more muted and much less cyclical absent aggregate price changes. This is particularly true for households in the top 70% of the wealth distribution, as evidenced by the modest variability in the blue line but near overlap of the red and black lines in Figure 13.

To provide quantitative summaries of the contribution from each return component to the wealth share cyclicality shown in Figure 13, we apply a Hodrick-Prescott filter to the actual and two counterfactual wealth share profiles for each group and then calculate the variance of the cyclical component. Table 3 presents the baseline variance of each wealth group's cyclical wealth profile in

[23] The relationship $\frac{Y_{t-1}^{NW,p}}{\sum_{p'} Y_t^{NW,p'}} > 0$ generally applies, as both the aggregate and the group specific net worth are positive. However, as we shown in Section 3.3, the net worth for the bottom 50% of the wealth distribution turned temporarily negative in the wake of the Global Financial Crisis. During those quarter, the reverse would thus apply. Namely, the bottom 50%'s wealth share would increase if $(E_t^p + U_t^p) < (\sum_{p'} Y_t^{NW,p'} / \sum_{p'} Y_{t-1}^{NW,p'} - 1)$.

Figure 13: Counterfactual Wealth Shares, No Aggregate Price and Group Specific Changes

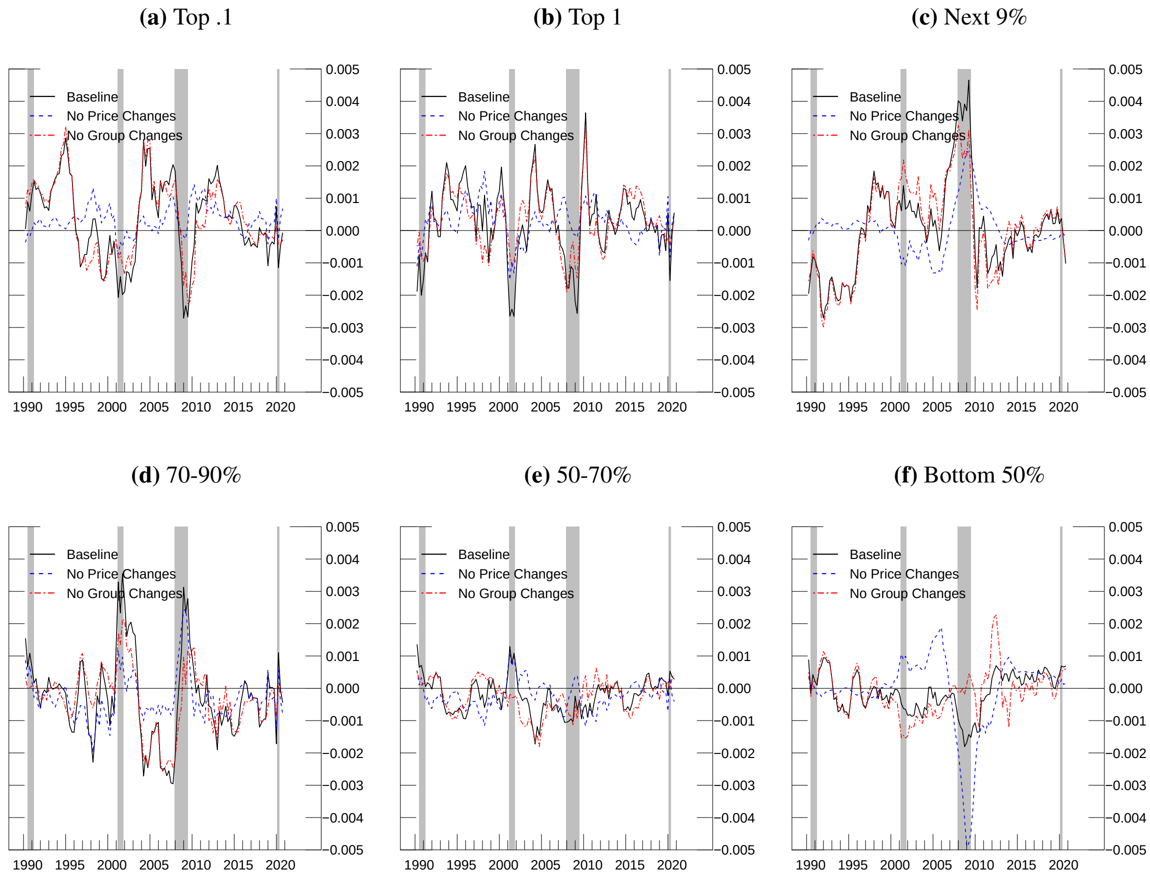

(a) Top .1 **(b)** Top 1 **(c)** Next 9% **(d)** 70-90% **(e)** 50-70% **(f)** Bottom 50%

Column 1, and the corresponding variance (as a percent of the baseline) for the two counterfactual profiles in Columns 2 and 3. As expected, the variance of the cyclical component of wealth shares is much lower absent aggregate price changes (Column 2). The average cyclical variance for households in the top half of the wealth distribution is only 50.4% as large as the baseline, with even larger reductions for households in the top 1 (30.8-100=-69.2%) and 90-99.9 percentiles (-53.7%). In contrast, average cyclical variance of wealth shares for the top half of the wealth distribution is 87.9% absent group specific factors, with variance reductions greatest for households in the 70-90th percentiles (-28.9%). These estimates confirm that the vast majority of cyclical wealth share variance is attributable to cyclical changes in aggregate prices.

Table 3: Decomposition into aggregate price changes and group-specific changes

	Baseline Variance (1)	No Aggregate Price Changes % of Baseline (2)	No Group Factors % of Baseline (3)
Top .1%	0.00344	30.8%	93.8%
99-99.9%	0.00317	46.3%	83.8%
90-99%	0.00357	52.5%	85.2%
70-90%	0.00419	49.8%	71.1%
50-70%	0.00143	72.5%	105.5%
0-50%	0.00144	223.7%	113.3%

4.4 Summary

Our analysis in this section provided several insights into the business cycle dynamics of the wealth distribution. First, wealthier households hold portfolios that are highly skewed towards assets with more cyclical returns. In addition however, group-specific factors that affect returns for wealthy households also tend to be more cyclical. These two patterns generate notable cyclicality in the distribution of overall wealth, with wealth shares shifting from the very wealthy to moderately wealthy during downturns and from moderately wealthy to very wealthy in expansions. Finally, although both aggregate and group-specific factors contribute to changes in the wealth distribution, aggregate price changes account for a much larger share of cyclical variance in wealth shares and therefore are the main driver of cyclical fluctuations in the distribution of wealth.

5 The Effect of Specific Shocks on the Wealth Distributions

Section 4 documented the cyclical fluctuations in the distribution of wealth. These changes—regardless whether they reflect aggregate price or group specific changes—reflect the realized set of all economic shocks and the consumption, investment, and saving decisions households make in response to these shocks. However, economists and policy makers are often interested in the effects of specific shocks. In lower-frequency distributional data, estimating the effects of specific shocks is complicated because many shocks and decisions occur between observations. However, our higher-frequency data offer more temporal variation that can help tease out the wealth distribution's response to economic shocks of interest. In this section, we therefore estimate impulse response functions of the Gini coefficient (a measure of overall inequality) to innovations in key economic variables.

Our impulse response functions are estimated via local projection methods (see Jordá (2005)) which summarize the effect of a specific shock (and subsequent household responses) at different horizons. Specifically, we estimate the following regression

$$y_{t+h} = \mu_h + \beta_h x_t + \sum_{i=1}^{I} \delta'_{h,i} w_{t-i} + \xi_{h,t} \tag{15}$$

where y_{t+h} is the outcome of interest at horizon h, μ_h is the regression constant, x_t is the macroeconomic variable projected onto y, w_{t-i} are lagged vectors of x and y, and $\xi_{h,t}$ is the projection residual. The local projection impulse response function of y_{t+h} with respect to x_t is given by $(\beta_h)_{h\geq 0}$.

When estimating Equation 15, we first consider the effect of key macroeconomic aggregates—the GDP gap, unemployment rate, and inflation—on the the Gini coefficient, before considering the effect of the Feds Fund rate on inequality. Equation 15 is estimated separately for each macroeconomic regressor x_t, meaning that IRFs are identified from time variation, although we also consider alternative identification strategies when examining monetary policy shocks. In all analyses, we control for 6 lags of the outcome and shock variable.

Our results are presented in Figure 14. First, Figure 14 (a) shows the effect of an output gap shock on inequality. Most of the work on the relationship between GDP growth and inequality focuses on long-run trends, with Kuznets (1955) famously predicting that economic growth first increases and later decreases inequality as countries develop, and subsequent empirical work (see, Barro (1999), Dollar and Kraay (2002), Sala-i Martin (2006), among others) largely supporting these predictions. In contrast, we focus on short-run dynamics and find that a short-run increase in GDP is associated with a statistically significant and persistent decline in wealth inequality. However, it is worth noting that the effect is only moderate in size, with a one standard deviation shock to the output gap reducing the Gini coefficient by .011, or about 2% of the change since 1989.

In Figure 14 (b) we consider an unemployment rate shock, which affects inequality through its effect on household income and saving. The theoretical predictions of the effect of unemployment on wealth inequality are relatively clear-cut, since 1) poorer households are disproportionately affected by changes in unemployment and 2) income for wealthier households is better buffered by a higher share of capital income (see Gornemann et al. (2016)). Consistent with these predictions, Figure 14 (b) shows that a 1 standard deviation increase in unemployment causes a a persistent, statistically significant, and again moderately sized increase in wealth inequality. Furthermore, the effect of an unemployment rate shock is slightly larger in magnitude than the effect of a GDP shock.

In Figure (c) we examine the effect of an inflation shocks on wealth inequality. Inflation shocks have two offsetting effects. First, low-wealth households tend to hold more liquid, cash-like assets which can be devalued relative to other assets like business equity following increases in inflation (Erosa and Ventura (2002)). Second, inflation lowers the real value of assets and liabilities, which helps borrowers and hurts savers, particularly those with longer duration positions. Doepke and Schneider (2006) show that this channel generally helps low and middle-class households who have a large share of long-duration fixed-rate mortgages, and hurts richer households with higher exposure to short-duration deposits and debt. The first channel therefore suggests an increase in inequality following an increase in inflation, while the second suggests a decrease. Our estimate implies a slight near-term decline and slight long-term increase in inequality following a positive inflation shock. However, our impulse estimate is never statistically significant, suggesting that the effects of inflation shocks are ambiguous empirically as well as theoretically.

We next consider the effect of monetary policy shocks on wealth inequality. The effect of monetary policy on inequality has received significant attention from policy makers in recent years due to concerns that accommodative monetary policy has contributed to the overall increase in inequality (see, e.g., Bernanke (2015), Bullard (2014), Mersch (2014), and Yellen (2014)). Theoretically, the effects of monetary policy shocks on inequality are ambiguous. As analyzed in Auclert (2019), accommodative monetary policy shocks might increase inequality through an interest rate expo-

Figure 14: Impulse response of Gini coefficient

A. Aggregate Macroeconomic Shocks

(a) GDP Output Gap **(b)** Unemployment **(c)** Inflation

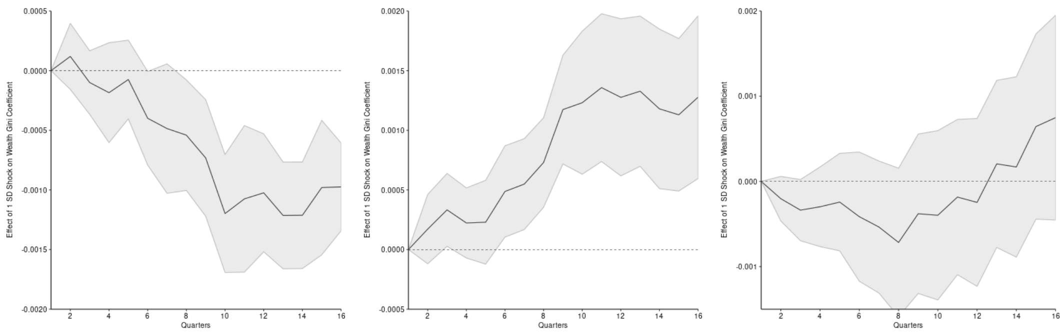

B. Accomodative Monetary Policy Shocks

(d) Time-Variation in FFR **(e)** Romer & Romer Policy Shocks

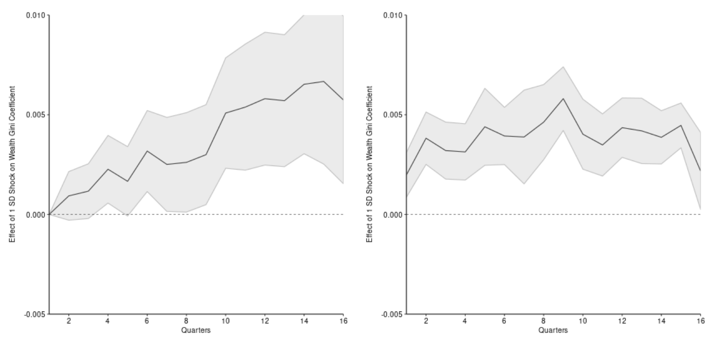

Note: Panel to change in the output gap (Congressional Budget Office), unemployment rate, or inflation (defined as the quarterly percent change in the GDP chain-weighed price index at an annual rate).

sure channel by increasing prices of stocks and other assets (e.g., Bernanke and Kuttner (2005), Rigobon and Sack (2004), Paul (2018)) which are primarily held by wealthier households. On the other hand, accomodative monetary policy is likely to boost earnings and wealth of households at the bottom of the distribution—as shown theoretically in Heathcote et al. (2010) and empirically in Coibion et al. (2012)—and therefore lower wealth inequality.

Figure 14 (d) shows the effect of a 1pp accomodative (negative) monetary policy shock on the wealth Gini. An accomodative monetary policy shock causes an increase in wealth inequality that increases for 10-12 quarters and persists for at least 16 quarters. This finding suggests that the effect of monetary policy shocks on asset prices (Paul (2020)) dominates the positive employment and income channel, leading to an increase in wealth inequality on net. Of course, significant changes in monetary policy often occur in volatile economic environments, so it is hard to say whether the change in inequality reflects the change in monetary policy or the change in economic circumstances that prompted the policy shift. To address this challenge, we repeat our analyses, but instrument for changes in the policy rate using monetary policy shocks obtained by extending the Romer and Romer (2004) through our sample period. These results, shown in Figure 14 (e), again imply that accomodative monetary policy shocks increase wealth inequality.[24]

In addition to being statistically significant, our results also suggest the effects of monetary policy on wealth inequality are quantitatively large. Since 1989, the wealth Gini coefficient—as implied by our data—has risen by 0.057 from 0.748 to .805. Our findings suggest that a 1pp accomodative monetary policy shock increases the Gini coefficient by about 0.005 after 16 quarters, or just under 10% of the 30 year increase. Thus, our results suggest that policy rate reductions designed to support the overall economy have the unintended effects of increasing wealth inequality, and furthermore, that these effects can be fairly large.

Although striking, these results come with several caveats. First, our estimates reflect the effect of short-term changes in the Fed funds rate, and do not speak to the effects of unconventional monetary policy or persistently low rates. Second, our results do not speak to the effects of macroeconomic shocks on earnings, consumption and overall welfare, which may be of primary interest to policy makers. Third, our estimates are specific to the period since 1989 covered by our data. As a result, the policy implications of these reduced-form IRFs are limited, and structural analysis is likely necessary to fully understand the effects of macroeconomic shocks on wealth inequality. Despite these caveats, the findings in this section are novel and provide a useful set of benchmarks for disciplining models featuring both a wealth distribution and aggregate risk.

[24] Additionally, we test our results using policy shocks from Paul (2020) Bu et al. (2021), and Nakamura and Steinsson (2018) and find qualitatively similar results.

6 Household Balance Sheets: A Tale of Two Recessions

Our analysis thus far has focused on describing how wealth shares vary across business cycles generally. However, the last two recessions–the Great Recession and the Covid-19 Pandemic Recession–were atypical in depth and duration, and specific features of household balance sheets have garnered a significant amount of attention in academic research. For example, Mian et al. (2013), Mian and Sufi (2014), and Mian et al. (2017) showed that the build-up in household leverage and subsequent deleveraging amplified the downturn during and slowed the recovery from the Great Recession. In contrast, the aggregate household balance sheet swelled during the pandemic recession as saving rates spiked due to a surge in fiscal transfers and collapse in consumption (see, e.g., Romer and Romer (2021)), suggesting that a similar balance-sheet hangover is unlikely. Our data are the high-frequency observations of the wealth distribution dynamics during these periods, which have and will likely continue to be a focus of economic research in the coming years. In this section we therefore compare and contrast balance sheet evolution through the past two recessions.

We first consider asset price dynamics during the two recessions in Figure 15. During the GFC, asset prices fell and recovered very slowly, with equity wealth (panel a) remaining below its pre-recession level—particularly for high wealth households—until 2012, and housing wealth (panel b) remaining depressed—particularly for middle-wealth households—until 2016. In the Covid-19 recession, an unprecedented set of fiscal and monetary policy actions to support businesses (Hanson et al. (2020)) led to quick improvements in investor sentiment (Cox et al. (2020)) and longer-term growth expectations (Gormsen and Koijen (2020)). This led business equity to surge after a brief decline, significantly boosting wealthy households' balances sheets.

Housing prices show even more divergent trends across the two recessions (Panel b). Housing prices increased dramatically in the years ahead of the GFC, largely due to a shift in beliefs regarding future price growth (Kaplan et al. (2019)). After the housing bubble burst, house prices declined for several years due to weak demand, depressed income, and household deleveraging due to tighter credit standards (see, e.g., Garriga and Hedlund (2019)). This created a long-term drag on wealth, particularly for the bottom 90% of households. In contrast, prices surged during the pandemic as shifts in preferences spurred moves from cities to suburbs (Ramani and Bloom (2021)) and raised housing demand. This rapid increase in demand could not be matched by new construction in the short term, causing housing markets to tighten considerably and home prices to appreciate rapidly (Anenberg et al. (2021)). For example, the Case-Shiller home price index rose by 25% from Jan 2020 to July 2021, an even faster rate of increase than in run-up in housing prices prior to the GFC. This acceleration in home price growth has led to a significant expansion in wealth, particularly for middle-wealth households.

In panel (c) we compare the accumulation of liquid assets—which include bank deposits, money market funds, and other cash-like holdings-during the two recessions. During the Great Recession, liquid assets for most groups moved sideways, as the large economic shock and slow recovery significantly slowed the accumulation in wealth. During the Covid-19 crisis, cash-like holdings increased sharply by almost $3tn with notable increases for all groups, with unprecedented cut-backs in spending leading to a large increase in liquid assets for wealthier households and large fiscal transfers supporting wealth accumulation among lower-wealth households. Furthermore, the aggregate increase in liquid assets is comparable in size to the total “excess savings" (Bilbiie et al. (2021)), suggesting that most extra savings during the pandemic are still being held in cash-like assets.

We next compare liability dynamics during the two recessions in Figure 16. During the Great Recession mortgage debt (panel a) fell—particularly for middle- and low-wealth households—due to defaults and forced deleveraging following the collapse in house prices. In the COVID Recession, mortgage debt rose—again, particularly for middle- and lower-wealth households—due to increased home purchases.

Panel (b) shows that patterns for non-mortgage debt (e.g., consumer credit) growth slowed after both downturns, but for very different reasons. After the Great Recession, tighter credit standards and borrowing conditions restricted credit to marginal borrowers, forcing lower-income households to delever and resulting in a large, persistent decline in non-mortgage debt (Bhutta (2015), Cooper (2012)). During the pandemic recession, many middle and lower-wealth households used their forced savings to pay down debts, resulting in a notable slowdown in consumer credit growth (Horvath et al. (2021), Coibion et al. (2020), Armantier et al. (2020)).

In Figure 17 we show how these dynamics have resulted in very different wealth dynamics for the different wealth groups. During the Great Recession, net worth for high wealth households

Figure 15: Asset Breakdown During the Great Recession and COVID-19 Crisis

(a) Business equity

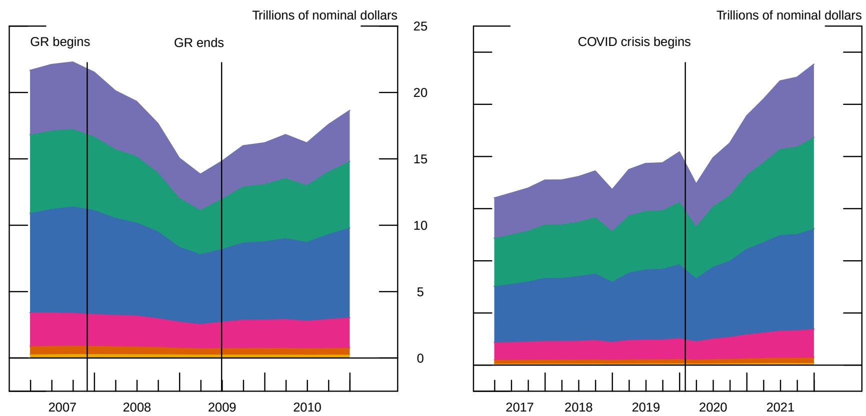

(b) Real estate

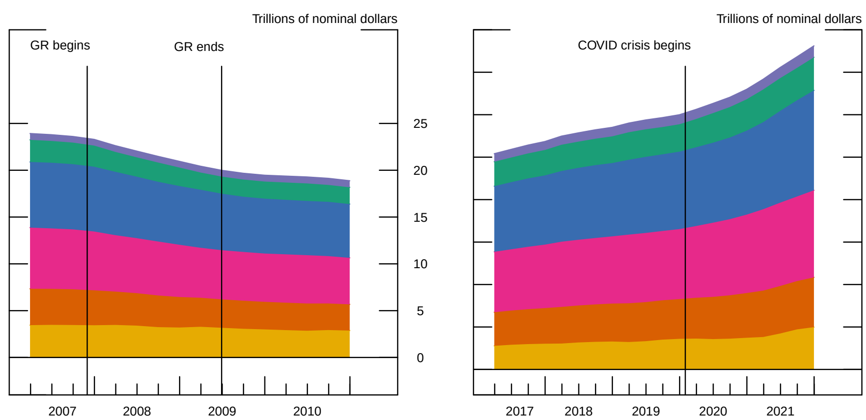

(c) Liquid assets

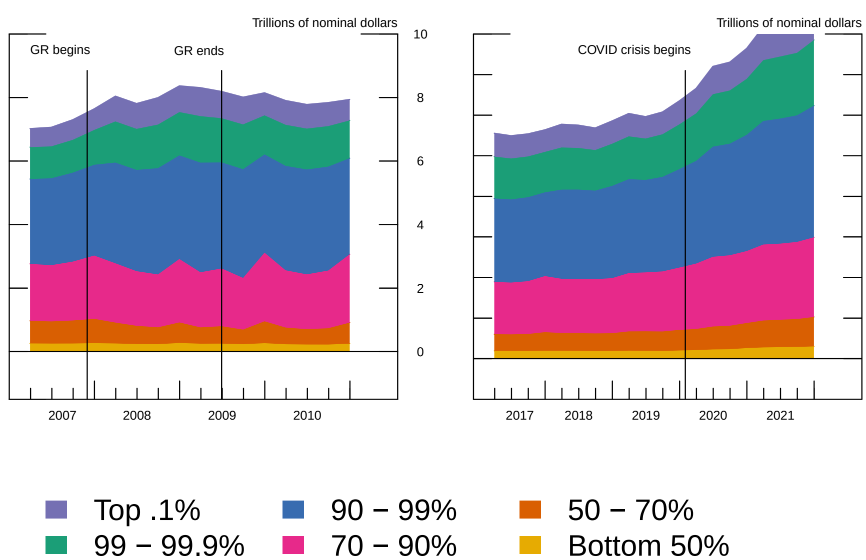

Figure 16: Debt Breakdown During the Great Recession and COVID-19 Crisis

(a) Mortgages

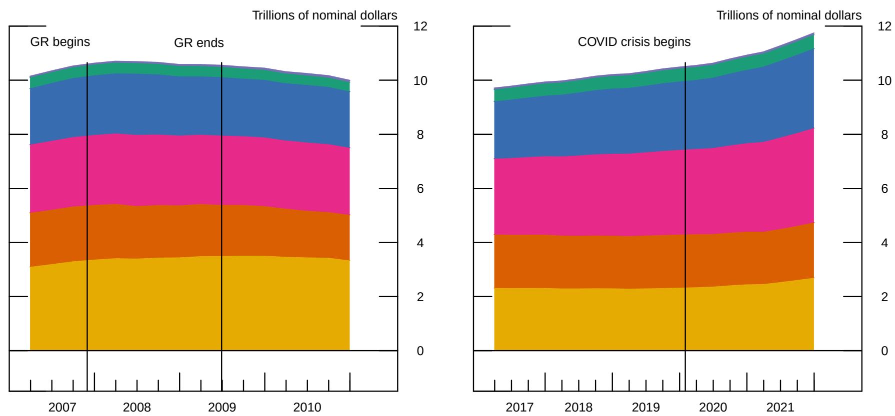

(b) Other liabilities

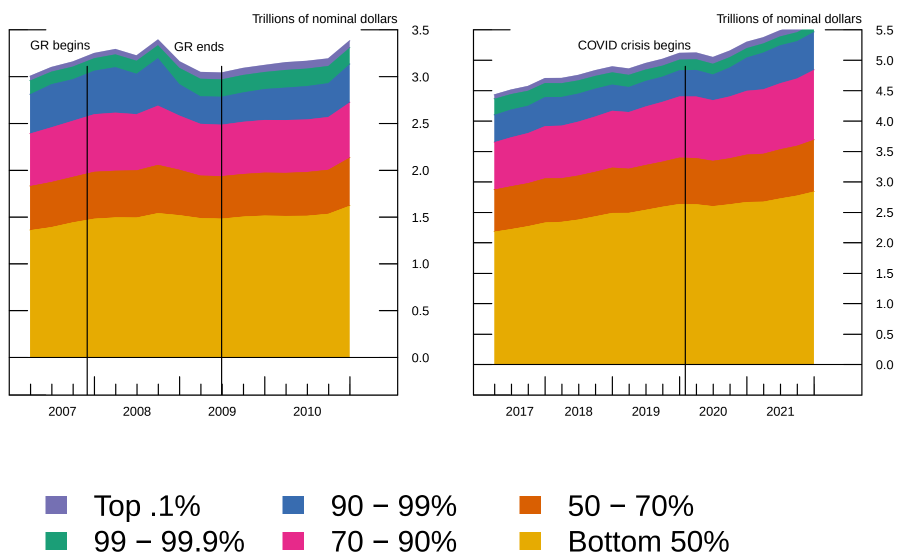

fell through early-2009 before starting to grow again as asset prices recovered, while net worth for the bottom 50% of households continued to decline through 2010 due to continued declines in real estate and a sluggish labor market recovery. In contrast, during the Covid-19 crisis, net worth for all groups recovered after a small decline in 2020Q1, with a rebound in business equity again driving gains among the top 10% wealthiest households, and housing wealth increases and increased savings due to forced spending cuts and increased fiscal transfers driving gains for the bottom 90% of households.

In addition to net worth, the distribution and dynamics of several other balance sheet measures that have attracted significant attention in recent applied macro research can be constructed from our data and their evolution tracked through the Great and Covid-19 pandemic recessions.

To demonstrate, in Panel (a) of Figure 18 we show the evolution of household leverage–the ratio of total liabilities to total assets—to document the comovement of asset and debt levels. Most importantly, leverage is a valuable predictor of financial fragility, since highly levered households might be forced to cut back consumption following a negative economic shock (Mian et al. (2013) and Justiniano et al. (2015)). Prior to the Great Recession, mortgage liabilities increased rapidly for households in the bottom 50% and 50-70th percentiles of the wealth distribution, with the level of liabilities for these groups peaking in 2009q3 and 2007q3, respectively (see Panel (a) in Figure 18). Real estate is the main asset on the balance sheet of these households (Figure 10), so the fall in house prices sparked a notable increase in leverage even as debt levels began to fall (Figure 18). Very different patterns emerged in the COVID recession. Leverage for all groups ticked up slightly in 2020Q1 as the value of business equity dropped, but has otherwise declined through the COVID-19 crisis—particularly for lower wealth groups—owing in large part to stable asset prices and use of forced savings to reduce debt.

A second useful measure, household liquity (the ratio of cash-like assets to total assets) is shown in Panel (b).[25] Recent research has also shown that the distribution of balance sheet liquidity is important to understanding consumption dynamics during downturns, with Kaplan and Violante (2014) documenting that a large share of households have high MPCS due to binding liquidity constraints despite having a large amount of positive wealth. During the Great Recession, liquidity as a share of wealth actually rose for wealthier households due to declines in other asset values, but declined or remained flat for the bottom 70% of households. During the Covid-19 crisis, liquidity rose for all wealth groups and particularly for those towards the bottom of the wealth distribution

[25] Specifically, we follow Kaplan and Violante (2014) and define liquidity as the sum of cash-line financial instruments (i.e, checking and savings accounts and money market funds), debt securities, mortgage and others loans on the asset side of the household balance sheet, and corporate equities, relative to household net worth.

Figure 17: Balance Sheets During the Great Recession and COVID periods

(a) Top 0.1%

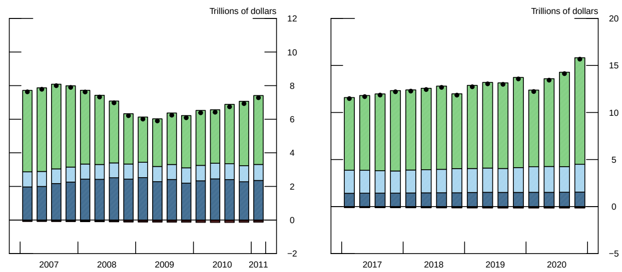

(b) 99-99.9%

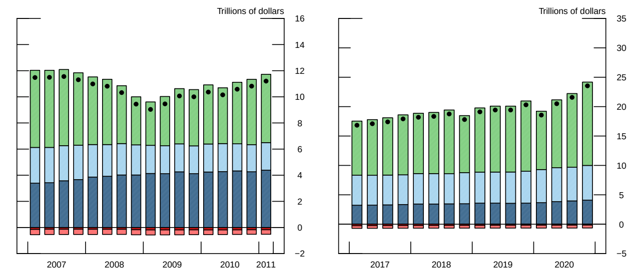

(c) 90 - 99%

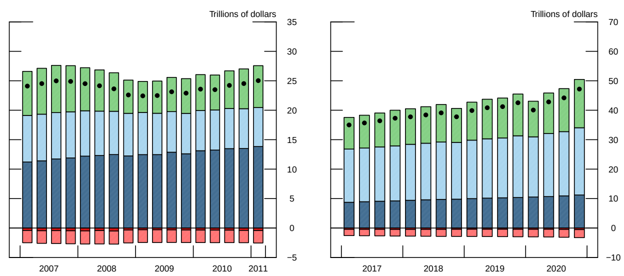

(d) 70 - 90%

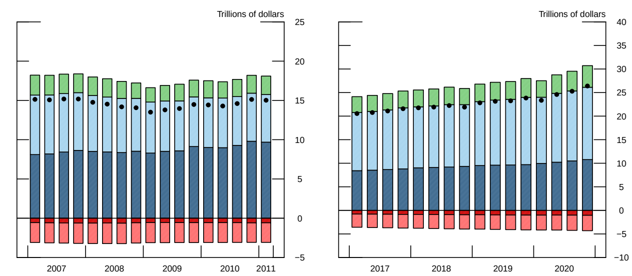

(e) 50 - 70%

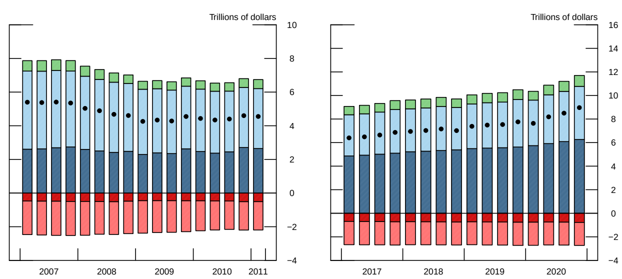

(f) Bottom 50

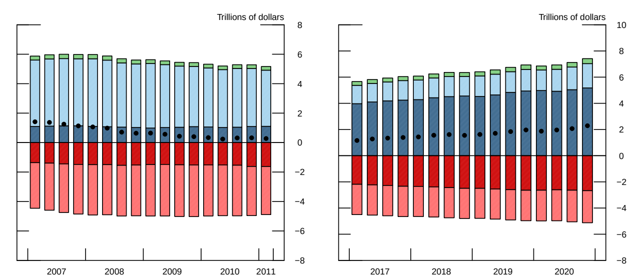

Other Financial Assets
Business Equity
Mortgages
Nonfinancial Assets
Other Liabilities
Net worth

due to the increase in liquid asset holdings we documented previously in Figure 15, panel c.

Figure 18: Balance Sheets During the Great Recession and COVID periods

(a) Leverage

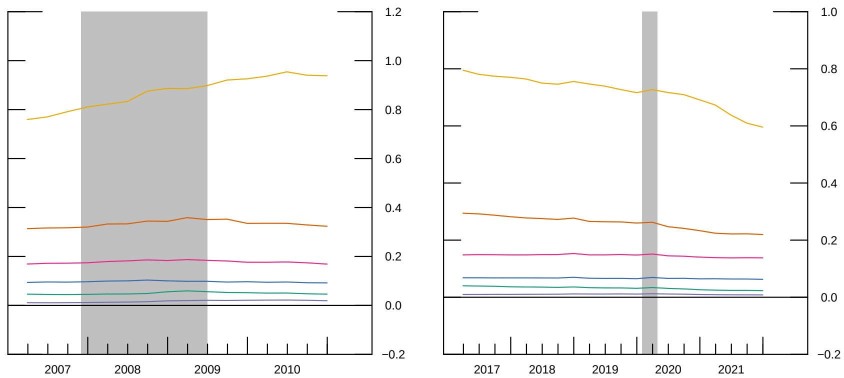

(b) Liquidity

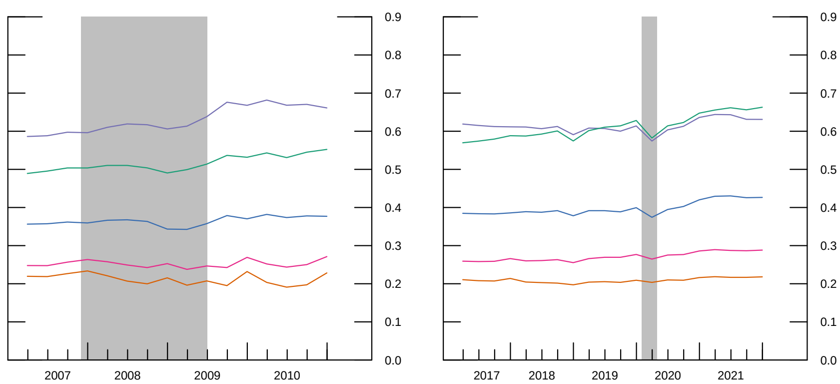

Taken altogether, our data document patterns during the Great Recession that are consistent with those proposed in a large applied macro literature (see, e.g., Corbae and Quintin (2015), Auclert (2019), Greenwald (2018), Jones et al. (2018), Kaplan et al. (2019), Garriga and Hedlund (2019)) that has studied how household balance sheet deterioration among specific parts of the wealth and earnings distribution prolonged the downturn and affected the transmission of macroeconomic shocks. However, opposite patterns have generally emerged in the COVID-19 crisis, as balance sheets for all groups and particularly lower- and middle-class households appear to have strengthened significantly. As a result, our data suggests that a balance sheet hangover is unlikely, and household balance sheets should provide a tailwind to the recovery from the pandemic.

The unique balance sheet dynamics in both these recessions—with unprecedented deterioration during the Great Recession and unprecedented strengthening in the COVID-19 recession—suggests that household balance sheet heterogeneity will remain an important research topic in the near- and medium-term, as applied and theoretical studies attempt to further our understanding of how heterogeneity interacts with aggregate and policy shocks. Our data, particularly the business cycle dynamics in recent recessions documented above, provide a useful set of moments that can be used to calibrate models featuring balance sheet heterogeneity, and will therefore be a key resource in advancing this agenda.

7 Conclusion

Most of the findings in this paper are fairly intuitive, especially, since heterogeneity in household portfolios—which drives much of the variation in the wealth distribution—is observable in lower frequency data. Additionally, most of our findings—for example, how various aggregate shocks affect the wealth distribution—are consistent with the predictions of an active academic literature. However, by providing the first empirical documentation of these patterns, this paper makes a significant contribution to economic research.

Most importantly, our findings have direct implications for macroeconomic models that aim to accurately capture wealth dynamics at business cycle frequencies. We find that different different segments of the wealth distribution hold different portfolios and are exposed to different shocks, which leads to systematic shifts in the distribution of wealth due to cyclical price changes and causes inequality to vary procyclically. However, we also show that group-specific factors also affect wealth accumulation and add to the procyclicality of returns for wealthy households. In contrast, the value of real estate holdings are less cyclical, and therefore dampen cyclical shifts in wealth for middle-class households that are highly exposed to this asset class.

Accurately modeling wealth distribution dynamics, then, requires a sufficiently rich portfolio choice problem to capture the differential exposure to shocks associated with different asset classes.

Our findings sugges that a model featuring three asset classes—riskless bonds, illiquid real estate, and risky capital—will be necessary to capture both the long-run and short-run dynamics of the wealth distribution and its effects on macroeconomic aggregates (see, e.g., Kaplan and Violante (2014)). Research along these lines is rapidly developing (see, e.g., Kekre and Lenel (2019)), and it is our hope that our patterns and data will provide a set of benchmarks and targets that help advance this research agenda and our understanding of how household balance sheets interact with aggregate shocks to determine aggregate outcomes.

References

Ahn, S., Kaplan, G., Moll, B., Winberry, T., and Wolf, C. (2017). When inequality matters for macro and macro matters for inequality. *NBER Macroeconomics Annual 2017*, 32.

Alvaredo, F., Atkinson, A., Chancel, L., Piketty, T., Saez, E., Zucman, G., et al. (2016). Distributional National Accounts (DINA) guidelines: Concepts and methods used in WID. world.

Álvarez-Peláez, M. J. and Díaz, A. (2005). Minimum consumption and transitional dynamics in wealth distribution. *Journal of Monetary Economics*, 52(3):633–667.

Anenberg, E., Ringo, D. R., et al. (2021). Housing market tightness during covid-19: Increased demand or reduced supply? Technical report, Board of Governors of the Federal Reserve System (US).

Antoniewicz, R. L. (1996). A comparison of the household sector from the Flow of Funds Accounts and the Survey of Consumer Finances. *Finance and Economics Discussion Series*, 1996(26):1–28.

Armantier, O., Goldman, L., Koşar, G., Lu, J., Pomerantz, R., and van der Klaauw, W. (2020). How have households used their stimulus payments and how would they spend the next? Liberty street economics, Federal Reserve Bank of New York.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–2367.

Auclert, A. and Rognlie, M. (2018). Inequality and aggregate demand. Working Paper 24280, National Bureau of Economic Research.

Avery, R. B., Elliehausen, G. E., and Kennickell, A. B. (1987). Changes in consumer installment debt: Evidence from the 1983 and 1986 Surveys of Consumer Finances. *Federal Reserve Bulletin*, 73:761.

Bach, L., Calvet, L. E., and Sodini, P. (2020). Rich pickings? risk, return, and skill in household wealth. *American Economic Review*, 110(9):2703–47.

Barro, R. J. (1999). Inequality, growth, and investment. Working Paper 7038, National Bureau of Economic Research.

Batty, M. M., Bricker, J., Briggs, J. S., Holmquist, E. B., McIntosh, S. H., Moore, K. B., Nielsen, E., Reber, S., Shatto, M., Sommer, K., Sweeney, T., and Alice (2019). Introducing the Distributional Financial Accounts of the United States. Finance and Economics Discussion Series 2019-017, Board of Governors of the Federal Reserve System (U.S.).

Benhabib, J., Bisin, A., and Luo, M. (2017). Earnings inequality and other determinants of wealth inequality. *American Economic Review*, 107(5):593–597.

Benhabib, J., Bisin, A., and Zhu, S. (2011). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica*, 79(1):123–157.

Benhabib, J., Bisin, A., and Zhu, S. (2016). The distribution of wealth in the BlanchardâYaari model. *Macroeconomic Dynamics*, 20(2):466â481.

Bernanke, B. S. (2015). Monetary policy and inequality.

Bernanke, B. S. and Kuttner, K. N. (2005). What explains the stock market's reaction to federal reserve policy? *The Journal of Finance*, 60(3):1221–1257.

Bhandari, A., Evans, D., Golosov, M., and Sargent, T. J. (2018). Inequality, business cycles, and monetary-fiscal policy. Working Paper 24710, National Bureau of Economic Research.

Bhutta, N. (2015). The ins and outs of mortgage debt during the housing boom and bust. *Journal of Monetary Economics*, 76:284–298.

Bilbiie, F., Eggertsson, G. B., Primiceri, G. E., and Tambalotti, A. (2021). Excess savings are not excessive. Liberty street economics, Federal Reserve Bank of New York.

Bricker, J., Henriques, A., Krimmel, J., and Sabelhaus, J. (2016). Measuring income and wealth at the top using administrative and survey data. *Brookings Papers on Economic Activity*, 2016(1):261–331.

Bu, C., Rogers, J., and Wu, W. (2021). A unified measure of fed monetary policy shocks. *Journal of Monetary Economics*, 118:331–349.

Bullard, J. B. (2014). Income inequality and monetary policy: a framework with answers to three questions. Speech 235, Federal Reserve Bank of St. Louis.

Calza, A., Monacelli, T., and Stracca, L. (2013). Housing finance and monetary policy. *Journal of the European Economic Association*, 11:101–122.

Carroll, C. D. (2014). Representing consumption and saving without a representative consumer. In *Measuring Economic Sustainability and Progress*, NBER Chapters, chapter 5, pages 115–134. National Bureau of Economic Research, Inc.

Chatterjee, S. (1994). Transitional dynamics and the distribution of wealth in a neoclassical growth model. *Journal of Public Economics*, 54(1):97–119.

Chow, G. C. and Lin, A.-L. (1971). Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *The Review of Economics and Statistics*, 53(4):372.

Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2012). Innocent bystanders? Monetary policy and inequality in the U.S. Working Paper 18170, National Bureau of Economic Research.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). How did u.s. consumers use their stimulus payments? Working Paper 27693, National Bureau of Economic Research.

Cooper, D. (2012). Us household deleveraging: what do the aggregate and household-level data tell us? *Federal Reserve Bank of Boston Public Policy Brief*, (12-2).

Corbae, D. and Quintin, E. (2015). Leverage and the foreclosure crisis. *Journal of Political Economy*, 123(1):1–65.

Cox, J., Greenwald, D. L., and Ludvigson, S. C. (2020). What explains the covid-19 stock market? Technical report, National Bureau of Economic Research.

Deaton, A. (2005). Measuring poverty in a growing world (or measuring growth in a poor world). *Review of Economics and Statistics*, 87(2):395–395.

Dettling, L. J., Devlin-Foltz, S., Krimmel, J., Pack, S., and Thompson, J. P. (2015). Comparing micro and macro sources for household accounts in the United States: evidence from the Survey of Consumer Finances.

Doepke, M. and Schneider, M. (2006). Inflation and the redistribution of nominal wealth. *Journal of Political Economy*, 114(6):1069–1097.

Dollar, D. and Kraay, A. (2002). Growth is good for the poor. *Journal of Economic Growth*, 7(3):195–225.

Dowd, T., McClelland, R., and Mortenson, J. A. (2019). Tax policy center.

Eller, T. (1994). Household wealth and asset ownership: 1991. *U.S. Bureau of the Census, Current Population Reports*.

Erosa, A. and Ventura, G. (2002). On inflation as a regressive consumption tax. *Journal of Monetary Economics*, 49(4):761–795.

Fagereng, A., Guiso, L., Malacrino, D., and Pistaferri, L. (2020). Heterogeneity and persistence in returns to wealth. *Econometrica*, 88(1):110–170.

Fernandez, R. B. (1981). A methodological note on the estimation of time series. *The Review of Economics and Statistics*, 63(3):471.

Ferreira, F. H. (1995). Roads to equality: Wealth distribution dynamics with public-private capital complementarity.

Gallin, J. H., Sommer, K., Nielsen, E., Molloy, R. S., and Smith, P. A. (2018). Measuring aggregate housing wealth: New insights from an automated valuation model. Finance and Economics Discussion Series 2018-064, Board of Governors of the Federal Reserve System (U.S.).

Garriga, C. and Hedlund, A. (2019). Crises in the housing market: Causes, consequences, and policy lessons. *Oxford Research Encyclopedia of Economics and Finance*.

Garriga, C., Kydland, F. E., and Šustek, R. (2015). Mortgages and monetary policy.

Gormsen, N. J. and Koijen, R. S. (2020). Coronavirus: Impact on stock prices and growth expectations. *The Review of Asset Pricing Studies*, 10(4):574–597.

Gornemann, N., Kuester, K., and Nakajima, M. (2016). Doves for the rich, hawks for the poor? Distributional consequences of monetary policy. *International Finance Discussion Paper*, 2016(1167):1–40.

Greenwald, D. L. (2018). The mortgage credit channel of macroeconomic transmission. *SSRN Electronic Journal*.

Guerrieri, V. and Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics*, 132(3):1427–1467.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2019). What do data on millions of U.S. workers reveal about life-cycle earnings dynamics?

Guvenen, F., Ozkan, S., and Song, J. (2012). The nature of countercyclical income risk.

Hagedorn, M., Manovskii, I., and Mitman, K. (2019). The fiscal multiplier. Working Paper 25571, National Bureau of Economic Research.

Hanson, S. G., Stein, J. C., Sunderam, A., and Zwick, E. (2020). Business credit programs in the pandemic era. *Brookings Papers on Economic Activity*, 2020(3):3–60.

Heathcote, J., Perri, F., and Violante, G. L. (2010). Unequal we stand: An empirical analysis of economic inequality in the United States: 1967-2006. *Review of Economic Dynamics*, 13(1):15–51.

Henriques, A. W. and Hsu, J. W. (2014). Analysis of wealth using micro and macro data: A comparison of the Survey of Consumer Finances and Flow of Funds Accounts. *Measuring Economic Sustainability and Progress*, pages 245–274.

Horvath, A., Kay, B. S., and Wix, C. (2021). The covid-19 shock and consumer credit: Evidence from credit card data. *Available at SSRN 3613408*.

Jones, C., Midrigan, V., and Philippon, T. (2018). Household leverage and the recession. *IMF Working Papers*, 18(194):1.

Jordá, Ó. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Justiniano, A., Primiceri, G., and Tambalotti, A. (2015). Household leveraging and deleveraging. *Review of Economic Dynamics*, 18(1):3–20.

Kaplan, G., Mitman, K., and Violante, G. L. (2019). The housing boom and bust: Model meets evidence. *SSRN Electronic Journal*.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3):697–743.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Kaplan, G. and Violante, G. L. (2018). Microeconomic heterogeneity and macroeconomic shocks. *Journal of Economic Perspectives*, 32(3):167–94.

Kekre, R. and Lenel, M. (2019). Redistribution, risk premia, and the macroeconomy. *Working Paper*.

Kopczuk, W. and Saez, E. (2004). Top wealth shares in the United States, 1916-2000: Evidence from estate tax returns. *National Tax Journal*, 57(2):445–487.

Krueger, D., Mitman, K., and Perri, F. (2016). Macroeconomics and household heterogeneity. *Handbook of Macroeconomics*, pages 843–921.

Krusell, P. and Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Kuhn, M., Schularick, M., and Steins, U. I. (2018). Income and wealth inequality in America, 1949-2016.

Kuznets, S. (1955). Economic growth and income inequality. *American Economic Review*, 45(1):1–28.

Litterman, R. B. (1983). A random walk, Markov model for the distribution of time series. *Journal of Business Economic Statistics*, 1(2):169–173.

Luetticke, R. (2018). Transmission of monetary policy with heterogeneity in household portfolios.

Maki, D. M. and Palumbo, M. (2001). Disentangling the wealth effect: A cohort analysis of household saving in the 1990s. *FEDS Working Paper Series*.

Mckay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, 106(10):3133–3158.

McKay, A. and Reis, R. (2016). The role of automatic stabilizers in the U.S. business cycle. *Econometrica*, 84:141–194.

Meh, C. A., Ríos-Rull, J.-V., and Terajima, Y. (2010). Aggregate and welfare effects of redistribution of wealth under inflation and price-level targeting. *Journal of Monetary Economics*, 57(6):637–652.

Mersch, Y. (2014). Corporate credit conference.

Mian, A., Rao, K., and Sufi, A. (2013). Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics*, 128(4):1687–1726.

Mian, A. and Sufi, A. (2014). What explains the 2007-2009 drop in employment? *Econometrica*, 82(6):2197–2223.

Mian, A., Sufi, A., and Verner, E. (2017). Household debt and business cycles worldwide*. *The Quarterly Journal of Economics*, 132(4):1755â1817.

Mönch, E. and Uhlig, H. (2005). Towards a monthly business cycle chronology for the Euro area. *Journal of Business Cycle Measurement and Analysis*, 2005(1):43–69.

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: the information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330.

Paul, P. (2018). Monetary Policy Cycles and Financial Stability. *FRBSF Economic Letter*.

Paul, P. (2020). The time-varying effect of monetary policy on asset prices. *Review of Economics and Statistics*, 102(4):690–704.

Piketty, T. (1997). The dynamics of the wealth distribution and the interest rate with credit rationing. *The Review of Economic Studies*, 64(2):173.

Piketty, T. (2013). *Capital in the Twenty-First Century*. Harvard University Press.

Piketty, T., Saez, E., and Zucman, G. (2018). Distributional National Accounts: Methods and estimates for the United States. *The Quarterly Journal of Economics*, 133(2):553–609.

Proietti, T., Marczak, M., and Mazzi, G. (2017). Euromind-: A density estimate of monthly gross domestic product for the euro area. *Journal of Applied Econometrics*, 32(3):683–703.

Ramani, A. and Bloom, N. (2021). The donut effect: How covid-19 shapes real estate. *SIEPR Policy Brief, January*.

Rigobon, R. and Sack, B. (2004). The impact of monetary policy on asset prices. *Journal of Monetary Economics*, 51(8):1553–1575.

Ríos-Rull, J.-V. and Kuhn, M. (2016). 2013 update on the u.s. earnings, income, and wealth distributional facts: A view from macroeconomics. *Quarterly Review*, 37(1).

Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084.

Romer, C. D. and Romer, D. H. (2021). The fiscal policy response to the pandemic. *Brookings Papers on Economic Activity*.

Rubio, M. (2011). Fixed- and variable-rate mortgages, business cycles, and monetary policy. *Journal of Money, Credit and Banking*, 43(4):657–688.

Sabelhaus, J. and Volz, A. H. (2019). Are disappearing employer pensions contributing to rising wealth inequality? *FEDS Notes*, 2019(2308).

Saez, E. and Zucman, G. (2016). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, 131(2):519–578.

Sala-i Martin, X. (2006). The world distribution of income: Falling poverty and convergence, period. *The Quarterly Journal of Economics*, 121(2):351–397.

Smith, M., Zidar, O., and Zwick, E. (2019). Top wealth in the United States: New estimates and implications for taxing the rich.

Werning, I. (2015). Incomplete markets and aggregate demand. *NBER Working Paper*.

Wolff, E. N., Zacharias, A., and Masterson, T. (2012). Trends in American living standards and inequality, 1959-2007. *Review of Income and Wealth*, 58(2):197–232.

Yellen, J. L. (2014). Perspectives on inequality and opportunity from the survey of consumer finances.

A More on Data and Methodology

A.1 Reconciliation of Major Asset Classes and Other Challenges

Both the SCF and FA aim to provide comprehensive wealth measures of household balance sheets, but wealth concepts sometimes differ across the two data sets. Here we describe issues with, and solutions to, reconciling the four largest balance sheet items (which account for 78% of total assets) and briefly discuss key challenges in reconciling other asset categories. We refer readers interested in details on the mapping from SCF to the Financial Accounts for other balance sheet lines to Batty et al. (2019), a technical paper which details the data's construction.

Pension entitlements Pension entitlements make up the largest B.101.h asset category, accounting for nearly a quarter of aggregate household assets. This category includes the balances of defined contribution (DC) pension plans (such as 401(k) and 403(b) plans), accrued benefits to be paid in the future from defined benefit (DB) plans (including those for which life insurance companies have assumed the payment obligation), and annuities sold by life insurers directly to individuals.[26] These three asset classes account for about 30%, 60%, and 10% of total pension entitlements in the Financial Accounts, respectively.[27] The SCF captures DC balances using a method that is compatible with the one used in the construction of the Financial Accounts. The DC aggregates between the two data sources are generally close, with a historical ratio of 97%. However, the SCF does not directly measure accrued DB benefits or annuities. The DFAs consequently follow the method developed in Sabelhaus and Volz (2019) to distribute the DB component of the B.101.h aggregate across the SCF households. Specifically, we first split SCF households who are entitled to DB benefits into those currently receiving pension payments, those expecting future payments from a past job, and those expecting future payments from a current job. We then use the benefit amount to assign DB pension wealth to those currently collecting, the expected timing and amount of future pension benefits to assign DB pension wealth to those with pensions from past jobs, and then allocate the residual DB pension reserves to those with pensions on their current job that are

[26]The annuities component also includes annuities held in individual retirement accounts (IRAs). IRA investments in other instruments, such as mutual fund shares, are included in the corporate equity and mutual fund balance sheet line.

[27]The defined-benefit component includes total accrued benefits from private-sector, state-and-local government, and federal employment, whether fully funded or not. Notably, it does not include Social Security, which is not currently included in the Financial Accounts.

not collecting according to wage, years in the plan, and age. See Sabelhaus and Volz (2019) for a more detailed description of this imputation methodology.

Similarly to accrued DB pensions, measures of annuity reserves are not directly collected by the SCF in a manner compatible with B.101.h. However, the SCF does report information that can be used to impute the value of annuities for SCF households. Specifically, the SCF reports the amount of income received from annuities that are in the payout phase, as well as the cash value of deferred annuities (which differs from the reserve due to surrender penalties and other policy benefits not immediately payable in cash).[28] To reconcile the SCF and B.101.h annuity measures, the DFAs capitalize the payout annuity income reported by SCF households into a present value using a set of sample annuity policies (Batty et al. (2019) for details), and then distribute the B.101.h annuity reserves according to the sum of the cash value of deferred annuities and capitalized value of payout annuities reported in the SCF.

Real estate Aggregate real estate measures in the Financial Accounts and SCF align reasonably well until the mid-2000s, although the SCF measure has consistently exceeded the B.101.h values. Important methodological differences drive the divergence between the SCF and Financial Accounts measures of housing wealth during the mid-2000s housing cycle. Specifically, the SCF measure is based on owner-reported values, whereas the Financial Accounts measure applies an automated valuation model (AVM) from Zillow. Gallin et al. (2018)—and studies cited therein—show that owner self-reports values tend to lag the market during market turns and also tend to be overly optimistic, potentially explaining a portion of the discrepancy. The sizable, time-varying gap between the Financial Accounts and SCF measures of housing wealth is notable, but the key question for our purposes is whether it causes bias when we apply the SCF distribution to the Financial Accounts measures. Batty et al. (2019) assesses assess the sensitivity of the DFAs distributional measures to a different aggregate housing wealth series recently developed by Board staff and finds that using this series results in minimal effect on the distribution of housing wealth, therefore suggesting that the difference in the level of housing wealth between the SCF and B.101.h is relatively

[28]In contrast to traditional annuities, deferred annuities are savings products offered by insurance companies. The account balance of some of these products accumulate at a rate set by the insurer, usually subject to a minimum guarantee determined at the time of sale. Others offer equity market participation, often with some type of embedded return guarantee. These products are called annuities because the policyholder has the option to later annuitize the value of the policy into periodic payments, but exercising this option is not typical.

evenly distributed across wealth groups.

Equity in noncorporate business This B.101.h balance sheet item includes non-publicly traded businesses and real estate owned by households for renting out to others. There are substantial differences in its measurement between the SCF and Financial Accounts. The B.101.h measure is a hybrid of different accounting bases. Real estate (e.g., rental properties), which accounts for approximately 60% of this category, is recorded at market value. In contrast, other nonfinancial assets are recorded at cost basis, based on investment data collected by the BEA, while financial assets and liabilities are recorded at book value using tax data. In the SCF, rental properties are reported at market value, as they are in the Financial Accounts, but for other noncorporate business assets, the SCF captures owners' self-reports of both the market value and the cost basis of their businesses. When we compare these two measures to B.101.h, we find (unsurprisingly) that the market-value SCF measure exceeds the B.101.h measure (with an average ratio of approximately 150%), while the cost-basis SCF measure falls below the B.101.h measure (with an average ratio of 70%). However, the SCF's cost-basis and market-value measures are near identical. Because the shares implied by the reconciled SCF distribution are applied to the Financial Accounts levels in the final step of the DFAs, the similarity in distributions suggests that either measure is unlikely to bias our estimates, as they are distributed similarly. We therefore use the average of the two SCF noncorporate business valuations in the DFAs.

Corporate equities and mutual funds This category includes all public equity, private equity, and mutual fund holdings, except for equity and mutual funds held through DC pensions. The corresponding SCF measure is comprised of directly held stocks and mutual funds, in addition to the portion of other investment vehicles that are invested in equities (such as IRAs, trusts, managed investment accounts, 529 plans, and Health Savings Accounts). Historically, the SCF measure is quite close to the B.101.h measure, averaging about 102%, and is relatively consistent across years.

Other Assets and Liabilities There are other existing differences between the reconciled SCF and Financial Accounts' balance sheets in smaller asset and liability categories. For instance, life insurance reserves are generally unknown by policy holders and thus are unmeasured in the SCF. We assign the B.101.h measures of term policy reserves according to the death benefit recorded in the SCF, and permanent policy reserves by the death benefit and the cash surrender value. Additionally, consumer durables, which are sometimes excluded from household balance sheets (see, e.g., Wolff et al. (2012); Saez and Zucman (2016)) are only 60% as large in the SCF as on B.101.h. This likely occurs because the BEA measure utilized in the construction of B.101.h covers any item that has resale value, whereas the SCF questions encourage respondents to focus only on the most substantial assets.[29] In unshown analyses, we divided the SCF assets into the twenty-eight BEA consumer durable categories and find no evidence that the SCF more severely underreports consumer durable goods that are likely more evenly distributed (such as "window covering" or "sporting equipment") than items that are more likely concentrated among the wealthy (such as "jewelry and watches" or "pleasure aircraft"). Finally, our reconciliation is not always perfect. For example, time deposits and short-term investments are understated in the SCF relative to B.101.h, while checkable deposits and currency holdings are overstated, despite these asset classes aligning conceptually between the two data sets. One partial explanation for this pattern could be misclassification by SCF respondents, while another could be mismeasurement of the B.101.h level, likely due to the residual nature of the construction of these series. Additionally, smaller asset and liability categories, e.g., corporate and foreign bonds (.7% of total assets) and depository institution loans (1.6% of liabilities), neither match well empirically nor is the difference easily explainable. However, given their relatively small contribution to household balance sheets, it is unlikely that these discrepancies will meaningfully affect our findings about the wealth distribution.

A.2 Estimating Covariance Matrices of the Error Process in the Chow-Lin Methodology

This appendix describes in greater detail the estimation process and how the higher-frequency covariance matrix V is identified in Chow and Lin (1971), Fernandez (1981), and Litterman (1983).

The Chow-Lin method solves the multiple regression model specified by Equations 1 and 3 to obtain an estimate of $\hat{X}$ given observations Y and Z and covariance matrix V. Chow and Lin

[29] While the SCF question about consumer durables offers examples of items that fall into many of the BEA categories, its prompt begins with a list geared towards items that may have considerable value, as opposed to typical household goods: "for example, artwork, precious metals, antiques, oil and gas leases, futures contracts, future proceeds from a lawsuit or estate that is being settled, royalties, or something else?"

(1971) show that a linear unbiased estimate $\hat{X}$ is given by

$$\hat{X} = Z\hat{\beta} + VB(B'VB)^{-1}[Y - B'Z\hat{\beta}] \tag{16}$$

$$\hat{\beta} = [Z'B(B'VB)^{-1}B'Z]^{-1}Z'B(B'VB)^{-1}Y. \tag{17}$$

Here, $\hat{\beta}$ is a vector obtained from the generalized least squares regression specified in Equation 3 with Y as the dependent variable, $B'Z$ as the dependent variable, and residual covariance matrix $(B'VB)$.

Equation 16 shows that the estimate $\hat{X}$ can be expressed as the sum of two components. The first component, $Z\hat{\beta}$, represents the predicted values of the higher-frequency target series X given the higher-frequency observations of Z, i.e., $\mathbb{E}[X|Z]$. The second component, $VB(B'VB)^{-1}[Y - B'Z\hat{\beta}]$, reflects the estimate of the vector of higher-frequency residuals obtained by distributing the vector of lower-frequency residuals $[Y - B'Z\hat{\beta}]$ across periods where the target series is unobserved. The distributing matrix $VB(B'VB)^{-1}$ is determined by the assumed covariance matrix V. Note that $\hat{X} = Y$ by construction for the periods that Y is observed.

A key input into this method is the assumed error structure of the higher-frequency residuals, represented by V. This covariance matrix is not observed and must be estimated—any consistent estimate for V can then be used to obtain FGLS estimates $\hat{\beta}$ and $\hat{X}$.

Chow and Lin (1971) show how to recover the higher-frequency covariance matrix V under two different assumptions about the underlying error process: serial independence and first-order autocorrelation, which is the leading case we pursue in this paper. In particular, they show that if the residuals follow a simple AR(1) process such that

$$u_t = au_{t-1} + \epsilon_t, \tag{18}$$

where the ϵ_t are iid with constant variance σ^2 then

$$\begin{aligned} V &= \begin{bmatrix} 1 & a & a^2 & ... & a^{n-1} \\ a & 1 & a & ... & a^{n-2} \\ a^2 & a & 1 & ... & a^{n-3} \\ \vdots & ... & ... & ... & \vdots \\ a^{3n-1} & ... & & ... & 1 \end{bmatrix} \\ &= A \times \frac{\sigma^2}{1-a^2}. \end{aligned}$$

Substituting Equation 18 into Equations 16 and 17 reveals that a feasible estimate of $\hat{X}$ requires an estimate of a but not σ^2 (the scalar factor $\sigma^2/(1-a^2)$ cancels in all of the expressions). To estimate a, note that the first order autocorrelation of $\left[Y - B'Z\hat{\beta}\right]$ is a^3. Iteratively using Equation 18 and Equation 16 and solving for a^3 by calculating the autocorrelation coefficient of $\left[[Y - B'Z\hat{\beta}\right]$ until convergence therefore yields a consistent estimate of a, and, by extension, V.

This basic approach has been generalized and extended by several other studies. Notably, Fernandez (1981) and Litterman (1983) characterize solutions for non-stationary error processes of the form

$$\begin{aligned} u_t &= a u_{t-1} + v_t \\ v_t &= \rho v_{t-1} + \eta_t. \end{aligned}$$

Fernandez (1981) assumes $\rho = 0$, while Litterman (1983) assumes $0 < \rho < 1$. In each of these cases, the solution follows the familiar form specified in Equations 16 and 17 with covariance matrix V given by

$$V = \left[\Delta' H(\rho)' H(\rho) \Delta\right]^{-1} \times \sigma_\eta^2,$$

where Δ is an $n \times n$ difference matrix with 1 on its diagonal, -1 on its subdiagonal, and zero elsewhere, $H(\rho)$ is an $n \times n$ matrix with 1 on its diagonal, $-\rho$ on its subdiagonal, and zero elsewhere, and σ_η^2 is the variance of the innovations η_t. In particular, Litterman (1983) shows that autoregres-

Figure 19: Alternative Error Assumptions

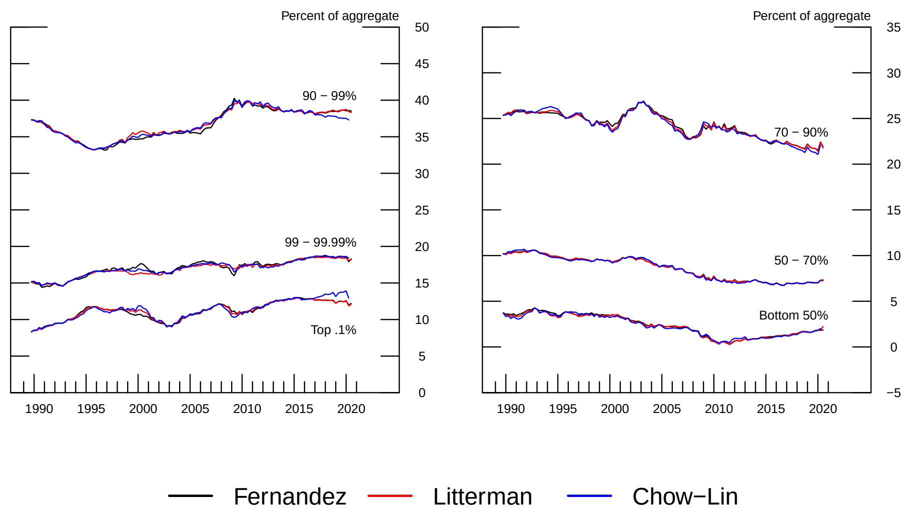

sive parameter ρ may be estimated by an iterative procedure similar to that proposed in Chow-Lin (1971) using Equations 16 and 17 and the first-order autocorrelation of the first difference of the residuals $\left[[Y - B'Z\hat{\beta}\right]$.

A.3 Alternative Error Processes

As noted in Section 2.2, our baseline estimation assumes an AR(1) error process. To explore whether this assumption affects our results, we reconstruct our data allowing for errors to follow a random walk as studied in Fernandez (1981) and Markov switching model as studied Litterman (1983).

The resulting data series are presented in Figure 19. Overall, we find little difference between our baseline data series and those allowing richer residual dynamics. Furthermore, the differences we do observe reflect tradeoffs among the top 1 and next 9 households, suggesting that these alternative error processes predominately affect the top of the wealth distribution. Still, the variation in wealth shares for the wealth groups considered in this study is relatively minimal, suggesting that our results are robust to alternative error assumptions.

A.4 The Data Overview: Zooming in on Major Asset Classes

Figure 20: Shares of major asset classes

(a) Real Estate **(b)** Pensions

(c) Corporate Equities and Mutual Funds **(d)** Noncorporate Business Equity

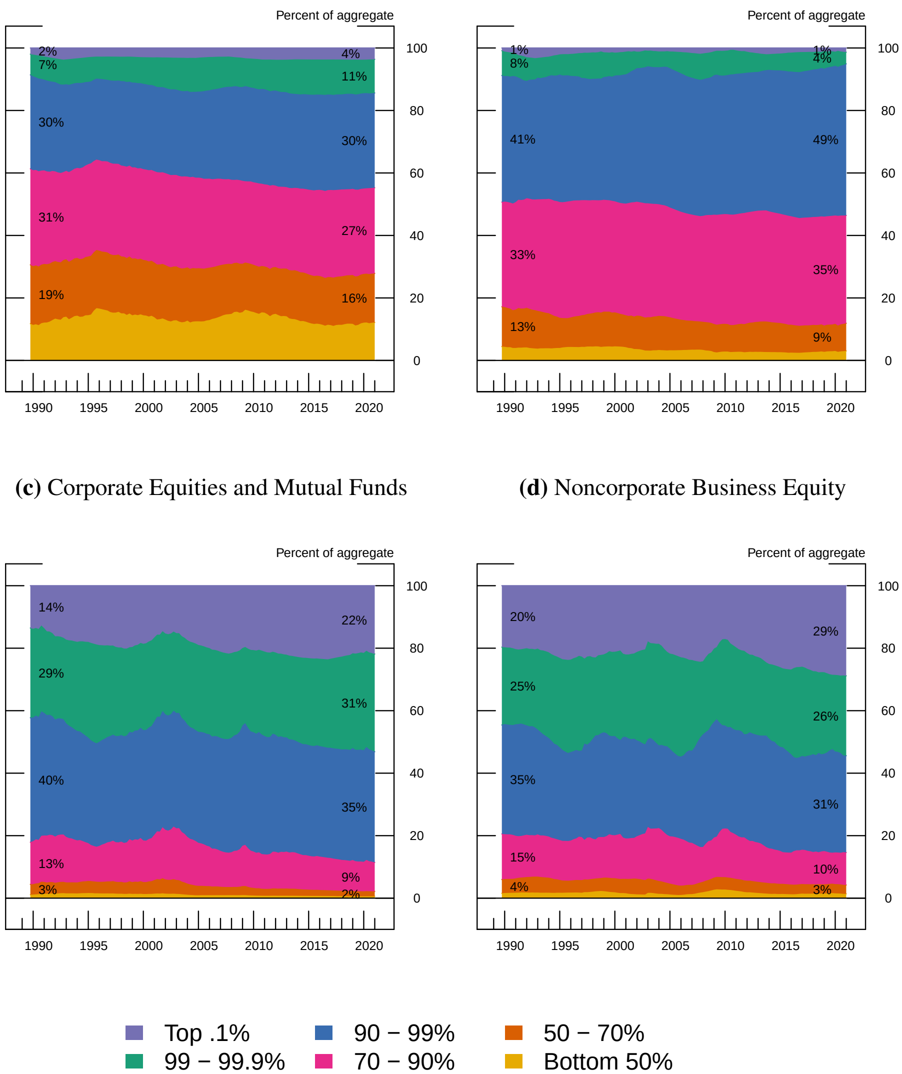

To complete our initial exploration of the data in Section 2.3, Figure 20 zooms in on the evolution of wealth and its distribution for four major asset categories: real estate, pensions, corporate equity, and non-corporate equity.

B Factor Pricing Model with a Full Estimation Sample

In Table 4, we repeat estimation from Table 2 but extend our sample to include the Covid-19 period (i.e., 2020q1-present). The patterns from these regressions are broadly the same, with the exception that the unemployment rate is no longer significant for low wealth households. This reflects the well documented pattern that wealth increased sharply, particularly for low wealth households, in the COVID-19 recession due to very high fiscal transfers.

Table 4: Full sample (including COVID-19) Fama French regression

	Dependent variable:					
	Top .1%	99-99.9%	90-99%	70-90%	50-70%	50%
	(1)	(2)	(3)	(4)	(5)	(6)
S&P 500	0.306***	0.295***	0.230***	0.160***	0.109***	0.177***
	(0.020)	(0.019)	(0.014)	(0.010)	(0.009)	(0.045)
Core Logic House Price Index	0.352***	0.354***	0.349***	0.421***	0.602***	2.344***
	(0.089)	(0.084)	(0.060)	(0.044)	(0.038)	(0.197)
FFR, 1 Quarter Change	−0.323	−0.297	−0.323	−0.279	−0.240	0.278
	(0.371)	(0.351)	(0.253)	(0.182)	(0.158)	(0.827)
5 Year Forward Rates	−0.483	−0.579	−0.212	0.073	0.404	3.345**
	(0.750)	(0.710)	(0.511)	(0.368)	(0.319)	(1.670)
Unemployment Rate	0.182	0.192	0.110	0.070	0.031	−0.164
	(0.146)	(0.138)	(0.099)	(0.072)	(0.062)	(0.325)
Constant	0.099	0.024	−0.148	−0.176	−0.010	0.494
	(0.227)	(0.215)	(0.155)	(0.111)	(0.097)	(0.505)
Observations	107	107	107	107	107	107
R^2	0.754	0.764	0.798	0.825	0.851	0.682
Adjusted R^2	0.742	0.752	0.789	0.816	0.844	0.666
Residual Std. Error (df = 101)	1.504	1.424	1.026	0.739	0.641	3.351
F Statistic (df = 5; 101)	61.941***	65.214***	80.038***	95.175***	115.603***	43.303***

Note: *p<0.1; **p<0.05; ***p<0.01

C Elasticity

Our analysis to this point establishes that exposure to asset price risk varies with wealth, and that this variation contributes to the dynamics of the wealth distribution. However, it does not quantify how this risk varies over time. To do so, we follow Kuhn et al. (2018) to derive the elasticity of group p's wealth share with respect to asset j's price:

$$\frac{\partial\left(\frac{\omega_t^{NW,p}}{\omega_{t-1}^{NW,p}}\right)}{\partial\left(\frac{p_t^j}{p_{t-1}^j}\right)} = (E_t + U_t)^{-1}\left(\alpha_{t,j}^i - \alpha_{t,j}\frac{\omega_t^{NW,p}}{\omega_{t-1}^{NW,p}}\right). \tag{19}$$

The above expression thus characterizes the effects of price changes on the distribution of wealth in each quarter. If a group's share of wealth in asset j is larger than the aggregate balance sheet's (i.e., $\alpha_{t,j}^i > \alpha_{t,j}$), the elasticity is larger if aggregate wealth growth ($E_t + U_t$) is smaller and if wealth growth if bigger, while the opposite patterns hold if $\alpha_{t,j}^i > \alpha_{t,j}$.

Figure 21 presents the four-quarter moving-average of the price elasticity for the four largest asset classes for our six wealth groups.[30] Households in the bottom half of the wealth distribution have a positive wealth share elasticity with respect to pensions and real estate and negative elasticity with respect to business equity. Furthermore, there is minimal cyclical trend in any of their elasticities, but their elasticities spiked during the Great Recession as leverage increased dramatically and their wealth share hovered at or below zero. The wealth share elasticities of households in the 50 to 90th percentiles of the wealth distribution confirm high exposure to changes in pensions and real estate wealth: a 1% increase in pension wealth increases their overall wealth share by .06 to.14% over this time period, while a 1% increase in real estate values would cause a .08 to .14% increase in wealth shares. Both of these elasticities exhibit interesting variation over time. For example, because this group owns a significant share of real estate overall, the wealth share elasticity with respect to real estate peaked at the height of the housing bubble in 2007 and fell sharply with with the collapse of house prices. With regards to corporate and noncorporate business equity, the

[30] Our elasticities are broadly consistent with those in Kuhn et al. (2018), and quantitative differences reflect conceptual differences between our wealth measure (i.e., Financial Accounts Table B.101.h) and that in their SCF-plus. Our measures complement the long-run trends they document by providing a more recent, higher frequency look at the the evolution of these elasticities for different segments of the wealth distribution (e.g., the top 1%).

Figure 21: Elasticities of Wealth Shares

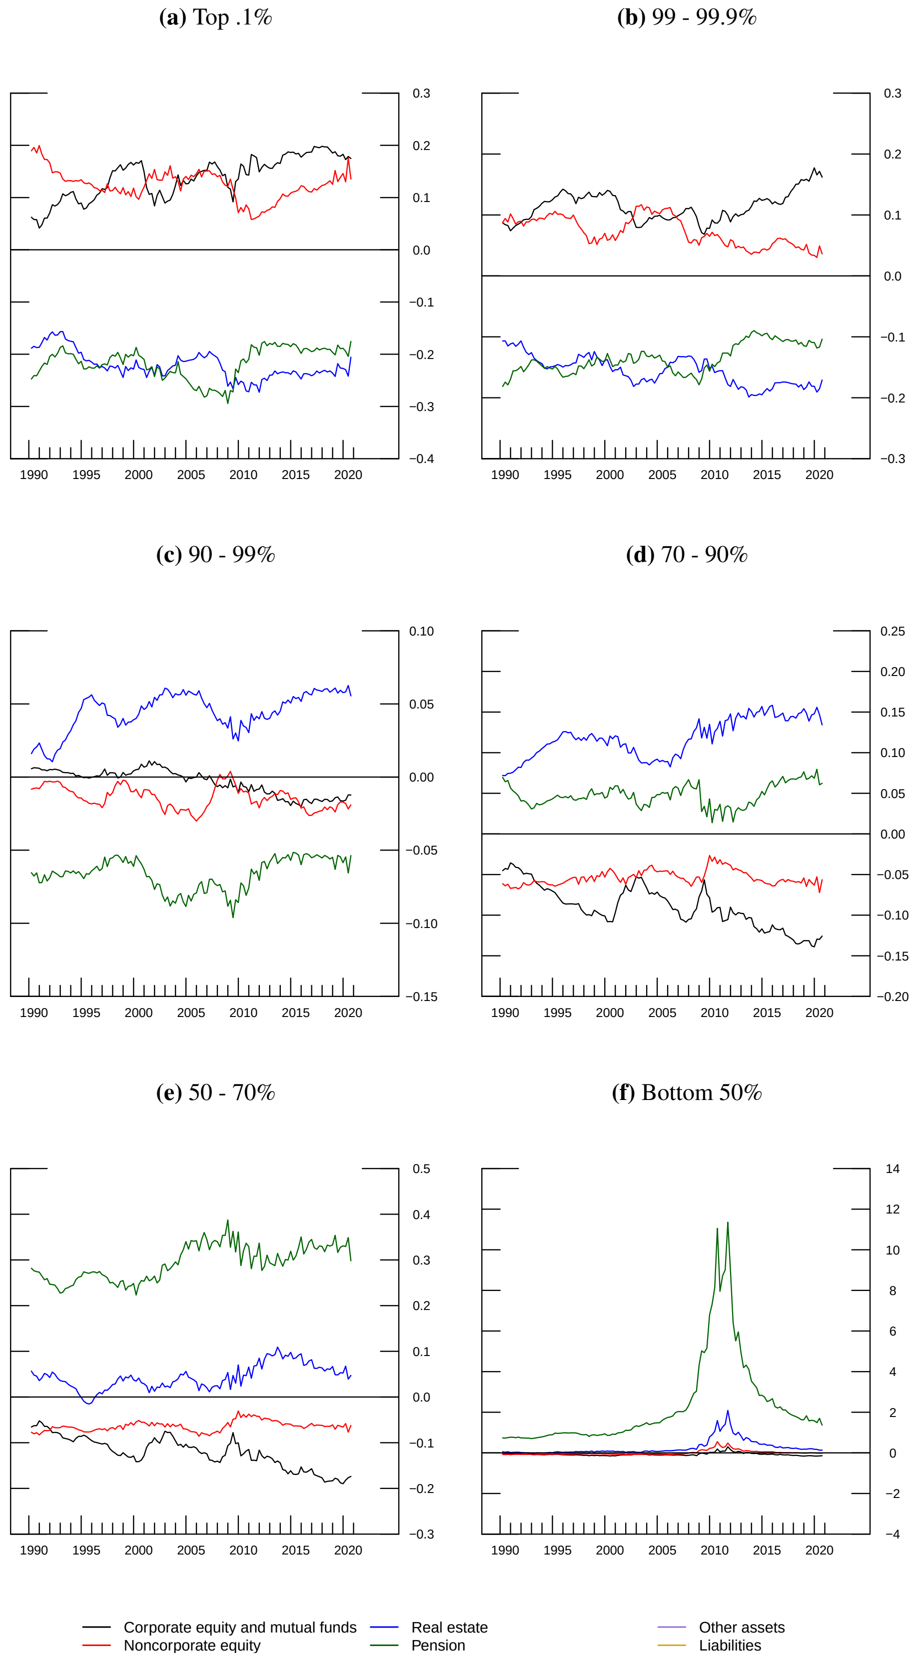

50-90th percentile households have a negative elasticity because $\alpha_{t,j}^{i} < \alpha_{t,j}$ for these households. Furthermore, these vary countercyclically because the gap between $\alpha_{t,j}^{i}$ and $\alpha_{t,j}$ decreases when equity prices fall during recessions, thus shrinking the second term in Equation 19.

As shown in Panel (d), variation in the top 1%'s elasticities are opposite those of the 50-90th percentile households. The top 1% has large negative elasticities with respect to pensions (a 1% price increase drops wealth shares by -.14 to -.22%) and real estate (-.13 to -.22%). The negative elasticity is driven by the top 1%'s underexposure to these assets relative to the aggregate households ($\alpha_{t,j}^{i} < \alpha_{t,j}$). Furthermore, the elasticity tends to vary countercyclically because the second term in Equation 19 becomes more negative during expansions due to the procyclicality of the top 1%'s overall wealth share (i.e., $\omega_{t}^{NW,p}/\omega_{t-1}^{NW,p}$ increases during recessions). The opposite holds for corporate and noncorporate business equity of the top 1%. Because for these assets $\alpha_{t,j}^{i} > \alpha_{t,j}$, this gap increases during expansions, and because $\omega_{t}^{NW,p}/\omega_{t-1}^{NW,p}$ is procyclical, these elastiticities are positive (.08-.16% and .06-.14%, respectively) and vary procyclically. For households in the 90-99th percentiles, the elasticity of wealth shares with respect to corporate equity and noncorporate equity are slightly negative but exhibit time variation similar to, but more muted than, the top 1%'s elasticities. Elasticities with respect to real estate are negative but also vary in a similar, but dampened manner, and the wealth elasticity with respect to pensions is slightly positive due to these households higher shares of wealth invested in pensions.

Our wealth share elasticities provide precise, quantitative measures of the wealth distributions exposure to asset price risk that confirm many of our earlier points. The large, positive elasticities of the top 1% with respect to corporate and noncorporate business equity, whose prices are very procyclical, suggest that these assets drive the procyclicality of the top 1%'s share. Furthermore, the effect of these price changes are amplified by the procyclicality of the elasticities themselves. In contrast, in the middle of the distribution wealth share elasticities with respect to pensions and real estate are much larger. While these elasticities do vary over time, they generally do so at a frequency distinct from business cycles (e.g., housing booms and busts).

D Decomposing Changes in Wealth Shares: Two Assets

Figure 22: Decomposition into aggregate price changes and group-specific changes with two-asset model price changes

(a) Top .1%

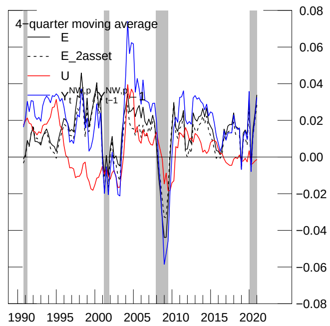

(b) 99 - 99.9%

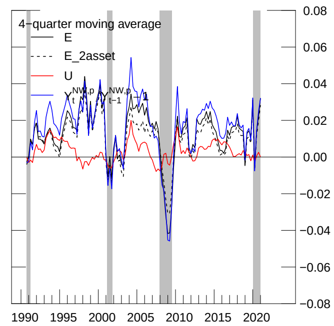

(c) 90 - 99%

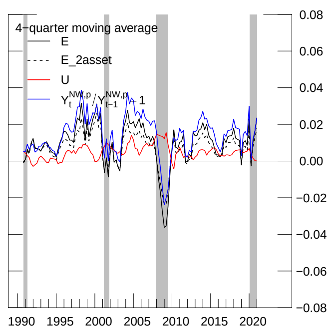

(d) 70 - 90%

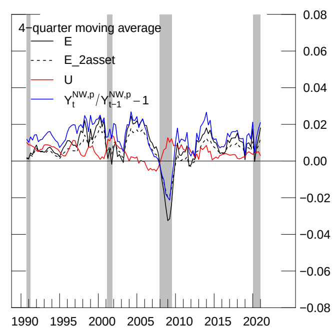

(e) 50 - 70%

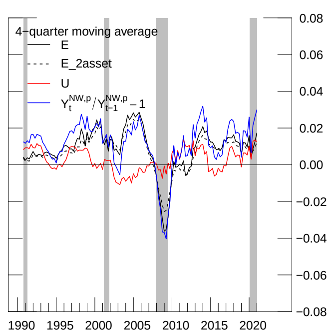

(f) Bottom 50%