The Rate of Return on Everything, 1870–2015[⋆]

Òscar Jordà[†] Katharina Knoll[‡] Dmitry Kuvshinov[§]
Moritz Schularick[¶] Alan M. Taylor[✠]

March 2019

Abstract

What is the aggregate real rate of return in the economy? Is it higher than the growth rate of the economy and, if so, by how much? Is there a tendency for returns to fall in the long-run? Which particular assets have the highest long-run returns? We answer these questions on the basis of a new and comprehensive dataset for all major asset classes, including housing. The annual data on total returns for equity, housing, bonds, and bills cover 16 advanced economies from 1870 to 2015, and our new evidence reveals many new findings and puzzles.

Keywords: return on capital, interest rates, yields, dividends, rents, capital gains, risk premiums, household wealth, housing markets.

JEL classification codes: D31, E44, E10, G10, G12, N10.

⋆This work is part of a larger project kindly supported by research grants from the Bundesministerium für Bildung und Forschung (BMBF) and the Institute for New Economic Thinking (INET). We are indebted to a large number of researchers who helped with data on individual countries. We are especially grateful to Francisco Amaral for outstanding research assistance, and would also like to thank Felix Rhiel, Mario Richarz, Thomas Schwarz, Lucie Stoppok, and Yevhenii Usenko for research assistance on large parts of the project. For their helpful comments we thank the editors and referees, along with Roger Farmer, John Fernald, Philipp Hofflin, David Le Bris, Clara Martínez-Toledano, Emi Nakamura, Thomas Piketty, Matthew Rognlie, Jón Steinsson, Johannes Stroebel, and Stijn Van Nieuwerburgh. We likewise thank conference participants at the NBER Summer Institute EFG Program Meeting, the Brevan Howard Centre for Financial Analysis at Imperial College Business School, the CEPR Housing Conference, and CEPR ESSIM at the Norges Bank, as well as seminar participants at Banca d'Italia, the Bank of England, Reserve Bank of New Zealand, Cornell University, New York University, University of Illinois at Urbana-Champaign, University of Chicago Booth School of Business, UC Berkeley, UCLA Anderson, Research Center SAFE, SciencesPo, and the Paris School of Economics. All errors are our own. The views expressed herein are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco, the Board of Governors of the Federal Reserve System, or the Deutsche Bundesbank.

†Federal Reserve Bank of San Francisco; and Department of Economics, University of California, Davis (oscar.jorda@sf.frb.org; ojorda@ucdavis.edu).

‡Deutsche Bundesbank (katharina.knoll@bundesbank.de).

§Department of Economics, University of Bonn (dmitry.kuvshinov@uni-bonn.de).

¶Department of Economics, University of Bonn; and CEPR (moritz.schularick@uni-bonn.de).

✠Department of Economics and Graduate School of Management, University of California, Davis; NBER; and CEPR (amtaylor@ucdavis.edu).

I. Introduction

What is the rate of return in an economy? It is a simple question, but hard to answer. The rate of return plays a central role in current debates on inequality, secular stagnation, risk premiums, and the decline in the natural rate of interest, to name a few. A main contribution of our paper is to introduce a large new dataset on the total rates of return for all major asset classes, including housing—the largest, but oft ignored component of household wealth. Our data cover most advanced economies—sixteen in all—starting in the year 1870.

Although housing wealth is on average roughly one-half of national wealth in a typical economy (Piketty, 2014), data on total housing returns (price appreciation plus rents) has been lacking (Shiller, 2000, provides some historical data on house prices, but not on rents). In this paper we build on more comprehensive work on house prices (Knoll, Schularick, and Steger, 2017) and newly constructed data on rents (Knoll, 2017) to enable us to track the total returns of the largest component of the national capital stock.

We further construct total returns broken down into investment income (yield) and capital gains (price changes) for four major asset classes, two of them typically seen as relatively risky—equities and housing—and two of them typically seen as relatively safe—government bonds and short-term bills. Importantly, we compute actual asset returns taken from market data and therefore construct more detailed series than returns inferred from wealth estimates in discrete benchmark years for a few countries as in Piketty (2014).

We also follow earlier work in documenting annual equity, bond, and bill returns, but here again we have taken the project further. We recompute all these measures from original sources, improve the links across some important historical market discontinuities (e.g., market closures and other gaps associated with wars and political instability), and in a number of cases we access new and previously unused raw data sources. In all cases, we have also brought in auxiliary sources to validate our data externally, and 100+ pages of online material documents our sources and methods. Our work thus provides researchers with the first broad noncommercial database of historical equity, bond, and bill returns—and the only database of housing returns—with the most extensive coverage across both countries and years.[1]

Our paper aims to bridge the gap between two related strands of the academic literature. The first strand is rooted in finance and is concerned with long-run returns on different assets. Dimson, Marsh, and Staunton (2009) probably marked the first comprehensive attempt to document and analyze long-run returns on investment for a broad cross-section of countries. Meanwhile, Homer and Sylla (2005) pioneered a multidecade project to document the history of interest rates.

The second related strand of literature is the analysis of comparative national balance sheets over time, as in Goldsmith (1985a). More recently, Piketty and Zucman (2014) have brought together data

[1]For example, our work complements and extends the database on equity returns by Dimson, Marsh, and Staunton (2009). Their dataset is commercially available, but has a shorter coverage and does not break down the yield and capital gain components.

from national accounts and other sources tracking the development of national wealth over long time periods. They also calculate rates of return on capital by dividing aggregate capital income in the national accounts by the aggregate value of capital, also from national accounts.

Our work is both complementary and supplementary to theirs. It is complementary as the asset price perspective and the national accounts approach are ultimately tied together by accounting rules and identities. Using market valuations, we are able to corroborate and improve the estimates of returns on capital that matter for wealth inequality dynamics. Our long-run return data are also supplementary to the work of Piketty and Zucman (2014) in the sense that we greatly extend the number of countries for which we can calculate real rates of return back to the late nineteenth century.

The new evidence we gathered can shed light on active research debates, reaching from asset pricing to inequality. For example, in one contentious area of research, the accumulation of capital, the expansion of capital's share in income, and the growth rate of the economy relative to the rate of return on capital all feature centrally in the current debate sparked by Piketty (2014) on the evolution of wealth, income, and inequality. What do the long-run patterns on the rates of return on different asset classes have to say about these possible drivers of inequality?

In many financial theories, preferences over current versus future consumption, attitudes toward risk, and covariation with consumption risk all show up in the premiums that the rates of return on risky assets carry over safe assets. Returns on different asset classes and their correlations with consumption sit at the core of the canonical consumption Euler equation that underpins textbook asset pricing theory (see, e.g., Mehra and Prescott, 1985). But tensions remain between theory and data, prompting further explorations of new asset pricing paradigms including behavioral finance. Our new data add another risky asset class to the mix, housing, and with it, new challenges.

In another strand of research triggered by the financial crisis, Summers (2014) seeks to revive the secular stagnation hypothesis first advanced in Alvin Hansen's (1939) AEA Presidential Address. Demographic trends are pushing the world's economies into uncharted territory as the relative weight of borrowers and savers is changing and with it the possibility increases that the interest rate will fall by an insufficient amount to balance saving and investment at full employment. What is the evidence that this is the case?

Lastly, in a related problem within the sphere of monetary economics, Holston, Laubach, and Williams (2017) show that estimates of the natural rate of interest in several advanced economies have gradually declined over the past four decades and are now near zero. What historical precedents are there for such low real rates that could inform today's policymakers, investors, and researchers?

The common thread running through each of these broad research topics is the notion that the rate of return is central to understanding long-, medium-, and short-run economic fluctuations. But which rate of return? And how do we measure it? For a given scarcity of funding supply, the risky rate is a measure of the profitability of private investment; in contrast, the safe rate plays an important role in benchmarking compensation for risk, and is often tied to discussions of monetary policy settings and the notion of the natural rate. Below, we summarize our main findings.

Main findings We present four main findings:

1. **On risky returns,** r^{risky}

   In terms of total returns, residential real estate and equities have shown very similar and high real total gains, on average about 7% per year. Housing outperformed equities before WW2. Since WW2, equities have outperformed housing on average, but had much higher volatility and higher synchronicity with the business cycle. The observation that housing returns are similar to equity returns, but much less volatile, is puzzling. Like Shiller (2000), we find that long-run capital gains on housing are relatively low, around 1% p.a. in real terms, and considerably lower than capital gains in the stock market. However, the rental yield component is typically considerably higher and more stable than the dividend yield of equities so that total returns are of comparable magnitude.

   Before WW2, the real returns on housing and equities (and safe assets) followed remarkably similar trajectories. After WW2 this was no longer the case, and across countries equities then experienced more frequent and correlated booms and busts. The low covariance of equity and housing returns reveals that there could be significant aggregate diversification gains (i.e., for a representative agent) from holding the two asset classes.

2. **On safe returns,** r^{safe}

   We find that the real safe asset return (bonds and bills) has been very volatile over the long-run, more so than one might expect, and oftentimes even more volatile than real risky returns. Each of the world wars was (unsurprisingly) a moment of very low safe rates, well below zero. So was the 1970s stagflation. The peaks in the real safe rate took place at the start of our sample, in the interwar period, and during the mid-1980s fight against inflation. In fact, the long decline observed in the past few decades is reminiscent of the secular decline that took place from 1870 to WW1. Viewed from a long-run perspective, the past decline and current low level of the real safe rate today is not unusual. The puzzle may well be why was the safe rate so high in the mid-1980s rather than why has it declined ever since.

   Safe returns have been low on average in the full sample, falling in the 1%–3% range for most countries and peacetime periods. While this combination of low returns and high volatility has offered a relatively poor risk-return trade-off to investors, the low returns have also eased the pressure on government finances, in particular allowing for a rapid debt reduction in the aftermath of WW2.

3. **On the risk premium,** $r^{risky} - r^{safe}$

   Over the very long run, the risk premium has been volatile. Our data uncover substantial swings in the risk premium at lower frequencies that sometimes endured for decades, and which far exceed the amplitudes of business-cycle swings.

In most peacetime eras, this premium has been stable at about 4%–5%. But risk premiums stayed curiously and persistently high from the 1950s to the 1970s, long after the conclusion of WW2. However, there is no visible long-run trend, and mean reversion appears strong. Interestingly, the bursts of the risk premium in the wartime and interwar years were mostly a phenomenon of collapsing safe returns rather than dramatic spikes in risky returns.

In fact, the risky return has often been smoother and more stable than the safe return, averaging about 6%–8% across all eras. Recently, with safe returns low and falling, the risk premium has widened due to a parallel but smaller decline in risky returns. But these shifts keep the two rates of return close to their normal historical range. Whether due to shifts in risk aversion or other phenomena, the fact that safe returns seem to absorb almost all of these adjustments seems like a puzzle in need of further exploration and explanation.

4. **On returns minus growth, $r^{wealth} - g$**

Piketty (2014) argued that, if investors' return to wealth exceeded the rate of economic growth, rentiers would accumulate wealth at a faster rate and thus worsen wealth inequality. Using a measure of portfolio returns to compute "r minus g" in Piketty's notation, we uncover an important finding. Even calculated from more granular asset price returns data, the same fact reported in Piketty (2014) holds true for more countries, more years, and more dramatically: namely "$r \gg g$."

In fact, the only exceptions to that rule happen in the years in or around wartime. In peacetime, r has always been much greater than g. In the pre-WW2 period, this gap was on average 5% (excluding WW1). As of today, this gap is still quite large, about 3%–4%, though it narrowed to 2% in the 1970s before widening in the years leading up to the Global Financial Crisis.

One puzzle that emerges from our analysis is that while "r minus g" fluctuates over time, it does not seem to do so systematically with the growth rate of the economy. This feature of the data poses a conundrum for the battling views of factor income, distribution, and substitution in the ongoing debate (Rognlie, 2015). The fact that returns to wealth have remained fairly high and stable while aggregate wealth increased rapidly since the 1970s, suggests that capital accumulation may have contributed to the decline in the labor share of income over the recent decades (Karabarbounis and Neiman, 2014). In thinking about inequality and several other characteristics of modern economies, the new data on the return to capital that we present here should spur further research.

II. A new historical global returns database

In this section, we will discuss the main sources and definitions for the calculation of long-run returns. A major innovation is the inclusion of housing. Residential real estate is the main asset in most household portfolios, as we shall see, but so far very little has been known about long-run

Table I: *Data coverage*

Country	Bills	Bonds	Equity	Housing
Australia	1870–2015	1900–2015	1870–2015	1901–2015
Belgium	1870–2015	1870–2015	1870–2015	1890–2015
Denmark	1875–2015	1870–2015	1873–2015	1876–2015
Finland	1870–2015	1870–2015	1896–2015	1920–2015
France	1870–2015	1870–2015	1870–2015	1871–2015
Germany	1870–2015	1870–2015	1870–2015	1871–2015
Italy	1870–2015	1870–2015	1870–2015	1928–2015
Japan	1876–2015	1881–2015	1886–2015	1931–2015
Netherlands	1870–2015	1870–2015	1900–2015	1871–2015
Norway	1870–2015	1870–2015	1881–2015	1871–2015
Portugal	1880–2015	1871–2015	1871–2015	1948–2015
Spain	1870–2015	1900–2015	1900–2015	1901–2015
Sweden	1870–2015	1871–2015	1871–2015	1883–2015
Switzerland	1870–2015	1900–2015	1900–2015	1902–2015
UK	1870–2015	1870–2015	1871–2015	1896–2015
USA	1870–2015	1871–2015	1872–2015	1891–2015

returns on housing. Our data on housing returns will cover capital gains, and imputed rents to owners and renters, the sum of the two being total returns.[2] Equity return data for publicly traded equities will then be used, as is standard, as a proxy for aggregate business equity returns.[3]

The data include nominal and real returns on bills, bonds, equities, and residential real estate for Australia, Belgium, Denmark, Finland, France, Germany, Italy, Japan, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, the United Kingdom, and the United States. The sample spans 1870 to 2015. Table I summarizes the data coverage by country and asset class.

Like most of the literature, we examine returns to national aggregate holdings of each asset class. Theoretically, these are the returns that would accrue for the hypothetical representative-agent investor holding each country's portfolio. An advantage of this approach is that it captures indirect holdings much better, although it leads to some double-counting thereby boosting the share of financial assets over housing somewhat. The differences are described in Online Appendix O.[4]

[2]Since the majority of housing is owner-occupied, and housing wealth is the largest asset class in the economy, owner-occupier returns and imputed rents also form the lion's share of the total return on housing, as well as the return on aggregate wealth.

[3]Moskowitz and Vissing-Jørgensen (2002) compare the returns on listed and unlisted U.S. equities over the period 1953–1999 and find that in aggregate, the returns on these two asset classes are similar and highly correlated, but private equity returns exhibit somewhat lower volatility. Moskowitz and Vissing-Jørgensen (2002) argue, however, that the risk-return tradeoff is worse for private compared to public equities, because aggregate data understate the true underlying volatility of private equity, and because private equity portfolios are typically much less diversified.

[4]Within-country heterogeneity is undoubtedly important, but clearly beyond the scope of a study covering nearly 150 years of data and 16 advanced economies.

II.A. The composition of wealth

Figure I shows the decomposition of economy-wide investable assets and capital stocks, based on data for five major economies at the end of 2015: France, Germany, Japan, UK and US.[5] Investable assets shown in the left panel of Figure I (and in Online Appendix Table A.23) exclude assets that relate to intra-financial holdings and cannot be held directly by investors, such as loans, derivatives (apart from employee stock options), financial institutions' deposits, and insurance and pension claims. Other financial assets mainly consist of corporate bonds and asset-backed securities. Other nonfinancial assets are other buildings, machinery and equipment, agricultural land, and intangible capital. The capital stock is business capital plus housing. Other capital is mostly made up of intangible capital and agricultural land. Data are sourced from national accounts and national wealth estimates published by the countries' central banks and statistical offices.[6]

Housing, equity, bonds, and bills comprise over half of all investable assets in the advanced economies today, and nearly two-thirds if deposits are included. The right-hand side panel of Figure I shows the decomposition of the capital stock into housing and various other nonfinancial assets. Housing is about one-half of the outstanding stock of capital. In fact, housing and equities alone represent over half of total assets in household balance sheets (see Online Appendix Figures A.5 and A.6).

The main asset categories *outside* the direct coverage of this study are: commercial real estate, business assets, and agricultural land; corporate bonds; pension and insurance claims; and deposits. But most of these assets represent claims of, or are closely related to, assets that we do cover. For example, pension claims tend to be invested in stocks and bonds; listed equity is a levered claim on business assets of firms; land and commercial property prices tend to comove with residential property prices; and deposit rates are either included in, or very similar to, our bill rate measure.[7]

Our data also exclude foreign assets. Even though the data on foreign asset holdings are relatively sparse, the evidence that we do have—presented in Online Appendix O.4—suggests that foreign assets have, through history, only accounted for a small share of aggregate wealth, and the return differentials between domestic and foreign asset holdings are, with few exceptions, not that large. Taken together, this means that our dataset almost fully captures the various components of the return on overall household wealth.

[5]Individual country data are shown in Online Appendix Tables A.23 and A.24.

[6]Both decompositions also exclude human capital, which cannot be bought or sold. Lustig, Van Nieuwerburgh, and Verdelhan (2013) show that for a broader measure of aggregate wealth that includes human capital, the size of human wealth is larger than of non-human wealth, and its return dynamics are similar to those of a long-term bond.

[7]Moreover, returns on commercial real estate are captured by the levered equity returns of the firms that own this real estate, and hence are indirectly proxied by our equity return data.

Figure I: *Composition of investable assets and capital stock in the major economies*

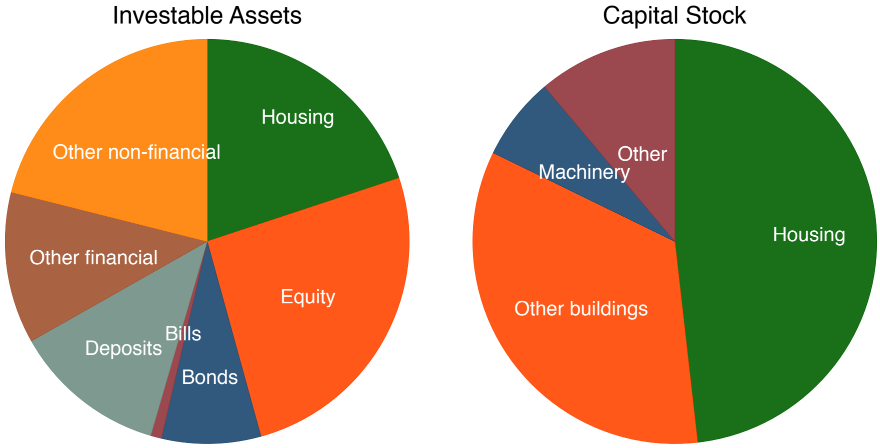

Note: Composition of total investable assets and capital stock. Average of the individual asset shares of France, Germany, Japan, UK, and US, as of end-2015. Investable assets are defined as the gross total of economy-wide assets excluding loans, derivatives, financial institutions' deposits, insurance, and pension claims. Other financial assets mainly consist of corporate bonds and asset-backed securities. Other nonfinancial assets are other buildings, machinery and equipment, agricultural land, and intangible capital. The capital stock is business capital plus housing. Other capital is mostly made up by intangible capital and agricultural land. Data are sourced from national accounts and national wealth estimates published by the countries' central banks and statistical offices.

II.B. Historical returns data

Bill returns The canonical risk-free rate is taken to be the yield on Treasury bills, i.e., short-term, fixed-income government securities. The yield data come from the latest vintage of the long-run macrohistory database (Jordà, Schularick, and Taylor, 2017).[8] Whenever data on Treasury bill returns were unavailable, we relied on either money market rates or deposit rates of banks from Zimmermann (2017). Since short-term government debt was rarely used and issued in the earlier historical period, much of our bill rate data before the 1960s actually consist of deposit rates.[9]

Bond returns These are conventionally the total returns on long-term government bonds. Unlike earlier cross-country studies, we focus on the bonds listed and traded on local exchanges and denominated in local currency. This focus makes bond returns more comparable with the returns of bills, equities, and housing. Moreover, this results in a larger sample of bonds, and on bonds that are more likely to be held by the representative household in the respective country. For some

[8] www.macrohistory.net/data

[9] In general, it is difficult to compute the total returns on deposits because of uncertainty about losses during banking crises, and we stick to the more easily measurable government bill and bond returns where these data are available. Comparisons with the deposit rate data in Zimmermann (2017), however, indicate that the interest rate differential between deposits and our bill series is very small, with deposit rates on average roughly 0.7 percentage points below bills—a return close to zero in real terms. The returns on government bills and deposits are also highly correlated over time.

countries and periods we have made use of listings on major global exchanges to fill gaps where domestic markets were thin, or local exchange data were not available (for example, Australian bonds listed in New York or London). Throughout the sample we target a maturity of around 10 years. For the second half of the 20th century, the maturity of government bonds is generally accurately defined. For the pre-WW2 period we sometimes had to rely on data for perpetuals, i.e., very long-term government securities (such as the British consol). Although as a convention we refer here to government bills and bonds as "safe" assets, both are naturally exposed to inflation and default risk, for example. In fact, real returns on these assets fluctuate substantially over time as we shall see (specifically, Sections V and VI).

Equity returns These returns come from a broad range of sources, including articles in economic and financial history journals, yearbooks of statistical offices and central banks, stock exchange listings, newspapers, and company reports. Throughout most of the sample, we aim to rely on indices weighted by market capitalization of individual stocks, and a stock selection that is representative of the entire stock market. For some historical time periods in individual countries, however, we also make use of indices weighted by company book capital, stock market transactions, or weighted equally, due to limited data availability.

Housing returns We combine the long-run house price series introduced by Knoll, Schularick, and Steger (2017) with a novel dataset on rents drawn from the unpublished PhD thesis of Knoll (2017). For most countries, the rent series rely on the rent components of the cost of living of consumer price indices constructed by national statistical offices. We then combine them with information from other sources to create long-run series reaching back to the late 19th century. To proxy the total return on the residential housing stock, our returns include both rented housing and owner-occupied properties.[10] Specifically, wherever possible we use house price and rental indices that include the prices of owner-occupied properties, and the imputed rents on these houses. Imputed rents estimate the rent that an owner-occupied house would earn on the rental market, typically by using rents of similar houses that are rented. This means that, in principle, imputed rents are similar to market rents, and are simply adjusted for the portfolio composition of owner-occupied as opposed to rented housing. Imputed rents, however, are not directly observed and hence are less precisely measured than market rents, and are typically not taxed.[11] To the best of our knowledge, we are the first to calculate total returns to housing in the literature for as long and as comprehensive a cross section of economies as we report.

Composite returns We compute the rate of return on safe assets, risky assets, and aggregate wealth, as weighted averages of the individual asset returns. To obtain a representative return from

[10]This is in line with the treatment of housing returns in the existing literature on returns to aggregate wealth—see, for example, Piketty et al. (2018) and Rognlie (2015).

[11]We discuss the issues around imputed rents measurement, and our rental yield series more generally in Section III.C.

the investor's perspective, we use the outstanding stocks of the respective asset in a given country as weights. To this end, we make use of new data on equity market capitalization (from Kuvshinov and Zimmermann, 2018) and housing wealth for each country and period in our sample, and combine them with existing estimates of public debt stocks to obtain the weights for the individual assets. A graphical representation of these asset portfolios, and further description of their construction is provided in the Online Appendix O.3. Tables A.28 and A.29 present an overview of our four asset return series by country, their main characteristics and coverage. The paper comes with an extensive data appendix that specifies the sources we consulted and discusses the construction of the series in greater detail (see the Online Data Appendix, Sections U, V, and W for housing returns, and Section X for equity and bond returns).

II.C. Calculating returns

The total annual return on any financial asset can be divided into two components: the capital gain from the change in the asset price P, and a yield component Y, that reflects the cash-flow return on an investment. The total nominal return R for asset j in country i at time t is calculated as:

$$\text{Total return:} \quad R^j_{i,t} = \frac{P^j_{i,t} - P^j_{i,t-1}}{P^j_{i,t-1}} + Y^j_{i,t}\,. \tag{1}$$

Because of wide differences in inflation across time and countries, it is helpful to compare returns in real terms. Let $\pi_{i,t} = (CPI_{i,t} - CPI_{i,t-1})/CPI_{i,t-1}$ be the realized consumer price index (CPI) inflation rate in a given country i and year t. We calculate inflation-adjusted *real returns* r for each asset class as,

$$\text{Real return:} \quad r^j_{i,t} = \frac{1 + R^j_{i,t}}{1 + \pi_{i,t}} - 1\,. \tag{2}$$

These returns will be summarized in period average form, by country, or for all countries.

Investors must be compensated for risk to invest in risky assets. A measure of this "excess return" can be calculated by comparing the real total return on the risky asset with the return on a risk-free benchmark—in our case, the government bill rate, $r^{bill}_{i,t}$. We therefore calculate the excess return ER for the risky asset j in country i as

$$\text{Excess return:} \quad ER^j_{i,t} = r^j_{i,t} - r^{bill}_{i,t}\,. \tag{3}$$

In addition to individual asset returns, we also present a number of weighted "composite" returns aimed at capturing broader trends in risky and safe investments, as well as the "overall return" or "return on wealth." Online Appendix O.3 provides further details on the estimates of country asset portfolios from which we derive country-year specific weights.

For safe assets, we assume that total public debt is divided equally into bonds and bills since there are no data on their market shares (only for total public debt) over our full sample. As a result, we compute the safe asset return as:

$$\text{Safe return:} \quad r_{i,t}^{safe} = \frac{r_{i,t}^{bill} + r_{i,t}^{bond}}{2}. \tag{4}$$

The risky asset return is calculated as a weighted average of the returns on equity and on housing. The weights w represent the share of asset holdings of equity and of housing stocks in the respective country i and year t, scaled to add up to 1. We use stock market capitalization and housing wealth to calculate each share and hence compute risky returns as:

$$\text{Risky return:} \quad r_{i,t}^{risky} = r_{i,t}^{equity} \times w_{i,t}^{equity} + r_{t}^{housing} \times w_{i,t}^{housing}. \tag{5}$$

The difference between our risky and safe return measures then provides a proxy for the aggregate risk premium in the economy:

$$\text{Risk premium:} \quad RP_{i,t} = r_{i,t}^{risky} - r_{i,t}^{safe}. \tag{6}$$

The "return on wealth" measure is a weighted average of returns on risky assets (equity and housing) and safe assets (bonds and bills). The weights w here are the asset holdings of risky and safe assets in the respective country i and year t, scaled to add to 1.[12]

$$\text{Return on wealth:} \quad r_{i,t}^{wealth} = r_{i,t}^{risky} \times w_{i,t}^{risky} + r_{i,t}^{safe} \times w_{i,t}^{safe}. \tag{7}$$

Finally, we also consider returns from a global investor perspective in Online Appendix I. There we measure the returns from investing in local markets in U.S. dollars (USD). These returns effectively subtract the depreciation of the local exchange rate vis-à-vis the dollar from the nominal return:

$$\text{USD return:} \quad R_{i,t}^{j,USD} = \frac{1 + R_{i,t}^{j}}{1 + \hat{s}_{i,t}} - 1, \tag{8}$$

where $\hat{s}_{i,t}$ is the rate of depreciation of the local currency versus the U.S. dollar in year t.

The real USD returns are then computed net of U.S. inflation $\pi_{US,t}$:

$$\text{Real USD return:} \quad r_{i,t}^{j,USD} = \frac{1 + R_{i,t}^{j,USD}}{1 + \pi_{US,t}} - 1. \tag{9}$$

II.D. Constructing housing returns using the rent-price approach

This section briefly describes our methodology to calculate total housing returns. We provide further details as needed later in Section III.C and in Online Appendix U. We construct estimates for total

[12]For comparison, Online Appendix P provides information on the equally-weighted risky return, and the equally-weighted rate of return on wealth, both calculated as simple averages of housing and equity, and housing, equity, and bonds respectively.

returns on housing using the rent-price approach. This approach starts from a benchmark rent-price ratio $\left(\frac{RI_0}{HPI_0}\right)$ estimated in a baseline year ($t = 0$). For this ratio we rely on net rental yields from the Investment Property Database (IPD).[13] We can then construct a time series of returns by combining separate information from a country-specific house price index series $\left(\frac{HPI_t}{HPI_0}\right)$ and a country-specific rent index series $\left(\frac{RI_t}{RI_0}\right)$. For these indices we rely on prior work on housing prices (Knoll, Schularick, and Steger, 2017) and new data on rents (Knoll, 2017). This method assumes that the indices cover a representative portfolio of houses. Under this assumption, there is no need to correct for changes in the housing stock, and only information about the growth rates in prices and rents is necessary.

Hence, a time series of the rent-price ratio can be derived from forward and back projection as

$$\frac{RI_t}{HPI_t} = \left[\left(\frac{RI_t}{RI_0}\right)\left(\frac{HPI_0}{HPI_t}\right)\right]\frac{RI_0}{HPI_0}. \tag{10}$$

In a second step, returns on housing can then be computed as:

$$R_{t+1}^{housing} = \frac{RI_{t+1}}{HPI_t} + \frac{HPI_{t+1} - HPI_t}{HPI_t}. \tag{11}$$

Our rent-price approach is sensitive to the choice of benchmark rent-price-ratios and cumulative errors from year-by-year extrapolation. We verify and adjust rent-price approach estimates using a range of alternative sources. The main source for comparison is the balance sheet approach to rental yields, which calculates the rent-price ratio using national accounts data on total rental income and housing wealth. The "balance sheet" rental yield RY_t^{BS} is calculated as the ratio of total net rental income to total housing wealth:

$$RY_t^{BS} = \frac{\text{Net rental income}_t}{\text{Housing Wealth}_t}, \tag{12}$$

This balance sheet rental yield estimate can then be added to the capital gains series in order to compute the total return on housing from the balance sheet perspective. We also collect additional point-in-time estimates of net rental yields from contemporary sources such as newspaper advertisements. These measures are less sensitive to the accumulated extrapolation errors in equation (10), but are themselves measured relatively imprecisely.[14] Wherever the rent-price approach estimates diverge from these historical sources, we make adjustments to benchmark the rent-price ratio estimates to these alternative historical measures of the rental yield. We also construct two additional housing return series—one benchmarked to all available alternative yield estimates, and

[13]These net rental yields use rental income net of maintenance costs, ground rent, and other irrecoverable expenditure. These adjustments are discussed exhaustively in the next section. We use net rather than gross yields to improve comparability with other asset classes.

[14]We discuss the advantages and disadvantages of these different approaches in Section III.C. Broadly speaking, the balance sheet approach can be imprecise due to measurement error in total imputed rent and national housing wealth estimates. Newspaper advertisements are geographically biased and only cover gross rental yields, so that the net rental yields have to be estimated.

another using primarily the balance sheet approach. The results of this exercise are discussed in Section III.C. Briefly, all the alternative estimates are close to one another, and the differences have little bearing on any of our results.

III. Rates of return: Aggregate trends

Our headline summary data appear in Table II and Figure II. The top panel of Table II shows the full sample (1870–2015) results whereas the bottom panel of the table shows results for the post-1950 sample. Note that here, and throughout the paper, rates of return are always annualized. Units are always expressed in percent per year, for raw data as well as for means and standard deviations. All means are arithmetic means, except when specifically referred to as geometric means.[15] Data are pooled and equally weighted, i.e., they are raw rather than portfolio returns. We will always include wars so that results are not polluted by bias from omitted disasters. We do, however, exclude hyperinflation years (but only a few) in order to focus on the underlying trends in returns, and to avoid biases from serious measurement errors in hyperinflation years, arising from the impossible retrospective task of matching within-year timing of asset and CPI price level readings which can create a spurious, massive under- or overstatement of returns in these episodes.[16]

The first key finding is that residential real estate, not equity, has been the best long-run investment over the course of modern history. Although returns on housing and equities are similar, the volatility of housing returns is substantially lower, as Table II shows. Returns on the two asset classes are in the same ballpark—around 7%—but the standard deviation of housing returns is substantially smaller than that of equities (10% for housing versus 22% for equities). Predictably, with thinner tails, the compounded return (using the geometric average) is vastly better for housing than for equities—6.6% for housing versus 4.7% for equities. This finding appears to contradict one of the basic tenets of modern valuation models: higher risks should come with higher rewards.

Differences in asset returns are not driven by unusual events in the early pre-WW2 part of the sample. The bottom panel of Table II makes this point. Compared to the full sample results in the top panel, the same clear pattern emerges: stocks and real estate dominate in terms of returns. Moreover, average returns post–1950 are similar to those for the full sample even though the postwar subperiod excludes the devastating effects of the two world wars. Robustness checks are reported in

[15]In what follows we focus on conventional average annual real returns. In addition, we often report period-average geometric mean returns corresponding to the annualized return that would be achieved through reinvestment or compounding. For any sample of years T, geometric mean returns are calculated as

$$\left(\prod_{t\in T}(1+r^j_{i,t})\right)^{\frac{1}{T}}-1.$$

Note that the arithmetic period-average return is always larger than the geometric period-average return, with the difference increasing with the volatility of the sequence of returns.

[16]Online Appendix G and Table A.12 do, however, provide some rough proxies for returns on different asset classes during hyperinflations.

Table II: *Global real returns*

	Real returns				Nominal Returns			
	Bills	Bonds	Equity	Housing	Bills	Bonds	Equity	Housing
Full sample:								
Mean return p.a.	1.03	2.53	6.88	7.06	4.58	6.06	10.65	11.00
Standard deviation	6.00	10.69	21.79	9.93	3.32	8.88	22.55	10.64
Geometric mean	0.83	1.97	4.66	6.62	4.53	5.71	8.49	10.53
Mean excess return p.a.	.	1.51	5.85	6.03				
Standard deviation	.	8.36	21.27	9.80				
Geometric mean	.	1.18	3.77	5.60				
Observations	1,767	1,767	1,767	1,767	1,767	1,767	1,767	1,767
Post-1950:								
Mean return p.a.	0.88	2.79	8.30	7.42	5.39	7.30	12.97	12.27
Standard deviation	3.42	9.94	24.21	8.87	4.03	9.81	25.03	10.14
Geometric mean	0.82	2.32	5.56	7.08	5.31	6.88	10.26	11.85
Mean excess return p.a.	.	1.91	7.42	6.54				
Standard deviation	.	9.21	23.78	9.17				
Geometric mean	.	1.51	4.79	6.18				
Observations	1,022	1,022	1,022	1,022	1,022	1,022	1,022	1,022

Note: Annual global returns in 16 countries, equally weighted. Period coverage differs across countries. Consistent coverage within countries: each country-year observation used to compute the statistics in this table has data for all four asset returns. Excess returns are computed relative to bills.

Figure II: *Global real rates of return*

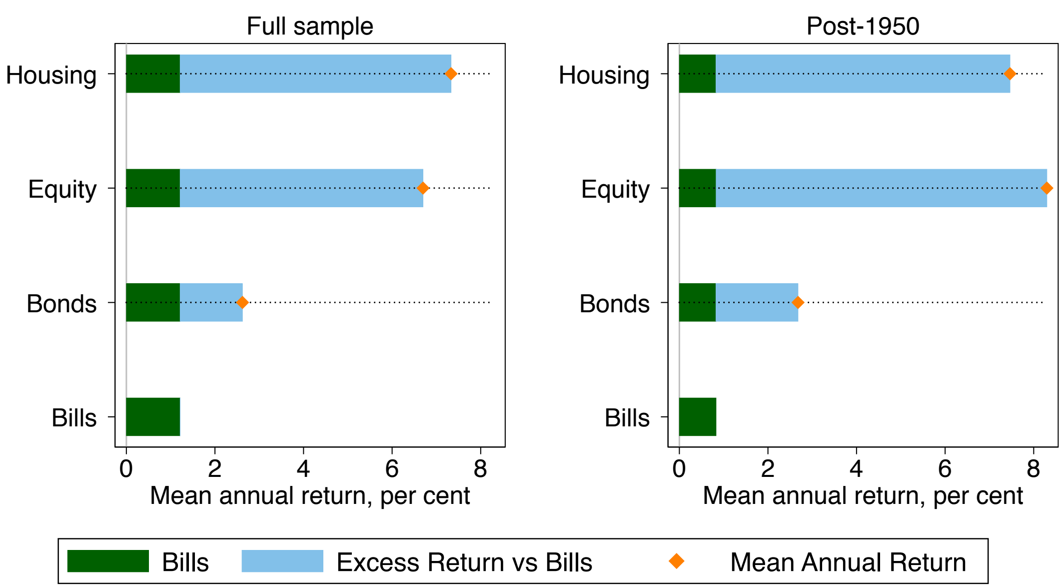

Notes: Arithmetic average real returns p.a., unweighted, 16 countries. Consistent coverage within each country: each country-year observation used to compute the average has data for all four asset returns.

Figures A.1, A.2, and A.3. Briefly, the observed patterns are not driven by the smaller European countries in our sample. Figure A.1 shows average real returns weighted by country-level real GDP, both for the full sample and post–1950 period. Compared to the unweighted averages, equity performs slightly better, but the returns on equity and housing remain very similar, and the returns and riskiness of all four asset classes are very close to the unweighted series in Table II.

The results could be biased due to the country composition of the sample at different dates given data availability. Online Appendix Figure A.2 plots the average returns for sample-consistent country groups, starting at benchmark years—the later the benchmark year, the more countries we can include. Again, the broad patterns discussed above are largely unaffected.

We also investigate whether the results are biased due to the world wars. Online Appendix Figure A.3 plots the average returns in this case. The main result remains largely unchanged. Online Appendix Table A.3 also considers the risky returns during wartime in more detail, to assess the evidence for rare disasters in our sample. Returns during both wars were indeed low and often negative, although returns during WW2 in a number of countries were relatively robust.

Finally, our aggregate return data take the perspective of a domestic investor in a representative country. Online Appendix Table A.14 instead takes the perspective of a global USD-investor, and assesses the USD value of the corresponding returns. The magnitude and ranking of returns are similar to those reported in Table II, although the volatilities are substantially higher. This is to be expected given that the underlying asset volatility is compounded by the volatility in the exchange rate. We also find somewhat higher levels of USD returns, compared to those in local currency.

What comes next in our discussion of raw rates of return? We will look more deeply at risky rates of return, and delve into their time trends and the decomposition of housing and equity returns into the capital gain and yield components in greater detail in Section IV. We will do the same for safe returns in Section V. But first, to justify our estimates, since these are new data, we have to spend considerable time to explain our sources, methods, and calculations. We next compare our data to other literature in Section III.A. We subject the equity returns and risk premium calculation to a variety of accuracy checks in Section III.B. We also subject the housing returns and risk premium calculation to a variety of accuracy checks in Section III.C. Section III.D then discusses the comparability of the housing and equity return series. For the purposes of our paper, these very lengthy next four subsections undertake the necessary due diligence and discuss the various quality and consistency checks we undertook to make our data a reliable source for future analysis—and only after that is done do we proceed with analysis and interpretation based on our data.

However, we caution that all these checks may be as exhausting as they are exhaustive and a time-constrained reader eager to confront our main findings may jump to the end of this section and resume where the analytical core of the paper begins at the start of Section IV on page 35.

III.A. Comparison to existing literature

Earlier work on asset returns has mainly focused on equities and the corresponding risk premium over safe assets (bills or bonds), starting with Shiller's analysis of historical US data (Shiller, 1981), later extended to cover post-1920 Sweden and the UK (Campbell, 1999), and other advanced economies back to 1900 (Dimson, Marsh, and Staunton, 2009), or back to 1870 (Barro and Ursúa, 2008). The general consensus in this literature is that equities earn a large premium over safe assets. The cross-country estimates of this premium vary between 7% in Barro and Ursúa (2008) and 6% in Dimson, Marsh, and Staunton (2009) using arithmetic means. Campbell (1999) documents a 4.7% geometric mean return premium instead.

We find a similarly high, though smaller, equity premium using our somewhat larger and more consistent historical dataset. Our estimate of the risk premium stands at 5.9% using arithmetic means, and 3.8% using geometric means (see Table II). This is lower than the estimates by Campbell (1999) and Barro and Ursúa (2008). The average risk premium is similar to that found by Dimson, Marsh, and Staunton (2009), but our returns tend to be slightly lower for the overlapping time period.[17] Details aside, our data do confirm the central finding of the literature on equity market returns: stocks earn a large premium over safe assets.

Studies on historical housing returns, starting with the seminal work of Robert Shiller (see Shiller, 2000, for a summary), have largely focused on capital gains. The rental yield component has received relatively little attention, and in many cases is missing entirely. Most of the literature pre-dating our work has therefore lacked the necessary data to calculate, infer, or discuss the total rates of return on housing over the long run. The few studies that take rents into account generally focus on the post-1970s US sample, and often on commercial real estate. Most existing evidence either places the real estate risk premium between equities and bonds, or finds that it is similar to that for equities (see Favilukis, Ludvigson, and Van Nieuwerburgh, 2017; Francis and Ibbotson, 2009; Ilmanen, 2011; Ruff, 2007). Some studies have even found that over the recent period, real estate seems to outperform equities in risk-adjusted terms (Cheng, Lin, and Liu, 2008; Shilling, 2003).

The stylized fact from the studies on long-run housing capital appreciation is that over long horizons, house prices only grow a little faster than the consumer price index. But again, note that this is *only* the capital gain component in (1). Low levels of real capital gains to housing was shown by Shiller (2000) for the US, and is also true, albeit to a lesser extent, in other countries, as documented in Knoll, Schularick, and Steger (2017). Our long-run average capital appreciation data for the US largely come from Shiller (2000), with two exceptions. For the 1930s, we use the more representative index of Fishback and Kollmann (2015) that documents a larger fall in prices during the Great Depression. From 1975 onwards, we use a Federal Housing Finance Agency index, which

[17] Our returns are substantially lower for France and Portugal (see Online Appendix Table A.18). These slightly lower returns are largely a result of more extensive consistency and accuracy checks that eliminate a number of upward biases in the series, and better coverage of economic disasters. We discuss these data quality issues further in Section III.B. Online Appendix L compares our equity return estimates with the existing literature on a country basis.

has a slightly broader coverage. Online Appendix M compares our series with Shiller's data and finds that switching to Shiller's index has no effect on our results for the US. See also the Online Appendix of Knoll et al. (2017) for further discussion.

However, our paper turns this notion of low housing returns on its head—because we show that including the yield component in (1) in the housing return calculation generates a housing risk premium roughly as large as the equity risk premium. Prior to our work on historical rental yields this finding could not be known. Coincidentally, in our long-run data we find that most of the real equity return also comes from the dividend yield rather than from real equity capital gains which are low, especially before the 1980s. Thus the post-1980 observation of large capital gain components in both equity and housing total returns is completely unrepresentative of the normal long-run patterns in the data, another fact underappreciated before now.

Data on historical returns for all major asset classes allow us to compute the return on aggregate wealth (see equation 7). In turn, this return can be decomposed into various components by asset class, and into capital gains and yields, to better understand the drivers of aggregate wealth fluctuations. This connects our study to the literature on capital income, and the stock of capital (or wealth) from a national accounts perspective. Even though national accounts and national balance sheet estimates have existed for some time (see Goldsmith, 1985b; Kuznets, 1941), it is only recently that scholars have systematized and compared these data to calculate a series of returns on national wealth.[18]

Piketty, Saez, and Zucman (2018) compute balance sheet returns on aggregate wealth and for individual asset classes using post-1913 US data. Balance sheet return data outside the US are sparse, although Piketty and Zucman (2014) provide historical estimates at benchmark years for three more countries, and, after 1970, continuous data for an additional five countries. Online Appendix R compares our market-based return estimates for the US with those of Piketty et al. (2018). Housing returns are very similar. However, equity returns are several percentage points above our estimates and those in the market-based returns literature more generally. Part of this difference reflects the fact that balance sheet returns are computed to measure income before corporate taxes, whereas our returns take the household perspective and are therefore net of corporate tax. Another explanation for the difference is likely to come from some measurement error in the national accounts data.[19] When it comes to housing, our rental yield estimates are broadly comparable and similar to those derived using the balance sheet approach, for a broad selection of countries and historical time

[18] The return on an asset from a national accounts perspective, or the "balance sheet approach" to returns, r_t^{BS} is the sum of the yield, which is capital income (such as rents or profits) in relation to wealth, and capital gain, which is the change in wealth not attributable to investment. See Online Appendix R and equation (13) for further details.

[19] See Online Appendix R for more detail. In brief, the main conceptual difference between the two sets of estimates, once our returns are grossed up for corporate tax, is the inclusion of returns on unlisted equities in the national accounts data. But existing evidence suggests that these return differentials are not large (Moskowitz and Vissing-Jørgensen, 2002), and the size of the unlisted sector not sufficiently large to place a large weight of this explanation, which leads us to conjecture that there is some measurement error in the national income and wealth estimates that is driving the remaining difference.

spans.[20]

Our dataset complements the market-based return literature by increasing the coverage in terms of assets, return components, countries, and years; improving data consistency and documentation; and making the dataset freely available for future research. This comprehensive coverage can also help connect the market-based return estimates to those centered around national accounts concepts. We hope that eventually, this can improve the consistency and quality of both market-based returns and national accounts data.

III.B. Accuracy of equity returns

The literature on equity returns has highlighted two main sources of bias in the data: weighting and sample selection. Weighting biases arise when the stock portfolio weights for the index do not correspond with those of a representative investor, or a representative agent in the economy. Selection biases arise when the selection of stocks does not correspond to the portfolio of the representative investor or agent. This second category also includes issues of survivorship bias and missing data bias arising from stock exchange closures and restrictions.

We consider how each of these biases affect our equity return estimates in this section. An accompanying Online Appendix Table A.29 summarizes the construction of the equity index for each country and time period, with further details provided in Online Appendix X.

Weighting bias The best practice when weighting equity indices is to use market capitalization of individual stocks. This approach most closely mirrors the composition of a hypothetical representative investor's portfolio. Equally weighted indices are likely to overweight smaller firms, which tend to carry higher returns and higher volatility.

The existing evidence from historical returns on the Brussels and Paris stock exchanges suggests that using equally weighted indices biases returns up by around 0.5 percentage points, and their standard deviation up by 2–3 percentage points (Annaert, Buelens, Cuyvers, De Ceuster, Deloof, and De Schepper, 2011; Le Bris and Hautcoeur, 2010). The size of the bias, however, is likely to vary across across markets and time periods. For example, Grossman (2017) shows that the market-weighted portfolio of UK stocks outperformed its equally weighted counterpart over the period 1869–1929.

To minimize this bias, we use market-capitalization-weighted indices for the vast majority of our sample (see Online Appendix Table A.29 and Online Appendix X). Where market-capitalization weighting was not available, we have generally used alternative weights such as book capital or transaction volumes, rather than equally weighted averages. For the few equally weighted indices that remain in our sample, the overall impact on aggregate return estimates ought to be negligible.

[20] See Section III.C and Online Appendix U for more detail.

Selection and survivorship bias Relying on an index whose selection does not mirror the representative investor's portfolio carries two main dangers. First, a small sample may be unrepresentative of overall stock market returns. And second, a sample that is selected ad hoc, and especially ex post, is likely to focus on surviving firms, or successful firms, thus overstating investment returns. This second bias extends not only to stock prices but also to dividend payments, as some historical studies only consider dividend-paying firms.[21] The magnitude of survivorship bias has generally been found to be around 0.5 to 1 percentage points (Annaert, Buelens, and De Ceuster, 2012; Nielsen and Risager, 2001), but in some time periods and markets it could be larger (see Le Bris and Hautcoeur, 2010, for France).

As a first best, we always strive to use all-share indices that avoid survivor and selection biases. For some countries and time periods where no such indices were previously available, we have constructed new weighted all-share indices from original historical sources (e.g., early historical data for Norway and Spain). Where an all-share index was not available or newly constructed, we have generally relied on "blue-chip" stock market indices. These are based on an ex ante value-weighted sample of the largest firms on the market. It is updated each year and tends to capture the lion's share of total market capitalization. Because the sample is selected ex ante, it avoids ex post selection and survivorship biases. And because historical equity markets have tended to be quite concentrated, "blue-chip" indices have been shown to be a good proxy for all-share returns (see Annaert, Buelens, Cuyvers, De Ceuster, Deloof, and De Schepper, 2011). Finally, we include non-dividend-paying firms in the dividend yield calculation.

Stock market closures and trading restrictions A more subtle form of selection bias arises when the stock market is closed and no market price data are available. One way of dealing with closures is to simply exclude them from the baseline return comparisons. But this implicitly assumes that the data are "missing at random"—i.e., that stock market closures are unrelated to underlying equity returns. Existing research on rare disasters and equity premiums shows that this is unlikely to be true (Nakamura, Steinsson, Barro, and Ursúa, 2013). Stock markets tend to be closed precisely at times when we would expect returns to be low, such as periods of war and civil unrest. Return estimates that exclude such rare disasters from the data will thus overstate stock returns.

To guard against this bias, we include return estimates for the periods of stock market closure in our sample. Where possible, we rely on alternative data sources to fill the gap, such as listings of other exchanges and over-the-counter transactions—for example, in the case of WW1 Germany we use the over-the-counter index from Ronge (2002) and for WW2 France we use the newspaper index from Le Bris and Hautcoeur (2010). In cases where alternative data are not available, we interpolate the prices of securities listed both before and after the exchange was closed to estimate the return (if

[21] As highlighted by Brailsford, Handley, and Maheswaran (2012), this was the case with early Australian data, and the index we use scales down the series for dividend-paying firms to proxy the dividends paid by all firms, as suggested by these authors.

Table III: *Geometric annual average and cumulative total equity returns in periods of stock market closure*

Episode	Real returns		Nominal returns		Real capitalization	
	Geometric average	Cumulative	Geometric average	Cumulative	Geometric average	Cumulative
Spanish Civil War, 1936–40	-4.01	-15.09	9.03	41.32	-10.22	-35.04
Portuguese Revolution, 1974–77	-54.98	-90.88	-44.23	-82.65	-75.29	-98.49
Germany WW1, 1914–18	-21.67	-62.35	3.49	14.72		
Switzerland WW1, 1914–16	-7.53	-14.50	-0.84	-1.67	-8.54	-16.34
Netherlands WW2, 1944–46	-12.77	-20.39	-5.09	-8.36		

Note: Cumulative and geometric average returns during periods of stock market closure. Estimated by interpolating returns of shares listed both before and after the exchange was closed. The change in market capitalization compares the capitalization of all firms before the market was closed, and once it was opened, and thus includes the effect of any new listings, delistings and bankruptcies that occured during the closure.

no dividend data are available, we also assume no dividends were paid).[22]

Even though this approach only gives us a rough proxy of returns, it is certainly better than excluding these periods, which effectively assumes that the return during stock market closures is the same as that when the stock markets are open. In the end, we only have one instance of stock market closure for which we are unable to estimate returns—that of the Tokyo stock exchange in 1946–1947. Online Appendix H further assesses the impact of return interpolation on the key moments of our data and finds that, over the full sample, it is negligible.

Table III shows the estimated stock returns during the periods of stock exchange closure in our sample. The first two columns show average and cumulative real returns, and the third and fourth columns show the nominal returns. Aside from the case of WW1 Germany, returns are calculated by comparing the prices of shares listed both before and after the market closure. Such a calculation may, however, overstate returns because it selects only those companies that "survived" the closure. As an additional check, the last two columns of Table III show the inflation-adjusted change in market capitalization of stocks before and after the exchange was closed. This serves as a lower bound for investor returns because it would be as if we assumed that all delisted stocks went bankrupt (i.e., to a zero price) during the market closure.

Indeed, the hypothetical investor returns during the periods of market closure are substantially below market averages. In line with Nakamura, Steinsson, Barro, and Ursúa (2013), we label these periods as "rare disasters." The average per-year geometric mean return ranges from a modestly negative –4% p.a. during the Spanish Civil War, to losses of roughly 55% p.a. during the three years after the Portuguese Carnation Revolution. Accounting for returns of delisted firms is likely to bring these estimates down even further, as evinced by the virtual disappearance of the Portuguese stock market in the aftermath of the revolution.

[22] For example, the Swiss stock exchange was closed between July 1914 and July 1916. Our data for 1914 capture the December 1913–July 1914 return, for 1915 the July 1914–July 1916 return, and for 1916 the July 1916–December 1916 return. For the Spanish Civil War, we take the prices of securities in end-1936 and end-1940, and apportion the price change in between equally to years 1937–1939.

Having said this, the impact of these rare events on the average cross-country returns (shown in Table II) is small, around –0.1 percentage points, precisely because protracted stock market closures are very infrequent. The impact on country-level average returns is sizeable for Portugal and Germany (around –1 percentage point), but small for the other countries (–0.1 to –0.4 percentage points). Online Appendix G provides a more detailed analysis of returns during consumption disasters. On average, equity returns during these times are low, with an average cumulative real equity return drop of 6.7% during the disaster years.

Lastly, Nakamura, Steinsson, Barro, and Ursúa (2013) also highlight a more subtle bias arising from asset price controls. This generally involves measures by the government to directly control transaction prices, as in Germany during 1943–47, or to influence the funds invested in the domestic stock market (and hence the prices) via controls on spending and investment, as in France during WW2 (Le Bris, 2012). These measures are more likely to affect the timing of returns rather than their long-run average level, and should thus have little impact on our headline estimates. For example, Germany experienced negative nominal and real returns despite the WW2 stock price controls; and even though the policies it enacted in occupied France succeeded in generating high nominal stock returns, the real return on French stocks during 1940–44 was close to zero. Both of these instances were also followed by sharp drops in stock prices when the controls were lifted.[23]

III.C. Accuracy of housing returns

The biases that affect equity returns—weighting and selection—can also apply to returns on housing. There are also other biases that are specific to housing return estimates. These include costs of running a housing investment, and the benchmarking of rent-price ratios to construct the historical rental yield series. We discuss each of these problematic issues in turn in this section. Our focus throughout is mainly on rental yield data, as the accuracy and robustness of the house price series has been extensively discussed in Knoll, Schularick, and Steger (2017) in their online appendix.

Maintenance costs Any homeowner incurs costs for maintenance and repairs which lower the rental yield and thus the effective return on housing. We deal with this issue by the choice of the benchmark rent-price ratios. Specifically, we anchor to the Investment Property Database (IPD) whose rental yields reflect net income—net of property management costs, ground rent, and other irrecoverable expenditure—as a percentage of the capital employed. The rental yields calculated using the rent-price approach detailed in Section II.D are therefore net yields. To enable a like-for-like comparison, our historical benchmark yields are calculated net of estimated running costs and depreciation. Running costs are broadly defined as housing-related expenses excluding interest, taxes, and utilities—i.e., maintenance costs, management, and insurance fees.

Applying the rent-price approach to net yield benchmarks assumes that running costs remain

[23]The losses in the German case are difficult to ascertain precisely because the lifting of controls was followed by a redenomination that imposed a 90% haircut on all shares.

Figure III: *Costs of running a housing investment*

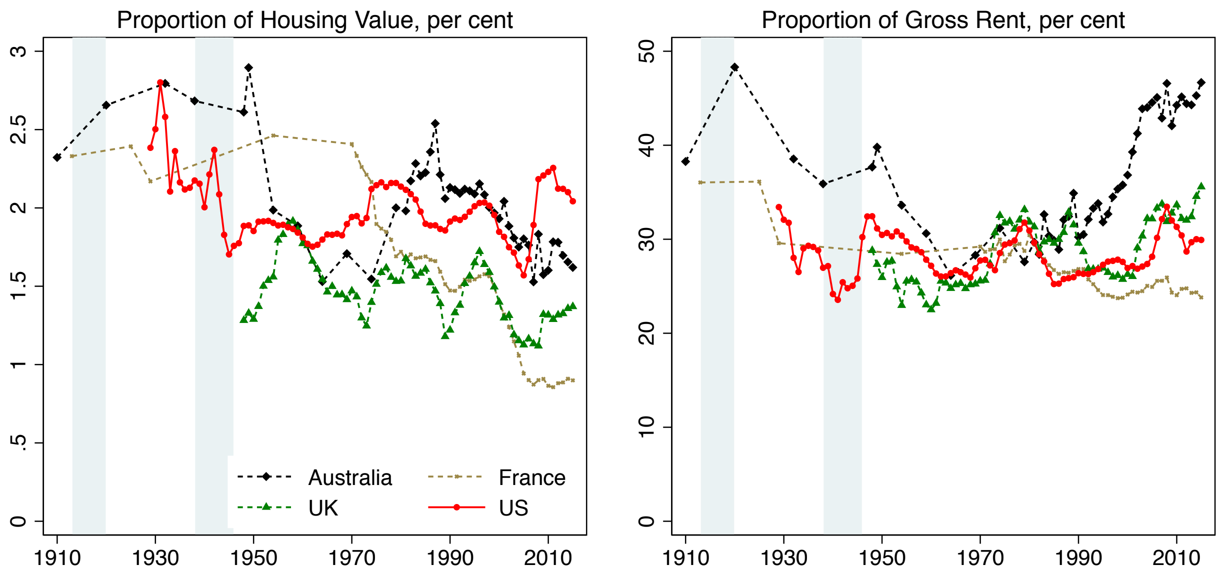

Note: Total costs include depreciation and all other housing-related expenses excluding interest, taxes and utilities (mainly maintenance and insurance payments). Costs are estimated as the household consumption of the relevant intermediate housing input, or fixed housing capital, in proportion to total housing wealth (left panel), or total gross rent (right panel).

stable relative to gross rental income over time within each country. To check this, Figure III presents historical estimates of running costs and depreciation for Australia, France, UK, and US, calculated as the sum of the corresponding housing expenditures and fixed capital consumption items in the national accounts. The left-hand panel presents these as a proportion of total housing value, and the right-hand panel as a proportion of gross rent. Relative to housing value, costs have been stable over the last 40 years, but were somewhat higher in the early to mid-twentieth century. This is to be expected since these costs are largely related to structures, not land, and structures constituted a greater share of housing value in the early twentieth century (Knoll, Schularick, and Steger, 2017). Additionally, structures themselves may have been of poorer quality in past times. When taken as a proportion of gross rent, however, as shown in the right-hand panel of Figure III, housing costs have been relatively stable, or at least not higher historically than they are today. This is likely because both gross yields and costs are low today, whereas historically both yields and costs were higher, with the two effects more or less cancelling out. This suggests that the historical rental yields that we have calculated using the rent-price approach are a good proxy for net yields.

Rental yield benchmarking To construct historical rental yield series using the rent-price approach, we start with a benchmark rent-price ratio from the Investment Property Database (IPD), and extend the series back using the historical rent and house price indices (see Section II.D).[24] This naturally implies that the level of returns is sensitive to the choice of the benchmark ratio. Moreover, past errors in rent and house price indices can potentially accumulate over time and may cause one to substantially over- or understate historical rental yields and housing returns. If the historical capital gains are overstated, the historical rental yields will be overstated too.

To try to avert such problems, we corroborate our rental yield estimates using a wide range of alternative historical and current-day sources. The main source of these independent comparisons comes from estimates using the balance sheet approach and national accounts data. As shown in equation 12, the "balance sheet" rental yield is the ratio of nationwide net rental income to total housing wealth. Net rental income is computed as gross rents paid less depreciation, maintenance and other housing-related expenses (excluding taxes and interest), with all data taken from the national accounts. The balance sheet approach gives us a rich set of alternative rental yield estimates both for the present day and even going back in time to the beginning of our sample in a number of countries. The second source for historical comparisons comes from advertisements in contemporary newspapers and various other contemporary publications. Third, we also make use of alternative current-day benchmarks based on transaction-level market rent data, and the rental expenditure and house price data from `numbeo.com`.[25] For all these measures, we adjust gross yields down to obtain a proxy for net rental yields.

Historical sources offer point-in-time estimates which avoid the cumulation of errors, but can nevertheless be imprecise. The balance sheet approach relies on housing wealth and rental income data, both of which are subject to potential measurement error. For housing wealth, it is inherently difficult to measure the precise value of all dwellings in the economy. Rental income is largely driven by the imputed rents of homeowners, which have to be estimated from market rents by matching the market rent to owner-occupied properties based on various property characteristics. This procedure can suffer from errors both in the survey data on property characteristics and market rents, and the matching algorithm.[26] Newspaper advertisements are tied to a specific location, and often biased towards cities. And transaction-level or survey data sometimes only cover the rental sector, rather than both renters and homeowners.

Given the potential measurement error in all the series, our final rental yield series uses data from both the rent-price approach and the alternative benchmarks listed above. More precisely, we use the following method to construct our final "best-practice" rental yield series. If the rent-price

---

[24] For Australia and Belgium, we instead rely on yield estimates from transaction-level data (Fox and Tulip (2014) and `Numbeo.com`, which are more in line with current-day and alternative historical estimates than IPD.

[25] The high-quality transaction level data are available for Australia and the US, from Fox and Tulip (2014) (sourced from RP Data) and Giglio, Maggiori, and Stroebel (2015) (sourced from Trulia) respectively. We use the Fox and Tulip (2014) yield as the benchmark for Australia. For the US, we use IPD because it is in line with several alternative estimates, unlike Trulia data which are much higher. See Online Appendix U for further details.

[26] For example, in the UK a change to imputation procedures in 2016 and the use of new survey data resulted in historical revisions which almost tripled imputed rents (see Office for National Statistics, 2016). We use a mixture of the old and new/revised data for our historical estimates.

Figure IV: *Comparison of the rent-price and balance-sheet approaches for historical rental yields*

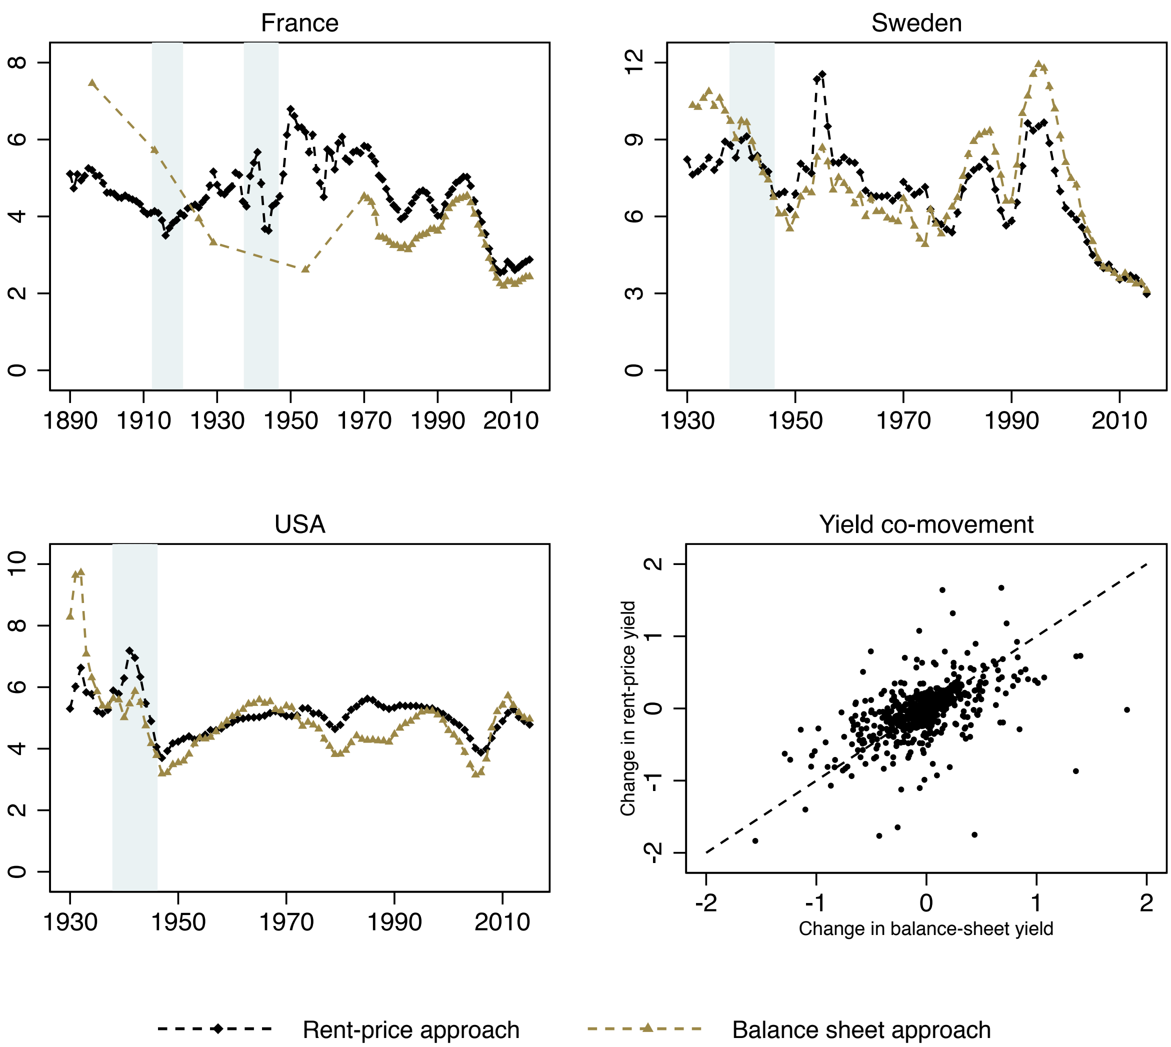

Note: The rent-price approach uses the baseline estimates in this paper. The balance sheet approach estimates the net yield in each year as total rental expenditure less housing running costs and depreciation, in proportion to total housing wealth.

Table IV: *Impact of using different rental yield benchmarks*

	Equity	Housing				
		Baseline	Low initial benchmark	High intial benchmark	Historical bench-marks	Balance sheet approach
Mean return p.a.	6.88	7.06	6.29	7.89	6.83	6.30
Standard deviation	21.79	9.93	9.89	10.03	9.93	9.95
Geometric mean	4.66	6.62	5.85	7.45	6.39	5.86
Observations	1,767	1,767	1,767	1,767	1,767	1,767

Note: Average total real returns across 16 countries, equally weighted.

approach estimates are close to alternative measures, we keep the rent-price approach data. This is the case for most historical periods in our sample. If there is a persistent level difference between the rent-price approach and alternative estimates, we adjust the benchmark yield to better match the historical and current data across the range of sources. This is the case for Australia and Belgium. If the levels are close for recent data but diverge historically, we adjust the historical estimates to match the alternative benchmarks. For most countries such adjustments are small or only apply to a short time span, but for Finland and Spain they are more substantial. Online Appendix U details the alternative sources and rental yield construction, including any such adjustments, for each country.

How large is the room for error in our final housing return series? To get a sense of the differences, Figure IV compares the rent-price approach of net rental yield estimates (black diamonds) with those using the balance sheet approach (brown triangles). The first three panels show the time series of the two measures for France, Sweden, and US, and the bottom-right panel shows the correlation between changes in rent-price and balance sheet yields in nine countries (Australia, Denmark, France, Germany, Italy, Japan, Sweden, UK, and US).[27] The level of the rent-price ratio using the two approaches is similar, both in the modern day and historically.[28] The two yield measures also follow a very similar time series pattern, both in the three countries depicted in panels 1–3, and the broader sample of countries summarized in the bottom-right panel.

Table IV provides a more comprehensive comparison. Columns 1 and 2 present the arithmetic and geometric mean, and the standard deviation, for the baseline measures of equity and housing annual real total returns in our sample (also shown in Table II). Column 3 instead uses the lowest possible initial benchmark for the housing series.[29] The resulting returns are around 0.8 percentage

[27]We limit our analysis to countries where the balance sheet approach data goes back at least several decades.

[28]For France, the historical data disagree somewhat, with balance sheet approach estimates both above and below the rent-price approach for some years. We further confirm the housing return series for France using returns on housing investment trusts, documented in the subsequent sections.

[29]For example, the balance sheet approach yield in 2013 Danish data is lower than the IPD yield; hence column 3 uses the 2013 balance sheet yield as the initial benchmark. For countries where we benchmark to historical rental yields, we use the same historical benchmark for all three series. For example, for Australia, we use a historical benchmark yield in 1949. So the "high" housing return series uses the high rental yield

points (henceforth, pps) lower, in both arithmetic and geometric mean terms. Column 4 instead uses the highest available benchmark, thus raising housing returns by 0.8 pps. Column 5 uses historical benchmarks for all rental yield series before 1980, i.e., we use these benchmarks as the main source for the yields, and only use the rent-price approach for interpolation.[30] This makes very little difference to the returns, lowering them by around 0.2 pps. The last column instead uses the balance sheet approach as the baseline estimate, both for the current and historical period. It then uses the rent-price approach to fill the gaps and interpolate between the balance sheet estimates.[31] Finally, we compute the total balance sheet return on housing as the sum of capital gains and the balance sheet yield.[32] The resulting return is 0.8 pps lower than our baseline estimates.

Taken together, this analysis suggests that the potential margins for error are small. Even under the more stringent checks, housing returns remain within a percentage point of our baseline estimates. The average return is always similar to equities in arithmetic mean terms, and always above equities when using the geometric mean.

Geographic coverage and selection biases Our data aim to approximate the return on a representative agent's housing portfolio. Selection bias means that the selection of properties in our dataset does not mirror the balance sheet of the representative agent. The main reason for this bias is selective geographical coverage. Housing returns can vary a lot by location, and our data are based on a sample of housing transactions.

To make our samples as representative as possible, we strive to attain a broad geographic coverage for both house price and rental data. Knoll, Schularick, and Steger (2017) discuss the potential location biases in house price data, but find that the house price trends in their, and hence our, dataset should be representative of country-level trends. When it comes to rents, the benchmark IPD yields are based on portfolios of institutional investors, which are slightly biased towards cities. This would lead to lower yields than the national average. On the other hand, investors may select higher-yielding properties within any given city. Comparisons with aggregate balance sheet approach data and alternative estimates indicate that, on average, IPD yields tend to be representative at country level. Further, IPD yields are capitalization weighted, which again better captures the yield on a representative portfolio. Finally, we aim for national coverage with the historical rental indices used for extrapolation, and historical balance sheet benchmarks.

Despite this, it is likely that both our house price and rental data are somewhat biased towards cities and urban areas, especially for historical periods—simply because urban housing data are more widely available and researched. Even though this would affect the level of capital gain and

benchmark for 1950–2015, and the historical benchmark for 1900–1949.

[30] For example, the series for Denmark is benchmarked to the lower balance sheet approach yield estimates for 1890–1910 and 1950–1970, and newspaper estimates for 1920–1940 (also see Online Appendix Figure A.14).

[31] Newspaper yield estimates are used as additional benchmarks for interpolation.

[32] This means that we use market-based house price data for capital gains, which is also common practice in the balance sheet approach computation, due to the large potential for error when estimating housing capital gains as a residual between wealth changes and investment. Piketty et al. (2018) use Shiller (2000) house price data for the balance sheet return computation.

Figure V: *Sensitivity of housing returns to a rent-price location correction*

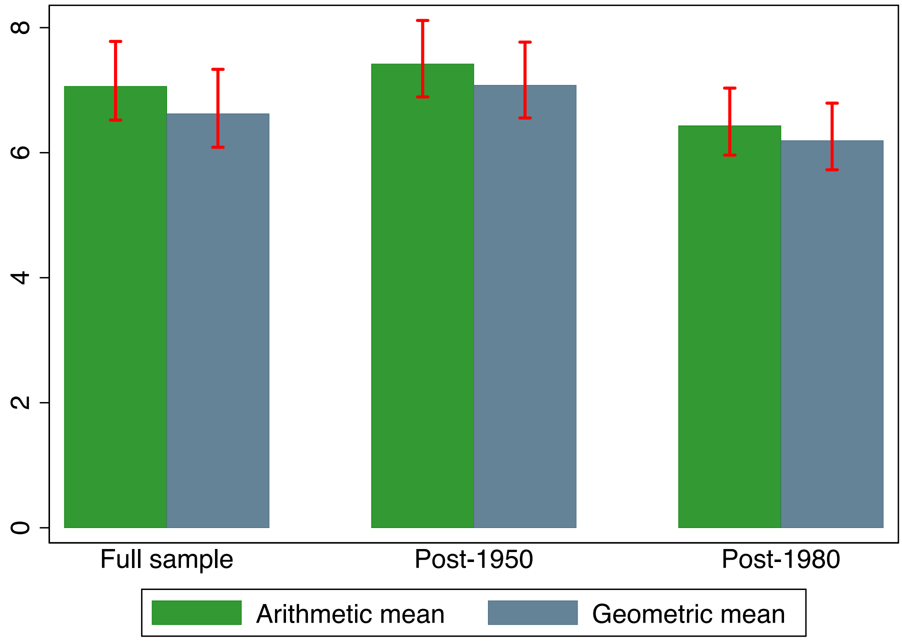

Note: Bars show the arithmetic- and geometric-average housing returns for selected subperiods. Error bars show the impact on historical returns of increasing or reducing the benchmark price/rent ratio by ± three, which broadly captures the difference between in- and out-of-city-center locations.

yield, it should have little influence on total returns, since cities tend to have higher capital gains, but lower rental yields.[33] Additionally, Knoll, Schularick, and Steger (2017) show that the rural-urban divide has a relatively small impact on capital gains. Relatedly, we can establish some bounds on how much our rental yields can vary with the choice of location. In 2013, `Numbeo.com` data suggest that price-rent ratios in and out of city centres differ by less than three times annual rent. The rental yield is the inverse of these price-rent ratios. This motivates us to construct a lower-bound rent-price ratio as $RP_{low} = \frac{1}{\frac{1}{RP_{actual}}+3}$ and an upper-bound rent-price ratio as $RP_{high} = \frac{1}{\frac{1}{RP_{actual}}-3}$ for each country in 2013 to estimate upper and lower bounds of our housing returns depending on the choice of location. Given the currently high price-rent ratios, these adjustments have a relatively small impact on our data. Figure V shows that increasing or reducing the price-rent ratio by three changes annual return estimates by about ±1 pps per year relative to our preferred baseline.

This suggests that the level of housing returns in our dataset should be representative of a country-wide portfolio. Still, it could be that returns on locations not included in our sample display higher volatility. For example, the post-1975 US indices are based on conforming mortgages and may exclude the more volatile segments of the market. To assess the likely magnitude of this bias, Table V compares the recent level and volatility of the US conforming mortgage based OFHEO

[33] Eisfeldt and Demers (2015) study the geographical distribution of returns on single-family rentals in the US from 1970s to today and find that lower capital gain areas tend to have much higher rental yields, and there is very little geographic variation in total returns.

Table V: *Level and volatility of real housing capital gains at different levels of coverage and aggregation*

	Baseline	Zillow			
	National	National	State	County	ZIP code
Mean real capital gain p.a.	1.42	0.79	1.07	0.53	0.92
Standard deviation	4.67	5.67	6.05	6.28	7.46

Note: US data, 1995–2015. Average annual real capital gain and standard deviation of house prices. Baseline data are sourced from the OFHEO index. Zillow data are sourced from the Zillow Home Value Index which covers around 95% of the US housing stock, and are averages of monthly values. National data are the returns and volatility of prices for a nationwide housing index, and the other figures cover a representative state, county or ZIP code level porftolio respectively.

house price indices with those that cover other segments of the market as well, which are sourced from Zillow.[34] Comparing columns 2 and 3 of Table V, the nationwide moments of the data are similar across the two measures—but, as expected, the OFHEO data display slightly higher real capital gains and slightly lower volatility, because they have a less comprehensive coverage of the areas that were hit hardest by the subprime crisis, which receives a relatively high weight in the 1995–2015 US sample used here.

Columns 3–5 of Table V also show that the volatility of the housing series increases as we move from the aggregate portfolio (column 2) to the subnational and local level. The standard deviation of ZIP code-level housing returns is roughly one-third higher than that in the national data. If investors owned one undiversified house whose price tracks the neighborhood index, the risk and return characteristics of this portfolio would be somewhat closer to those of the aggregate equity index, although the gap would still be large.

Of course, it is much more difficult to invest in a diversified housing portfolio than a well-diversified equity portfolio. That being said, Benhabib and Bisin (2016) show that most equity is also held in an undiversified manner. The data regarding returns on individual housing and private equity returns are, however, at this point in time, very sparse. To understand exactly how these risk-return characteristics play out at a disaggregated level, a more detailed study of individual portfolios and returns is necessary. This would be a worthy goal of future research.

Another selection bias comes about from the fact that rent data largely come from the rental market, whereas the majority of housing stock is held by owner-occupiers. To guard against this, we rely on rental indices that, whenever possible, account for the growth of imputed rents. Additionally, we benchmark our series to the balance sheet yield data that are constructed to cover the total housing stock. Still, imputed rents are measured with error, and may not reflect the cost that the homeowner would pay on the rental market. If owning is relatively cheaper than renting—for example, due to tax exemptions or long-run house price appreciation—homeowners would purchase

[34] As we show later in Section IV.C, almost all the volatility in housing returns comes about from house prices. Therefore for the analysis of volatility, we focus on house prices rather than rental yields.

larger or better houses than they would rent, and imputed rents would overstate the value of housing services accruing to homeowners. On the other hand, buying a house is subject to credit constraints, which means that renters can afford better houses than homeowners all else equal. In this case, imputed rents would understate the flow value of housing services. Overall, the direction of any potential bias is unclear and leaves much scope for future work.

Finally, the portfolio selection in the price and rent series is subject to survivorship bias. In general, our price and rental yield estimates aim to capture transaction or appraisal values, and rental costs on a broad and impartially selected portfolio of properties. Some survivorship bias may, however, enter the series for the following reasons. First, indices that rely on an ex post selection of cities may inadvertently choose the more "successful" cities over the less successful ones. Second, houses that decline in value are likely to lose liquidity and be sold less frequently, hence carrying a lower weight in the index. And third, chain-linking historical house price and rent indices to compute annual returns will generally ignore the impact of large destructions of the housing stock, in particular those occurring during wartime.

Several factors suggest that the impact of survivorship bias on housing returns should be limited. First, Figure V and Knoll, Schularick, and Steger (2017) show that any location-specific bias in our estimates is likely to be small. Second, if the magnitude of survivorship bias is similar to that in equity markets (Section III.B), then the bias is also unlikely to be large. Third, the low liquidity and weight of houses with declining prices is in some ways similar to the documented negative returns on delisted equities (Shumway, 1997; Shumway and Warther, 1999), which in general cannot be incorporated into the stock return series due to the lack of data. Therefore this bias should be less of a concern when comparing housing and equity returns. Finally, similarly to the stock market closures discussed in Section III.B, even though capital stock destruction during wars can have a substantial impact on returns in specific years, it is unlikely to profoundly affect cross-country long-run returns due to the rarity of such events.[35] And as Figure IX later shows, the main facts in the data are similar for countries that experienced major war destruction on their own territory versus countries that did not (e.g., Australia, Canada, Sweden, Switzerland, and US). Further, Online Appendix Table A.5 shows that housing offers a similar return relative to equity on average even after wars are excluded.

Returns on real estate investment trusts Another way to check our housing returns is to compare them to the historical returns on housing investment trusts. These trusts offer independent estimates of returns. Real estate investment trusts, or REITs, are investment funds that specialize in the purchase and management of residential and commercial real estate. Many of these funds list their shares on the local stock exchange. The return on these shares should closely track total real estate returns. Differences will arise because the REIT portfolio is more geographically concentrated,

[35]As a reasonable upper bound, existing estimates suggest that around 33%–40% of the German housing stock was destroyed by Allied bombing during WW2 (Akbulut-Yuksel, 2014; Diefendorf, 1993), but this would lower the full-sample country-specific average annual real total return by only about 0.30 pps per year.

Figure VI: *Returns on housing compared to real estate investment funds*

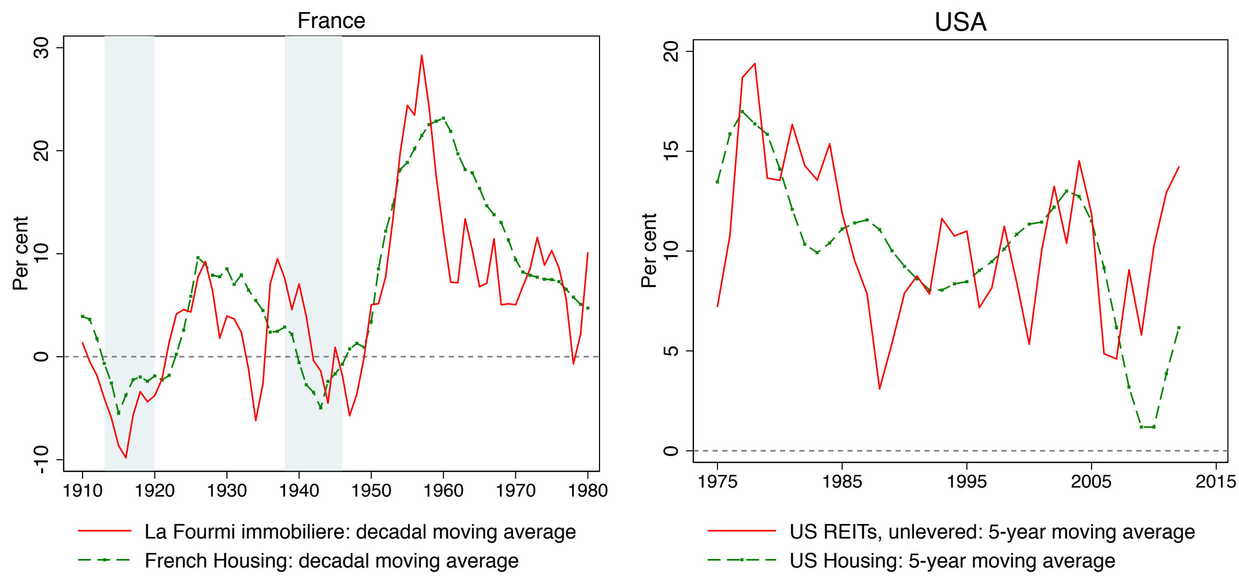

Note: Total real return on housing, and shares of housing investment firms in France and US. Moving averages. Following Giacomini, Ling, and Naranjo (2015), we assume a 45% leverage ratio for US REITs.

its assets often contain nonresidential property, and share price fluctuations may reflect expectations of future earnings and sentiment, as well as underlying portfolio returns. Further, the REIT portfolio returns should be net of taxes and transaction costs as well as housing running costs, and may therefore be somewhat lower than our housing series. Still, returns on the REIT portfolio should be comparable to housing and can be used to check the general plausibility of our return series.

Figure VI compares our historical housing returns (dashed line) with those on investments in REITs (solid line) in France and US, two countries for which longer-run REIT return data are available. The REIT returns series for France refers to shares of the fund "La Fourmi Immobilière" (see Simonnet, Gallais-Hamonno, and Arbulu, 1998). The fund acquired a portfolio of 15 properties in Paris between 1900 and 1913, worth around 36 million euros at 2015 prices, and its shares were listed on the Paris stock exchange between 1904 and 1997. We exclude the period after 1985, when "La Fourmi Immobilière" was taken over by AGF. For the US, we use the FTSE NAREIT residential total return index after 1994, and the general FTSE equity NAREIT before. REIT returns have to be unlevered to capture the returns on the REIT housing portfolio. "La Fourmi Immobilière" had an unlevered balance sheet structure, hence we do not adjust their returns. We assume a REIT leverage of 45% for the U.S. following Giacomini, Ling, and Naranjo (2015). Returns for France are presented as decadal moving averages, and for the US as five-year moving averages, given the shorter span of the US data.

Comparing the solid and dashed lines in Figure VI, the long-run levels of unlevered REIT and housing returns are remarkably similar. The time trend also follows a similar pattern, especially in France. The REIT returns, however, tend to be somewhat more volatile—most likely because they reflect changes in the market's valuations of future earnings, as well as the current portfolio performance. The REIT returns also seem to be affected by the general ups and downs of the stock market: for example, the 1987 "Black Monday" crash and dot-com bust in the U.S., as well as the 1930s Great Depression and 1960s stock crises in France. This suggests that the valuations of the funds' housing portfolios may have been affected by shifts in general stock market sentiment, possibly unrelated to housing market fundamentals.

Overall, the returns on real estate investment funds serve to confirm the general housing return level in our dataset. The comparison also suggests that returns in housing markets tend to be smoother than those in stock markets. The next section examines various factors that can affect the comparability of housing and equity returns more generally.

III.D. Comparability of housing and equity returns

Even if the fundamentals driving housing and equity returns (expected dividend/profit, and rental flows) are similar, investor returns for the two asset classes may differ for a number of reasons including taxes, transaction costs, and the financial structure of the investment claim. In this subsection we consider such comparability issues.

Transaction costs The conventional wisdom is that while bonds and equities can be purchased with low transaction costs and at short notice, the seller of a house typically incurs significant costs. We provide a rough estimate of how transaction costs affect our return estimates for housing. We perform a simple back-of-the-envelope calculation to do this using contemporary data on average holding periods of residential real estate and average transaction costs incurred by the buyer. According to the (OECD, 2012), average round-trip transaction costs across 13 of the 16 countries in our sample amount to about 7.7 percent of the property's value.[36]

However, these simple cost ratios need to be adjusted for the typical trading frequency of each asset. According to the American Community Survey of 2007, more than 50 percent of U.S. homeowners had lived in their current home for more than 10 years. Current average holding periods are similar in, e.g., the U.K., Australia, and the Netherlands. Another way to estimate housing turnover is using housing sales data, which for the U.S. gives us an average holding period of close to 20 years.[37] Either way, accounting for transaction costs would thus lower the average annual return to housing to less than 100 basis points (e.g., 77 basis points per year based on a 7.7% cost incurred every 10 years).

[36]Data are available for Australia, Belgium, Switzerland, Germany, Denmark, Finland, France, the U.K., Japan, the Netherlands, Norway, Sweden, and the U.S. Transaction costs are highest in Belgium amounting to nearly 15 percent of the property value and lowest in Denmark amounting to only 1 percent of the property value.

[37]Between April 2017 and March 2018, 5.5 million existing homes were sold in the U.S. at an average price of $250,000, which amounts to roughly one-twentieth of the total U.S. housing stock.

For equities, the cost of each individual transaction is much smaller, but the number of transactions is much higher. Jones (2002) estimates that at the New York Stock exchange over the period 1900–2001, the average transaction cost was around 80 bps (half bid-ask spread of 30 bps plus commission rate of 50 bps), while turnover was roughly 60% per year, resulting in the annual average equity transaction costs of 40 bps. Comparing this number to the back-of-the-envelope housing transaction cost estimates reported above, it seems that even though equity transaction costs are probably somewhat lower, the difference between two asset classes is likely to be small—and no more than 0.5 pps per year.

The fact that housing faces much higher costs per each transaction, however, means that the realized housing transaction costs may understate the "shadow" utility cost which would include the suboptimal allocation choices from staying in the same house and not moving, for example. It might also reduce the volatility of housing returns, making them react more sluggishly to shocks. This means that the relatively modest short-run volatility of housing returns could mask more pronounced fluctuations at lower frequencies. Online Appendix K and Table A.17 compare equity and housing return volatility over longer horizons of up to 20 years. It turns out that the standard deviation of housing returns is always around one-half that of equity returns, regardless of the time horizon, which suggests that housing returns not only have lower short-run volatility, but also less pronounced swings at all longer horizons.

Leverage Household-level returns on real estate and equity will be affected by the structure of the household balance sheet, and how these investments are financed. Jordà, Schularick, and Taylor (2016) show that advanced economies in the second half of the 20th century experienced a boom in mortgage lending and borrowing. This surge in household borrowing not only reflected rising house prices, it also reflected substantially higher household debt levels relative to asset values (and relative to household incomes). The majority of households in advanced economies today hold a leveraged portfolio in their local real estate market. As with any leveraged portfolio, this significantly increases both the risk and the return associated with the investment. And today, unlike in the early twentieth century, houses can be levered much more than equities. The benchmark rent-price ratios from the IPD used to construct estimates of the return to housing refer to rent-price ratios of unleveraged real estate. Consequently, the estimates presented so far constitute only unlevered housing returns of a hypothetical long-only investor, which is symmetric to the way we (and the literature) have treated equities.

However, computing raw returns to housing and equity indices neglects the fact that an equity investment contains embedded leverage. The underlying corporations have balance sheets with both debt and equity liabilities. Thus, reconciliation is needed, and two routes can be taken. For truly comparable raw unlevered returns, equity returns could be delevered. Alternatively, for truly comparable levered returns, housing returns would have to be levered up to factor in the actual leverage (using mortgages) seen on household balance sheets. Is this a big deal in practice? We argue that it does not bias our conclusions significantly based on some elementary calculations.

Consider, for example, the second reconciliation of levering up housing returns. Let the real long-term mortgage borrowing rate be r_0, and α be the leverage of the average house proxied by total mortgages divided by the value of the housing stock. Then we can solve for levered real housing returns TR' as a function of unlevered real housing returns TR using the formula $TR' = \frac{TR - \alpha r_0}{1-\alpha}$. In our data, $TR \approx 7.0\%$ and $\alpha \approx 0.2$. Using a long bond return as a proxy for r_0 of around 2.5% p.a., this would imply $TR' = 8.1\%$.[38] In other words, for the representative agent the levered housing return is about 110 bps higher than the unlevered housing return (8.1% versus 7%), a small difference. Such adjustments appear to be inconsequential for the main conclusions we present in this paper. In fact, they would bolster one of our central new claims which is that real housing returns at least match or even exceed real equity returns in the long run when the two are compared on an equal footing.

Taxes When computing equity and housing returns we do not account for taxes. From an investor's perspective accounting for taxes is clearly important. Typically, equity capital gains—and, for some countries and periods, even dividend income—have been subject to a capital gains tax. When dividends are not taxed as capital gains, they tend to be taxed as income. In some countries, housing capital gains are subject to capital gains taxes, but owner-occupied houses in particular have been granted exemptions in many cases. Imputed rents of homeowners are, unlike dividend income, almost never taxed. Additionally, housing tends to be subject to asset-specific levies in the form of property taxes, documented extensively in Online Appendix Y.

For both equities and housing, the level and applicability of taxes has varied over time. For housing, this variation in treatment also extends to assessment rules, valuations, and tax band specifications. As a ballpark estimate, the impact of property taxes would lower real estate returns by less than 1.0 pps per year relative to equity (see Online Appendix Y for further details). The various exemptions for homeowners make the impact of capital gains taxes on real estate returns even harder to quantify but such exemptions also imply that differential tax treatment is unlikely to play an important role in explaining differences in the return between equities and housing.[39]

Since quantifying the time- and country-varying effect of taxes on returns with precision is beyond the scope of this study, throughout this paper we focus on pretax returns from an investor perspective. Importantly, these pretax returns are net of corporate profit taxes, which are netted out before the cashflow payment to the investor. Studies of returns from an aggregate wealth perspective such as Piketty, Saez, and Zucman (2018) typically compute business equity returns gross of corporate tax. Online Appendix S discusses the impact of adding back corporate taxes on

[38]For evidence on α, the average economy-wide housing leverage measured by total mortgages divided by the value of the housing stock, see Jordà, Schularick, and Taylor (2016). If one preferred to use the mortgage rate rather than the long bond in this calculation, the evidence in Zimmermann (2017) points to an average real mortgage rate r_m of around 3% p.a. This would imply $TR' = 8\%$, only slightly lower than the figure quoted in the main text.

[39]Note that whilst this is true for aggregate or owner-occupied housing, the tax burden on landlords is likely to be somewhat higher than that on holders of listed equity, because landlords do not benefit from the homeowner exemptions to property taxes, and their rental income is taxed.

Table VI: *Impact of using end-of-year versus yearly-average asset prices*

	Equity (MSCI index)		Housing (this paper)
	End-of-year	Yearly average	Yearly average
Mean return p.a.	8.70	7.51	6.55
Standard deviation	27.56	22.00	7.45
Observations	694	694	694

Note: Annual global real returns in 16 countries, equally weighted, 1970–2015. End-of-year returns are computed using the return index value for the last day of the year. Yearly average returns are computed using the average index value throughout the year.

our return data. Equity returns before corporate tax would be roughly one percentage point higher than our baseline estimates (Table A.27). This adjustment is, however, very much an upper bound on the housing-equity return differential for the following reasons. First, as noted above, a true like-for-like comparison should also delever equity returns and compare the returns on business and housing wealth. Online Appendix S Table A.27 estimates that first adding back corporate taxes, and then delevering equity returns leaves them approximately equal to the baseline estimates that we report. Second, the total tax burden on the pre-corporate-tax equity returns is likely to be higher than on housing, since in light of the various homeowner exemptions, the post-corporate-tax burden on the two assets appears to be roughly similar. Third, the returns on the two asset classes are similar before 1920, when the corporate tax rate was close to zero.

Temporal aggregation and return averaging The way house and equity price indices are constructed is likely to influence the volatility of the return series. The house price indices used for return calculations (e.g., indices from national statistical agencies) tend to be an *average* of all transactions in a given year, or use a sample of transactions or appraisal values throughout the year. But the equity prices used for return calculations, by the usual convention followed here, on the contrary, compare *end-of-year* prices of shares. The use of end-of-year rather than yearly-average prices mechanically makes equity returns more volatile.

We can assess the magnitude of this effect by constructing an equity return index based on annual averages of daily data, to mimic the way housing returns are computed, and then comparing it to a "normal" return using end-of-year index values. For this robustness exercise we use daily MSCI equity returns data for 1970–2015. Table VI presents the end-of-year and yearly-average real equity returns in the first two columns, and our yearly-average housing returns for the same time period in the third column. Comparing the first two columns shows that yearly averaging lowers the standard deviation of returns by around one-fifth, or 5 pps. It also lowers the average return by around 1 ppt, because the return series is a transformation of the raw price data, and lowering the variance reduces the mean of the return. But the standard deviation of the smoothed yearly-average equity series is still almost three times that of housing over the same time period.

Because historical house price data sometimes rely on relatively few transactions, they are likely to be somewhat less smooth than averages of daily data. Therefore Table VI provides an upper bound of the impact of averaging on our return series. Even taking this upper bound at face value, the averaging of house price indices is likely to explain some, but far from all, of the differences in volatility of equity and housing returns.

IV. Risky rates of return

At this waystation the lengthy pilgrimage of Section III ends: the numerous details of how we compiled our data; the many important, but somewhat technical, aspects of data construction; the extensive accuracy checks. In these next sections the pace picks up and the focus turns to analysis and interpretation of the data. We examine broad trends and explore their implications for how we think about macroeconomics and finance, confronting the four big themes laid out in the introduction: the long-run behavior of risky returns, safe returns, risk premia, and "r minus g." Readers who skipped the better part of Section III are welcomed to rejoin the flow here.

IV.A. Global returns

We first turn in Figure VII to a more detailed inspection of the returns on the risky assets, equity and housing. The global returns are GDP-weighted averages of the 16 countries in our sample. Although we do not show the unweighted data, the corresponding figure would look very similar. We smooth the data using decadal moving averages as explained earlier. For example, the observation reported in 1900 is the average of data from 1895 to 1905. Figure VII shows the trends in decadal-average real returns on housing (solid line) and equity (dashed line) for our entire sample. In addition, Figure VIII displays the average rolling decadal correlation of annual risky returns between asset classes, across countries, and with inflation.

Risky returns were high in the 1870s and 1880s, fell slowly at first, but then sharply after 1900, with the decade-average real equity returns turning negative during WW1. Risky returns recovered quickly in the 1920s, before experiencing a drop in the the Great Depression, especially for equities. Strikingly, after WW2 the trajectories of returns for the two risky asset classes had similar long-run means but over shorter periods diverged markedly from each other.

Equity returns have experienced many pronounced global boom-bust cycles, much more so than housing returns, with average real returns as high as +16% and as low as –4% over entire decades. Equity returns fell in WW2, boomed in the post-war reconstruction, and fell off again in the climate of general macroeconomic instability in the 1970s. Equity returns bounced back following a wave of deregulation and privatization in the 1980s. The next major event was the Global Financial Crisis, which exacted its toll on equities and to some extent housing, as we shall see.

Housing returns, on the other hand, have remained remarkably stable over almost the entire post-WW2 period. As a consequence, the correlation between equity and housing returns, depicted

Figure VII: *Trends in real returns on equity and housing*

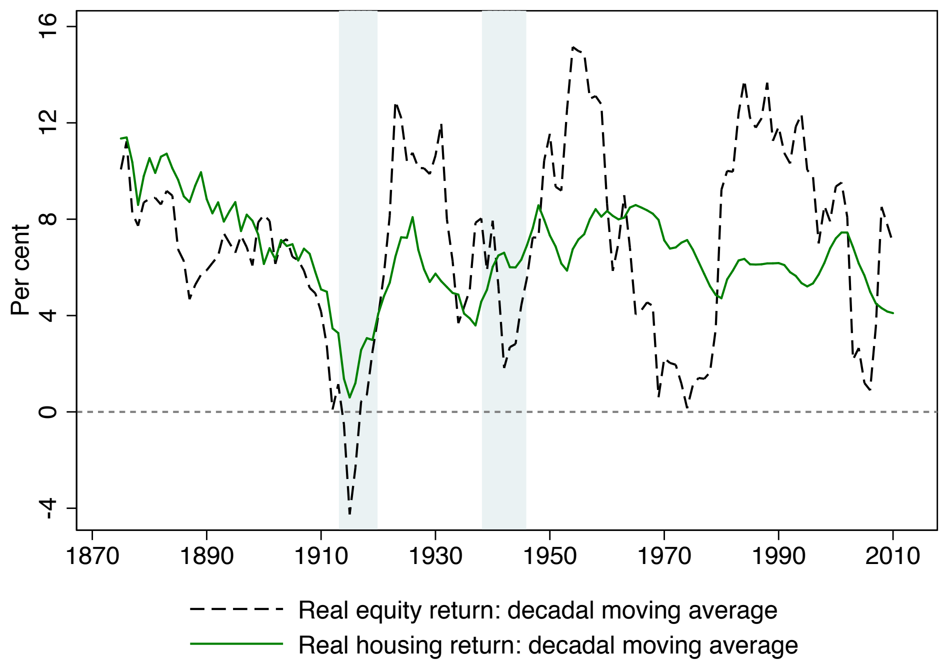

Note: Mean returns for 16 countries, weighted by real GDP. Decadal moving averages.

in the top left panel of Figure VIII, was highly positive before WW2, but has all but disappeared over the past five decades. The low covariance of equity and housing returns over the long run reveals potential attractive gains from diversification across these two asset classes that economists, up to now, have been unable to measure or analyze.

In terms of relative returns, we see that housing persistently outperformed equity up until WW1, even though both of these asset returns followed a broadly similar temporal pattern. In recent decades, equities have slightly outperformed housing in (arithmetic, not geometric) average, but with much higher volatility and cyclicality. Furthermore, upswings in equity prices have generally not coincided with times of low growth or high inflation, when standard asset pricing theory would say high returns would have been particularly valuable.

The top-right panel of Figure VIII examines the correlation between risky rates of return and inflation. It shows that equity comoved negatively with inflation in the 1970s, while housing provided a more robust hedge against an unusually rapid surge in consumer prices. In fact, apart from the interwar period, when the world was gripped by a broad deflationary bias, we find that equity returns have comoved negatively with inflation in almost all eras. Moreover, the big downswings in equity returns in the two world wars and the 1970s coincided with periods of generally poor economic performance.

In the past two decades equity returns have also become highly correlated across countries, as shown by the sharp rise in the degree of cross-country comovement in the bottom-left panel

Figure VIII: *Correlations across risky asset returns*

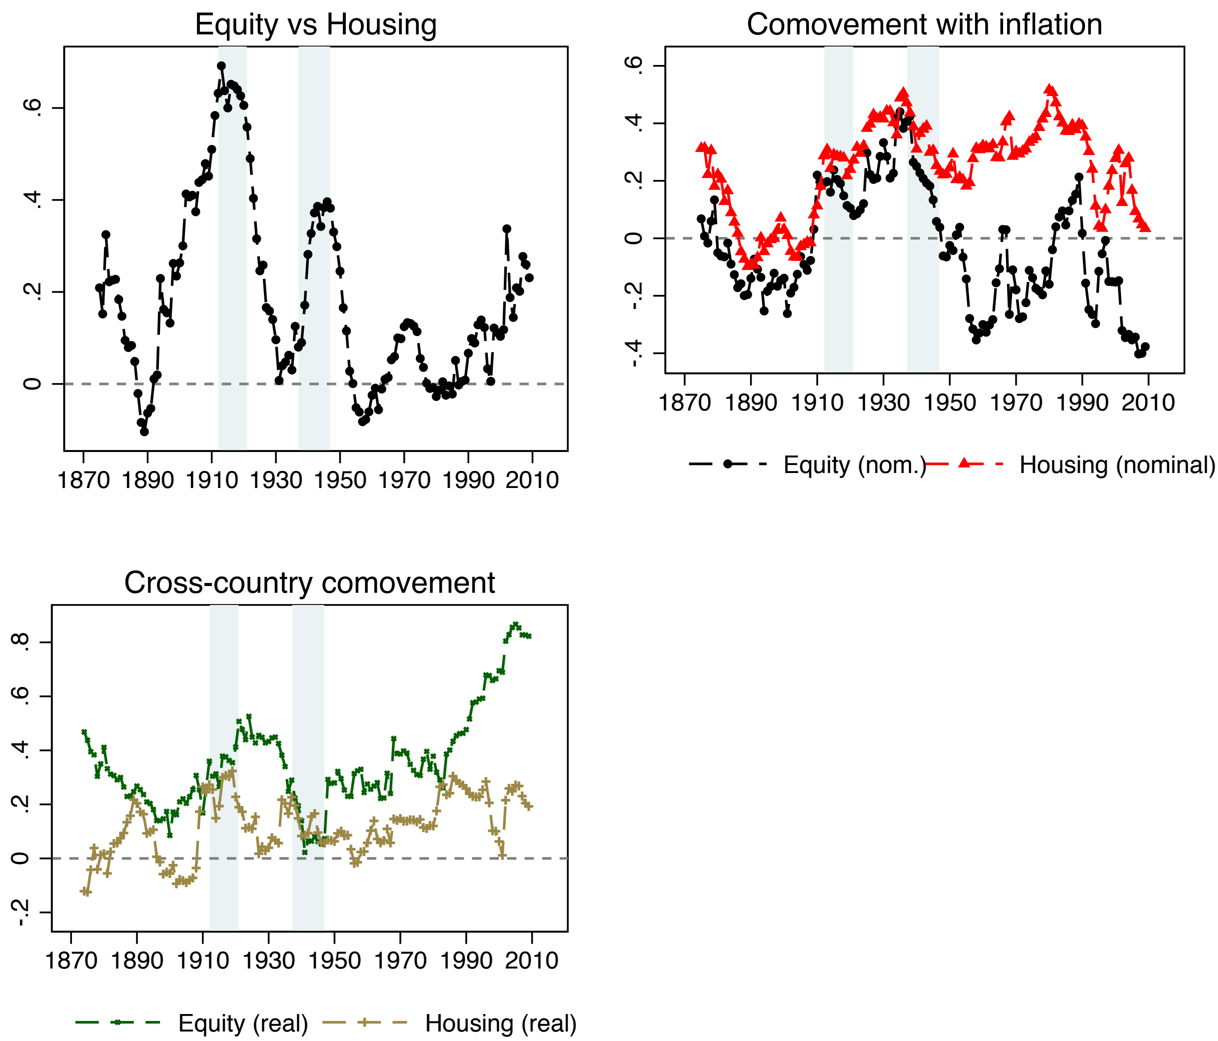

Note: Rolling decadal correlations. The global correlation coefficient is the average of individual countries for the rolling window. Cross-country correlation coefficient is the average of all country pairs for a given asset class. Country coverage differs across time periods.

of Figure VIII, measured as the average of all country-pair correlations for a given window.[40] A well-diversified global equity portfolio has thus become less of a hedge against country-specific risk (Quinn and Voth, 2008). As is a matter of debate, this may reflect greater freedom to arbitrage and

[40]We report the average of all country-pair combinations for a given window, calculated as

$$Corr_{i,t} = \frac{\sum_j \sum_{k \neq j} Corr(r_{i,j,t\in T}, r_{i,k,t\in T})}{\sum_j \sum_{k \neq j} 1}$$

for asset i (here: equities or housing), and time window $T = (t-5, t+5)$. Here j and k denote the country pairs, and r denotes real returns, constructed as described in Section II.C.

Table VII: *Real rates of return on equity and housing*

Country	Full Sample		Post 1950		Post 1980	
	Equity	Housing	Equity	Housing	Equity	Housing
Australia	7.79	6.37	7.53	8.29	8.70	7.16
Belgium	6.23	7.89	9.65	8.14	11.49	7.20
Denmark	7.49	8.22	9.73	7.04	13.30	5.14
Finland	10.03	9.58	12.89	11.18	16.32	9.47
France	3.21	6.39	6.01	9.68	9.61	5.78
Germany	7.11	7.82	7.53	5.30	10.07	4.13
Italy	7.25	4.77	6.09	5.55	9.45	4.57
Japan	6.00	6.54	6.21	6.74	5.62	3.58
Netherlands	6.96	7.28	9.19	8.53	11.51	6.41
Norway	5.67	8.03	7.33	9.10	12.22	9.82
Portugal	4.51	6.31	4.84	6.01	8.60	7.15
Spain	5.83	5.21	7.75	5.83	11.96	4.62
Sweden	8.02	8.30	11.37	8.94	15.87	9.00
Switzerland	6.51	5.63	8.37	5.64	9.29	6.19
UK	6.83	5.44	9.10	6.57	9.11	6.81
USA	8.46	6.10	8.89	5.76	9.31	5.86
Average, unweighted	6.67	7.26	8.30	7.47	10.78	6.43
Average, weighted	7.12	6.72	8.19	6.40	9.08	5.50

Note: Average annual real returns. Period coverage differs across countries. Consistent coverage within countries: each country-year observation used to compute the statistics in this table has data for both real housing and equity returns. The average, unweighted and average, weighted figures are respectively the unweighted and real-GDP-weighted arithmetic averages of individual country returns.

trade across equity markets globally, or an increase in the global shocks to which firms, especially those in the typical equity index, are increasingly exposed. In contrast to equities, cross-country housing returns have remained relatively uncorrelated, perhaps because housing assets remain less globally tradable than equities, or because they are more exposed to idiosyncratic country-level shocks.

IV.B. Country returns

Next we explore risky returns in individual countries. Table VII shows returns on equities and housing by country for the full sample and for the post-1950 and post-1980 subsamples. Long-run risky asset returns for most countries are close to 6%–8% per year, a figure which we think represents a robust and strong real return to risky capital. Still, the figures also show an important degree of heterogeneity among countries. Many of the countries that experienced large political shocks show lower equity returns. This is the case for Portugal and Spain which both underwent prolonged civil strife, and France which undertook a wave of nationalizations in the aftermath of WW2. French equity returns are also negatively affected by the fallout from the world wars, and the fallout from an oil crisis in the 1960s (for more detail, see Blancheton, Bonin, and Le Bris, 2014; Le Bris and Hautcoeur, 2010). In contrast, real equity returns in Finland have been as high as 10%, on average throughout the sample. Housing returns also show considerable heterogeneity. Returns on housing

Figure IX: *Risk and return of equity and housing*

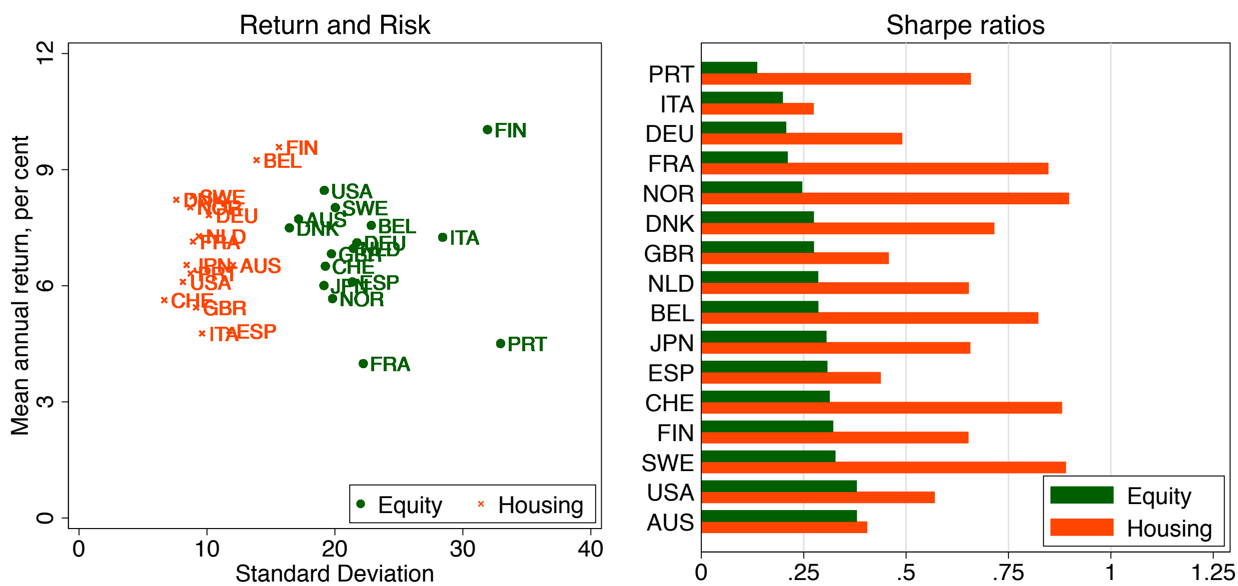

Note: Left panel: average real return p.a. and standard deviation. Right panel: Sharpe ratios, measured as $\frac{\bar{r}_i - \bar{r}_{bill}}{\sigma_i}$, where i is the risky asset with $\bar{r}_i$ mean return and σ_i standard deviation. 16 countries. Consistent coverage within each country.

have been high on average in the Nordic countries, but low in Italy and Spain. US risky asset returns fall roughly in the middle of the country-specific figures, with equity returns slightly above average, and housing returns slightly below. Our estimates of post-WW2 US housing returns are in line with those in Favilukis, Ludvigson, and Van Nieuwerburgh (2017).[41] The degree of heterogeneity and the relative ranking of returns is broadly similar when comparing the full sample to the post-1950 period.

This country-level evidence reinforces one of our main findings: housing has been as good a long-run investment as equities, and possibly better. Housing has offered a similar return to equity in the majority of countries and time periods. In the long-run, housing outperformed equities in absolute terms in six countries, and equities outperformed housing in five. Returns on the two assets were about the same in the remaining five countries. After WW2, housing was the best-performing asset class in three countries, and equities in nine.

However, although aggregate total returns on equities exceed those on housing for certain countries and time periods, equities do not outperform housing in simple risk-adjusted terms. Figure IX compares the risk and returns of housing and equities for each country. The left panel

[41] Favilukis, Ludvigson, and Van Nieuwerburgh (2017) estimate a gross nominal return on US housing of 9%–11%, based on three data sources going back to the 1950s and 1970s. This implies a net real return of around 5%–7% (once inflation, maintenance and running costs are subtracted), in line with our estimates in Table VII.

plots average annual real returns on housing and equities against their standard deviation. The right panel shows the Sharpe ratios for equities and housing for each country in the sample.[42] Housing provides a higher return per unit of risk in each of the 16 countries in our sample, with Sharpe ratios on average more than double those for equities.

IV.C. Decomposition of returns

To further look into the underlying drivers of housing and equity returns, we decompose them into the capital gain (price) and yield (dividend or rent) components. To be consistent with the data in Section III and Table II, we decompose *real* total return into *real* capital gain—that is, the price change net of inflation—and dividend or rental yield—that is, the nominal yield as proportion of the previous year's share or house price.[43] Yet caveats arise. In principle, it is not entirely clear whether inflation should be subtracted from the capital gain or yield component. Moreover, firms may buy back shares to generate low-tax capital gains instead of paying out higher-taxed dividends; thus, the manner of distribution of total returns may not be invariant to circumstances.

Table VIII decomposes equity and housing returns into capital gains and dividends or rents, for the full cross-country sample and the period after 1950. Over the full sample, most of the real return is attributable to the yield. Dividends account for roughly 60% of real equity returns, and rents for roughly 80% of real housing returns. In terms of geometric means (Table VIII, row 3), almost all of both equity and housing returns are attributable to, respectively, dividend and rental income. After 1950, capital gains become more important for both equities and housing. For equities, real capital gains account for the majority of the total return after 1950, and for housing for roughly one-third.

The importance of dividends and rents is partly a matter of convention. Online Appendix N and Online Appendix Table A.20 computes the equivalent decomposition for nominal returns, and finds that the capital gain versus dividend/rental income split is then closer to roughly 50/50. Nevertheless, without dividends or rents, the real returns on both assets would be low, especially in geometric mean terms. This is consistent with the existing literature on real house prices: Shiller (2000) documents that house prices in the US moved in line with inflation before the 2000s bubble, and Knoll, Schularick, and Steger (2017) show that real house prices in advanced economies were more or less flat before 1950. This is also true in our data: the pre-1950 annual real housing capital gains are just 0.5%. Post-1950 capital gains are somewhat higher at 2.5%, but still only half the magnitude of the rental yields. Adding rents to the equation radically changes the picture, and brings the long-run housing total return close to 7%. Interestingly, the broad picture is similar for equities: the real equity capital gain before 1950 is, on average, just 0.4%, compared to 4.7% per year after 1950. However, the contribution of dividend and rental income means that housing and equity

[42] The Sharpe ratio is calculated as $\frac{\bar{r}_i - \bar{r}_{bill}}{\sigma_i}$, where i is the risky asset (housing or equity) with $\bar{r}_i$ mean return and σ_i standard deviation.

[43] The small residual difference between combined capital gain and dividend income, and the equity total return, accounts for gain and loss from capital operations such as stock splits or share buybacks, and income from reinvestment of dividends.

Table VIII: *Total return components for equity and housing*

	Equity			Housing		
	Real capital gain	Dividend income	Real total return	Real capital gain	Rental income	Real total return
Full sample:						
Mean return p.a.	2.78	4.17	6.82	1.61	5.50	6.92
Standard deviation	21.37	1.74	21.89	9.87	2.05	10.40
Geometric mean	0.57	4.16	4.58	1.15	5.48	6.43
Observations	1,707	1,707	1,707	1,707	1,707	1,707
Post 1950:						
Mean return p.a.	4.73	3.80	8.36	2.39	5.22	7.38
Standard deviation	23.70	1.81	24.24	8.59	1.93	8.95
Geometric mean	2.03	3.79	5.62	2.06	5.21	7.04
Observations	995	995	995	995	995	995

Note: Average annual real capital gain, dividend, or rental income, and total return across 16 countries, unweighted. Period coverage differs across countries. Consistent coverage within countries: each country-year observation used to compute the statistics in this table has data for both equity and housing returns, capital gains, and yields. Dividend and rental income are expressed in percent of the previous year's asset price.

returns were high both before and after 1950.

While most of the return can be attributed to dividends and rents, almost all of the volatility comes from equity and house prices, i.e., the capital gains component. Table VIII row 2 shows that both dividends and rents are very stable, with a standard deviation of dividend/rental yields of about 2%. Prices, on the contrary, move around much more, with a standard deviation roughly equal to that of total returns (21.4% for equities and 9.9% for housing). The higher variability of equity returns, and the superior risk-adjusted performance of housing can, therefore, largely be attributed to the lower volatility of house prices compared to those of equities. An additional factor is that capital gains—the more volatile component—account for a relatively larger share of equity returns.

Since aggregate equity prices are subject to large and prolonged swings, a representative investor would have to hold on to his equity portfolio for longer in order to ensure a high real return. Aggregate housing returns, on the other hand, are more stable because swings in aggregate house prices are generally less pronounced. National aggregate housing portfolios have had comparable real returns to national aggregate equity portfolios, but with only half the volatility.

Table IX examines the relative importance of capital gains versus dividends/rents at the country level (figures are arithmetic means and standard deviations). Additionally we report the share of real total return accounted for by capital gains. The fact that the majority of housing returns come

Table IX: *Total return components for equity and housing by country*

	Equity				Housing			
	Real capital gain	Dividend income	Real total return	Capital gain share	Real capital gain	Rental income	Real total return	Capital gain share
Australia	3.06	4.90	7.79	0.38	2.53	3.99	6.37	0.24
	(16.30)	(1.08)	(16.94)		(11.94)	(0.92)	(11.92)	
Belgium	2.53	3.83	6.23	0.40	1.95	6.15	7.89	0.14
	(22.92)	(1.64)	(23.61)		(15.05)	(1.46)	(15.51)	
Denmark	2.71	4.95	7.49	0.35	1.26	7.13	8.22	0.08
	(16.14)	(2.09)	(16.45)		(7.02)	(2.42)	(7.60)	
Finland	5.19	5.08	10.03	0.51	2.82	7.14	9.58	0.17
	(31.18)	(1.95)	(31.93)		(14.61)	(2.86)	(15.62)	
France	-0.37	3.73	3.21	0.09	1.55	5.09	6.39	0.13
	(21.57)	(1.33)	(22.14)		(9.39)	(1.14)	(10.03)	
Germany	2.74	4.08	7.11	0.40	1.86	6.03	7.82	0.13
	(20.99)	(1.58)	(21.72)		(9.24)	(2.61)	(10.16)	
Italy	3.78	3.61	7.25	0.51	1.45	3.49	4.77	0.18
	(27.99)	(1.34)	(28.42)		(9.28)	(1.59)	(9.61)	
Japan	3.12	2.65	6.00	0.54	2.00	4.70	6.54	0.18
	(18.94)	(1.77)	(19.15)		(7.99)	(1.24)	(8.41)	
Netherlands	3.38	4.87	8.10	0.41	1.75	5.96	7.51	0.13
	(19.21)	(1.57)	(19.61)		(8.22)	(1.68)	(8.76)	
Norway	1.61	4.21	5.67	0.28	1.49	6.72	8.03	0.10
	(19.33)	(1.60)	(19.82)		(8.26)	(1.19)	(8.70)	
Portugal	2.92	2.28	5.11	0.56	1.13	4.47	5.21	0.12
	(34.34)	(1.22)	(34.73)		(9.26)	(1.98)	(9.37)	
Spain	1.80	4.53	5.83	0.28	1.26	4.16	5.21	0.13
	(20.48)	(2.30)	(21.15)		(11.59)	(1.60)	(12.00)	
Sweden	4.08	4.12	8.02	0.50	1.39	7.12	8.30	0.09
	(19.54)	(1.05)	(20.03)		(8.46)	(1.61)	(8.88)	
Switzerland	3.17	3.20	6.27	0.50	0.81	4.54	5.24	0.08
	(20.61)	(1.46)	(20.73)		(6.50)	(0.62)	(6.74)	
UK	2.48	4.53	6.83	0.35	1.63	3.94	5.44	0.17
	(19.12)	(1.39)	(19.73)		(8.94)	(0.86)	(9.15)	
USA	4.19	4.38	8.46	0.49	0.90	5.33	6.10	0.08
	(18.90)	(1.57)	(19.17)		(7.84)	(0.75)	(8.12)	

Note: Arithmetic average of annual real capital gain, dividend or rental income, and total return, full sample. Standard deviation in parentheses. Period coverage differs across countries. Consistent coverage within countries: each country-year observation used to compute the statistics in this table has data for both equity and housing returns, capital gains, and yields. Dividend and rental income are expressed as a percentage of the previous year's asset price. Capital gain share is the proportion of real total return attributable to real capital gains.

from yields is true across all countries. The lowest real capital gains are observed in Switzerland, and the highest in Finland, but none exceed 3% per year in the full sample. Rents are relatively more important in the US, accounting for roughly 90% of returns, but this is not unusual: Denmark, Sweden, and Switzerland have the same share of capital gains as the US. For equities, the picture is more mixed. Seven countries, including the US, have a roughly 50/50 split between real capital gain and dividend yield shares. Other countries record low or negative real capital gains over the full sample, and especially so in geometric mean terms (see Online Appendix Table A.22).

V. Safe rates of return

Turning to safe asset returns, Figure X shows the trends in real returns on government bonds (solid line) and bills (dashed line) since 1870. Again, returns are GDP-weighted averages of the 16 countries in our sample; the corresponding unweighted figure would look very similar. We smooth the data using decadal moving averages as explained earlier.

Three striking features of Figure X deserve comment. First, low real rates and, in fact, negative real rates have been relatively common during modern financial history. Second, for the most part, returns to long-term and short-term safe assets have tracked each other very closely—with a premium of about 1% that has widened considerably since the well-documented decline of the mid-1980s (see, e.g., Holston, Laubach, and Williams, 2017). Third, a major stylized fact leaps out once we compare the safe rates of return in Figure X to the risky rates of return in Figure VII above. Prior to WW2, real returns on housing, safe assets, and equities followed remarkably similar trajectories. After WW2 this was no longer the case.

Safe rates are far from stable in the medium-term. There is enormous time-series, as well as cross-country, variability. In fact, real safe rates appear to be as volatile as real risky rates (sometimes more volatile), a topic we return to in the next section. Considerable variation in the risk premium often comes from sharp changes in safe real returns, not from real returns on risky assets.

Two four-decade-long declines in real rates stand out: (1) from 1870 to WW1 (with a subsequent further collapse during the war); and (2) the well-documented decline that started in the mid-1980s. We could add to this list the briefer, albeit more dramatic decline that followed the Great Depression into WW2. Some observers have therefore interpreted the recent downward trend in safe rates as a sign of a new era of "secular stagnation" (see, e.g., Summers, 2014).

However, in contrast to 1870–1913 and the 1930s, the more recent decline is characterized by a much higher term premium—a feature with few precedents in our sample.[44] There are other periods in which real rates remained low, such as in the 1960s. They were pushed below zero, particularly for the longer-tenor bonds, during the 1970s inflation spike, although here too term

[44] One important qualification is that this is the ex post, not ex ante term premium. It therefore includes any unexpected shocks that affect either the short rate or the long-run bond return series. Furthermore, because the long-run bond return measure includes capital gains, and the short-term rate measure is the yield only (since the security matures within one year), most of the post-1980 increase in the term premium is driven by higher capital gains on long-term government bonds.

Figure X: *Trends in real returns on bonds and bills*

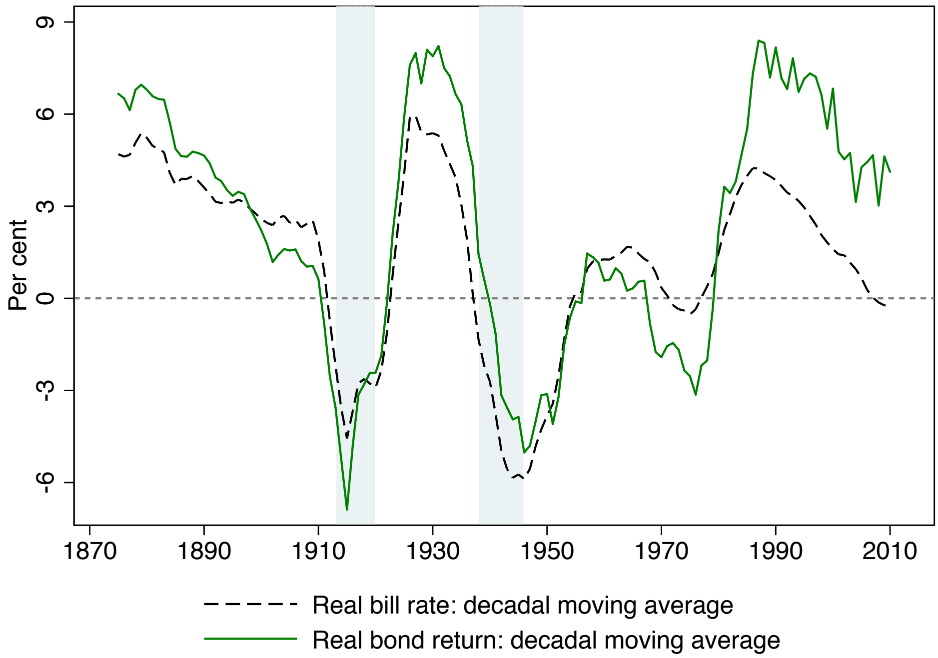

Note: Mean returns for 16 countries, weighted by real GDP. Decadal moving averages.

premiums remained relatively tight. Returns also dipped dramatically during both world wars. This is perhaps to be expected: demand for safe assets spikes during disasters although the dip may also reflect periods of financial repression and high inflation that usually emerge during times of conflict, and which often persist into peacetime. Thus, from a broad historical perspective, high rates of return on safe assets and high term premiums are more the exception than the rule.

Summing up, over more than 140 years from the late 19th to the 21st century, real returns on safe assets have been low—on average 1% for bills and 2.5% for bonds—relative to alternative investments. Although the return volatility—measured as annual standard deviation—is lower than that of housing and equities, these assets offered little protection during high-inflation eras and during the two world wars, both periods of low consumption growth.

Again moving on to examine correlations, Figure XI explores additional key moments of the data. The top-left panel plots the correlation between real bond and real bill returns, again using decadal rolling windows and computed as the cross-sectional average of correlations. In parallel to our discussion of the term premium, real returns on bonds and bills have been highly correlated for most of the sample up until the 1960s. From the 1970s onwards, the era of fiat money and higher average inflation, this correlation has become much weaker, and near zero at times, coinciding with a widening term premium.

The top right panel of Figure XI displays the correlation between nominal safe asset returns and inflation, both for real bond and real bill returns. The figure shows that safe assets provided more

Figure XI: *Correlations across safe asset returns*

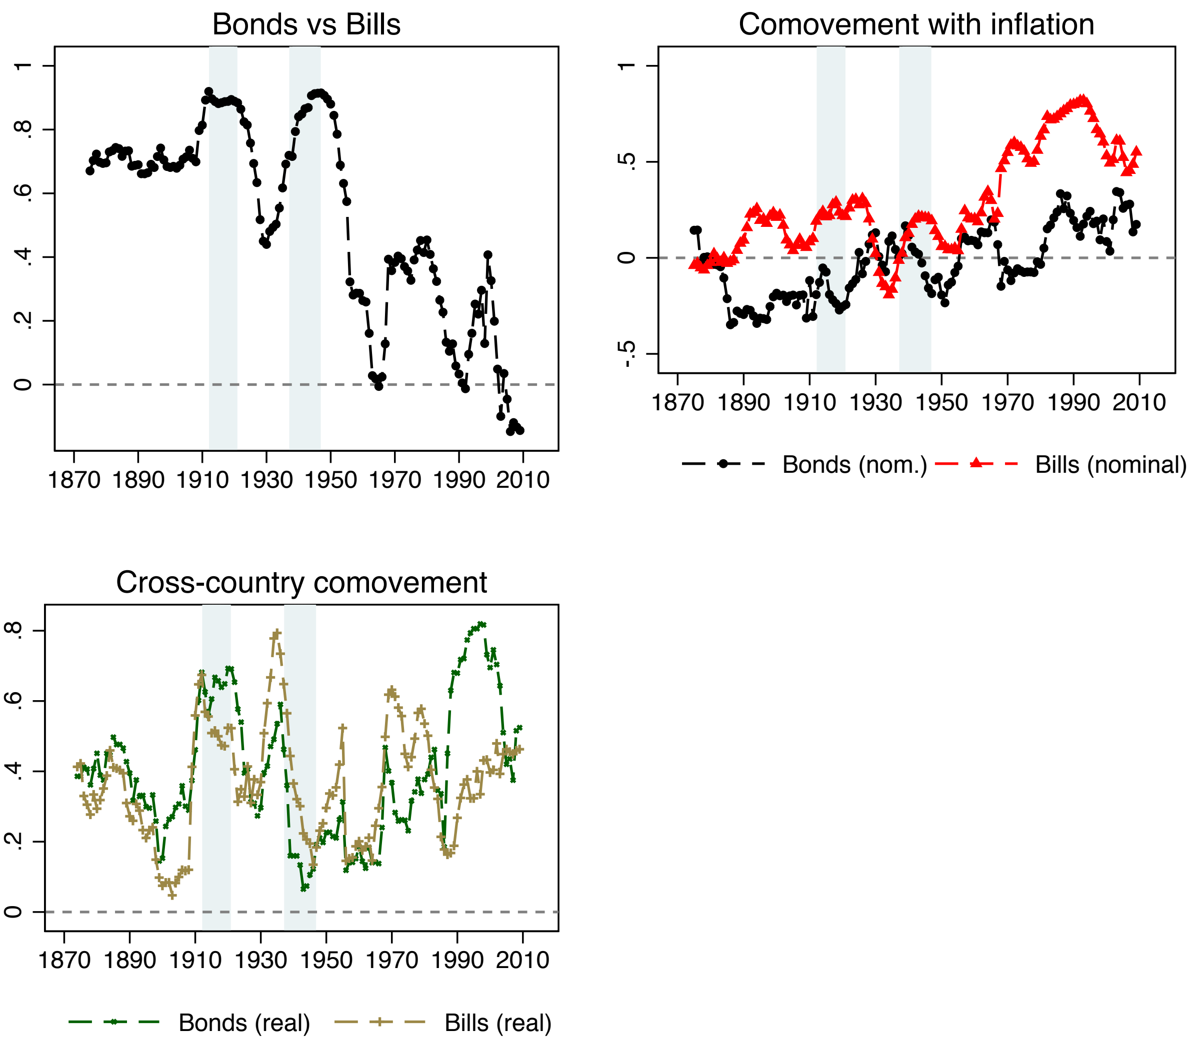

Note: Rolling decadal correlations. The global correlation coefficient is the average of individual countries for the rolling window. The cross-country correlation coefficient is the average of all country pairs for a given asset class. Country coverage differs across time periods.

of an inflation hedge starting in the 1970s, around the start of the era of modern central banking. However, as Figure X showed, both bonds and bills have experienced prolonged periods of negative real returns—both during wartime inflations, and in the high-inflation period of the late 1970s. Although safe asset rates usually comove positively with inflation, they do not always compensate the investor fully.

The bottom panel of Figure XI displays the cross correlation of safe returns over rolling decadal windows, averaged for all country-pair combinations, to examine how much risk can be diversified with debt instruments. Cross-country real safe returns have exhibited positive comovement throughout history. The degree of comovement shows a few marked increases in WW1 and in the 1930s.

Table X: *Real rates of return on bonds and bills*

Country	Full Sample		Post 1950		Post 1980	
	Bills	Bonds	Bills	Bonds	Bills	Bonds
Australia	1.29	2.24	1.32	2.45	3.23	5.85
Belgium	1.21	3.01	1.61	3.86	2.51	6.24
Denmark	3.08	3.58	2.18	3.50	2.80	7.13
Finland	0.64	3.22	0.63	4.86	2.61	5.76
France	-0.47	1.54	0.96	2.97	2.24	6.96
Germany	1.51	3.15	1.86	3.70	1.97	4.23
Italy	1.20	2.53	1.30	2.83	2.42	5.85
Japan	0.68	2.54	1.36	2.83	1.48	4.53
Netherlands	1.37	2.71	1.04	2.14	2.08	5.59
Norway	1.10	2.55	-0.26	1.94	1.50	5.62
Portugal	-0.01	2.23	-0.65	1.59	0.65	6.25
Spain	-0.04	1.41	-0.32	1.21	2.20	5.72
Sweden	1.77	3.25	0.82	2.71	1.52	6.60
Switzerland	0.89	2.41	0.12	2.33	0.33	3.35
UK	1.16	2.29	1.14	2.63	2.70	6.67
USA	2.23	2.85	1.43	2.77	1.91	5.90
Average, unweighted	1.14	2.61	0.91	2.77	2.01	5.77
Average, weighted	1.34	2.51	1.23	2.70	1.98	5.64

Note: Average annual real returns. Period coverage differs across countries. Consistent coverage within countries: each country-year observation used to compute the statistics in this table has data for both real bill and bond returns. The average, unweighted and average, weighted figures are respectively the unweighted and real-GDP-weighted arithmetic averages of individual country returns.

The effect of these major global shocks on individual countries seems to have resulted in a higher correlation of cross-country asset returns.

Turning to cross-sectional features, Table X shows country-specific safe asset returns for three samples: all years, post 1950, and post 1980. Here the experiences of a few countries stand out. In France, real bill returns have been negative when averaged over the full sample. In Portugal and Spain, they have been approximately zero. In Norway, the average return on bills has been negative for the post-1950 sample. However, most other countries have experienced reasonably similar returns on safe assets, in the ballpark of 1%–3%.

Aside from the investor perspective discussed above, safe rates of return have important implications for government finances, as they measure the cost of raising and servicing government debt. What matters for this is not the level of real return *per se*, but its comparison to real GDP growth, or $r^{safe} - g$. If the rate of return exceeds real GDP growth, $r^{safe} > g$, reducing the debt/GDP ratio requires continuous budget surpluses. When r^{safe} is less than g, however, a reduction in debt/GDP is possible even with the government running modest deficits. Existing evidence points to $r^{safe} < g$ being the norm rather than the exception, both in recent years and broader historical data (Ball, Elmendorf, and Mankiw, 1998; Mehrotra, 2017).

Figure XII: *Trends in the real return on safe assets and GDP growth*

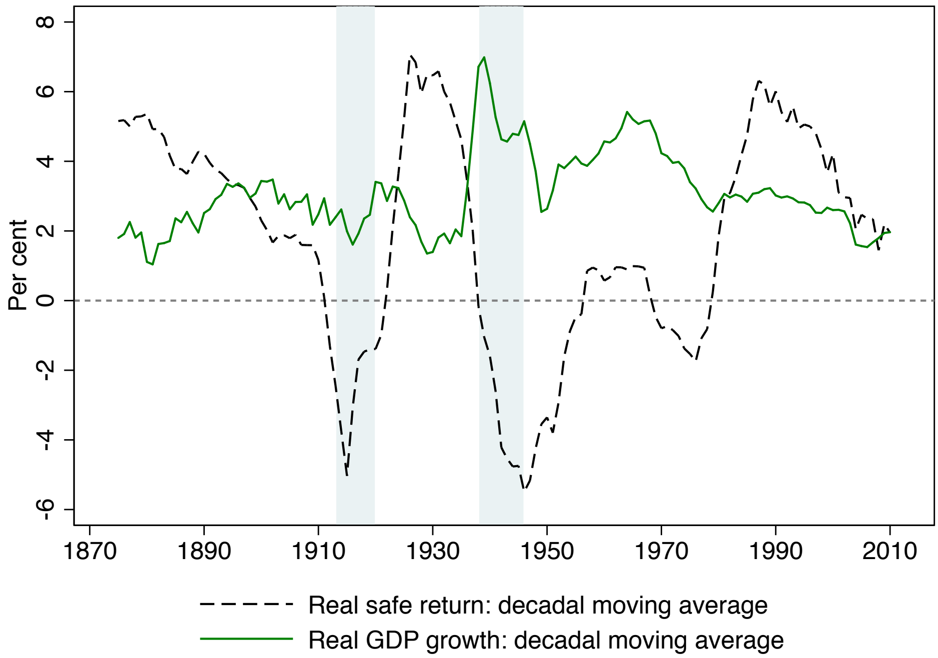

Note: Mean returns and GDP growth for 16 countries, weighted by real GDP. Decadal moving averages. The safe rate of return is an arithmetic average of bonds and bills.

Figure XII plots the representative "safe rate of return" as the arithmetic average of bond and bill returns (dashed line) alongside real GDP growth (solid line), again as decadal moving averages. Starting in the late 19th century, safe rates were higher than GDP growth, meaning that any government wishing to reduce debt had to run persistent budget surpluses. Indeed, this was the strategy adopted by Britain to pay off the debt incurred during the Napoleonic War (Crafts, 2016). The two world wars saw low real returns, but nevertheless a large debt accumulation to finance the wartime effort. The aftermath of these two wars, however, offered vastly different experiences for public finances. After WW1, safe returns were high and growth low, requiring significant budgetary efforts to repay the war debts. This was particularly difficult for many countries given the large interlocking reparations imposed by the Treaty of Versailles, and the turbulent macroeconomic environment at the time. After WW2, on the contrary, high growth and inflation helped greatly reduce the value of national debt, creating $r^{safe} - g$ gaps as large as –10 percentage points.

More recently, the Great Moderation saw a reduction in inflation rates and a corresponding increase in the debt financing burden, whereas the impact of $r^{safe} - g$ in the aftermath of the Global Financial Crisis remains broadly neutral, with the two rates roughly equal. On average throughout our sample, the real growth rate has been around 1 percentage point higher than the safe rate of return (3% growth versus 2% safe rate), meaning that governments could run small deficits without increasing the public debt burden.

In sum, real returns on safe assets have been quite low across the advanced countries over the last

150 years. In fact, for some countries, these returns have often been persistently negative. Periods of unexpected inflation, in war and peace, have often diluted returns, and flights to safety may have depressed safe returns even further in the more turbulent periods of global financial history. The low return for investors has, on the flipside, implied a low financing cost for governments, which was particularly important in reducing the debts incurred during WW2.

VI. Risky versus safe returns

Having established the general trends in each risky and safe asset class, we now turn to examine broader patterns of returns across the different asset classes. We start by comparing returns on risky and safe assets. Figure XIII depicts the trends in global safe and risky asset returns, again using decadal moving averages of GDP-weighted global return series.

The risky return in each country is a weighted average of housing and equity returns, with weights corresponding to equity market capitalization and housing wealth in each respective country. The safe return is a simple unweighted average of bonds and bills.[45] The left panel of Figure XIII shows the risky and safe asset returns, and the right panel depicts the risk premium, calculated as the risky minus safe difference.

As in Sections IV and V, the data presented in this Section measure ex post returns and risk premiums, inclusive of capital gains. For some of the debates that we touch on, however, a forward-looking expected return measure would have been preferable. Realized returns are likely to fall below ex ante return expectations during periods of large negative shocks, such as the two world wars, and rise above them in times of high capital gains, such as that between 1980 and today. Long-run data on expected returns are, however, difficult to obtain. Our focus on long-run trends, to an extent, allows us to look through some of the unexpected shocks that drive a wedge between ex ante and ex post returns. Nevertheless, the discussion in this section should be treated with a note of caution.

Both risky and safe returns were high during the 19th century but had been gradually declining in the run up to WW1, after which they declined sharply, as is to be expected. After the war, returns were recovering during the 1920s. From 1930 onwards, the risky return stayed high and relatively stable, whereas the safe return dropped sharply and remained low until the late 1970s, before increasing and falling back again during the past three decades. These findings have implications for current debates around secular stagnation and the pricing, or mispricing, of risk.

Secular stagnation is associated with low rates of return, driven by an excess of savings or a general unwillingness to borrow and invest. These in turn reflect a variety of potential factors, including: (1) lower rates of productivity growth; (2) lower fertility and mortality rates; (3) a decline in the relative price of investment goods; (4) greater firm-level market power; and (5) higher income

[45]For details on the construction of the weighted returns and the asset weights, see Section II.C and Online Appendix O.3. Online Appendix P further compares the portfolio-weighted returns to equally weighted returns, i.e., a simple average of housing and equity.

Figure XIII: *Global real risky vs. real safe return*

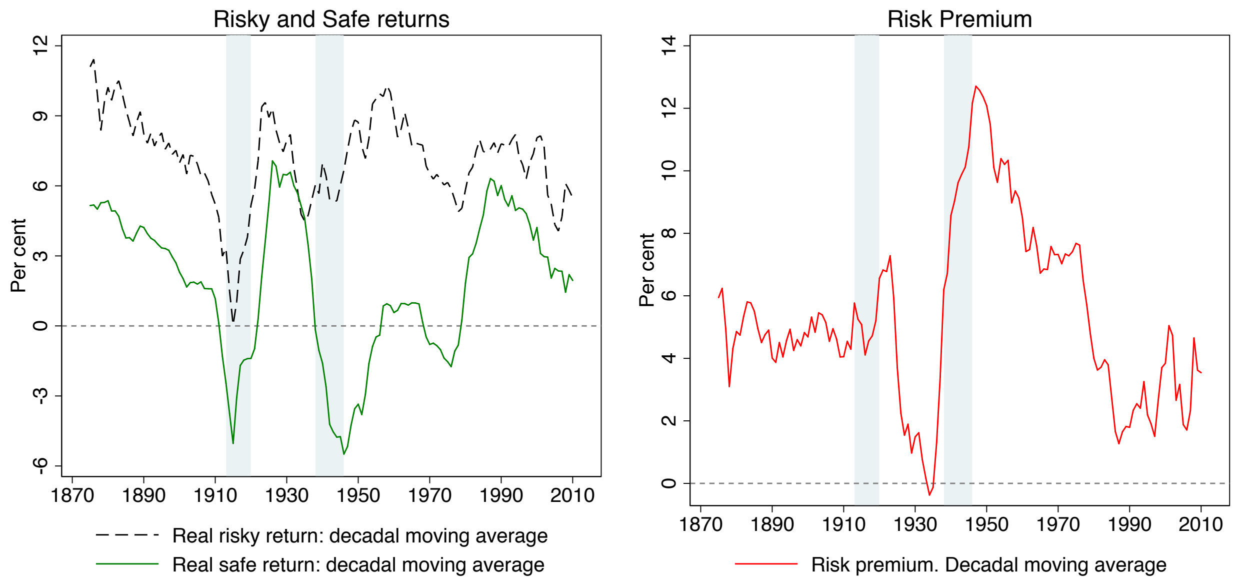

Note: Mean returns for 16 countries, weighted by real GDP. Decadal moving averages. Within each country, the real risky return is a weighted average of equities and housing, and safe return is a weighted average of bonds and bills. The within-country weights correspond to the shares of the respective asset in the country's wealth portfolio. Risk premium = risky return - safe return.

inequality (Eggertsson, Mehrotra, and Robbins, 2017; Rachel and Smith, 2015; Thwaites, 2015).

Indeed, we can see that the safe return fell sharply during the 1930s, when Hansen (1939) originally proposed the secular stagnation hypothesis. That time also coincided with a demographic bust and was preceded by a big rise in income inequality in the run-up to the Great Depression. The safe return has been falling again since the mid-1980s as many have noted.[46] Understandably, this has led some observers to suggest that advanced economies are again in danger of entering secular stagnation, e.g., Summers (2014), and Eggertsson and Mehrotra (2014).

But the picture changes radically when we consider the trend in risky returns in addition to safe returns. Unlike safe returns, risky returns have remained high and broadly stable through the best part of the last 100 years, and show little sign of a secular decline. Turning back to the trend for safe assets, even though the safe return has declined recently, much as it did at the start of our sample, it remains close to its historical average. These two observations call into question whether secular stagnation is quite with us. The high and stable risky return coupled with falling safe rates could also be consistent with the notion of a "safety trap" brought about by the relative shortage of safe assets (Caballero and Farhi, 2017). However with risk premiums still not far off their historical averages, the evidence for a safety trap is thus far also not clear-cut.

[46]Note that the safe interest rate—i.e. the component of the safe return that excludes capital gains, and is more relevant for the secular stagnation and safety trap debates, has also fallen sharply since 1980. However, like the bill rate in Figure X, it remains close to its historical average.

We now turn to examine the long-run developments in the ex post risk premium, i.e., the spread between safe and risky returns (right panel of Figure XIII). This spread was low and stable at around 5 percentage points before WW1. It rose slightly after WW1, before falling to an all-time low of near zero by around 1930. The decades following the onset of WW2 saw a dramatic widening in the risk premium, with the spread reaching its historical high of around 14 percentage points in the 1950s, before falling back to around its historical average.

Interestingly, the period of high risk premiums coincided with an era of few systemic banking crises. In fact, not a single such crisis occurred in our advanced-economy sample between 1946 and 1973. By contrast, banking crises appear to have been relatively more frequent when risk premiums were low. This finding speaks to the recent literature on the mispricing of risk around financial crises. Among others, Krishnamurthy and Muir (2017) argue that when risk is underpriced, i.e., risk premiums are excessively low, severe financial crises become more likely.

The long-run trends in risk premiums presented here seem to confirm this hypothesis. Online Appendix F further examines how these long-run movements in the risk premium, and the returns on the individual risky and safe asset classes, correspond to the changing monetary regimes, and finds, in accordance with Figure XIII, that the risk premium during the Bretton-Woods fixed exchange rate era was unusually high by historical standards, driven largely by the low, even negative, safe asset returns, but also by reasonably high housing returns.

Table XI zooms in to examine the evolution of safe and risky asset returns across different countries, as well as time periods. To enable a comparison with the aggregate trends in Figure XIII, we split the post-WW2 period into two subperiods: 1950–1980, when global risk premiums were high and global safe returns low, and post-1980, which saw an initial recovery, and subsequent decline in global safe returns.

The vast majority of countries in our sample follow similar patterns. The risky return is largely stable across time, even though it varies somewhat across countries: from just over 5% in Italy and Spain to 11% in Finland. Risk premiums were at or near their highest level in almost every country during the period 1950–1980, largely due to low returns on safe assets. The real safe rate of return was close to zero or negative for the majority of the countries in the sample, with the lowest level of –3.5% observed in Spain and Portugal, and only Belgium, Finland, and Germany experiencing robustly positive real returns. Meanwhile, risky returns were also somewhat above their long-run level in a number of countries, but the differences are relatively smaller than those for safe rates. The post-1980 period saw a recovery in safe returns across the board, with the recent downward trend not yet apparent in these longer-run period averages. Risky returns, meanwhile, were close to their historical levels in most countries, with only Japan experiencing a strong decline following the bursting of its asset price bubble in the 1990s.

We now turn to examine the correlations between risky and safe returns, which are displayed in Figure XIV. The top-left panel of this figure shows the rolling decadal correlation between the risky and safe returns, calculated as the average of rolling correlations in individual countries in a similar fashion to the calculations in Figure VIII. Throughout most of the historical period under

Table XI: *Real risky and safe asset returns across countries and time*

Country	Full Sample		1950–1980		Post 1980	
	Risky return	Safe return	Risky return	Safe return	Risky return	Safe return
Australia	6.96	1.77	6.51	-1.34	7.71	4.54
Belgium	8.31	1.82	9.68	1.05	7.99	4.38
Denmark	8.02	3.05	8.57	0.49	6.84	4.97
Finland	10.87	2.16	13.47	1.28	13.06	4.18
France	6.54	0.54	12.33	-1.15	6.61	4.60
Germany	7.90	3.34	7.00	1.77	5.20	3.10
Italy	5.32	2.28	7.08	-0.83	5.21	4.14
Japan	6.79	1.29	10.86	0.05	4.81	3.00
Netherlands	7.30	1.31	10.26	-0.89	7.42	3.83
Norway	7.96	1.59	7.75	-2.34	10.65	3.56
Portugal	6.46	0.45	5.19	-3.30	7.41	3.45
Spain	5.39	0.68	7.27	-3.56	5.46	3.96
Sweden	8.52	2.35	8.67	-1.12	11.42	4.06
Switzerland	6.51	1.57	6.07	0.25	7.76	1.84
UK	6.35	1.51	8.33	-1.36	7.66	4.69
USA	7.12	1.92	6.44	-0.32	7.28	3.91
Average, unweighted	7.44	1.88	8.48	-0.81	7.65	3.89
Average, weighted	7.22	1.89	7.88	-0.56	6.66	3.81

Note: Average annual real returns. Real risky return is a weighted average of equity and housing, and safe return is a weighted average of bonds and bills. The weights correspond to the shares of the respective asset in the country's wealth portfolio. Period coverage differs across countries. Coverage is consistent within countries: each country-year observation used to compute the statistics in this table has data for both the risky and safe return. The average, unweighted and average, weighted figures are respectively the unweighted and real-GDP-weighted arithmetic averages of individual country returns.

consideration, risky and safe returns had been positively correlated. In other words, safe assets have not generally provided a hedge against risk since safe returns were low when risky returns were low—in particular during both world wars—and vice versa. This positive correlation has weakened over the more recent decades, and turned negative from the 1990s onwards. This suggests that safe assets have acted as a better hedge for risk during both the Great Moderation and the recent Global Financial Crisis.

The top-right panel of Figure XIV shows the comovement of risky and safe nominal returns with inflation. Mirroring our findings presented in the preceding sections, safe returns have tended to comove more strongly with inflation, particularly during the post-WW2 period. Moving to cross-country correlations depicted in the bottom two panels of Figure XIV, historically safe returns in different countries have been more correlated than risky returns. This has reversed over the past decades, however, as cross-country risky returns have become substantially more correlated. This seems to be mainly driven by a remarkable rise in the cross-country correlations in risk premiums, depicted in the bottom-right panel of Figure XIV. This increase in global risk comovement may pose

Figure XIV: *Correlations across risky and safe asset returns*

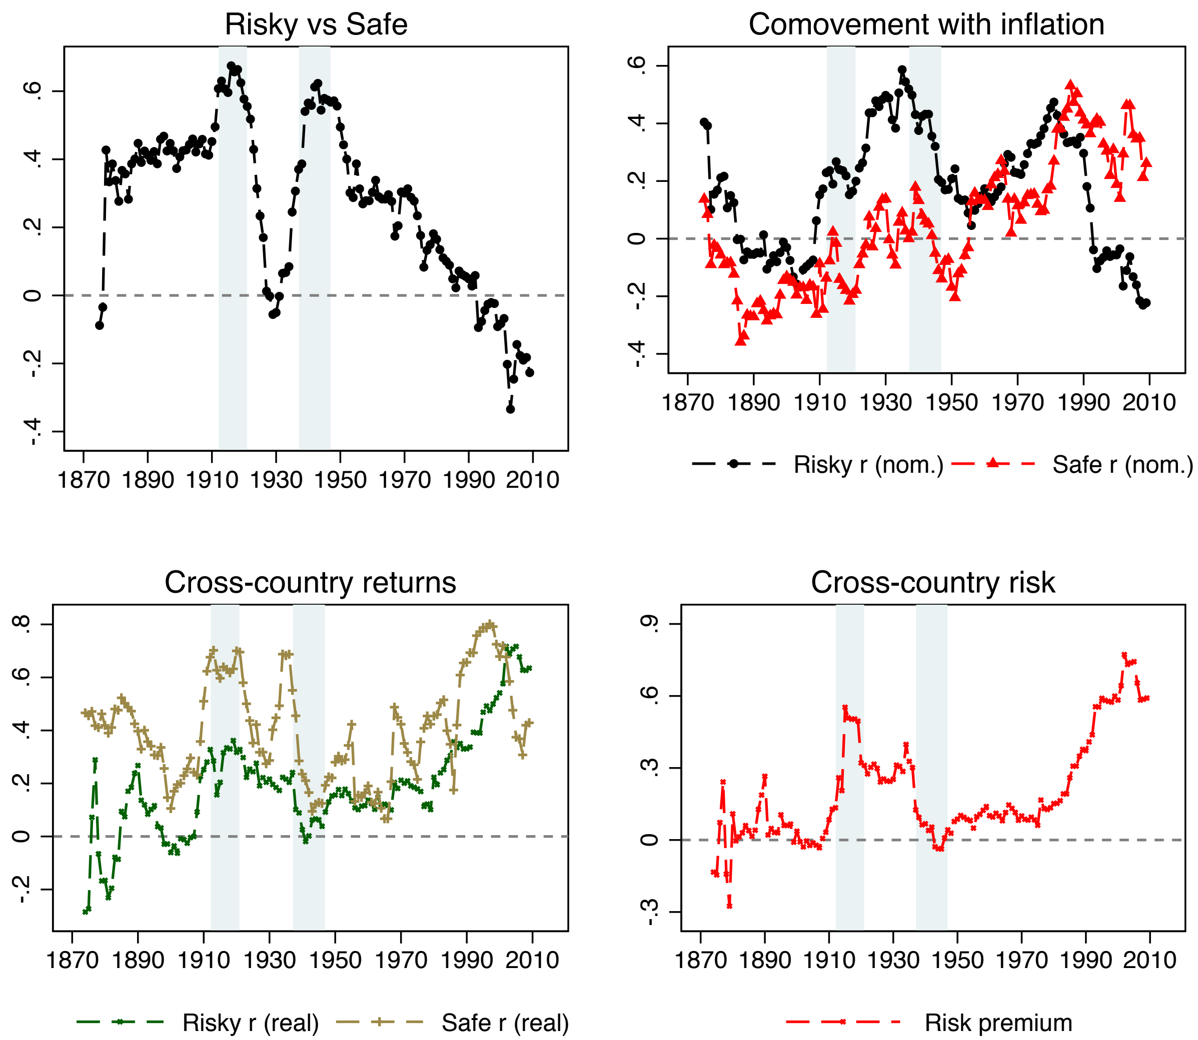

Note: Rolling decadal correlations. The global correlation coefficient is the average of individual countries for the rolling window. Cross-country correlation coefficient is the average of all country pairs for a given asset class. Country coverage differs across time periods.

new challenges to the risk-bearing capacity of the global financial system, a trend consistent with other macro indicators of risk-sharing (Jordà, Schularick, and Taylor, 2017).

VII. *r* versus *g*

Our analysis also provides insights into the debate on inequality. Piketty (2014) and Piketty and Zucman (2014) argue that inequality and wealth-to-income ratios in advanced economies have followed a U-shaped pattern over the past century and a half. They further hypothesize that wealth inequality may continue to rise in the future, along with a predicted decline in the rate of economic

Figure XV: *The real return on wealth and real GDP growth*

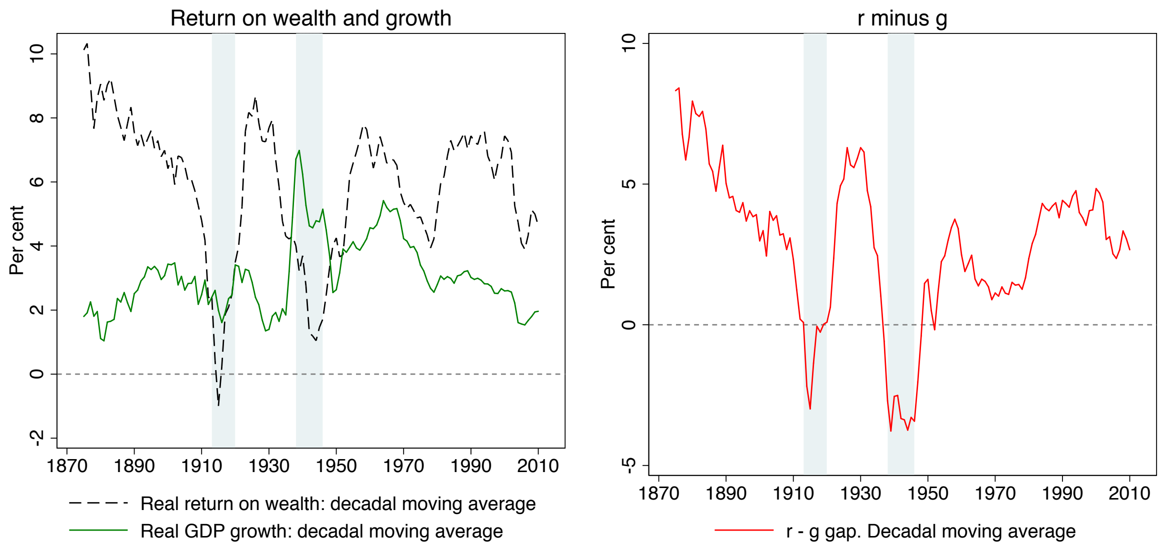

Note: Mean returns and real GDP growth for 16 countries, weighted by real GDP. Decadal moving averages. Within each country, the real return on wealth is a weighted average of bonds, bills, equity, and housing. The within-country weights correspond to the shares of the respective asset in each country's wealth portfolio.

growth. The main theoretical argument for this comes about from a simple relation: $r > g$. In their approach, a higher spread between the real rate of return on wealth, denoted r, and the rate of real GDP growth, g, tends to magnify the steady-state level of wealth inequality. Benhabib and Bisin (2016) show that in a wide class of models featuring stochastic returns to wealth, a higher gap between r and g increases the Pareto index of the steady-state wealth distribution, making it more unequal.

Of course, this is not the only channel through which rates of return can impact the wealth distribution. Rate of return differentials between asset classes can affect the wealth distribution if there are systematic differences in the portfolio composition between rich and poor households as Kuhn, Schularick, and Steins (2017) show, or if rates of returns vary with portfolio size as stressed by Piketty (2014). Studying administrative Swedish data, Bach, Calvet, and Sodini (2016) find that wealthy households earn higher returns on their portfolios, and Fagereng, Guiso, Malacrino, and Pistaferri (2016) use Norwegian tax data to document substantial heterogeneity in wealth returns. Rates of return on wealth are beginning to receive attention in the theoretical literature. For instance, Benhabib and Bisin (2016) point to return differences of assets as one potential channel to explain diverging trends between income and wealth inequality, and Garbinti, Goupille-Lebret, and Piketty (2017) show that asset price effects played an important role in shaping the French wealth distribution over the past 200 years. Further, wealth inequality may depend not only on the magnitude of r in relation to g, but also on return volatility. Higher return volatility can increase the dispersion of wealth outcomes, and make the distribution of wealth more unequal.

To bring our data to bear on these debates, we construct a measure of the world's real return on wealth as a weighted average of real returns on bonds, equities, and housing—reflecting the typical portfolio of a private household end-investor. We then compare this measure to the rate of real GDP growth of economies over the long-run. Importantly, our approach differs from Piketty (2014) in that we rely on annual returns from observed market prices and yields for each individual asset class, rather than implicit returns derived from aggregate balance sheet data at selected benchmark dates. This, we think, is more robust and provides a vital cross check for the core argument.

Similarly to the risky returns in Section VI, we weight the individual returns by the size of the respective asset portfolio: stock market capitalization, housing wealth, and public debt (divided equally between bonds and bills).[47] Figure XV displays the long-run trends in the global real rate of return on wealth (dashed line) and the global real GDP growth rate (solid line) since the late 19th century, again using decadal moving averages of GDP-weighted data.

Our data show that the trend long-run real rate of return on wealth has consistently been *much* higher than the real GDP growth rate. Over the past 150 years, the real return on wealth has substantially exceeded real GDP growth in 13 decades, and has only been below GDP growth in the two decades corresponding to the two world wars. That is, in peacetime, r has always exceeded g. The gap between r and g has been persistently large. Since 1870, the weighted average return on wealth (r) has been about 6.0%, compared to a weighted average real GDP growth rate (g) of 3.0%, with the average $r - g$ gap of 3.0 percentage points, which is about the same magnitude as the real GDP growth rate itself. The peacetime gap between r and g has been larger still, averaging around 3.8 percentage points.

Table XII shows how the rate of return on wealth and the GDP growth rate have varied across different countries and time periods. Despite some variation, the positive gap between r and g is a persistent feature of the data: r is bigger than g in every country and every time period that we consider. The last few decades prior to the Global Financial Crisis saw a general widening of this gap, mirroring the aggregate pattern shown in Figure XV.

As previously discussed, returns on housing play an important part in this story—but with scant data until now, their exact role was unclear. The high level of housing returns that we have uncovered serves to push up the level of r, and thus, potentially, wealth inequality. But what is the counterfactual? We need to bear in mind that housing wealth is more equally distributed than, for instance, equities (see, e.g., Kuhn et al., 2017), and returns on housing are less volatile than those on shares—with both of these factors serving to flatten the distribution of wealth changes, making the overall impact of housing returns on wealth inequality less clear-cut and offering substantial scope for further research.

Rognlie (2015) notes that recent trends in wealth and income could be influenced primarily by what has happened in housing. Real house prices have experienced a dramatic increase in the past

[47]For details on the construction of the weighted returns and the asset weights, see Section II.C and Online Appendix O.3. Online Appendix P further compares the portfolio-weighted returns to equally weighted returns, with the equally weighted return on wealth a simple average of equity, housing, and bonds.

Table XII: *The return on wealth and GDP growth across countries and time*

Country	Full Sample		Post 1950		Post 1980	
	Return on wealth	GDP growth	Return on wealth	GDP growth	Return on wealth	GDP growth
Australia	5.91	3.51	7.39	3.73	7.53	3.19
Belgium	6.38	2.32	7.29	2.68	6.90	2.17
Denmark	7.37	2.70	7.21	2.51	6.62	1.60
Finland	9.76	3.49	11.92	3.16	11.81	2.16
France	4.92	2.55	7.76	3.17	6.29	1.92
Germany	7.07	2.81	5.26	2.80	4.72	2.40
Italy	5.08	3.82	5.07	3.30	5.01	1.37
Japan	5.59	4.18	6.35	4.20	4.23	2.09
Netherlands	5.33	3.16	6.67	3.21	6.71	2.29
Norway	6.86	3.06	7.67	3.45	9.35	2.80
Portugal	5.87	3.39	5.65	3.48	6.99	2.13
Spain	4.58	3.21	5.50	4.03	5.34	2.56
Sweden	7.41	2.89	8.69	2.86	9.87	2.36
Switzerland	5.63	2.33	5.98	2.69	7.03	1.95
UK	4.75	2.09	5.90	2.49	7.23	2.45
USA	6.03	3.38	5.91	3.33	6.58	2.82
Average, unweighted	6.30	2.86	6.92	3.23	7.01	2.26
Average, weighted	5.98	3.04	6.09	3.33	6.08	2.48

Note: Average annual real returns. Real return on wealth is a weighted average of bonds, bills, equity, and housing. The weights correspond to the shares of the respective asset in each country's wealth portfolio. Period coverage differs across countries. Coverage is consistent within countries: each country-year observation used to compute the statistics in this table has data for both the real return on wealth and the real GDP growth rate. The average, unweighted and average, weighted figures are respectively the unweighted and real-GDP-weighted arithmetic averages of individual country returns.

40 years, coinciding with the rapid expansion of mortgage lending (Jordà, Schularick, and Taylor, 2015, 2016; Knoll, Schularick, and Steger, 2017). This is very much evident from Table IX. Measured as a ratio to GDP, rental income has been growing, as Rognlie (2015) argues. However, the rental yield has declined slightly—given the substantial increase in house prices—so that total returns on housing have remained pretty stable, as we have discussed. In this sense, recent housing trends have diverged little.

Our data allow us to more formally examine whether movements in the $r - g$ gap are more closely related to return fluctuations or movements in the real GDP growth rate. Online Appendix Q and Table A.26 document the correlations between $r - g$ and g, and $r - g$ and r across different time horizons, for the full sample and the period after 1950. Overall, the correlation between $r - g$ and g is negative, and somewhat stronger at longer horizons, with the correlation coefficients ranging between –0.2 and –0.6 depending on the historical subperiod and time window. At the same time, the $r - g$ gap is even more robustly related to changes in the return on wealth r, with a positive correlation between the two and a correlation coefficient of around 0.9, both over the short and long run. This suggests that both falling GDP growth and higher returns would tend to increase the $r - g$ gap, although historically much of the changes in $r - g$ have come about from movements in the return on wealth. During peacetime r has been quite stable, and so has been the $r - g$ gap.

Since the 1970s, the stable and high levels of the rate of return on wealth have coincided with high and rising wealth-to-income ratios (see Piketty and Zucman, 2014, and Online Appendix Figure A.7). Together, these two facts have meant that the capital share of GDP has increased across advanced economies (Karabarbounis and Neiman, 2014). A large part of these high returns, and of the increase in wealth ratios, can be attributed to high capital gains on risky assets, namely housing and equity. Rognlie (2015) argues that house prices have played an important role in the evolution of wealth-to-income ratios in the US. Kuvshinov and Zimmermann (2018) show that most of the recent increase in the value of listed firms in our cross-country sample is accounted for by higher equity valuations.

These high capital gains in recent decades have allowed the stock of measured wealth to rise without running into diminishing returns. Understanding the drivers behind these long-run trends in returns and valuations seems key to disentangling the underlying causes behind the recent upsurge in wealth, inequality, and the capital share of income.

VIII. Conclusion

In this paper we provide an investigation of the long history of advanced economy asset returns for all the major categories of the investable wealth portfolio. Our work brings new stylized facts to light and rigorously documents many broad patterns that have stimulated so much research in core areas of economics and finance over the last two centuries.

The returns to risky assets and risk premiums have been high and stable over the past 150 years. Substantial diversification opportunities exist between risky asset classes, and across countries. Arguably the most surprising result of our study is that long-run returns on housing and equity look remarkably similar. Yet while returns are comparable, residential real estate is less volatile on a national level, opening up new and interesting risk premium puzzles.

Our research speaks directly to the relationship between r, the rate of return on wealth, and g, the growth rate of the economy, that figures prominently in the current debate on inequality. One robust finding in this paper is that $r \gg g$: globally, and across most countries, the weighted rate of return on capital was twice as high as the growth rate in the past 150 years.

These and other discoveries can provide a rich agenda for future research, by us and by others. Many issues remain to be studied, among them determining the particular fundamentals that drive the returns on each of the asset classes in typical economies. For now, we hope our introduction of this new compilation of asset return data can provide the evidentiary basis for new lines of exploration in years to come.

Supplementary Material

An Online Appendix for this article can be found at The Quarterly Journal of Economics online. Data and code replicating tables and figures in this article can be found in Jordà et al. (2019), in the Harvard Dataverse, doi: 10.7910/DVN/GGDQGJ.

References

Akbulut-Yuksel, Mevlude. 2014. Children of War. The Long-Run Effects of Large-Scale Physical Destruction and Warfare on Children. *Journal of Human Resources* 49(3): 634–662.

Annaert, Jan, Frans Buelens, Ludo Cuyvers, Marc De Ceuster, Marc Deloof, and Ann De Schepper. 2011. Are Blue Chip Stock Market Indices Good Proxies for All-Shares Market Indices? The Case of the Brussels Stock Exchange 1833–2005. *Financial History Review* 18(3): 277–308.

Annaert, Jan, Frans Buelens, and Marc De Ceuster. 2012. New Belgian Stock Market Returns: 1832–1914. *Explorations in Economic History* 49(2): 189–204.

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2016. Rich Pickings? Risk, Return, and Skill in the Portfolios of the Wealthy. CEPR Discussion Paper 11734.

Ball, Laurence, Douglas W. Elmendorf, and N. Gregory Mankiw. 1998. The Deficit Gamble. *Journal of Money, Credit and Banking* 30(4): 699–720.

Barro, Robert J., and José Ursúa. 2008. Macroeconomic Crises since 1870. *Brookings Papers on Economic Activity* 39(1): 225–350.

Benhabib, Jess, and Alberto Bisin. 2016. Skewed Wealth Distributions: Theory and Empirics. NBER Working Paper 21924.

Blancheton, Bertrand, Hubert Bonin, and David Le Bris. 2014. The French Paradox: A Financial Crisis During the Golden Age of the 1960s. *Business History* 56(3): 391–413.

Brailsford, Tim, John C. Handley, and Krishnan Maheswaran. 2012. The Historical Equity Risk Premium in Australia: Post-GFC and 128 Years of Data. *Accounting and Finance* 52(1): 237–247.

Caballero, Ricardo J., and Emmanuel Farhi. 2017. The Safety Trap. *Review of Economic Studies.* Forthcoming.

Campbell, John Y. 1999. Asset Prices, Consumption, and the Business Cycle. volume 1 of *Handbook of Macroeconomics*, chapter 19, 1231–1303. Elsevier.

Cheng, Ping, Zhenguo Lin, and Yingchun Liu. 2008. The Real Estate Risk Premium Puzzle: A Solution. Working paper.

Crafts, Nicholas. 2016. Reducing High Public Debt Ratios: Lessons from UK Experience. *Fiscal Studies* 37(2): 201–223.

Diefendorf, Jeffry M. 1993. *In the Wake of War: The Reconstruction of German Cities After World War II*. Oxford: Oxford University Press.

Dimson, Elroy, Paul Marsh, and Mike Staunton. 2009. *Triumph of the Optimists: 101 Years of Global Investment Returns*. Princeton, N.J.: Princeton University Press.

Eggertsson, Gauti B., and Neil R. Mehrotra. 2014. A Model of Secular Stagnation. NBER Working Paper 20574.

Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins. 2017. A Model of Secular Stagnation: Theory and Quantitative Evaluation. NBER Working Paper 23093.

Eisfeldt, Andrea, and Andrew Demers. 2015. Total Returns to Single Family Rentals. NBER Working Paper 21804.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2016. Heterogeneity and Persistence in Returns to Wealth. NBER Working Paper 22822.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh. 2017. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy* 125(1): 140–223.

Fishback, Price V., and Trevor Kollmann. 2015. Hedonic Housing Indexes During the Great Depression. Unpublished manuscript.

Fox, Ryan, and Peter Tulip. 2014. Is Housing Overvalued? RBA Research Discussion Paper 2014-06.

Francis, Jack Clark, and Roger G. Ibbotson. 2009. Contrasting Real Estate with Comparable Investments, 1978 to 2008. *The Journal of Portfolio Management* 36(1): 141–155.

Garbinti, Bertrand, Jonathan Goupille-Lebret, and Thomas Piketty. 2017. Accounting for Wealth Inequality Dynamics: Methods, Estimates and Simulations for France (1800–2014). CEPR Discussion Paper 11848.

Giacomini, Emanuela, David C. Ling, and Andy Naranjo. 2015. Leverage and Returns: A Cross-Country Analysis of Public Real Estate Markets. *Journal of Real Estate Finance and Economics* 51(2): 125–159.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel. 2015. Very Long-Run Discount Rates. *Quarterly Journal of Economics* 130(1): 1–53.

Goldsmith, Raymond W. 1985a. *Comparative National Balance Sheets: A Study of Twenty Countries, 1688–1978*. Chicago, Ill.: University of Chicago Press.

Goldsmith, Raymond W. 1985b. *Comparative National Balance Sheets: A Study of Twenty Countries, 1688–1978*. Chicago, Ill.: University of Chicago Press.

Grossman, Richard S. 2017. Stocks for the Long Run: New Monthly Indices of British Equities, 1869–1929. CEPR Discussion Paper 12042.

Hansen, Alvin H. 1939. Economic Progress and Declining Population Growth. *American Economic Review* 29(1): 1–15.

Holston, Kathryn, Thomas Laubach, and John C. Williams. 2017. Measuring the Natural Rate of Interest: International Trends and Determinants. *Journal of International Economics* 108(S1): 59–75.

Homer, Sidney, and Richard E. Sylla. 2005. *A History of Interest Rates*. Hoboken, N.J.: Wiley, 4th edition.

Ilmanen, Antti. 2011. *Expected Returns: An Investor's Guide to Harvesting Market Rewards*. Hoboken, N.J.: Wiley.

Jones, Charles M. 2002. A Century of Stock Market Liquidity and Trading Costs. Working paper, New York.

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor. 2019. Replication Data for: 'The Rate of Return on Everything, 1870-2015'. Harvard Dataverse, doi:10.7910/DVN/GGDQGJ.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2015. Betting the House. *Journal of International Economics* 96(S1): 2–18.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2016. The Great Mortgaging: Housing Finance, Crises and Business Cycles. *Economic Policy* 31(85): 107–152.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2017. Macrofinancial History and the New Business Cycle Facts. In *NBER Macroeconomics Annual 2016, Volume 31*, edited by Martin Eichenbaum, Jonathan A. Parker, pp. 213–263. Chicago, Ill.: University of Chicago Press.

Karabarbounis, Loukas, and Brent Neiman. 2014. The Global Decline of the Labor Share. *Quarterly Journal of Economics* 129(1): 61–103.

Knoll, Katharina. 2017. *Our Home in Days Gone By: Housing Markets in Advanced Economies in Historical Perspective*. Ph.D. thesis, Free University of Berlin. Chapter 3. As Volatile As Houses: Return Predictability in International Housing Markets, 1870–2015.

Knoll, Katharina, Moritz Schularick, and Thomas M. Steger. 2017. No Price like Home: Global

House Prices, 1870–2012. *American Economic Review* 107(2): 331–352.
Krishnamurthy, Arvind, and Tyler Muir. 2017. How Credit Cycles Across a Financial Crisis. NBER Working Paper 23850.
Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins. 2017. Income and Wealth Inequality in America , 1949–2013. CEPR Discussion Paper 20547.
Kuvshinov, Dmitry, and Kaspar Zimmermann. 2018. The Big Bang: Stock Market Capitalization in the Long Run. EHES Working Paper 136.
Kuznets, Simon. 1941. *National Income and its Composition*. Cambridge, Mass.: National Bureau of Economic Research.
Le Bris, David. 2012. Wars, Inflation and Stock Market Returns in France, 1870–1945. *Financial History Review* 19(3): 337–361.
Le Bris, David, and Pierre-Cyrille Hautcoeur. 2010. A Challenge to Triumphant Optimists? A Blue Chips Index for the Paris Stock Exchange, 1854–2007. *Financial History Review* 17(2): 141–183.
Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan. 2013. The Wealth-Consumption Ratio. *Review of Asset Pricing Studies* 3(1): 38–94.
Mehra, Rajnish, and Edward C. Prescott. 1985. The Equity Premium: A Puzzle. *Journal of Monetary Economics* 15(2): 145–161.
Mehrotra, Neil. 2017. Debt Sustainability in a Low Interest Rate World. Hutchins Center Working Paper 32.
Moskowitz, Tobias J., and Annette Vissing-Jørgensen. 2002. The Returns to Entrepreneurial Investment: A Private Equity Premium Puzzle? *American Economic Review* 92(4): 745–778.
Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa. 2013. Crises and Recoveries in an Empirical Model of Consumption Disasters. *American Economic Journal: Macroeconomics* 5(3): 35–74.
Nielsen, Steen, and Ole Risager. 2001. Stock Returns and Bond Yields in Denmark, 1922–1999. *Scandinavian Economic History Review* 49(1): 63–82.
OECD. 2012. *OECD Economic Surveys: European Union 2012*. Paris: OECD Publishing.
Office for National Statistics. 2016. Changes to National Accounts: Imputed Rental. https://www.ons.gov.uk/economy/nationalaccounts/uksectoraccounts/articles/changestonationalaccounts/imputedrental.
Piketty, Thomas. 2014. *Capital in the Twenty-First Century*. Cambridge, Mass.: Harvard University Press.
Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. Distributional National Accounts: Methods and Estimates for the United States. *Quarterly Journal of Economics* 133(2): 553–609.
Piketty, Thomas, and Gabriel Zucman. 2014. Capital is Back: Wealth-Income Ratios in Rich Countries 1700–2010. *Quarterly Journal of Economics* 129(3): 1255–1310.
Quinn, Dennis P., and Hans-Joachim Voth. 2008. A Century of Global Equity Market Correlations. *American Economic Review* 98(2): 535–540.
Rachel, Lukasz, and Thomas Smith. 2015. Secular Drivers of the Global Real Interest Rate. Bank of England Working Paper 571.
Rognlie, Matthew. 2015. Deciphering the Fall and Rise in the Net Capital Share. *Brookings Papers on Economic Activity* 46(1): 1–69.
Ronge, Ulrich. 2002. *Die Langfristige Rendite Deutscher Standardaktien: Konstruktion eines Historischen Aktienindex ab Ultimo 1870 bis Ultimo 1959*. Frankfurt am Main: Lang.
Ruff, Jon. 2007. Commercial Real Estate: New Paradigm or Old Story? *Journal of Portfolio Management* 27.
Shiller, Robert J. 1981. Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? *American Economic Review* 71(3): 421–436.

Shiller, Robert J. 2000. *Irrational Exuberance*. Princeton, N.J.: Princeton University Press.
Shilling, James D. 2003. Is There a Risk Premium Puzzle in Real Estate? *Real Estate Economics* 31(4): 501–525.
Shumway, Tyler. 1997. The Delisting Bias in CRSP Data. *Journal of Finance* 52(1): 327–340.
Shumway, Tyler, and Vincent A. Warther. 1999. The Delisting Bias in CRSP's Nasdaq Data and Its Implications for the Size Effect. *Journal of Finance* 54(6): 2361–2379.
Simonnet, François, Georges Gallais-Hamonno, and Pedro Arbulu. 1998. Un Siècle de Placement Immobilier. L'exemple de La Fourmi Immobilière. *Journal de la Société Française de Statistique* 139(2): 95–135.
Summers, Lawrence H. 2014. US Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound. *Business Economics* 49(2): 65–73.
Thwaites, Gregory. 2015. Why are Real Interest Rates So Low? Secular Stagnation and the Relative Price of Investment Goods. Bank of England Working Paper 564.
Zimmermann, Kaspar. 2017. Breaking Banks? Bank Profitability and Monetary Policy. Unpublished.