

NBER WORKING PAPER SERIES

THE TOTAL RISK PREMIUM PUZZLE

Òscar Jordà
Moritz Schularick
Alan M. Taylor

Working Paper 25653
<http://www.nber.org/papers/w25653>

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

This work is part of a larger project kindly supported by research grants from the European Research Council (ERC) and the Institute for New Economic Thinking, and we are grateful for this support. We are particularly thankful to Johannes Matschke and Sarah Quincy for outstanding research assistance. We received helpful comments from Thomas Mertens. 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 National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w25653.ack>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Òscar Jordà, Moritz Schularick, and Alan M. Taylor. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Total Risk Premium Puzzle
Òscar Jordà, Moritz Schularick, and Alan M. Taylor
NBER Working Paper No. 25653
March 2019
JEL No. E44,G12,G15,N20

ABSTRACT

The risk premium puzzle is worse than you think. Using a new database for the U.S. and 15 other advanced economies from 1870 to the present that includes housing as well as equity returns (to capture the full risky capital portfolio of the representative agent), standard calculations using returns to total wealth and consumption show that: housing returns in the long run are comparable to those of equities, and yet housing returns have lower volatility and lower covariance with consumption growth than equities. The same applies to a weighted total-wealth portfolio, and over a range of horizons. As a result, the implied risk aversion parameters for housing wealth and total wealth are even larger than those for equities, often by a factor of 2 or more. We find that more exotic models cannot resolve these even bigger puzzles, and we see little role for limited participation, idiosyncratic housing risk, transaction costs, or liquidity premiums.

Òscar Jordà
Economic Research, MS 1130
Federal Reserve Bank of San Francisco
San Francisco, CA 94105
and University of California, Davis
oscar.jorda@sf.frb.org

Moritz Schularick
University of Bonn
Department of Economics
Adenauerallee 24-42
53113 Bonn
Germany
and CEPR
moritz.schularick@uni-bonn.de

Alan M. Taylor
Department of Economics and
Graduate School of Management
University of California
One Shields Ave
Davis, CA 95616-8578
and CEPR
and also NBER
amtaylor@ucdavis.edu

1. INTRODUCTION

If economics is galvanized by the raw conflict between pure theories and brute facts, then the battle over the equity premium puzzle has arguably energized as much intellectual warfare as almost any other empirical anomaly. Over four decades it has stood as a rebuke to many macro-financial theories, and consumption-based asset pricing in particular, triggering an avalanche of research.¹

Thus far, the near-universal reaction has been to come up with more elaborate theories taking the magnitude of the puzzle as given. This paper is different. We bring new evidence from historical data to this research agenda along three dimensions. First, international evidence from a wider spectrum of advanced economies. Second, a longer time span that includes wars, financial and currency crises, and other rare but dramatic events that greatly influence how one should think about consumption risk. Finally, we bring in returns to housing, a large category of investable assets for the representative consumer, getting us closer to a more accurate measure of the investor's total wealth portfolio returns and its relation to consumption growth. We focus on measurement and how alternative theories fare against this new evidence.

Methodologically, the paper follows a long lineage. According to textbook theory, a risky asset should earn a risk premium, an excess return relative to a safe asset, which depends on the statistical properties of the risky return and the investor's discount factor—more specifically the covariance of the two, as we recap below. As Shiller (1982) demonstrated, by deriving volatility bounds, if the stochastic discount factor is derived from a standard power utility function, then the high excess return on equities combined with a plausible risk aversion parameter would imply a discount factor that was inadmissible as it was many times more volatile than could be generated by actual U.S. post-WW2 time-series data on consumption. Mehra and Prescott (1985) worked in a two-state model, and reached a formulation which cast the result in a different light: given a plausible risk aversion parameter and the measured consumption data, the excess return was inexplicably large, and they attached to this conundrum its enduring name, the equity premium puzzle.²

A huge literature has tried a variety of approaches to “solve” the puzzle. One major approach has been to alter the consumption and return data fed into the model. For example, some narrow the excess return or increase the covariance by various means such as choosing longer-duration safe assets (U.S. bonds rather than bills), or considering longer spans of pre-WW2 U.S. economic history where negative skew from “rare disasters” (e.g., the 1930s) change key statistical moments; some introduce non-U.S. time series data to further address the arguably quite serious small-sample and survivorship bias derived from studying only the equity market of the global capitalist hegemon in its heyday. The other major approach has been to alter the model itself, specifically by adding

¹Articles too numerous to list address the point. For major surveys see Breeden et al. (2015ab), Campbell (1999; 2005; 2017), Cochrane (2008), Cochrane and Hansen (1992), and Mehra (2008).

²On the history of thought, Cochrane (2008, 261–266) notes that many papers from that era found unrealistically high risk aversion was needed to fit consumption-based models to the data, e.g., Grossman and Shiller (1981), Shiller (1982), Hansen and Singleton (1983), as well as Mehra and Prescott (1985). Cochrane also offers an interesting discussion of why the latter paper got much more credit than the others.

“exotic” preference assumptions to derive the stochastic discount factor, and to amplify the necessary moments to achieve a better fit with a plausible level of risk aversion. For example, some impose habits in consumption, so that changes are penalized more heavily and underlying marginal utilities become more volatile; some impose recursive preferences and long-run risks to growth.

The central contribution of this paper is to bring insights from the workhorse models to bear again once we re-measure returns to the hypothetical representative investor’s total portfolio including both equities and housing. Note that to keep all else equal, our benchmark is the standard representative agent model. In that model, the equilibrium for the agent is, ipso facto, a buy-and-hold, no-trade equilibrium. We take no stand on whether this is the right model of the world, which is a different discussion. Of course, in reality key assumptions fail: assets trade, transactions costs intrude, etc. But this is true for equities as well as housing, and we shall take some time to ponder these issues even if they must be left to a different paper. Rather, our aim is to ask how well this widely-used canonical model works once returns to wealth are properly measured to include a proxy for total wealth that includes the hitherto ignored role of housing.

We make that critical but arduous correction and arrive at a striking conclusion. We rely on a new database for the U.S. and 15 other advanced economies, from 1870 to the present, including total returns on housing as well as equities (Jordà, Knoll, Kuvshinov, Schularick, and Taylor, forthcoming). We do standard calculations on returns and consumption and show that housing returns are comparable to those of equities, and yet housing returns have even lower covariance with consumption growth than equities. The same also holds for a total wealth portfolio, and over a range of horizons. In other words, the total risk premium puzzle is even bigger: The implied risk aversion parameters for housing wealth and total wealth are considerably larger than those for equities alone, often by a factor of 2 or more. The total risk premium puzzle is even more pronounced than the equity risk premium puzzle, given the implausibly large degree of risk aversion implied by the data.

The outline of the paper is as follows. Section 2 presents summary data to fix facts and set the scene. Section 3 reviews the key literature on equity risk premiums, while bringing in studies of housing risk premiums that use the standard asset pricing framework. Section 4 reviews our new data on housing and equity returns. Section 5 presents our key results: asset pricing moments for equities, housing, and the total portfolio are computed, and implied risk aversion parameters are derived. Section 6 shows that several well-known “exotic” variants of the benchmark consumption-based asset pricing model, which partly resolve the equity premium puzzle, struggle to explain the housing and total risk premium puzzles uncovered here. This section also discusses potential deviations from the canonical representative agent model such as limited participation, transaction costs, idiosyncratic risk, and liquidity, how these factors might not much mitigate the puzzles we uncover. Section 7 concludes, and poses a question for future research: if asset pricing research has struggled for decades with equity-return implied risk aversion parameters too large for comfort, how can we proceed now, confronted with housing-return and total-return implied risk aversion parameters that are 2 or 3 times larger still?

Our paper opens a new front in the risk premium conflict, with challenging findings for asset

pricing research based on a new dataset with unprecedented spatial and temporal coverage for both key asset returns, equity and housing. A key contribution is not just to highlight that housing returns are likely to worsen long-standing puzzles in the literature, but to actually quantify how much worse things get using the near universe of alternative explanations that have been explored in the literature. Whatever future scholarship decides, our new findings and data will be available to those who aim to better understand the behavior of returns to wealth.

2. TOTAL WEALTH AND RETURNS: EQUITY VERSUS HOUSING AT A GLANCE

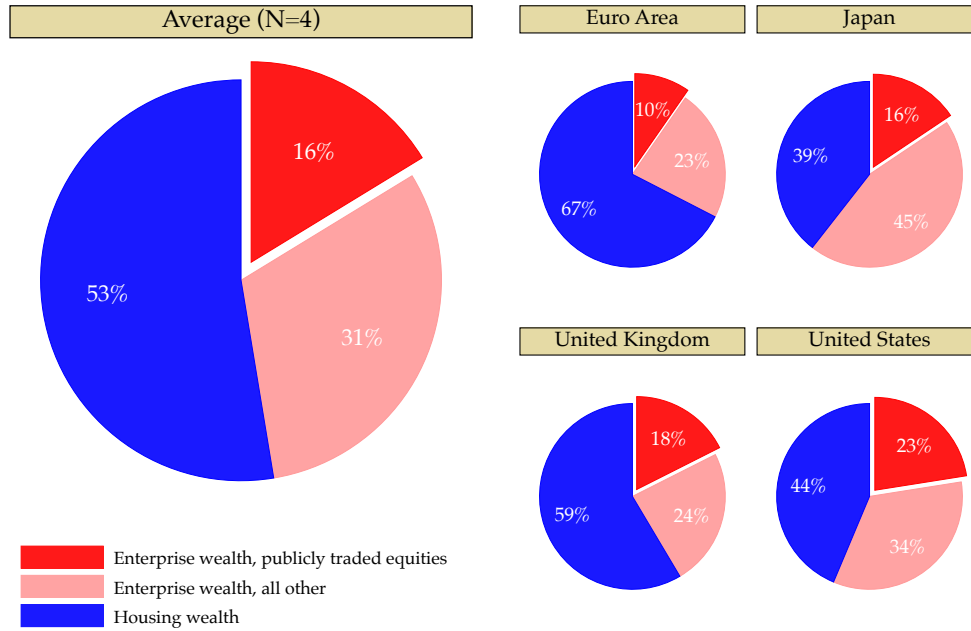
The core argument of this paper can be summed up in two statements. First, modern asset pricing research has been unduly focused on traded equities. This is unfortunate not only because traded equities are a small and possibly unrepresentative fraction of the total capital invested in enterprises. The larger problem is that this capital is itself only a fraction of total wealth in the economy, given the large share of capital consisting of housing wealth. Second, by extension, traded equity returns may be a poor guide to total wealth returns. Whether or not they are an acceptable proxy for returns on equity in enterprise capital, returns to housing may have very different statistical properties.

Figure 1 shows that in advanced economies the share of traded equities in total wealth is about one sixth, and represent about one third of total equity capital invested in enterprises. There is also a large share of wealth not in enterprises, roughly one half, which is invested in housing. The putative representative agent owns *all* of this wealth, not just the small slice in traded equities—but the mainstream asset pricing literature basically ignores this.

Does this matter? As Table 1 shows, we can now see that this is the case. In the table, R^e and R^h denote the annual gross total real returns on equities and housing, respectively. Total returns are capital gains plus dividends for equities, and capital gains plus rent (net of depreciation) for housing, with the real returns then deflated by consumer price inflation. The statistics displayed are means and standard deviations of net annual returns (in percent), the correlation of the two returns, and the sample size. The rows correspond to 4 samples, the U.S. and a pooled world sample of 16 advanced economies, for both the full 1870–2015 sample (including wars) and a balanced-panel post-WW2 sample (starting in 1963).

The key message is that of the two risky assets, housing has been a strong and safe investment, comparable to equities. In the World full sample, equities have returned 6.7% pa, and housing 6.9% pa in real terms. In the U.S. full sample, equities have returned 8.5% pa, and housing 6.1% pa. The standard deviation of equity returns has been large, 19% for the U.S. (22% for the World); but the standard deviation of housing returns has only been less than half as large, 8% for the U.S. (10% for the World). Similar findings apply to the Post-WW2 sample. The striking conclusion is that if the equity premium on its own is a puzzle that is hard to explain, introducing into the representative investor's portfolio a second asset class, of equal weight, and with higher returns, lower volatility, and low correlation to the first, is likely to make the puzzle even bigger. The core of this paper is devoted to robustly proving this point in the canonical consumption-based asset pricing framework

Figure 1: Composition of total wealth in selected advanced economies in 2005



Notes: This figure shows estimates of the breakdown of total wealth into housing and enterprise wealth for four regions (U.S., U.K., Euro Area, Japan) in 2005, based on national estimates of total wealth, the value of housing, and the capitalization of the tradable equity market. Other enterprise wealth is computed as the residual. Data on housing and total wealth relative to the size of the economy are taken from IMF, *World Economic Outlook*, April 2009, Box 2.1.1. Data on equity market wealth relative to the size of the economy are taken from World Bank, World Development Indicators, online database.

Table 1: Summary table: Key moments for annual real total returns on equities and housing

	$E(R) - 1$		$\sigma(R)$		$\rho(R_E, R_H)$	N
	Equities	Housing	Equities	Housing		
U.S. Full Sample	8.46	6.10	19.2	8.12	0.34	125
U.S. Post-WW2 Balanced	7.35	5.83	16.3	3.56	0.22	53
World Full Sample	6.73	6.93	21.9	10.3	0.18	1790
World Post-WW2 Balanced	7.92	6.89	25.0	7.46	0.07	837

Notes: This table shows the means and standard deviations of annual real total returns for housing and equities in the U.S. and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). The moments are computed for raw returns, and the units are in percent for means and standard deviations. Also shown is the correlation of the total returns for housing and equities, and the sample size. See text.

and its most familiar exotic variants.

Even more significantly, a portfolio composed, as in the real world, of 50-50 equities and housing would offer significant diversification benefits. The correlation of equity and housing returns is below 0.34 in the U.S. (0.18 world), and in the Post-WW2 sample it is smaller still, virtually uncorrelated. This ostensibly low correlation between equity and housing markets has an important implication. If returns in the equity and housing markets do not co-move closely, the properties of the diversified market portfolio consisting of the weighted sum of investable assets will be very different from that of the equity market alone. The CAPM-betas that quantify the systematic risk of an asset in comparison to the market portfolio will look very different when they are calculated with the true market portfolio that includes housing as opposed to using the equity market alone.

3. THEORETICAL AND EMPIRICAL BACKDROP

3.1. Notation and Statistical Preliminaries

Total real returns will be defined as nominal returns deflated appropriately for consumer price inflation. For any country, let $R_{i,t+h}$ be the gross h -period (depending on the periodicity of the data and the investor's horizons considered) total real return on a risky asset of type i . Formally,

$$\text{Return} = R_{i,t+h} = \frac{p_{i,t+h} + y_{i,t+h}}{p_{i,t}} \frac{1}{1 + \pi_{i,t+h}},$$

where y_t is the asset's nominal dividend in period t and p_t its nominal price at end of the period. The h period returns are then cumulated accordingly. We also denote by $R_{f,t+h}$ the gross h -period return on a safe asset if type f . Unless otherwise noted we restrict attention to the case $h = 1$, and drop the time subscript for clarity. The risky asset types i will be equities (E), housing (H), and a total (T) consisting of a 50-50 weight of each asset, given the above evidence on the weights of each asset type in total wealth. The safe assets will be government bills (typically 3-month, or a proxy short-term instrument) and bonds (typically 5–10 year).

Several summary statistics for raw returns will be of interest and allow a preliminary understanding of the properties of the data. We begin with a focus on a few key moments, namely,

$$\begin{aligned} \text{Expected return, asset } i: & E(R_i), \\ \text{S.d. of return, asset } i: & \sigma(R_i), \\ \text{Correlation of returns, asset } i, j: & \rho(R_i, R_j), \\ \text{Risk premium or excess return, asset } i: & E(R_i - R_f), \\ \text{Sharpe ratio, asset } i: & SR_i = \frac{E(R_i - R_f)}{\sigma(R_i - R_f)}. \end{aligned} \tag{1}$$

3.2. Consumption-Based Asset Pricing

To explain some of these moments, we seek to explore how well, or how badly, the standard framework for consumption-based asset pricing can perform, applied now to housing as compared to its familiar use for equities. We follow the simplest textbook approach, assuming a power utility with lognormal returns for tractability. (e.g., Abel 1988; Campbell 1999, 2004; Mehra 2008; Cochrane 2008).

The representative household at any time (e.g., normalized to $t = 0$), maximizes expected time-separable utility from per capita consumption c_t given by

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t U(c_t) \right\}, \quad 0 < \beta < 1,$$

where, for now, we assume a constant relative risk aversion (CRRA) or power utility function with curvature or *risk aversion parameter* γ , so that

$$U(c, \gamma) = \frac{c^{1-\gamma}}{1-\gamma}, \quad 0 < \gamma < \infty.$$

Intertemporal optimization requires that the discounted marginal utility tomorrow of $R_{i,t+1}$ units should equal the marginal utility of one unit today,

$$1 = \beta E_t \left\{ \frac{U'(c_{t+1})}{U'(c_t)} R_{i,t+1} \right\}.$$

Applying the same formula to a safe asset return $R_{f,t+1}$ (riskless one-period bond) we obtain

$$1 = \beta E_t \left\{ \frac{U'(c_{t+1})}{U'(c_t)} \right\} R_{f,t+1}.$$

Combining these last two expressions with some tedious algebra one can solve for the risk premium.

Specifically, given a CRRA utility, defining gross real per capita consumption growth as $g_t = c_{t+1}/c_t$, and dropping time subscripts, the theory yields an expression for the risk premium as the product of the risk aversion parameter and the covariance of returns and consumption growth,

$$\underbrace{\ln E(R_i) - \ln E(R_f)}_{\text{Risk premium or excess return (log form)}} = \underbrace{\gamma}_{\text{Risk aversion parameter}} \times \underbrace{\text{Cov}(\ln R_i, \ln g)}_{\text{Covariance of returns and consumption growth}}$$

The intuition for this expression is that a risk averse agent will be less willing to hold an asset whose payoffs more tightly covary with consumption, since this asset will tend to have low payoffs in bad states of the world. As a result, in equilibrium, they must be compensated with a higher excess return for bearing this risk.

The well-trodden path to the risk premium puzzle is taken by going to the data, computing the moments on either side of this expression, and solving for the implied risk aversion parameter γ . Along the way, some additional informative moments can be computed, namely,

$$\begin{aligned}
 \text{S.d. of consumption growth: } & \sigma(g), \\
 \text{Correlation of return and consumption growth, asset } i: & \rho(R_i, g), \\
 \text{Consumption beta, asset } i: & \frac{\text{Cov}(R_i, g)}{\text{Var}(g)}, \\
 \text{Hansen-Jagannathan bound, asset } i: & \gamma \geq \frac{SR_i}{\sigma(g)}, \\
 \text{Implied risk aversion parameter, asset } i: & \gamma = \frac{\ln E(R_i) - \ln E(R_f)}{\text{Cov}(\ln R_i, \ln g)}.
 \end{aligned} \tag{2}$$

When we turn to our empirical work in Section 5, we will discuss the relevance of moments (1)–(2) and present their empirical values for equity, housing, and total returns using our new panel dataset.

Before that we first review important prior work, which comprises a massive literature devoted to U.S. and world equity returns, and then discuss the much smaller literature which has applied the same framework to housing returns. We then briefly discuss the sources and construction of our new returns database.

3.3. The Evidence on the U.S. Equity Return Premium

The empirical moments just discussed in expressions (1) and (2) for U.S. postwar equity data have by now been computed many times and the stylized facts are well known. The basic conundrum presented in Shiller (1982) and Mehra and Prescott (1985) remains, and is a core feature of textbook treatments (Cochrane 2000; Campbell 2017).

According to the survey by Campbell (1999) the average U.S. real total equity return from 1947 to 1996 was high at 7.6% pa, and the average riskless rate was low at 0.8% pa based on 3-month Treasury bills; the implied equity risk premium was thus 6.8% pa. Similar estimates appear in Campbell (2003). For a more historical comparison, Mehra (2008) reports a range of equity premiums based on longer run annual data for the U.S. In the longest span sample, 1802 to 2004 due to Ibbotson and Siegel (1984), the value is 5.4% pa. In the shortest span sample, 1926 to 2004 due to Ibbotson and Siegel (1984), the value is 8.6%. The Shiller (1982) sample from 1871 to 2005 gives a premium 8.3% and the Mehra-Prescott sample from 1889 to 2005 gives 6.4%. Campbell (1999) reports an annualized value of 6.7% for 1891 to 1995. A reasonable long-run value of the U.S. annual equity risk premium in both postwar and longer samples is therefore perhaps around 6.5% pa.

Turning to the correlation of real consumption growth and real equity returns, Campbell (1999) finds a postwar value of 0.22 in quarterly data, and 0.33 in annual data. Expanding to a U.S. historical sample from 1891 to 1994 he reports an annual value of 0.45. This higher value presumably reflects some of the more synchronized stock-consumption episodes in the more turbulent pre-WW2

period, especially in the 1920s and 1930s. In addition, he shows that the annualized standard deviation of real equity returns is 15.5% and that of real consumption growth is 1.1% in quarterly postwar data. In annual U.S. historical data these figures rise to 18.6% and 3.3%, respectively.

Some results quickly follow. The implied Sharpe ratio for postwar U.S. equities is roughly $(0.068/0.155)$ or 0.44, annualized. Since the standard deviation of real consumption growth is 0.011, the Hansen-Jagannathan lower bound for the implied risk aversion parameter is about 40. But this lower bound is attained if and only if real consumption growth and real returns have a perfect correlation of 1; in reality, they do not. Empirically, the covariance of real equity returns and real consumption growth is small: Campbell (1999) shows that it is approximately $0.22 \times 0.155 \times 0.011$ or about 0.0004 in the quarterly postwar data, and appears to be somewhat higher, at $0.45 \times 0.186 \times 0.033$ or about 0.003, in the longer-run annual data. Plugging these values into equation (10) gives an implied risk aversion parameter of about 150 for U.S. postwar data, and around 20 for the longer-run data (see his Table 5). The former value is far above a plausible value, which is considered to be something below 10 (Mehra 2008); even the latter value is still quite high, although it forcefully illustrates that the puzzle abates when a broader and more volatile period of economic growth is included in the U.S. sample.

3.4. International Evidence on the Equity Return Premium

Brown, Goetzmann, and Ross (1995) suggested that results based on U.S. data alone may suffer from selection bias. A more recent literature (e.g., Jorion and Goetzmann 1999; Campbell 1999, 2000; Dimson, Marsh, and Staunton 2002, 2008; Mehra and Prescott 2003; Mehra 2012) discusses the international experience in more detail in the context of equity returns alone, but finds that international differences are few—the equity premium puzzle endures on a global level, and is not simply a U.S. phenomenon.

Here, research has consistently found a smaller puzzle in the wider international sample, with notably smaller values for the equity risk premium in long-run samples. The reason is, perhaps, obvious: the U.S. experience is exceptional, reflecting roughly 150 years without any major destructive wartime conflict on the U.S. mainland, and no episodes of major political or economic cataclysm such as coups, revolutions, hyperinflation, or sovereign default. In contrast, many of today's so-called Advanced Economies have faced such disruptions within living memory, and these episodes have in general been very adverse for returns to risky (or even "safe") assets. Thus, the empirical moments computed from U.S. data alone might be criticized as being subject to sample selection, or a type of survivorship bias. The same concerns might also apply to a lesser degree to the use of U.S. postwar data, where economic growth has been especially tranquil and even in the pre-2008 samples widely used.

Some typical postwar international results for equity risk premiums can be found in the survey by Campbell (2003). In his Table 4, some advanced economies, e.g., Canada (3.9%), Italy (4.7%), and Japan (5.0%), do exhibit lower equity risk premiums than the United States (6.3%) from 1970

to the 1990s. But conversely, some other economies have even higher equity risk premiums, e.g., Netherlands (11.4%), Sweden (11.5%) and Switzerland (14.9%). Thus, compared to the postwar U.S. case, the puzzle is similar in international data—sometimes worse, sometimes not.

Over the longer run back to the 19th century, the issue of sample selection bias is more salient, and the U.S. emerges as a more exceptional country. According to Dimson, Marsh, and Staunton (2008), over the period 1900 to 2005, the real equity returns for the “World ex-U.S.” have been 5.23% (16 countries) compared to 6.52% in the U.S. alone. Only 3 out of the 16 other countries attained higher average annual equity returns in this sample: South Africa, Australia, and Sweden. Zooming in on turbulent periods like 1914 to 1945 also emphasizes the much greater volatility of equity returns once the sample period extends outside the post-WW2 era. Converting these to excess returns over bills, these authors found annualized equity risk premiums of 4.2% for the “World ex-U.S.” and 5.5% for the U.S. over the same period. Thus perhaps one quarter of the risk premium puzzle might be explained away by sample selection bias when looking only at the U.S. data.

3.5. The Evidence on the Housing Return Premium

Although housing constitutes a significant share of wealth, empirical work on the housing risk premium has been held back by lack of widely available data on returns. Most references refer to U.S. returns on residential and commercial real estate. Starting with the early studies for the U.S. by Ibbotson and Siegel (1984), Case and Shiller (1989), and Goetzmann and Ibbotson (1990), a common theme of this literature has been that the U.S. real estate market seems to offer attractive risk-adjusted returns.

More precisely, diversified real estate investments have average returns that are only slightly lower than common stocks, but much less volatile, so that real estate investments have much higher Sharpe ratios. On the basis of the Case-Shiller housing return data, Goetzman (1993) discusses how diversified real estate portfolios can help investors achieve superior risk-adjusted returns. Shilling (2003) concluded that, since the 1980s, U.S. real estate investments carried a puzzlingly large risk premium of about 6% annually, yet with only a fraction of the volatility of the stock market.

Real estate investments thus seem to offer very attractive risk-return trade-offs. As a result, many have questioned whether the publicly available house price indices really capture the true volatility of housing assets and to what extent the attractive risk-return characteristics of real estate hold up to further scrutiny (Cheng, Liu, and Lin 2008). Moreover, from a household finance perspective, it is the volatility of the individual house that matters, since households typically do not invest in a diversified housing portfolio, but rather take on local idiosyncratic risk. Though tangential to the representative agent approach to asset pricing, these issues may deserve more study.

For example, Flavin and Yamashita (2002), using PSID micro-level data from 1968 to 1992, studied optimal portfolios including housing for mean-variance investors. In the data they found annual real after-tax returns of 6.6% for housing and 8.2% for equities. Standard deviations were, respectively, 14.2% and 24.2%. (Note that the after-tax basis of these returns depress them relative to

the pre-tax values we report in this paper.) Hence, their implied after-tax Sharpe ratios were about 0.3 for equities, but closer to 0.5 for housing in this sample; and their implied Hansen-Jagannathan lower bound for the implied risk aversion parameter would therefore be about 1.6 times larger for housing compared to equities. In addition, this study found housing returns to be negatively correlated with all other asset classes, indicating substantial diversification benefits to a mixed portfolio. However, studies such as Goetzmann (1993), Landvoigt, Piazzesi, and Schneider (2015), as well as Eisfeldt and Demers (2015), find considerably lower city-wide house price volatilities in the range of 7%–9%, depending on city, suggesting that self-assessed house prices in the PSID could suffer from excessive high-frequency noise, leading to the volatility estimate being upward biased.

In our aggregate data, we will find analogous results, not just for a 25-year window of data for the U.S., but for all advanced economies, and over long-run going back to 1870. Of course, macro and micro data may deliver different results in reality as there is a significant asymmetry in practice between equities, where the representative agent can hold the index, and housing, where such diversification is not yet possible. The representative agent model used since Mehra and Prescott (1985) ignores this constraint: in a no-trade equilibrium the agent buys and holds the market, the full equity bundle and the full housing bundle. Yet, as Davis and Van Nieuwerburgh (2014) caution that “the literature still misses high-quality data to pin down the return correlation matrix between stocks, bonds, and individual houses”, so the study of individual-level housing investments is painstakingly difficult. Nonetheless, even with our aggregate focus motivated by a literal and brute application of the standard representative agent model, we shall discuss the sensitivity of our results to micro-to-macro housing portfolio aggregation issues in the later part of this paper.

4. NEW DATA ON HOUSING AND EQUITY RISK PREMIUMS

In this section, we will discuss a major new source for the calculation of long-run returns and risk premiums in 16 countries for 145 years, for four asset classes: equities, housing, bills and bonds. This section draws on Jordà, Knoll, Kuvshinov, Schularick, and Taylor (forthcoming) [abbreviated JKKST] which discusses the data sources and construction in great detail.

A major innovation in JKKST is the inclusion of housing returns. These new data on housing returns cover capital gains, and also the net rents paid renters and owners (i.e., imputed). Equity return data for publicly-traded equities will then be used, as is standard, as a proxy for aggregate business equity returns.

The JKKST 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 2 summarizes the data coverage by country and asset class.

Like most of the literature, JKKST examine returns to national aggregate holdings of each asset class. Theoretically, as in the benchmark asset pricing models in this paper, these are the returns that accrue for the hypothetical representative-agent investor holding each country’s portfolio.

Table 2: 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

Bill returns In JKKST, the canonical risk-free rate is taken to be the yield on Treasury bills, i.e., short-term, fixed-income government securities, typically 3-month maturity. The yield data come from the latest vintage of the long-run macrohistory database (Jordà, Schularick, and Taylor 2016).³ Whenever data on Treasury bill returns were unavailable, JKKST 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 the bill rate data before the 1960s actually consist of deposit rates.

Bond returns In JKKST, these are conventionally the total return on long-term government bonds. Unlike earlier cross-country studies, JKKST 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 countries and periods JKKST 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 JKKST target a maturity of around 10 years.

³The data are online at www.macrohistory.net/data.

Equity returns In JKKST, 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, JKKST 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, JKKST also make use of indices weighted by company book capital, stock market transactions, or weighted equally, due to limited data availability.

Housing returns In JKKST, these data 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. JKKST 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, the returns include both rented housing and owner-occupied properties. These series also include an adjustment for maintenance and depreciation. To the best of our knowledge, JKKST are the first to calculate housing returns in the literature for as long and comprehensive a cross section of economies.

5. A BIGGER PUZZLE: RETURNS ON EQUITIES AND HOUSING

The consumption-based asset pricing model is disarmingly simple. In equilibrium, the returns of an asset, any asset, should reflect how we value consumption today against that in the uncertain future, and the insurance the asset provides against bad consumption states. In aggregate, little else matters. Realistic frictions faced by individual investors play a secondary role in the aggregate. The quantitative assessment of this model to date has relied primarily on easily available data on equity returns, with a focus on the U.S. postwar experience. The equity-premium puzzle documented by Mehra and Prescott (1985) highlighted how implausibly high the implied price of risk seems to be.

Tables 3 and 4 provide a bird’s eye view of the new total return data. As discussed in Section 2, equity returns are capital gains adjusted with dividends, and reported in real raw terms deflated by CPI inflation. Similarly, housing returns are capital gains adjusted with unique data on rents from Knoll (2017), and also reported in real raw terms and net of maintenance and depreciation. The tables calculate real and excess returns for equities, housing and the 50-50 portfolio alongside the return on two “safe assets” short- and long-term government securities, usually 3-months and 10-years in duration respectively, where we label these “bills” and “bonds,” respectively.

We report our results here, and in subsequent tables, using two samples. A pooled unbalanced panel of all available cross-country data from 1870 to 2015 (but excluding a window of 2 years around both world wars), and a balanced panel based on a post-WW2 subsample that begins in 1963 and ends in 2015.

Table 3: Total returns on safe and risky assets

	$E(R) - 1$					$\sigma(R)$		
	Bills	Bonds	Equities	Housing	Total	Equities	Housing	Total
(a) Full Sample								
USA	1.52	2.10	8.46	6.10	7.28	19.17	8.12	11.60
World (pooled)	1.04	2.03	6.73	6.93	6.83	21.92	10.31	12.92
(b) Post-WW2								
USA	1.66	2.55	7.35	5.83	6.59	16.29	3.56	8.71
World (pooled)	1.11	2.26	7.92	6.89	7.41	25.03	7.46	13.29

Notes: This table shows the means of annual real total returns $E(R) - 1$ for bills, bonds, housing, equities, and total wealth, and standard deviations $\sigma(R)$, for the U.S. and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). Calculations for each of the 16 countries in our sample provided in Appendix table A.1. The moments are computed for raw returns, and the units are percent. See text.

Table 4: Excess returns

	$E(R - R_f^{bills})$			$E(R - R_f^{bonds})$		
	Equities	Housing	Total	Equities	Housing	Total
(a) Full Sample						
USA	6.95	4.58	5.76	6.37	4.01	5.19
World (pooled)	5.86	6.03	5.94	4.82	5.00	4.91
(b) Post-WW2 Balanced						
USA	5.69	4.17	4.93	4.81	3.28	4.05
World (pooled)	6.81	5.78	6.29	5.66	4.63	5.14

Notes: This table shows the means of annual real excess returns $E(R - R_f)$ for housing, equities, and total wealth, over bills and bonds, for the U.S. and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). Calculations for each of the 16 countries in our sample provided in Appendix table A.2. The moments are computed for raw returns, and the units are percent. See text.

Table 3 (and individually for each country in Table A.1) shows that real total returns on safe assets hover in the 1%–2% range across countries in the full sample. For the U.S., bills and bonds returned 1.5% and 2.0%, respectively, about the same as the average for the World (here, “World” refers to the average for the 16 advanced economies in our sample). These returns are about the same in the post-WW2 sample, being higher in the 1970s and 1980s, before falling with the well-documented secular decline in safe rates observed internationally over the past 30 or so years

(see, e.g., Holston, Laubach and Williams 2017). Cross-country variation is rather limited in both samples, regardless of whether one examines bills or bonds.

Equity and housing returns are similar across countries. In the World samples, returns are about 7% although equities are 2 to 3 times as volatile (measured by their standard deviation) as housing. For the U.S., housing returns are about 6% compared with about 7% for equities in the post-WW2 era, and yet they are about one-quarter as volatile (3.5% versus 16%). Not surprisingly, when we look at excess returns reported in Table 4 (and for individual countries in Table A.2), we find that these are sizable—around 6% on average for the World when using bills as the safe asset, and about 1 percentage point lower when using bonds as the safe asset. The U.S. displays a higher equity risk premium and lower housing risk premium in the full sample, but is well within the global range of variation shown.

5.1. Sharpe Ratios and Hansen-Jagannathan Bounds

The ratio of the mean of the excess return to its standard deviation is known as the Sharpe ratio, and it provides a good summary statistic of the basic risk-return trade-offs of an investment. In simple settings—e.g., the mean-variance optimizing investor—the higher this ratio, the higher the returns for a given level of risk, and the more desirable the investment is, all else equal.

Of course, what is meant by “risk” in calculating the Sharpe ratio can be an imperfect guide. Investors that favor stability in their consumption streams—that is, risk-averse investors—will favor investments that insure them against bad consumption states. Thus, the volatility of the investment is as important as how the returns on the investment covary with consumption growth. As we discuss in a moment, an investment with higher pay-offs in a downturn should perhaps command a higher premium than an investment that fares poorly in recession.

The Hansen-Jagannathan lower bound (Hansen and Jagannathan 1991) combines these two concepts of risk into a measure that sets a floor to the volatility of the pricing kernel. Remember that the pricing kernel reflects the utility trade-offs between consuming today versus consuming tomorrow given that the future is uncertain and given that consumers dislike uncertainty. (In our baseline setting the pricing kernel is derived from power utility, as seen below.)

Information on Sharpe ratios and consumption growth volatility are provided in Table 5 for both of the samples (Full versus Post-WW2) that we have been considering (expanded per country results appear in Table A.3). Several results deserve comment. In line with the results reported earlier, housing offers considerably better investment opportunities than equities as measured by the Sharpe ratio. In some countries and depending on the sample, the Sharpe ratio rises above the fabled threshold of 1. More generally, averaging across the countries in our sample, the Sharpe is about two to three times larger for housing than it is for equities. Given $\sigma(g)$, the Hansen-Jagannathan lower bound will also then be twice as large too. The housing puzzle starts to loom larger.

Absent additional information, these returns seem wildly out of proportion compared to the volatility of consumption. After all, if the premiums demanded by consumers on risky investments

Table 5: *Sharpe ratios and standard deviations of consumption growth*

	(a) Full Sample				(b) Post-WW2 Balanced			
	SR				SR			
	Equities	Housing	Total	$\sigma(g)$	Equities	Housing	Total	$\sigma(g)$
USA	0.38	0.54	0.52	0.03	0.35	1.00	0.57	0.02
World (pooled)	0.28	0.62	0.49	0.05	0.28	0.74	0.48	0.02

Notes: This table shows the Sharpe Ratios $SR = E(R - R_f) / \sigma(R - R_f)$ for excess returns for housing, equities, and total wealth, over bills only, for the U.S. and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). Calculations for each of the 16 countries in our sample provided in Appendix table A.3. Also shown are standard deviations of real per capita consumption growth. The moments are computed for raw returns, and the units are raw, not percent. See text.

Table 6: *Correlations of returns and consumption growth*

	(a) Full Sample			(b) Post-WW2 Balanced		
	$\rho(R, g)$			$\rho(R, g)$		
	Equities	Housing	Total	Equities	Housing	Total
USA	0.60	0.37	0.62	0.47	0.33	0.50
World (pooled)	0.11	0.13	0.14	0.26	0.30	0.33

Notes: This table shows the correlation of returns and consumption growth $\rho(R, g)$ for housing, equities, and total wealth, for 16 countries and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). Calculations for each of the 16 countries in our sample provided in Appendix table A.4. The moments are computed for raw returns. See text.

reflect—at least in part—distaste for volatility in consumption over time, it must be because consumption itself is volatile. That is clearly not the case. The volatility of consumption for any country, regardless of sample, is an order of magnitude smaller than one would expect.

One possibility not accounted for in Table 5 is that risky returns could be highly correlated with consumption growth. Or, how else would we explain such high excess returns? The less an investment insures against fluctuations in consumption the higher a premium a consumer would require as compensation for holding such an investment. Table 6 (and Table A.4 in the Appendix) therefore calculates the correlation between returns and consumption growth.

In the U.S., where a great deal of the literature has focused, the correlation between equities and consumption growth is 0.47 (0.60 postwar). This number holds promise but appears to be on the high side when looked at in a comparative international context. The pooled average correlation between equity returns and consumption growth across countries in our sample is less than half that of the U.S., at around 0.11 (0.26 postwar). However, a new wrinkle appears when we look at

housing returns. In the U.S., housing returns, in addition to having a lower volatility than equities, also have much lower annual correlation with consumption growth than equities (about 0.4 versus 0.6). Again, the correlation is lower still around the world (just 0.13 full sample). Indeed, the Table A.4 shows that no other risky return in any country is as highly correlated with that country's consumption growth as it is for U.S. equity. Rather than abating, the puzzle grows.

5.2. Consumption Betas and Implied Risk Aversion

The analysis thus far has described basic properties of the returns on equities and housing with scant reference to the statistical properties of consumption growth and the economic implications of their interaction. It is time to extend the analysis in this direction. We explore two distinct but closely related summary calculations based on the consumption-capital asset pricing model (C-CAPM).

If one is willing to make the further assumption that the investor's preferences are well captured by the ubiquitous power utility function, Breeden, Gibbons, and Litzenberger (1989) show that the average gross excess real returns of an asset can be expressed in terms of the (constant) coefficient of relative risk aversion, γ , and the gross growth rate of real consumption per capita between periods t and $t + h$ and denoted $g_{c,t+h}$,

$$E[R_{i,t+h} - R_{f,t+h}] = \lambda_h \beta_{i,h}, \quad (3)$$

with

$$\lambda_h = \frac{\gamma \text{Var}(g_{c,t+h})}{1 - \gamma E(g_{c,t+h} - 1)}; \quad \beta_{i,h} = \frac{\text{Cov}(g_{c,t+h}, R_{i,t+h} - R_{f,t+h})}{\text{Var}(g_{c,t+h})}. \quad (4)$$

Here λ_h is the strictly positive market price of consumption risk. Similarly, $\beta_{i,h} > 0$ for most assets.

The expressions (3) and (4) require linear approximation assumptions, assumptions about the utility function underlying the investor's preferences, and an assumption that the risk aversion parameter remain constant over time.

Estimating λ_j would require further assumptions, but the parameter $\beta_{i,h}$ can be calculated directly from the data for various horizons h . Based on this approach, Table 7 reports estimates of $\beta_{i,1}$ for $i = E, H, T$, and Figure 2 displays estimates of $\beta_{i,h}$ for $h = 1, \dots, 5$ for the same three possible investments.

Finally, under lognormal return assumptions and with approximations, an estimate of the coefficient of relative risk aversion (RRA) can be easily calculated as:

$$\text{RRA} \equiv \gamma = \frac{\ln E(R_{i,t+h}) - \ln E(R_{f,t+h})}{\text{Cov}(\ln R_{i,t+h}, \ln g_{c,t+h})} \quad (5)$$

Table 8 reports such an estimate for $h = 1$ and for each of the three investments considered $i = E, H, T$. Meanwhile, Figure 3 shows the estimates of γ using $h = 1, \dots, 5$.

Turning to first to the consumption beta estimates, consider the values of $\beta_{i,h}$ based on expression (4) first, as reported in Table 7 and Figure 2b. Perhaps the most useful comparison is of $\beta_{E,h}$ and $\beta_{H,h}$. Recall that the model implies $\beta_{i,h} > 0$ but otherwise, there is little guidance on interpreting

the specific values that this parameter may take for a given asset, but a helpful metric is to compare the equity β against the housing β . Table 7 reports the results for $h = 1$. Nearly uniformly across samples (full and post-WW2 samples) and across countries (see Table A.5 in the Appendix), it is almost always the case that $\beta_{E,h} > \beta_{H,h}$. For the U.S., the implied beta for equities is 5 to 10 times that for housing; in the World samples it is 2 to 3 times.

Turning to Figure 2b, which focuses on the post-WW2 sample only, we display the values of $\beta_{i,h}$ for the U.S. in panel (a) and World in panel (b). The figure suggests that differences in the betas are largest for $h = 1$ and somewhat dissipate as h increases (although a more sizable gap remains for the U.S. as shown in panel (a)). The figure also shows that the differences between equity and housing betas are more extreme in the U.S. than in the World sample at all horizons shown. These results confirm much of what we found in previous sections, namely, that the standard consumption-based asset price model has greater difficulty in explaining the observed historical returns to housing, even more than was the case for equities.

We end this section by returning to the calculations of the implied coefficient of relative risk aversion implied by power utility. The objective is less to come with a precise estimate and more to show that conventional estimates are off by almost an order of magnitude for equities and even more so for housing. Begin with Table 8. The estimated γ based on expression (5) is between 18 and 40 for the U.S. and for equities (comparing full versus post-WW2 samples) but shoots up to between 48 to 203 for housing. This is not just a U.S. phenomenon. The same figures for the World sample are pretty stable over time, hovering around 40–50 for equities and 80–100 for housing in both samples. For the total portfolios the implied World coefficient of relative risk aversion is 60 in both samples, about halfway between the implied equity and housing coefficients.

The picture improves a little bit when considering estimates of γ over extended horizons, as displayed in Figure 3. This figure only shows the post-WW2 balanced sample results over a maximum horizon of 5 years. Panel (a) corresponds to the U.S. and shows that the γ for equities fluctuates in the range 40–60, similar to the numbers reported in Table 8. The housing γ becomes more muted over time as it declines from about 200 at the 1-year horizon to a low of about 100, although it is still far above the benchmark value of 10 or lower, the high end of what is considered to be a plausible value of the risk aversion parameter.

Panel (b) corresponds to the World pooled sample. Here too, we see a similar temporal pattern where the housing γ declines quickly beyond the 1-year horizon and in fact quickly becomes very similar to that for equities, both fluctuating in value at around 40. That said, these values are again still much too high relative to the benchmark level of 10.

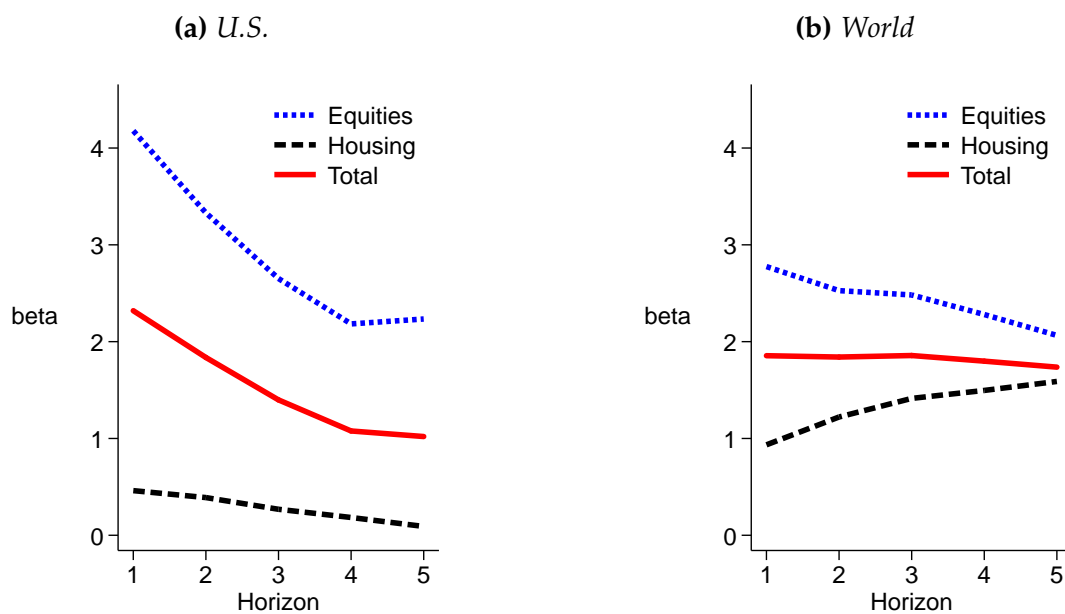
Finally, note that the low correlation of equity and housing returns means that the puzzle remains serious even when we look at the 50-50 total wealth portfolio. In panel (a), for the U.S. post-WW2 balanced sample, the implied coefficient of relative risk aversion varies between about 50 and 75 at all horizons, and comes in at an average of about 65. In panel (b), for the World post-WW2 balanced sample, the implied coefficient of relative risk aversion varies between 30 and 60 at all horizons, and comes in at an average of about 50.

Table 7: Consumption betas

	(a) Full Sample			(b) Post-WW2 Balanced		
	$\beta = \frac{\text{Cov}(R - R_f, g)}{\text{Var}(g)}$			$\beta = \frac{\text{Cov}(R - R_f, g)}{\text{Var}(g)}$		
	Equities	Housing	Total	Equities	Housing	Total
USA	3.17	0.64	1.91	4.18	0.46	2.32
World (pooled)	0.43	0.23	0.33	2.78	0.93	1.85

Notes: This table shows the consumption beta $\beta = \text{Cov}(R - R_f, g) / \text{Var}(g)$ for housing, equities, and total wealth, for the U.S. and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). Calculations for each of the 16 countries in our sample provided in Appendix table A.5. The moments are computed for raw returns. See text.

Figure 2: Consumption betas at longer horizons



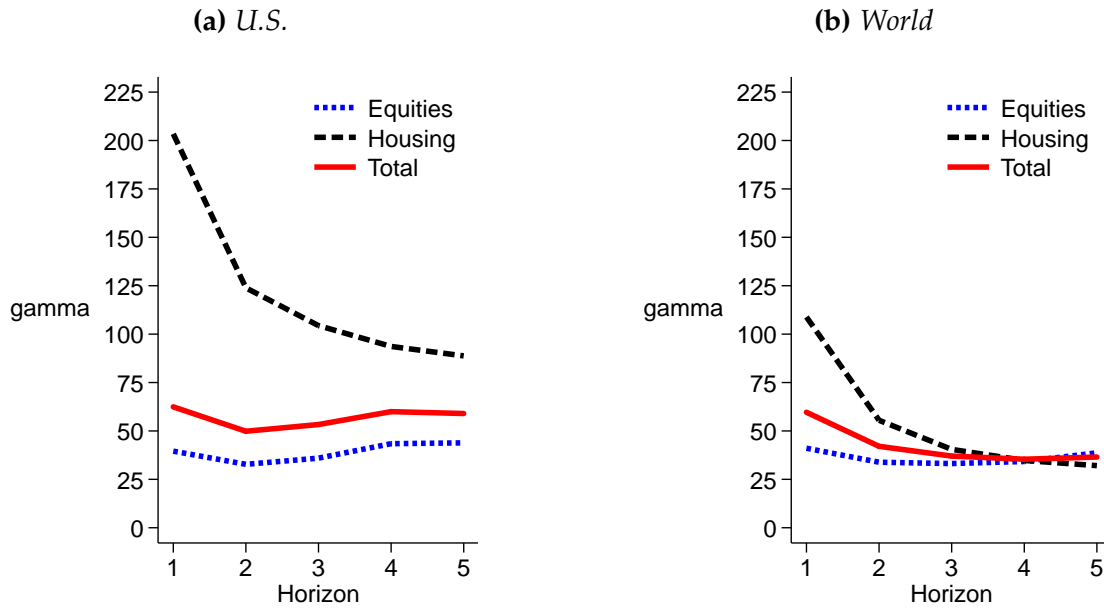
Notes: This figure shows the consumption beta $\beta = \text{Cov}(R, g) / \text{Var}(g)$ over horizons $h = 1, \dots, 5$ for housing, equities, and total wealth, for the U.S. and pooled World samples, for the post-WW2 balanced sample (1963–2015). The moments are computed for raw returns. See text.

Table 8: Implied risk aversion parameters

	(a) Full Sample			(b) Post-WW2 Balanced		
	γ			γ		
	Equities	Housing	Total	Equities	Housing	Total
USA	18	48	24	40	203	62
World (pooled)	46	80	60	41	109	60

Notes: This table shows the implied risk aversion parameter $\gamma = [\ln E(R) - \ln E(R_f)] / \text{Cov}(\ln R, \ln g)$ for housing, equities, and total wealth, for the U.S. and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). Calculations for each of the 16 countries in our sample provided in Appendix table A.6. See text.

Figure 3: Implied risk aversion parameters at longer horizons



Notes: This figure shows the implied risk aversion parameter $\gamma = [\ln E(R) - \ln E(R_f)] / \text{Cov}(\ln R, \ln g)$ over horizons $h = 1, \dots, 5$ for housing, equities, and total wealth, for the U.S. and pooled World samples, for the post-WW2 balanced sample (1963–2015). See text.

Table 9: Summary table: Consumption correlations and risk aversion parameters

	$\rho(R, g)$			γ		
	Equities	Housing	Total	Equities	Housing	Total
USA Full Sample	0.60	0.37	0.62	18	48	24
USA Post-WW2 Balanced	0.47	0.33	0.50	40	203	62
World Full Sample	0.11	0.13	0.14	46	80	60
World Post-WW2 Balanced	0.26	0.30	0.33	41	109	60

Notrs: This table shows the correlation of returns and consumption growth and the implied risk aversion parameter for housing, equities, and total wealth, for the United States and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). See text.

5.3. The Total Risk Premium Puzzle

A concise way to summarize what we have learned in this section is presented in Table 9. This table reports the annual correlation of equity, housing and the 50-50 portfolio with consumption growth in the left hand block of columns, both for the U.S. and for the 17-country world average, calculated as in Table 6. We report these results for both the full and the post-WW2 samples that we have considered throughout. The right-hand side block of columns report the model-implied estimates of the γ coefficient of relative risk aversion calculated as in Table 8. The correlation of returns with consumption growth is similar across categories yet the differences in the risk aversion parameter are vast.

This potentially opens new doors for macro-finance research. Moving from equity to the full wealth portfolio of the representative agent aggravates the challenges of the model to match the data. In the context of the canonical representative-agent asset pricing model, we cannot speak of just an equity risk premium puzzle any more. We are also confronted with a housing risk premium puzzle and a total risk premium puzzle.

5.4. Further Checks and Adjustments

Before we infer that the inclusion of housing in consumption-based asset pricing aggravates the problems of the canonical model, it is important to think about a number of objections that have been raised with regard to the standard model. Specifically, as noted in the introduction, for *ceteris paribus* reasons we have strictly adhered to the logic of the representative agent model and its no-trade equilibrium. But some deviations from that setup might be particularly salient in the case of housing, so we need to consider how they might affect our core results presented above.

Three issues stand out, and we will discuss them here as they are potentially important when it comes to housing returns: limited participation, transaction costs, idiosyncratic risks, and liquidity.

Do any of these adjustment have the potential to overturn our key finding that the puzzles of the consumption based asset price model get worse when one includes housing into the total wealth portfolio? We will argue that the answer is likely negative.

5.4.1 Limited Participation

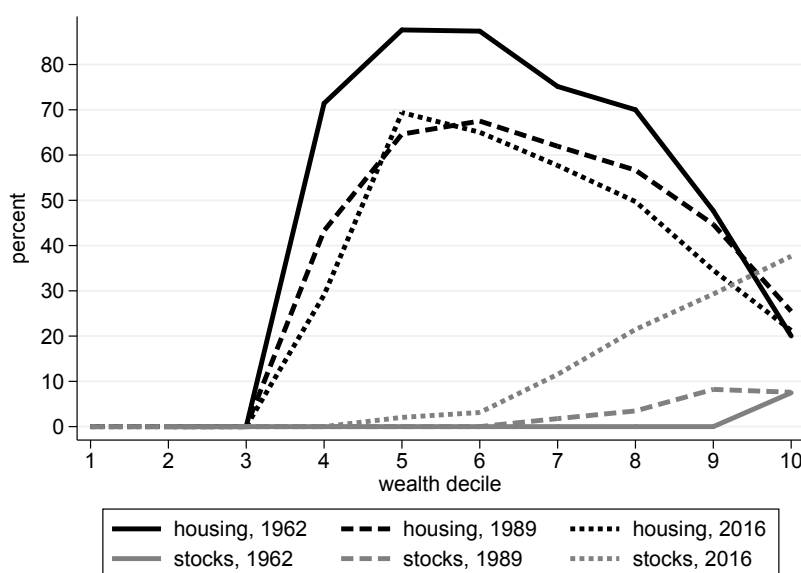
We start with limited participation as a potential reason for the observed failure of the standard consumption based model. With respect to equity returns, papers such as Mankiw and Zeldes (1991) and Vissing-Jorgensen (2002) argue that the consumption stream of stockholders is more volatile and more highly correlated with excess returns to equities. In Brav, Constantinides, and Geczy (2002), as the participation threshold is raised (based on value of assets owned), the correlation of the per capita consumption growth with the equity premium rises, and so again the implied RRA parameter falls. Hence, the RRA parameter implied by the consumption of luxury goods falls relative to national accounts data. Ait-Sahalia, Parker, and Yogo (2004) find a similar result when the consumption stream is restricted to a set luxury goods which form a more significant share of rich households' spending.⁴

However, limited participation arguments are much less likely to work in the context of housing as compared to equities. And the reason is simple: housing is a much more broadly held asset than equities. According to Mankiw and Zeldes (1991) only 25% of U.S. households own stock. In contrast, U.S. Census Bureau data shows that for more than 50 years, going back to 1962, the U.S. home ownership rate (share owner-occupied) has been between 63% and 69% in all years. Only in the uppermost deciles do equities start to take up a large fraction of household portfolios. Thus, while limited participation is a serious problem for equities, and could strongly bias the results of consumption-based capital asset pricing models applied to aggregate data, the bias should be significantly smaller in the case of housing, since a much larger fraction of households own a house.

A graphical view of the long-run cross-section patterns in U.S. real estate and equity holdings by households can be gleaned from Survey of Consumer Finances, using newly collected data by Kuhn, Schularick, and Steins (2017). Some representative snapshots for the 1960s, 1980s, and today are shown in Figure 4, and these clearly makes the point that exposure to housing investment risk runs very far down the wealth distribution, but material exposure to equity investment risk is only a matter for the top quintile. For confirmation using a different source, PSID data by wealth decile, Barras and Betermier (2016) have shown similar patterns in today's data.

⁴The point may well apply more broadly. An alternative disaggregated approach would be to look at other less aggregate consumption streams, which can be more volatile than the averaged national accounts data would suggest. E.g., Attanasio and Weber (1993) and Beaudry and van Wincoop (1996) have found higher values for the IES (or lower RRA) using disaggregated cohort-level and state-level consumption data.

Figure 4: Share of household assets in corporate equity and real estate by wealth percentile



Notes: Based on data from Kuhn, Schularick, and Steins (2017).

5.4.2 Transaction Costs

Transaction costs for assets might potentially explain the problems of the consumption-based asset pricing model. Moreover, such concerns might loom especially large in the case of housing. So far the results seem mixed, in that plain symmetric transaction costs can do very little to help bring down the implausibly high values of the implied RRA, or its Hansen-Jagannathan lower bound (He and Modest 1995; Luttmer 1996). Instead, more complicated transaction costs, specifically on short-sales might be required to make any difference to the puzzle.

In Table 10 and what follows we explore whether this approach could have promise if applied to housing premium puzzle. We argue that it likely does not. For the equity market, typical transaction costs values used applied to the U.S. are 1.5 bps and 75 bps for the Treasury bill and value-weighted equity returns, respectively. The 75 bps value may be a little low for the equity market over the longer run, as costs have fallen in recent decades on U.S. bourses. Jones (2002) finds a round-trip transaction cost of 100 bps. For less frequently traded stocks spreads could be as high or higher (and likewise, of course, for nontraded equity), and they could well be higher in overseas markets and in more distant historical epochs.

However, these simple cost ratios need to be adjusted for the typical trading frequency of each asset. Jones (2002) also shows that equity market turnover, at least post-WW2, has been at a minimum of 25% annually on the NYSE, rising rapidly in recent years. Over a longer horizon NYSE turnover has been at least 50% on average implying annualized round-trip transaction costs of at least 50 bps (100 bps times 50%) over a century or so.

Table 10: *Transaction costs for equities and housing*

	Equities	Housing
Cost per transaction, two-way	100 bps	700 bps
Annualized turnover rate	50 %	15 %
Annualized transaction costs, two-way	50 bps	105 bps

Notes: This table shows transaction cost computations for U.S. equities and housing as detailed in the text.

One’s first thought might be that in the context of our risk premium puzzle applied to housing, the transaction costs explanation might bear more fruit. After all, on any individual transaction the costs are much higher for the trade of a house, than for an equity trade. For example, in the U.S. standard realtor fees are typically 600 bps and an additional 100 bps might be absorbed by other fees such as title, recording, and closing costs, for a total round-trip transaction cost of 700 bps.

Perhaps surprisingly, it turns out the analogous calculations of transaction costs for U.S. housing generate a not so different result. A recent study by Bachmann and Cooper (2014) finds that annual turnover in the housing market is 15%, using PSID data (see their Figure 2). That is, the typical house is bought and sold roughly once every 7 years. Given the round-trip transaction cost of 700 bps for housing, this works out to annual round-trip transaction costs of 105 bps for this asset class too, only about 50 bps higher than equities. Thus, *prima facie*, it would seem that using transaction costs to provide an explanation of the housing premium puzzle will be about as successful, or unsuccessful, as using transaction costs to provide an explanation of the equity premium puzzle. That is, not very.

As a final caveat, we could worry that it is possibly misleading to base this argument entirely on moments computed from short holding periods, such as annual data, when most houses are not being traded, but equities are. However, this is where our decision to explore the puzzle at longer horizons provides extra assurances. We showed about that at horizons all the way out to $h = 5$ years, the puzzles remain intact, and the implied RRAs remain ludicrously large. As we extend to ever longer holding periods, the differences in transaction frequency, possibly due to different liquidity and other characteristics of the two assets, will matter much less, but as we have seen earlier the puzzle still remains strong at longer horizons.

5.4.3 Idiosyncratic Risks

Idiosyncratic risks could be especially salient for the housing component of wealth. It is, of course, relatively easy for a U.S. investor to “hold the market” by investing in a low-cost mutual fund or ETF which tracks the total return to a broad stock index like the S&P500. (Of course, this argument does not apply to the large share of enterprise wealth that is not publicly traded as shares on stock markets, and which is often held in concentrated forms, e.g., family businesses.)

Table 11: *Level and volatility of real housing capital gains at different levels of coverage and aggregation*

	Baseline		Zillow		
	National	National	State	County	Zipcode
<i>Mean real capital gain p.a.</i>	1.42	0.79	1.07	0.53	0.92
<i>Standard deviation</i>	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, and the other figures cover a representative state, county or zipcode level portfolio respectively. See Jordà, Knoll, Kuvshinov, Schularick, and Taylor (forthcoming)

Having said that, it is even more difficult, nay, impossible, for a U.S. investor to buy an index that tracks the total return to aggregate U.S. housing. Rather, if we are a typical homeowner, we own our own house, and that house alone. It will suffer regional, state, county, and possibly also city and neighborhood risks which are not present in the national house price index.

The question is, are these disaggregated local housing risks large enough to matter, relative to the national, to change the picture substantially? We argue that although they go in the right direction, these adjustments are unlikely to solve the problem entirely, simply because we find the housing premium puzzle is so big in our data.

Drawing on Jordà, Knoll, Kuvshinov, Schularick, and Taylor (forthcoming), Table 11 compares the recent level and volatility of the US conforming mortgage based OFHEO national house price indices with those that cover sub-national markets, where the latter are sourced from Zillow. Comparing columns 2 and 3 of Table 11, 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 11 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. But not by much. Any individual house will track Zillow Zipcode-level price movements closely, and the standard deviation of those housing returns is roughly one-third (a factor $4/3$) higher than that in the national data, so the Sharpe Ratio and related puzzles revealed by the moments our housing return data might be mitigated by 25% (i.e., by a factor of $3/4$).

This pattern is consistent with a very long tradition of work. In a seminal paper, Case and Shiller (1988) documented that prices of individual homes are twice as volatile as city-level house price indexes. Similar patterns appear in Flavin and Yamashita (2002) and Piazzesi et al. (2007). Flavin and Yamashita (2002) looked at individual-level PSID data from 1968 to 1992 they found the annual standard deviation of housing returns to be 14%, versus 24% for stocks. This compared with Case-Shiller city wide annual standard deviation of housing returns of 7%. In that same period the

FHFA U.S. price index has standard deviation of about 3.3%. However, as noted earlier, one concern might be the potential excess volatility of self-reported PSID housing wealth data.

In addition, as noted in a recent work by Giacoletti (2016), all these works except the first, assumed price shocks were i.i.d., but Case and Shiller (1987) and Goetzmann (1993) show that at short horizons the volatility does not scale to zero, indicating that care must be taken with regard to the holding period being studied. Giacoletti (2016) looks across horizons and, using CoreLogic data, focuses narrowly on a few real estate markets with big price action (LA, San Francisco, San Diego). He shows that the fraction of individual house price volatility determined by idiosyncratic risk is decreasing as the holding period increases. Idiosyncratic risk makes up to 60% of capital gains volatility of home resales after one year, but it accounts for less than 20% after five years.

Therefore, a putative resolution of the housing puzzle via this mechanism appears incomplete; it might only be promising at short horizons, and unlikely to change the story at long horizons where the puzzle remains strong. It is again fortunate that we have computed values of the implied RRA, or its Hansen-Jagannathan lower bound at a range of horizons. Once we get out to five years, the housing puzzle still remains large, but the idiosyncratic volatility is less than 20%, and this would be too small of an amplification to materially change the key asset pricing moments, as it would not do much to raise the consumption beta or depress the implied RRA for housing. Even at the one year horizon, a 75% increase in housing price volatility would leave the implied RRAs in Table 8 still at a level of about 40 or higher—still much higher than the equity-implied RRAs, and well above the a priori reasonable range.

5.4.4 Liquidity

From the point of view of an individual investor, real estate is a heterogeneous asset held over over relative long periods of time that is usually traded in decentralized markets far more infrequently than equities, as we just showed. Uncertainty on how much time the seller will need to close the transaction is considered an important reason why real estate returns carry a liquidity premium. When prices are rising, selling times are short and the volume of transactions are high. When prices are falling, all these features go into reverse.

It has been long recognized that for an individual investor, housing must carry a liquidity premium. Theoretical asset pricing models on housing have incorporated this and other features of housing markets (see, e.g. Krainer 2001; Piazzesi and Schneider 2009; Ngai and Tenreyro 2014). However, empirical estimates of the magnitude of this liquidity premium are hard to come by.

One notable exception is Lin and Vandell (2007), who introduce a model of the liquidity premium and estimates based on U.S. data from the Office of Federal Housing Enterprise (OFHEO) from 1981Q1 to 2004Q4. Table 2 of their paper in particular, reports a calculation informative for our purposes. The table is organized as a function of the expected time the house will be in the market before it is sold, and adjusted for the holding period.

For example, data from the National Association of Realtors in the U.S.⁵ reports that recently sold homes were on the market a median of 4 weeks (one month). Using those figures, Lin and Vandell (2007) indicate the variance would need to be scaled by a factor of 1.02 based on a holding period of 10 years. Holding periods in the U.S. appear to be in line with those in other countries. Data from various sources⁶ indicate holding periods in 2015 were 10 years in Australia, 9 years in New Zealand (not in our dataset), 9 years in Canada, and 11 years in the U.K., all very close to the U.S. median. Prior to the Great Recession, the median tenure in the U.S. had been in the neighborhood of 6 to 7 years historically, bringing the scaling factor closer to 1.05. These numbers increase to 1.10 and 1.19 respectively if we quadruple to 4 months the time it takes to sell a house. Even under these adverse scenarios, the increase in the variance of housing returns suggested by these scaling factors would fall well short to explain the housing premium that we report above.

An alternative options-based framework for gauging liquidity premia is provided by Longstaff (2018). Here, illiquidity takes the form of a lockup or “gate” which prevents sale in a given window. Theoretically, the *upper bound* on the opportunity cost of not being able to trade is equal to forgoing an option to sell at the asset’s expected high price in the window, the *best-case* scenario. The value of selling at this optimal stopping time equals the value of a put option at the exercise date with a strike equal to the expected high, which has a closed-form solution. Using, Table 2 of Longstaff (2018), for an asset with 10% annualized volatility and a holding period of 10 years (comparable moments to housing found in our data) the required liquidity premium is 126 basis point or 1.25 ppt. Since the excess return to housing is about 6 ppt this correction amounts to a downscaling of the risk premium of only 125/600 or about 20%. Since the Sharpe ratio of housing is double that of equities, and since this is very much an upper bound, this correction also does little to explain away the even larger risk premium puzzle we have uncovered for housing as compared to equities.

6. “EXOTIC” MODELS AND THE TOTAL RISK PREMIUM PUZZLE

In this section we study how more “exotic” variants of the benchmark consumption-based asset pricing model cope with the addition of housing to the aggregate portfolio. More precisely, we ask to what extent well-known extensions of the standard model can resolve the total risk premium puzzle that we discussed in the previous sections. The literature is vast, and here we focus on two important approaches, both of which respond to the fundamental problem that in the benchmark model the stochastic discount factor derived from the second moment of the consumption claim has very low volatility compared to the asset price volatility which it is supposed to explain.

The first approach is to expand the sample, dispense with lognormal assumptions, and admit “rare disaster” events: unusually low consumption values mostly associated with wars or major financial crises. Their inclusion increases both the volatility and negative skewness in the consump-

⁵2016 Profile of Home Buyers and Sellers. <https://www.nar.realtor/sites/default/files/reports/2016/2016-profile-of-home-buyers-and-sellers-10-31-2016.pdf>

⁶See Figure 2 in <https://www.corelogic.com/blog/2016/08/us-economic-outlook-august-2016.aspx>

tion growth process, making the asset less attractive, i.e., boosting the risk premium. An asset that loses its value in disasters provides poor insurance and an investor may therefore demand a high risk premium on such an asset. This is the argument developed by Rietz (1988) and advanced further by Barro (2006) and subsequent papers, to explain why there is a large equity premium. A key idea is that the sample sizes used in much previous research may have been too short (or too focused on the U.S.) to include any or many such rarely observed disasters, where consumption declines by dramatic amounts. In related research, Muir (2017) expands on this idea, using historical data similar to ours, and finds risk premium spikes in disasters: that is, asset prices fall by significantly more than dividends. But while this is true in financial crises, it is not very evident in deep recessions or even wars—suggesting still another conundrum for the traditional consumption based asset pricing model even allowing for a disaster-based approach.

The other major approach, instead of boosting the volatility of the claim, changes the underlying utility functions of agents to amplify the volatility of the model's stochastic discount factor. The two influential examples we consider are the habits model of Campbell and Cochrane (1999) and the long-run risk model of Bansal and Yaron (2004). In the habits model, preferences are such that the relevant consumption level is not absolute, but the level relative to some inherited habitual baseline that originates from the agent's past history; by construction, volatility around this sluggishly-moving level is far more pronounced when consumption is persistent and the habitual baseline lags behind actual consumption. Thus, the stochastic discount factor gets more variable.

In the long-run risks model, preferences are such that utility derives not just from today's consumption but also from a recursive term incorporating future expectations. Where habits add a backward-looking term, this model adds a forward-looking term, but it serves a similar purpose here when combined with a volatile long-run growth. If utility puts weight on this recursive term, small changes in the agent's expectations of long-run growth can deliver large changes in utility, and hence in the stochastic discount factor. As with rare disasters, this model makes the consumption path change, but it does this via a small growth rate change along the entire future path, rather than a large front-loaded drop in levels, and recursion then makes the change bite on asset prices today.

Looking at rare disasters, habits, and long-run risk approaches, we now ask how these models stand up to inclusion of housing into the asset portfolio. How successful are such extensions in explaining the total risk premium puzzle? To address this question we now augment and refine some of the models, implement them with two risky assets, rather than just one, and discuss the details of how we implement these extensions.

Still, when all is said and done, the bottom line is not that surprising. Although all of these approaches have been shown to make some progress in addressing the equity risk premium puzzle, they are, like the benchmark model, faced with a much harder challenge when seeking to explain the housing risk premium puzzle and the total risk premium puzzle. For these latter two puzzles, as we have seen, the risk premium is just as high as in the case of equities, but the key correlation between the asset return and the consumption process is even lower. In most models, this also leads to an even bigger problem to match model to data with "reasonable" parameters.

6.1. Barro (2006) with Risky Assets

A series of influential papers building on Rietz (1988) and Barro (2006) attributes the risk premium on equity to the existence of rare disasters. Infrequent but severe disasters, such as World Wars, the Great Depression and major financial crises, generate dramatic contractions in output and consumption, which in conjunction with diminishing marginal utilities imply large risk premiums despite their infrequent nature. The rarity of disasters, of course, implies a “peso problem”: as Barro (2006) notes, the mere *potential* for disasters has major effects on required rates of returns despite prolonged stability in most Western countries during the second half of the 20th century.

In the spirit of the Barro (2006) rare disasters model, we augment the model to incorporate two risky assets and investigate whether the distinct behavior of equity and housing claims during consumption disasters can work in this framework to explain the total risk premium puzzle. In a nutshell, we show empirically that rare disasters trigger larger risk premiums for equity claims, but not housing claims. Rare disasters hence cannot resolve the issue of similar observed equity and housing returns and risk premiums given the much more severe equity price declines seen during consumption disasters.

One contribution vis-à-vis the standard Barro (2006) model relates to deriving a closed form solution for the expected risky return of any arbitrary asset.⁷ Full details are given in the Appendix and in Barro (2006) for the initial setup, yet we briefly highlight our procedure. In essence, we assume that the dividend stream follows a similar exogenous process as the consumption process, i.e., it is characterized by uncertainty during normal times as well as rare disasters which unfold over one period. We allow for co-movement during calm periods by introducing a joint normal process for regular fluctuations. With respect to disaster periods, we impose simultaneity among consumption and assets disasters. As such, we restrict attention to dividend declines during consumption disaster periods and not vice versa. We believe this assumption is justified given that the risk premium crucially depends on high marginal utilities during consumption disasters. In our setting, assets are by construction only imperfectly correlated with consumption. As such, required risk premiums for assets are lower than the risk premium on the consumption claim itself.

We present the results of our exercise in Table 12. The risk premium for the consumption claim equals 4.7% and is about 1.1% higher than the risk premium in Barro (2006) for output claims. Most of the gap can be attributed to a higher disaster probability of 2.1% versus 1.7% (in our data) for consumption versus output disasters. More importantly though, the model clearly predicts lower risk premiums for housing returns. The associated premium equals 0.5% while the risk premium for equity returns amounts to 2.9%. The gap is substantial and partially due to a higher correlation between equity and consumption claims during normal times. However, the majority of this difference can be attributed to the behavior during consumption disasters simply because

⁷Barro and Ursúa (2008) extend the CRRA framework and introduce Epstein-Zin preferences. Epstein-Zin preferences do not affect the risk premium per se, but mitigate some counterfactual asset pricing implications. Since we focus on differences in risk premiums among equity and housing claims in the presence of rare disasters, we review the CRRA case which we also impose in the main body of this paper.

Table 12: Barro (2006) model with risky assets

Model Asset	Barro (2006) Y claim	JST this paper C claim	JST this paper E claim	JST this paper H claim
<i>Parameters</i>				
θ (relative risk aversion)	4	4	4	4
σ (s.d. of cons. claim, no disasters)	0.02	0.043	0.043	0.043
γ (growth of cons claim, no disasters)	0.025	0.024	0.024	0.024
ρ (time preference)	0.03	0.03	0.03	0.03
p (disaster probability)	0.017	0.021	0.021	0.021
q (bill default probability in disaster)	0.4	0	0	0
$\tilde{\sigma}$ (s.d. of asset claim, no disasters)			0.272	0.071
κ (corr. between asset and cons. claim, no disaster)			0.105	0.077
<i>Implied Returns</i>				
Risky Return	0.071	0.071	0.054	0.030
Bill Return	0.035	0.025	0.025	0.025
Risk Premium	0.036	0.047	0.029	0.005

Notes: This table shows the extended Barro (2006) model with risky assets applied to our data. The first column reproduces the Barro (2006) estimates and calibration for comparison. The remaining columns are based on our data. Y=Output, C=Consumption, E=Equity, H=Housing. Consumption disaster definitions are based on Barro and Ursúa (2008) and include peak to trough declines of at least 10%. Barro (2006) defines a threshold of 15% which he subsequently revised to 10% in follow-up papers. The partial default rate for the "risk-free" asset is set to zero to more closely align with observed returns.

owning houses seem to be a superior hedge—that is, during the rare disasters, drops in prices in the housing market are statistically less associated with (utility weighted) consumption declines than equity. We shall expand on this point when we turn to our application of the Muir (2017) methodology below.

6.2. Nakamura, Steinsson, Barro, and Ursúa (2013) with Risky Assets

We just highlighted that the Barro (2006) rare disasters model is unable to explain the total risk premium puzzle of roughly equal housing and equity returns. Yet as Gourio (2008) and Constantinides (2008) among others argue, the Barro (2006) model features two assumptions that are not supported by data and could potentially undermine any conclusions, namely instant disasters and irreversibility. Intuitively, fast recoveries reduce the persistence of disasters on dividend streams and therefore reduce the drop in prices when disasters occur. Multi-period disasters, everything else equal, make disasters less risky.

In subsequent influential work, Nakamura, Steinsson, Barro, and Ursúa (2013) addressed this critique and consequently adjust the rare disaster model of Barro (2006), again focusing only on the consumption claim. They find strong support for partial recoveries and disasters that last for multiple periods. Their associated (consumption claim) risk premium therefore ranks in between the Barro (2006) model and the plain vanilla asset pricing model of Mehra and Prescott (1985).

Before diving into details about why we think including partial recoveries and multi-period disasters makes the total risk premium puzzle more puzzling, it is worth stressing the conditions under which the new features may help to bridge the gap in risk premiums and hence resolve the puzzle proposed in this paper. First, equity disasters may unfold over a considerably longer stretch of time and losses could hence be distributed over multiple periods. Second, despite the severity of peak to trough equity declines, stock markets might recover more substantially than housing markets, thereby reducing the persistence of equity disasters.

In the remainder of this section, we provide evidence against these possibilities and conclude that the total risk premium puzzle gets worse when the model expands to include recoveries and multi-period disasters. Our strategy, following Nakamura, Steinsson, Barro, and Ursúa (2013), is to develop a method to re-estimate the model for equity and housing claims. We then show that short-run per period declines are stronger and recoveries weaker for equity compared to housing claims. We further highlight that housing and not equity disasters unfold over a longer period of time. Last but not least we provide evidence of a positive co-movement between disaster probabilities for equity and consumption claims if a country finds itself in a consumption disaster, but we find no such pattern for housing claims. The latter breaks the link between a high marginal utility of consumption during consumption disasters and housing price drops, which in turn ultimately and counter-factually implies a much lower risk premium for housing.

We briefly highlight the empirical model and focus on three parameters of interest: the temporary drop and permanent decline in assets or consumption due to a disaster, as well as the persistence of disasters. The model itself is estimated using Bayesian MCMC methods. We use the same prior calibration as in Nakamura, Steinsson, Barro, and Ursúa (2013) and refer to their paper for full details. In particular, the prior on the long-run impact of disasters is uninformative in the sense that we are agnostic about whether disasters have any long-run effect at all and allow for positive long-run effects.

Formally, we denote log dividends or consumption $d_{i,j,t}$, in country i in year t , where the claim type is $j \in \{\text{Consumption, Equity, Housing}\}$. These are modeled as the sum of three unobserved components: Potential "output" ($x_{i,j,t}$), a disaster gap ($z_{i,j,t}$) and an i.i.d. distributed normal shock ($\epsilon_{i,j,t}$). Equations (7) and (8) highlight the evolution of potential consumption and the disaster gap.

$$d_{i,j,t} = x_{i,j,t} + z_{i,j,t} + \epsilon_{i,j,t}, \quad (6)$$

$$\Delta x_{i,j,t} = \mu_{i,j,t} + \theta_{i,j,t} I_{i,j,t} + \eta_{i,j,t}, \quad (7)$$

$$z_{i,j,t} = \rho_j z_{i,j,t-1} - \theta_{i,j,t} I_{i,j,t} + \phi_{i,j,t} I_{i,j,t} + v_{i,j,t}. \quad (8)$$

Table 13: *Disaster Parameters for Nakamura, Steinsson, Barro and Ursúa (2013) Extended Model*

Model	NBSU (2013)	JST this paper	JST this paper	JST this paper
Asset	C claim	C claim	E claim	H claim
$1 - p_e$	0.835	0.7517	0.7102	0.8003
ϕ	-0.111	-0.0771	-0.1479	-0.0426
θ	-0.025	-0.0190	-0.1284	0.0031

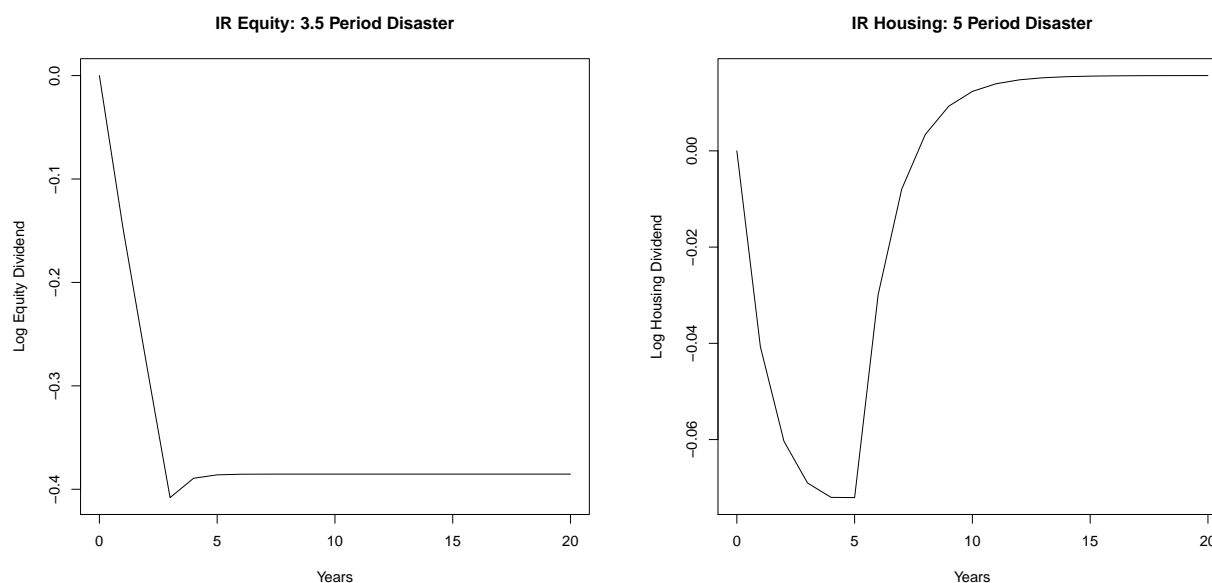
Notes: This table shows the NSBU and Extended NSBU model parameters. Estimates are posterior means. In line with Nakamura, Steinsson, Barro, and Ursúa (2013), we assume that disaster parameters are common across countries and time. The first column refers to the original estimates in Nakamura, Steinsson, Barro, and Ursúa (2013) for comparison, and the remaining columns to results based on our data. C=Consumption, H=Housing, E=Equity. The estimates are based on 40 runs with 50,000 iterations each.

The disaster gap follows an AR(1) structure with persistence given by ρ_j , while potential consumption follows a random walk with drift. The terms $v_{i,j,t}$ and $\eta_{i,j,t}$ are i.i.d. shocks similar to $\epsilon_{i,j,t}$, and $\mu_{i,j,t}$ is the country-specific average growth rate. The term $I_{i,j,t}$ is a Bernoulli random variable and indicates disaster periods. The occurrence of disasters itself follows a Markov process. It is characterized by an exit probability in each period (p_e), based on which we are able to compute the average duration of disasters. The term $\phi_{i,j,t}$ captures the temporary drop and only enters the disaster gap equation. In contrast, $\theta_{i,j,t}$ captures the permanent loss. As in Nakamura, Steinsson, Barro, and Ursúa (2013), we assume that the permanent decline does not affect assets or consumption during the disaster, hence the opposing signs in equations (7) and (8).

We estimate our empirical version of the NSBU model for our consumption, equity, and housing data separately. The sample consists of countries for which we have data prior to World War 1 (Australia, Belgium, Denmark, France, Norway, Sweden, UK, USA) in order to ensure sufficient disasters in the sample and hence allow the model to properly estimate disaster parameters. Since we directly compare estimation results for assets and consumption, it is furthermore crucial to have identical data availability among all three models. Table 13 presents the estimates of our main disaster parameters.

The first row displays the probability of remaining in a disastrous state if a country is currently exposed to a disaster. Comparing the different assets, we conclude that equity disasters are less likely to continue followed by consumption and housing disasters. Based on these numbers we compute the average disaster duration, which equals 4 years for consumption claims, and roughly 3.5 years and 5 years for equity and housing disasters, respectively. Turning attention to the second row, i.e., short-run drops, we conclude that equity disasters are characterized by much larger drops than housing, whereas consumption is in between. Finally, in the third row, we see that equity disasters feature almost no reversal, while housing disasters are associated with a more than complete recovery. Hence, based on this exercise, we can conclude that—compared to housing disasters—equity disasters are of shorter duration, feature more significant within-disaster drops, and experience less recovery.

Figure 5: *Impulse Response of a Typical Disaster*



Notes: This figure shows the median percentage change in D/P ratios in a ± 4 -year disaster window. We set parameter estimates for long-run and short-run drops equal to the posterior means as displayed in Table 13. The average disaster duration equals 3.5 years for equity and 5 years for housing. All idiosyncratic shocks are shut-off and the mean growth rate is set to zero.

We can visualize results via impulse response functions for “typical disasters” which, following Nakamura, Steinsson, Barro, and Ursúa (2013), we define as disasters with average duration and with the short-run and long-run effects equal to their respective posterior means, as in Table 13. All remaining parameters, in particular the idiosyncratic shocks, are set equal to zero.

These impulse response functions are shown in Figure 5 and they resemble the parameter estimates highlighted in Table 13. Specifically, we observe much larger within-disaster drops for equity as compared to housing claims. Further, the graphs are very illustrative with respect to the long-term impact. Housing disasters are more than offset by subsequent recoveries, while reversal is very small for equity disasters. Finally, the graphs highlight a stronger peak-to-trough decline for equity, consistent with the above replication of Barro (2006) in the previous subsection.

The above evidence clearly suggests that equity disasters are worse than housing disasters. However we still lack one important puzzle piece before we can infer a larger equity premium. The missing piece relates to the co-movement of asset disasters with consumption disasters. Risk premiums are ultimately determined by the riskiness of a particular asset in conjunction with its co-movement with consumption. If equity disasters were totally unrelated to consumption disasters while housing disasters were not, high marginal utilities due to a consumption disaster would not align with drops in stock prices and the risk premium would be small or even zero, despite all of the above.

However, we can provide evidence against this hypothesis too, based on correlations between

disaster probabilities for each country. We restrict the sample to observations with a consumption disaster probability of at least 50% since we are primarily interested in correlations during consumption disaster periods. The numbers are however robust to different specifications. The results are striking: The correlation between consumption and housing disaster probabilities is negative (-0.04), but for consumption and equity disasters it is positive (0.19). Consumption disasters tend to coincide with equity disasters in terms of likelihood, while there is no such pattern for housing disasters, which strengthens our previous argument.⁸

To sum up, we are confident that the Nakamura, Steinsson, Barro, and Ursúa (2013) extensions, i.e., partial recovery and disasters that unfold over multiple periods, will likely increase the equity premium vis-à-vis the baseline Barro (2006) model, but it does not work for housing in the same way. Equity declines during disasters are more severe than housing declines and recovery is essentially non-existent. Moreover, the probability of equity disasters co-moves with consumption disasters, which is also consistent with a sizable risk premium for equity. Yet this is not the case for housing.

6.3. Muir (2017) with Two Risky Assets

The rare disaster approach teaches us that examining the performance of assets during consumption disasters could be a useful starting point for evaluating risk premiums. Having reviewed leading asset pricing models which feature particularly severe disasters, i.e., the rare disaster models of Barro (2006) and Nakamura, Steinsson, Barro and Ursúa (2013), we now provide empirical facts regarding the performance of equity and housing dividend yields during different types of disasters, a granular approach suggested by Muir (2017).

Consumption disasters are inherently volatile and by definition they are associated with drops in consumption as shown in Table 14. Generally speaking, assets that perform well when marginal utilities are high should be associated with a low required return. Further, since households care more about periods in which the marginal utility is high, the risk premium should be largely determined by asset behavior during disasters.

We find most notably that housing yields fail to respond to disasters at their onset and display heterogeneous behavior afterwards. Equity dividend yields in contrast, co-move well with the arrival of a disasters, independent of the type of disaster, and level-off afterwards. Based on these observations, we would expect a higher premium for equity returns compared to housing returns; however as we documented in the main body of the paper, this does not seem to be the case.

In line with Muir (2017), we define four different types of disasters—recessions, deep recessions, financial crises, and wars—and compare dividend price ratios of equity and housing claims before and after these disasters. Recessions are described as consumption declines smaller than 2%, deep recessions as collapses of more than 2% in a given year. Wars are characterized as consumption

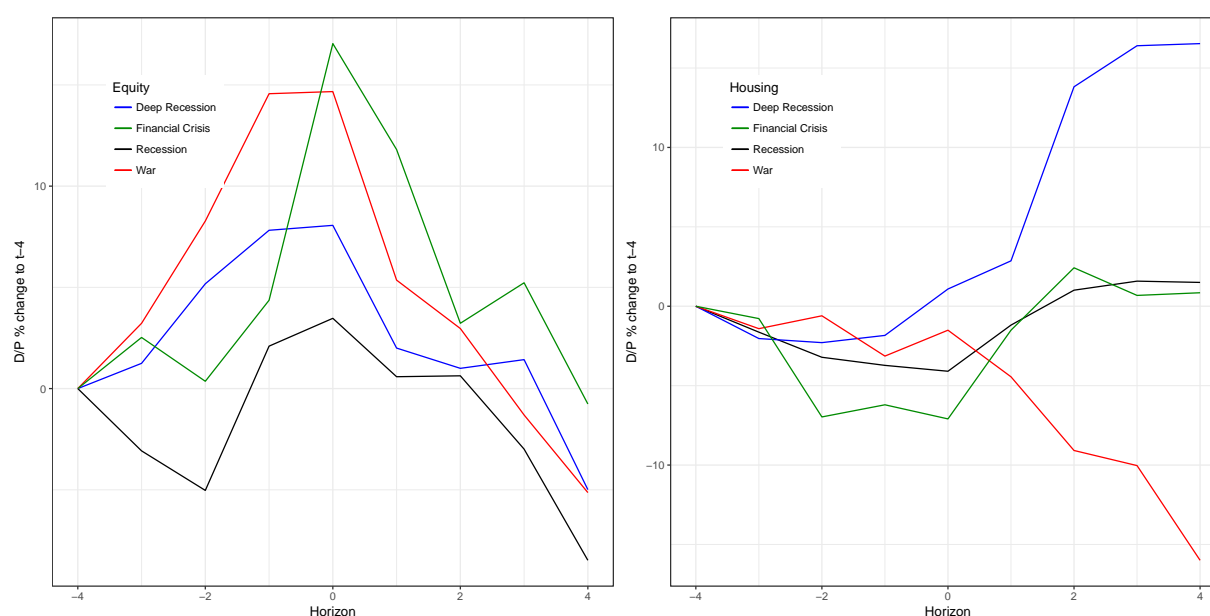
⁸Values are computed separately for each country and then averaged across all 8 countries. Since we are predominately interested in correlations during consumption disasters, correlations are based on a subset in which the posterior mean consumption disaster probability is at least 50%.

Table 14: Disaster characteristics following Muir (2017)

Event Type	Recession	Deep Recession	War	Financial Crisis
Peak to Trough Decline	1.20	5.34	16.24	1.59
σ_c	2.81	4.04	9.48	3.91
N	129	89	32	79

Notes: Based on our data. The median peak to trough decline is defined as the median of the maximum consumption growth declines over a five year window around the disaster date. σ_c represents the median standard deviation of consumption growth during 10 years after the disaster date.

Figure 6: Change in Dividend Price Ratios around Consumption Disasters



Notes: Median disaster percentage change in D/P ratios around a four period window. Recessions are described as consumption declines smaller than 2%, deep recessions as collapses of more than 2% in a given year. Wars are characterized as consumption drops during war periods and financial crises are defined as in Jordà, Schularick, and Taylor (2017). Consecutive disaster events up to five years are aggregated into one single event to avoid double counting.

drops during war periods and financial crises are defined as in Jordà, Schularick, and Taylor (2017), who follow Laeven and Valencia (2012) in determining financial crises as “events during which a country’s banking sector experiences bank runs, sharp increases in default rates accompanied by large losses of capital that result in public intervention, bankruptcy, or forced merger of financial institutions.” We use dividend yields as a well-established proxy for expected risk premiums, since dividend yields strongly predict future stock returns and only weakly forecast future dividend growth (Shiller 1981; Campbell 1988; Campbell and Cochrane 1999).

Further insight can be gleaned from Figure 6 which displays the median dividend price change with reference to 4 years prior to a consumption disaster event split by asset type, with equity on the left and housing on the right. The left panel highlights that changes in equity dividend yields are well aligned with the arrival of consumption disaster at period $t = 0$. Deep recessions are worse than regular recessions in terms of dividend yields as well, while both are topped by war and financial crises. Financial crises tend to be the worst, exactly in line with the finding of Muir (2017). But, as he noted, this finding is puzzling, since wars are associated with the largest peak to trough consumption declines as well as the largest volatility following a disaster as is evident in Table 14.

We repeat the same exercise for housing dividend yields in the right panel of Figure 6. Most notably, in contrast to equity claims, housing claims do not contemporaneously respond to consumption disasters and they show a heterogeneous response afterwards. Due to the well established link between dividend yields and the risk premium, we would hence expect a larger risk premium for equity versus housing. However, as discussed in the main body of the paper, we observe roughly equal risk premiums in the data.

6.4. Campbell and Cochrane (1999) with Two Risky Assets

In an alternative approach to resolving the equity risk premium puzzle, Campbell and Cochrane (1999) augment the standard Mehra and Prescott (1985) asset pricing model with a slow moving external habit process following Abel (1990).⁹ In contrast to the Barro (2006) rare disasters model, risk premiums are internally propagated and thus don't rely on large, infrequent exogenous consumption drops. Intuitively, habits penalize varying consumption growth rates and households care about their relative performance vis-a-vis the habit level rather than declines in the actual absolute consumption level.

We re-evaluate the original Campbell and Cochrane (1999) habits model and fit it to our post-WW2 balanced sample (1963–2015) for equity, housing, and total returns. We find that habit formation is able to explain observed equity risk premiums but not housing premiums: that is, as in the original study, we are able to match the equity risk premium with a coefficient of relative risk aversion of about 1, but meanwhile an implausibly high risk aversion parameter, way beyond 10, is then required to match the housing risk premium.

The intuition is that this result is driven by the quite different serial correlation properties of dividend/price and rent/price ratios. The autocorrelation is around 0.85 for equity dividend/price ratios but is a much higher 0.96 for housing rent/price ratios. This means that housing rent/price ratios are going to be much more sluggish in their adjustments, since they have much more powerful persistence. Note that we also considered the entire sample, and re-ran the autocorrelation estimates. This resulted in a first order auto-correlation of 0.52 for equity and 0.96 for housing, so this result is robust. Our results based on the post-WW2 period instead of the entire sample are therefore quite

⁹Uhlig (2007) and Grishchenko (2010), among others, present more recent evidence on habits as a way of explaining risk premiums.

Table 15: *Campbell and Cochrane (1999) with Two Risky Assets: Risk Premiums versus Risk Aversion*

	Risk premium (log times 100)					
	$\gamma = 2$			$\gamma = 10$		
	Equity	Housing	Total	Equity	Housing	Total
Model, $\Phi = 0.85$	5.63	6.52	6.49	7.63	7.23	7.88
Model, $\Phi = 0.90$	3.07	4.10	3.98	4.46	4.50	4.91
Data	3.74	5.38	5.34	3.74	5.38	5.34

Notes: γ is the CRRA parameter and Φ governs the persistence of habits. A large value is associated with volatile habits and hence a muted impact on the risk premium. Calibration: $g_c=0.02$, $\sigma_c=0.024$, $r_f=1.03$, $\sigma_{eq}=0.25$, $\rho(\Delta c, \Delta d_{eq})=0.10$, $\sigma_{hou}=0.043$, $\rho(\Delta c, \Delta d_{hou})=0.20$, $\sigma_{tot}=0.13$, $\rho(\Delta c, \Delta d_{tot})=0.14$. Post-WW2 balanced sample (1963-2015). The model is estimated using a log scale, and units are log times 100.

conservative.

However, for any asset, a more persistent dividend process will require a more smooth “effective risk aversion.” However, as we now explain, a smooth “effective risk aversion” is then associated with weak or quickly adjusting habits. In detail, in the Campbell and Cochrane (1999) habits model, the effect of external habits on the curvature of the utility function can be described by the “effective relative risk aversion” (η_t), which is defined as and equal to

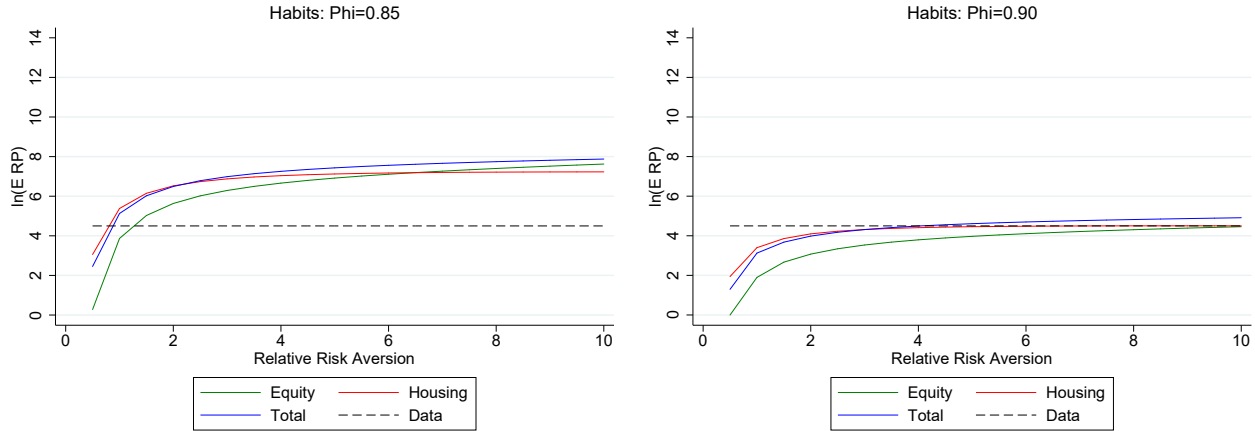
$$\eta_t \equiv \frac{C_t U_{cc}}{U_c} = \frac{\gamma}{S_t}. \quad (9)$$

where U_c (U_{cc}) is the first (second) derivative of a standard CRRA utility function with respect to consumption and S_t represents the surplus consumption ratio defined as $S_t = \frac{C_t - X_t}{C_t}$. Here, X_t characterizes the external habit, and γ indicates risk aversion absent habits. However with habits, curvature is enhanced for low values of S_t , that is, when actual consumption C_t is close to habits X_t . As a result of this property of the model, Campbell and Cochrane (1999) are able to generate reasonable risk premiums on consumption claims even for modest levels of γ .

A crucial target for fitting this model to the data is the serial correlation of equity and housing dividend price ratios, which—as noted above—is significantly higher for housing (0.96) than for equity (0.85) in the post-WW2 balanced sample. Persistence in Campbell and Cochrane (1999) is mainly driven by the adjustment process of the consumption habit characterized by Φ . A large value of Φ implies slowly moving surplus consumption ratios, which is equivalent to swift adjustments in habits, i.e., weak habits. Hence, the observed serial correlation of equity streams is associated with stronger habits than housing dividend yields, which in turn corresponds to a large premium for equity returns. Table 15 quantifies this verbal argument, where note that returns and premiums are now measured using a log scale, given this is the correct definition for the habits model setup.

A few features are worth highlighting. Foremost, the results are very sensitive to the degree of habit persistence. A change of Φ by just 0.05 corresponds to a 2–3 log point change in risk premium

Figure 7: *Campbell and Cochrane (1999) with Two Risky Assets: Risk Premiums versus Risk Aversion*



Notes: Model implied risk premiums for varying degrees of relative risk aversion (γ). The dashed line represents an equally weighted risk premium between equity and housing for comparison. Calibration: $g_c=0.02$, $\sigma_c=0.024$, $r_f=1.03$, $\sigma_{eq}=0.25$, $\rho(\Delta c, \Delta d_{eq})=0.10$, $\sigma_{hou}=0.043$, $\rho(\Delta c, \Delta d_{hou})=0.20$, $\sigma_{tot}=0.13$, $\rho(\Delta c, \Delta d_{tot})=0.14$. Post-WW2 balanced sample (1963-2015). The model is estimated using a log scale, and units are log times 100.

(where 1 log point = 0.01 times 100). We are able to explain the observed equity risk premium (3.74%) with a risk aversion parameter of less than 2, as is clear in row one. However, housing dividend yields are more persistent, hence the appropriate row is $\Phi=0.90$ (or higher still). We see here that a relative risk aversion of 10, which is considered the upper bound of reasonable values in the literature, is still insufficient to match observed risk premiums (4.50 versus 5.38 log points).

The sensitivity of the model with respect to habits is further illustrated in Figure 7. The plots are an extension to the previous table and provide model implied risk premiums (vertical axis) for a range of “reasonable” relative risk aversion parameters (horizontal axis). The left figure relates to stronger habits and should be consulted to gauge the risk premium on equity, while the right figure features weaker habits and is hence associated with housing returns. Nevertheless, we plot lines for equity, housing and total returns in both graphs. The graphs visualize the substantial impact of habits on the risk premium. Risk premiums are significantly shifted down by roughly 2%-3% despite a modest change in habit strength. Furthermore, the implied risk premiums with weak habits level off below 5% at higher relative risk aversion. Thus, it is very difficult to reconcile the observed housing risk premium with any reasonable choice of relative risk aversion.

6.5. Bansal and Yaron (2004) with Two Risky Assets

A different approach to resolving the equity premium puzzle proceeds by incorporating highly persistent long-run risk and news about consumption growth. This idea was pioneered by Bansal and Yaron (2004) and they showed that both components in combination with Epstein-Zin-Weil

preferences (Epstein and Zin 1989; Weil 1989) can yield a risk premium of roughly 7% for a market portfolio based on CRSP with reasonable choices for the relative risk aversion (10) and the intertemporal elasticity of substitution (1.5).

In this model, the mechanism depends on high persistence in both components such that a small shock is able to create a long-lasting almost permanent effect on dividend growth and volatility which ultimately manifests in large required risk premiums. We will argue that this framework is not suitable to resolve the housing or total risk premium puzzles.

The intuition for our claim is as follows. Since observed equity and housing returns are roughly in line, the model implies similar dividend processes with respect to the impact of the risk and news component. Yet we point out that with an almost identical calibration, model based moments are inconsistent with data.

Specifically, the Bansal and Yaron (2004) model induces a tight link between correlation of the dividend and consumption processes and the resulting model-implied risk premium. Since we observe higher correlations between equity returns and consumption as compared to housing returns and consumption, it is inevitable that the model cannot generate similar risk premiums for equity and housing.

The previously mentioned link between risk premium and correlation manifests in the news component, which is the only source of co-movement among assets and consumption in the model. At the same time news about future growth is crucial for deriving a large risk premium even more than the risk component. Consequently, the model implicitly ties co-movement with risk premiums for reasonably calibrated values.

In the data, equity returns are more volatile than housing returns. A proper calibration therefore requires us to either increase the dependence on uncertainty and/or news, as both, *ceteris paribus*, increase volatility. Raising the dependence on the predictable growth component generates the aforementioned problem, but increasing the impact of uncertainty generates another problem: It decreases the correlation between consumption and dividend stream. Hence both options, necessary for a good calibration, cause the model to fail when attempting to reconcile the housing and total risk premium puzzles.

In what follows we provide theoretical and empirical support by expanding the Bansal and Yaron (2004) model. The dynamics of the our expanded model are characterized by four equations: a consumption and dividend growth process as well as processes guiding the news component and uncertainty. x_t represents news, g_t and $g_{j,t}$ the consumption and asset growth rate where $j \in \{\text{Equity}, \text{Housing}\}$. σ_t^2 refers to time-varying economic uncertainty. All remaining parameters correspond to unconditional means, loadings to uncertainty and news, and the persistence thereof.

Table 16: *Bansal and Yaron (2004) with Two Risky Assets: Moments*

	Consumption		Housing		Equity	
	Data	Model	Data	Model	Data	Model
σ	0.06	0.05	0.08	0.11	0.28	0.15
AC(1)	0.12	0.26	0.45	0.39	-0.02	0.59
AC(2)	0.04	0.19	0.24	0.29	-0.06	0.44
$corr(g, g_j)$			0.18	0.39	0.24	0.48

Notes: The model is simulated in monthly frequency. Model parameters are based on the process characterized by Equations (1) - (4). The parameters are $\mu = 0.00134$, $\mu_e = 0.000143$, $\mu_h = 0.000775$, $\sigma = 0.0078$, $\sigma_\omega = 0.23 \times 10^{-5}$, $\chi_e = 4.5$, $\chi_h = 4.5$, $\phi_e = 5$, $\phi_h = 3$, $\rho = 0.979$, $\nu = 0.987$, Intertemporal Elasticity of Substitution=1.5, Relative Risk Aversion = 10. $AC(n)$ refers to the auto-correlation of lag n . The model statistics are based on 1,000 simulations each with 1,740 monthly observations to resemble the 145 years of data availability. The values display the annualized average over all samples and periods.

$e_{t+1}, \eta_{t+1}, v_{j,t+1}, \omega_{t+1}$ are i.i.d. normal distributed shocks with mean zero and variance one.

$$x_{t+1} = \rho x_t + \chi \sigma_t e_{t+1} \quad (10)$$

$$g_{t+1} = \mu + x_t + \sigma_t \eta_{t+1} \quad (11)$$

$$g_{j,t+1} = \mu_j + \phi_j x_t + \chi_j \sigma_t v_{j,t+1} \quad (12)$$

$$\sigma_{t+1}^2 = \sigma^2 + \nu(\sigma_t^2 - \sigma^2) + \sigma_\omega \omega_{t+1} \quad (13)$$

With idiosyncratic error terms, the co-variance between dividend growth rates and consumption growth rates is pinned down by ϕ_j and the volatility for each asset by a combination of ϕ_j and χ_j . Further, distinct characteristics of equity and housing returns are characterized by the two parameters ϕ_j and χ_j only.

We calibrate the model using values in Bansal and Yaron (2004) for all parameters except the mean growth rates, which are straightforward to compute from our raw data, with one exception. The key difference relates to coefficient governing the loading for the news component. To account for the stronger correlation between equity and consumption streams plus the higher volatility of the former, we set the value to $\phi_e = 5$ for equity and $\phi_h = 3$ for housing. As we just explained, modifying the loading towards uncertainty generates a problematic trade-off between volatility and co-movement, hence we there decided to keep the value proposed in Bansal and Yaron (2004).

Table 16 highlights the impact of distinct news loading on the correlation with consumption and volatility. The large value for equity clearly leads to both higher volatility as well as a stronger correlation with consumption growth. The model fit is reasonable with respect to the standard deviation and correlation, but poor with respect to the autocorrelation. This finding is in line with Bansal and Yaron (2004) and due to equation (12), which suggests that comovement with consumption can only be achieved via higher autocorrelation.

Table 17: *Bansal and Yaron (2004) with Two Risky Assets: Risk Premiums*

	Consumption	Housing	Equity
$E_t(r_{t+1} - r_{f,t+1})$	2.99%	7.98%	13.94%

Notes: The model statistics are based on 1,000 simulations each with 1,740 monthly observations to resemble the 145 years of data availability. The values display the annualized average over all samples and periods.

The model predictions for risk premiums are summarized in Table 17. Most importantly, the annualized risk premium for equity is significantly higher than for housing. The difference in the risk premium equals 6% and purely results from a distinct loading towards news. Thus fitting the dividend growth equation to data yields a much larger risk premium for equity returns than for housing returns, whereas in the data these are equal.

To conclude, we are not able match the observed risk premiums for our equity and housing data. While the overall level of excess returns could be lowered by, for example, slightly lowering the persistence of the predictable consumption component, the qualitative difference between the risk premiums remains. Similar risk premiums require a similar calibration of the processes governing the dividend growth rates, which ultimately yields inconsistent simulated moments.

7. CONCLUSIONS

The risk premium puzzle is worse than you think. The unadorned, benchmark, representative-agent, consumption-based asset pricing model implies an absurdly high risk-aversion parameter, because equity returns have low risk (low covariance with consumption-growth) but high excess returns (relative to safe assets). A gigantic research program has sought to ameliorate or explain away this puzzle, but the bar has been set too low: by ignoring half of total wealth invested outside firm equity, specifically in housing, the literature has hitherto avoided a more serious challenge.

Using a new database for the U.S. and 15 other advanced economies, from 1870 to the present, including housing as well as equities, we perform standard calculations using returns to equity, housing, and total wealth, and consumption growth, to show that the total risk premium puzzle is even bigger. Housing returns are comparable to those of equities and yet housing returns have even lower covariance with consumption growth than equities. The same also holds for a total-wealth portfolio, and over a range of horizons.

The implied risk aversion parameters for housing wealth and total wealth are even larger than those for equities, often by a factor of 2 or more. A variety of standard adjustments to the benchmark representative-agent model also appear to be unable to resolve these even bigger puzzles, since in all models, whether standard or exotic, the time series properties of housing and equity returns imply that the former should have a much lower risk premium than the latter, yet this is contradicted by the historical data.

REFERENCES

- Abel, Andrew. 1990. Asset Prices under Habit Formation and Catching Up with the Joneses. *American Economic Review* 80(2):38–42.
- Ait-Sahalia, Yacine, Jonathan A. Parker, and Motohiro Yogo. Luxury Goods and the Equity Premium *Journal of Finance* 59(6): 2959–3004
- Attanasio, Orazio P., and Guglielmo Weber. 1995. Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy* 103(6): 1121–1157.
- Bachmann, Rüdiger, and Cooper, Daniel. 2014. The Ins and Arounds in the U.S. Housing Market. CEPR Discussion Paper 10041.
- Bansal, Ravi, and Amir Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59(4): 1481–1509.
- Barras, Laurent, and Betermier, Sebastien. 2016. The Predominance of Real Estate in the Household Portfolio. SSRN. <https://ssrn.com/abstract=2776350>.
- Barro, Robert J. 2006 Rare Disasters and Asset Markets in the Twentieth Century. *Quarterly Journal of Economics* 121(3): 823–866.
- Barro, Robert J., and José F. Ursúa. 2008. Macroeconomic Crises since 1870. *Brookings Papers on Economic Activity* 39(1): 255–350.
- Beaudry, Paul, and Eric van Wincoop. 1996. The Intertemporal Elasticity of Substitution: An Exploration Using a US Panel of State Data. *Economica* 63(251): 495–512.
- Brav, Alon, George M. Constantinides, and Christopher C. Geczy. 2002. Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence. *Journal of Political Economy* 110(4): 793–824.
- Breeden, Douglas T., Robert H. Litzenberger, and Tingyan Jia. 2015a. Consumption-Based Asset Pricing, Part 1: Classic Theory and Tests, Measurement Issues, and Limited Participation. *Annual Review of Financial Economics* 7(1): 35–83.
- Breeden, Douglas T., Robert H. Litzenberger, and Tingyan Jia. 2015b. Consumption-Based Asset Pricing, Part 2: Habit Formation, Conditional Risks, Long-Run Risks, and Rare Disasters. *Annual Review of Financial Economics* 7(1): 85–131.
- Brown, Stephen J., William N. Goetzmann, and Stephen A. Ross. 1995. Survival. *Journal of Finance* 50 (3): 853–873.
- Campbell, John Y. 1999. Asset Prices, Consumption, and the Business Cycle. In *Handbook of Macroeconomics*, volume 1, edited by J. B. Taylor and M. Woodford. Amsterdam: Elsevier, pp. 1231–1303.
- Campbell, John Y. 1988. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies* 1(3): 195–228.
- Campbell, John Y. 2000. Asset Pricing at the Millennium. *Journal of Finance* 55(4): 1515–1567.
- Campbell, John Y. 2003. Consumption-Based Asset Pricing. In *Handbook of the Economics of Finance*, volume 1, edited by G.M. Constantinides, M. Harris, and R. M. Stulz. Amsterdam: Elsevier, pp. 803–887.
- Campbell, John Y. 2017. *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton, N.J.: Princeton University Press.
- Campbell, John Y., and John H. Cochrane. 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107(2): 205–251.
- Campbell, John Y., and John H. Cochrane. 2000. Explaining the Poor Performance of Consumption-Based Asset Pricing Models. *Journal of Finance* 55(6): 2863–2878.

- Case, Karl E., and Robert J. Shiller. 1987. Prices of Single Family Homes Since 1970: New Indexes for Four Cities. NBER Working Paper 2393.
- Case, Karl E., and Robert J. Shiller. 1989. The Efficiency of the Market for Single-Family Homes. *American Economic Review* 79(1): 125–137.
- Cheng, Ping, Zhenguang Lin, and Yingchun Liu. 2008. The Real Estate Risk Premium Puzzle: A Solution. Unpublished.
- Cochrane, John H. 2001. *Asset Pricing*. Princeton, N.J.: Princeton University Press.
- Cochrane, John H. 2008. Financial Markets and the Real Economy. In *The Handbook of the Equity Risk Premium* edited by R. Mehra. Amsterdam: Elsevier, pp. 237–325.
- Constantinides, George M. 2008. Macroeconomic Crises Since 1870: Comments and Discussion. *Brookings Papers on Economic Activity* 39(1): 336–350.
- Davis, Morris A., and Robert F. Martin. 2013. Housing, House Prices, and the Equity Premium Puzzle. Federal Reserve Board, Finance and Economics Discussion Series 2005-13.
- Davis, Morris A., and Stijn Van Nieuwerburgh. 2014. Housing, Finance and the Macroeconomy. NBER Working Paper No. 20287.
- Dimson, Elroy, Paul Marsh, and Mike Staunton. 2002. *Triumph of the Optimists: 101 Years of Global Investment Returns*. Princeton, N.J.: Princeton University Press.
- Dimson, Elroy, Paul Marsh, and Mike Staunton. 2008. The Worldwide Equity Premium: A Smaller Puzzle. In *The Handbook of the Equity Risk Premium* edited by R. Mehra. Amsterdam: Elsevier, pp. 467–514.
- Eisfeldt, Andrea, and Andrew Demers. 2015. Total Returns to Single Family Rentals. NBER Working Paper 21804.
- Epstein, Larry G., and Stanley E. Zin. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57(4): 937–969.
- Flavin, Marjorie, and Xuan Liang. 2013. The Housing CCAPM with Adjustment Costs and Heterogeneous Agents. UCSD. Photocopy.
- Flavin, Marjorie, and Shinobu Nakagawa. 2008. A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence. *American Economic Review* 98(1): 474–495.
- Flavin, Marjorie, and Takashi Yamashita. 2002. Owner-Occupied Housing and the Composition of the Household Portfolio. *American Economic Review* 92(1): 345–362.
- Giacoletti, Marco. 2016. Idiosyncratic Risk in Housing Markets. Stanford University. Unpublished. <https://www.gsb.stanford.edu/sites/gsb/files/phd-job-market-paper/marcogiacolettijmp.pdf>.
- Goetzmann, William N. 1993. The Single Family Home in the Investment Portfolio. *Journal of Real Estate Finance and Economics* 6: 201–222.
- Goetzmann, William N., and Roger G. Ibbotson. 1990. The Performance of Real Estate as an Asset Class. *Journal of Applied Corporate Finance* 3(1): 65–76.
- Gourio, Francois. 2008. Disasters and Recoveries. *American Economic Review* 98(2):68–73.
- Grishchenko, Olesya V. 2010. Internal vs. External Habit Formation: The Relative Importance for Asset Pricing. *Journal of Economics and Business* 62(3): 176–194.
- Hansen, Lars Peter, and Ravi Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99(2): 225–262.
- He, Hua, and David M. Modest. 1995. Market Frictions and Consumption-Based Asset Pricing. *Journal of Political Economy* 103(1): 94–117.
- Iacoviello, Matteo M. 2011. Housing Wealth and Consumption. Federal Reserve Board, International Finance Discussion Paper 1027.
- Ibbotson, Roger G. and Laurence B. Siegel. 1984. Real Estate Returns: A Comparison with Other Investments. *Journal of the American Real Estate and Urban Economics Association* 12(3): 219–42.

- Jagannathan, Ravi, and Yong Wang, 2007. Lazy Investors, Discretionary Consumption, and the Cross-Section of Stock Returns. *Journal of Finance* 62(4): 1623–1661.
- Jones, Charles M. 2002. A Century of Stock Market Liquidity and Trading Costs. SSRN <https://ssrn.com/abstract=313681>.
- Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2010. Financial Crises, Credit Booms, and External Imbalances: 140 Years of Lessons. *IMF Economic Review* 59(2): 340–378.
- Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2016. 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.
- Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor. Forthcoming. The Rate of Return on Everything, 1870–2015. *Quarterly Journal of Economics*.
- Jorion, Philippe, and William N. Goetzmann. Global Stock Markets in the Twentieth Century. *Journal of Finance* 54 (3): 953–980.
- 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.
- Krainer, John. 2001. A Theory of Liquidity in Residential Real Estate Markets. *Journal of Urban Economics* 49(1): 32–53.
- Kuhn, Moritz, Moritz Schularick, and Ulrike Steins. 2017. Income and Wealth Inequality in America, 1949–2013. CEPR Discussion Papers 12218.
- Landvoigt, Tim, Monika Piazzesi, and Martin Schneider. 2015. The Housing Market(s) of San Diego. *American Economic Review* 105(4): 1371–1407.
- Liang, Xuan. 2015. Implications of Homeownership for Endogenous Risk Aversion, Asset Pricing and Portfolio Composition. Ph.D. dissertation. UC San Diego.
- Lin, Zhenguang, and Kerry D. Vandell. 2007. Illiquidity and Pricing Biases in the Real Estate Market. *Real Estate Economics* 35(3): 291–330.
- Longstaff, Francis A. 2018. Valuing Thinly Traded Assets. *Management Science*. 64(8): 3868–3878.
- Lustig, Hanno, and Stijn Van Nieuwerburgh. 2006. Can Housing Collateral Explain Long-Run Swings in Asset Returns? NBER Working Papers 12766.
- Luttmer, Erzo G. J. 1996. Asset Pricing in Economies with Frictions. *Econometrica* 64(6): 1439–1467.
- Mankiw, N. Gregory, and Stephen Zeldes. 1991. The Consumption of Stockholders and Nonstockholders. *Journal of Financial Economics* 29(1): 97–112.
- Mehra, Rajnish. 2007. The Equity Premium Puzzle: A Review. *Foundations and Trends in Finance* 2(1): 1–81.
- Mehra, Rajnish. 2012. Consumption-Based Asset Pricing Models. *Annual Review of Financial Economics* 4(1): 385–409.
- Mehra, Rajnish, and Edward C. Prescott. 1985. The Equity Premium: A Puzzle. *Journal of Monetary Economics* 15(2): 145–161.
- Mehra, Rajnish, and Edward C. Prescott. 2003. The Equity Premium in Retrospect. In *Handbook of the Economics of Finance*, Volume 1, edited by G. M. Constantinides, M. Harris, and R. M. Stulz. Amsterdam: Elsevier, pp. 889–938.
- Muir, Tyler. 2017. Financial Crises and Risk Premia. *Quarterly Journal of Economics* 132(2): 765–809.
- Nakamura, Emi, Jón Steinsson, Robert J. Barro, and José Ursúa. 2013. Crises and Recoveries in an Empirical Model of Consumption Disasters. *American Economic Journal: Macroeconomics* 5(3): 35–74.

- Ngai, L. Rachel, and Silvana Tenreyro. 2014. Hot and Cold Seasons in the Housing Market. *American Economic Review* 104(12): 3991–4026.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2007. Housing, Consumption and Asset Pricing. *Journal of Financial Economics* 83(3): 531–569.
- Piazzesi, Monika, and Martin Schneider. 2009. Momentum Traders in the Housing Market: Survey Evidence and a Search Model. *American Economic Review, Papers and Proceedings*, 99(2): 406–411.
- Rietz, Thomas A. 1988. The equity risk premium a solution. *Journal of Monetary Economics* 22(1): 117–131.
- Shiller, Robert J. 1981. Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? *American Economic Review* 71(3): 421–36.
- Shilling, James D. 2003. Is There a Risk Premium Puzzle in Real Estate? *Real Estate Economics* 31(4): 501–525.
- Tracy, Joseph, Henry Schneider, and Sewin Chan. 1999. Are Stocks Overtaking Real Estate in Household Portfolios? *Current Issues in Economics and Finance* vol. 5, no. 5. New York, N.Y.: Federal Reserve Bank of New York, pp. 1–6.
- Uhlig, Harald. 2007. Explaining Asset Prices with External Habits and Wage Rigidities in a DSGE Model. *American Economic Review* 97(2): 239–243.
- Vissing-Jørgensen, Annette. 2002. Limited Asset Market Participation and the Elasticity of Intertemporal Substitution. *Journal of Political Economy* 110(4): 825–853.
- Weil, Philippe. 1989. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics* 24(3): 401–421.
- Yoshida, Jiro. 2008. Technology Shocks and Asset Price Dynamics: The Role of Housing in General Equilibrium. University of Tokyo. Unpublished.
- Zimmermann, Kaspar. 2017. Breaking Banks? Bank Profitability and Monetary Policy. University of Bonn. Unpublished.

APPENDICES

A. Derivation of the Extended Barro (2006) model with Risky Assets

The original Barro (2006) model features a standard Lucas-Tree type model with an exogenous endowment, as in Mehra and Prescott (1985), augmented for one period rare disasters. We modify the economy and introduce risky assets. To keep things tractable and derive a closed form solution, we assume that asset disasters coincide with consumption disasters. Hence, a consumption crisis automatically triggers an asset crises, whereat the magnitude of the asset crisis is determined based on empirical grounds. To fix ideas, Barro (2006) assumes the following exogenous evolution of the endowment:

$$\log(A_{t+1}) = \log(A_t) + \gamma + u_{t+1} + v_{t+1} \quad (14)$$

where A_t denotes endowment which is equal to consumption in equilibrium, γ its long-run growth rate and u_t is i.i.d. normal with mean zero and variance σ^2 . The term v_t captures rare disasters and equals $\log(1 - b)$ in bad times with probability p and zero with probability $1 - p$. b in turn is a random variables and captures the fractional consumption decline during disasters, which we derive based on peak to trough consumption declines. We assume dividends follow a similar process and denote variables related to assets with a tilde.

$$\log(\tilde{A}_{t+1}) = \log(\tilde{A}_t) + \tilde{\gamma} + \tilde{u}_{t+1} + \tilde{v}_{t+1} \quad (15)$$

Each of the parameters associated with the dividend process is calibrated on asset data and hence distinct from the parameters chosen for the consumption claim. The only exception pertains to asset disaster probabilities. In order to derive an analytical solution we assume that asset disasters coincide with consumption disasters and calibrate this probability based on consumption data only. However, we do not make an assumption on the relationship between the fractional declines of consumption (b) and assets (d) during disasters. We further introduce correlation between consumption and dividend claims during normal times, hence we allow for a non-zero correlation (κ) between u_{t+1} and \tilde{u}_{t+1} .

In the remainder of the appendix, we derive a closed form solution for the asset return $\log E_t[\tilde{R}_{t+1}]$ which we combine with the formula for the risk free rate in Barro (2006) to derive the risk premium presented in the main body of the paper. As in Barro (2006) we assume a standard CRRA utility function. $\rho \geq 0$ deontes the rate of time preference and $\theta > 0$ the coefficient of relative risk aversion. The Euler Equation for asset holdings corresponds to:

$$A_t^{-\theta} = e^{-\rho} E_t[A_{t+1}^{-\theta} \tilde{R}_{t+1}] \quad (16)$$

which is equivalent to

$$\tilde{P}_t = e^{-\rho} E_t[A_{t+1}^{-\theta} \tilde{A}_{t+1}] A_t^\theta \quad (17)$$

due to the definition of the risky asset return $\tilde{R}_{t+1} = \frac{\tilde{A}_{t+1}}{\tilde{P}_t}$. We can rewrite the expectations operator as

$$E_t[A_{t+1}^{-\theta} \tilde{A}_{t+1}] = E_t[A_t^{-\theta} e^{-\theta\gamma} e^{-\theta u_{t+1}} e^{-\theta v_{t+1}} \tilde{A}_t e^{\tilde{\gamma}} e^{\tilde{u}_{t+1}} e^{\tilde{v}_{t+1}}] \quad (18)$$

Since we condition on period t we pull out constants as well as period t variables, hence:

$$E_t[A_{t+1}^{-\theta} \tilde{A}_{t+1}] = A_t^{-\theta} e^{-\theta\gamma} \tilde{A}_t e^{\tilde{\gamma}} E_t[e^{-\theta u_{t+1}} e^{-\theta v_{t+1}} e^{\tilde{u}_{t+1}} e^{\tilde{v}_{t+1}}] \quad (19)$$

We once again focus on the expectations operator and decompose it into:

$$E_t[e^{-\theta u_{t+1}} e^{-\theta v_{t+1}} e^{\tilde{u}_{t+1}} e^{\tilde{v}_{t+1}}] = E_t[e^{-\theta u_{t+1}} e^{\tilde{u}_{t+1}}] E_t[e^{-\theta v_{t+1}} e^{\tilde{v}_{t+1}}] \quad (20)$$

Equivalence holds since idiosyncratic shocks are independent of crises. u and \tilde{u} follow a normal distribution with correlation κ . Further, the sum of normal distributed variables is normal distributed.

Hence,

$$E_t[e^{-\theta u_{t+1}} e^{\tilde{u}_{t+1}}] = E_t[e^{\tilde{u}_{t+1} - \theta u_{t+1}}] = e^{\frac{\tilde{\sigma}^2}{2} + \frac{\theta^2 \sigma^2}{2} - \theta \kappa \sigma \tilde{\sigma}} \quad (21)$$

The second term involving expectations in Equation (7) depends on three random variables. A Bernoulli distributed variable on disaster occurrence independent of the fractional decline in consumption and the asset, and the fractional declines in consumption and asset claims, b and d . We can thus rewrite the expression as

$$E_t[e^{\tilde{v}_{t+1}} e^{-\theta v_{t+1}}] = pE[(1-b)^{-\theta}(1-d)] + 1-p \quad (22)$$

Technically, the expectations operator is conditioning on period t , however disaster declines are assumed to be independently distributed across time, hence we omit indexation. Once we plug the previous equations in Equation (4), the asset price can be determined as

$$\tilde{P}_t = \tilde{A}_t e^{-\rho - \theta\gamma + \tilde{\gamma} + \frac{\tilde{\sigma}^2}{2} + \frac{\theta^2 \sigma^2}{2} - \theta \kappa \sigma \tilde{\sigma}} [pE[(1-b)^{-\theta}(1-d)] + 1-p] \quad (23)$$

Last but not least, we plug the asset price into the gross return formula. The derivation of $E_t[\tilde{A}_{t+1}]$ is analogous to Barro (2006) and omitted.

$$E_t[\tilde{R}_{t+1}] = \frac{E_t[\tilde{A}_{t+1}]}{\tilde{P}_t} = \frac{e^{\tilde{\gamma} + \frac{\tilde{\sigma}^2}{2}} [pE[1-d] + 1-p]}{e^{-\rho - \theta\gamma + \tilde{\gamma} + \frac{\tilde{\sigma}^2}{2} + \frac{\theta^2 \sigma^2}{2} - \theta \kappa \sigma \tilde{\sigma}} [pE[(1-b)^{-\theta}(1-d)] + 1-p]} \quad (24)$$

Taking logs results in the following approximation

$$\log E_t[\tilde{R}_{t+1}] \approx \rho + \theta\gamma - \frac{\theta^2 \sigma^2}{2} + \theta \kappa \sigma \tilde{\sigma} - p[E[(1-b)^{-\theta}(1-d)] + E[d] - 1] \quad (25)$$

Equation (12) nests the risky return on consumption claims in Barro (2006) with $d = b$ and $\kappa = 1$.

B. Our MS-VAR Version of the Nakamura, Steinsson, Barro, and Ursúa (2013) Model with Risky Assets

We standardize, pool all observations and estimate a Markov Switching VAR with two regimes, a crisis and a non-crisis regime.

We estimate an MS-VAR correlation matrix for the crisis regime as follows:

MS-VAR: Correlation Matrix Crisis Regime

	C	H	E
C	1.0000	0.1615	0.2542
H	0.1615	1.0000	0.1755
E	0.2542	0.1755	1.0000

The correlation matrix is based on standardized pooled data and two regimes.

We only report here the correlation matrix for the crisis regime. The striking feature of the correlation matrix is the difference between the correlation between consumption and housing (0.16), and consumption and equity (0.25).

Thus equity crises tend to be more in line with consumption crises than housing crises, which consequently strengthens our previous findings.

Despite the limitations, the VAR exercise hints towards a larger risk premium for equity.

C. Detailed Tables

Table A.1: Total returns on safe and risky assets

	$E(R) - 1$					$\sigma(R)$		
	Bills	Bonds	Equities	Housing	Total	Equities	Housing	Total
(a) Full Sample								
AUS	1.29	2.13	7.79	6.37	7.08	16.94	11.92	10.87
BEL	1.37	2.07	6.23	7.89	7.06	23.61	15.51	16.42
CHE	2.90	3.04	7.49	8.22	7.86	16.45	7.60	9.66
DEU	0.08	3.87	9.98	9.58	9.78	31.92	15.62	19.51
DNK	-0.47	0.23	3.24	6.39	4.81	22.15	10.03	13.14
ESP	2.65	3.69	7.11	7.82	7.47	21.72	10.16	12.48
FIN	1.37	2.62	7.32	4.77	6.04	28.75	9.61	15.40
FRA	0.39	1.39	6.09	6.54	6.31	19.25	8.41	11.34
GBR	0.78	1.88	7.09	7.28	7.18	21.25	9.38	12.18
ITA	0.90	2.13	5.55	8.03	6.79	19.48	8.70	11.15
JPN	-0.48	0.99	4.37	6.31	5.34	33.47	8.73	18.52
NLD	-0.23	0.50	5.46	5.21	5.34	20.58	12.00	12.52
NOR	1.56	2.35	7.98	8.30	8.14	20.12	8.88	11.25
PRT	0.72	1.65	6.71	5.63	6.17	19.48	6.66	10.45
SWE	1.15	2.03	6.90	5.44	6.17	19.90	9.15	11.04
USA	1.52	2.10	8.46	6.10	7.28	19.17	8.12	11.60
World (pooled)	1.04	2.03	6.73	6.93	6.83	21.92	10.31	12.92
(b) Post-WW2 Balanced								
AUS	1.98	2.71	7.76	7.15	7.46	21.09	5.50	10.78
BEL	2.36	3.53	10.10	8.26	9.18	22.48	6.00	11.34
CHE	2.30	3.50	10.20	5.88	8.04	23.17	7.90	13.13
DEU	0.87	2.24	12.63	9.56	11.09	35.68	8.89	19.53
DNK	1.63	2.69	5.23	7.48	6.35	23.70	6.46	11.80
ESP	1.71	3.25	7.53	5.30	6.41	21.60	4.32	10.79
FIN	1.32	2.51	4.14	5.75	4.95	27.27	10.12	13.42
FRA	0.98	1.52	5.97	5.54	5.75	20.29	6.53	11.17
GBR	1.29	2.54	8.74	7.63	8.18	21.19	9.51	11.73
ITA	0.26	2.04	8.39	9.55	8.97	27.38	7.57	14.16
JPN	-1.04	0.39	4.65	5.63	5.14	37.42	8.51	20.49
NLD	0.13	0.50	6.45	5.15	5.80	24.81	8.29	13.91
NOR	1.09	2.68	11.79	9.13	10.46	25.32	6.93	13.33
PRT	-0.11	1.22	7.48	5.60	6.54	22.03	4.49	10.74
SWE	1.63	2.60	8.79	7.10	7.94	23.29	9.09	11.74
USA	1.66	2.55	7.35	5.83	6.59	16.29	3.56	8.71
World (pooled)	1.11	2.26	7.92	6.89	7.41	25.03	7.46	13.29

Notes: This table shows the means of annual real total returns $E(R) - 1$ for bills, bonds, housing, equities, and total wealth, and standard deviations $\sigma(R)$, for 16 countries and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). The moments are computed for raw returns, and the units are percent. See text.

Table A.2: Excess returns

	$E(R - R_f^{bills})$			$E(R - R_f^{bonds})$		
	Equities	Housing	Total	Equities	Housing	Total
(a) Full Sample						
AUS	6.44	5.25	5.84	5.66	4.23	4.95
BEL	6.19	7.88	7.04	4.77	6.52	5.64
CHE	4.60	5.32	4.96	4.45	5.18	4.82
DEU	9.89	9.50	9.70	7.79	7.44	7.62
DNK	4.49	7.61	6.05	3.02	6.16	4.59
ESP	4.46	5.17	4.81	3.42	4.13	3.78
FIN	5.95	3.39	4.67	4.70	2.15	3.42
FRA	5.70	6.15	5.93	4.69	5.14	4.92
GBR	6.31	6.50	6.41	5.21	5.40	5.30
ITA	4.65	7.13	5.89	3.42	5.90	4.66
JPN	4.85	6.80	5.82	3.39	5.33	4.36
NLD	5.95	5.07	5.51	5.18	4.52	4.85
NOR	6.42	6.74	6.58	5.63	5.95	5.79
PRT	5.99	4.90	5.44	5.06	3.98	4.52
SWE	5.75	4.30	5.02	4.87	3.41	4.14
USA	6.95	4.58	5.76	6.37	4.01	5.19
World (pooled)	5.86	6.03	5.94	4.82	5.00	4.91
(b) Post-WW2 Balanced						
AUS	5.79	5.17	5.48	5.05	4.44	4.75
BEL	7.74	5.90	6.82	6.57	4.72	5.65
CHE	7.90	3.59	5.74	6.70	2.38	4.54
DEU	11.77	8.69	10.23	10.39	7.31	8.85
DNK	3.60	5.84	4.72	2.55	4.79	3.67
ESP	5.82	3.59	4.71	4.28	2.05	3.17
FIN	2.82	4.43	3.63	1.63	3.23	2.43
FRA	4.99	4.56	4.77	4.45	4.02	4.24
GBR	7.44	6.34	6.89	6.20	5.10	5.65
ITA	8.12	9.28	8.70	6.34	7.50	6.92
JPN	5.69	6.67	6.18	4.26	5.24	4.75
NLD	6.32	5.02	5.67	5.95	4.65	5.30
NOR	10.71	8.04	9.38	9.11	6.45	7.78
PRT	7.59	5.71	6.65	6.26	4.38	5.32
SWE	7.16	5.47	6.32	6.19	4.50	5.34
USA	5.69	4.17	4.93	4.81	3.28	4.05
World (pooled)	6.81	5.78	6.29	5.66	4.63	5.14

Notes: This table shows the means of annual real excess returns $E(R - R_f)$ for housing, equities, and total wealth, over bills and bonds, for 16 countries and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). The moments are computed for raw returns, and the units are percent. See text.

Table A.3: Sharpe ratios and standard deviations of consumption growth

	(a) Full Sample				(b) Post-WW2 Balanced			
	SR			$\sigma(g)$	SR			$\sigma(g)$
	Equities	Housing	Total		Equities	Housing	Total	
AUS	0.39	0.41	0.54	0.05	0.28	0.87	0.51	0.01
BEL	0.28	0.82	0.57	0.09	0.34	0.88	0.59	0.01
CHE	0.27	0.67	0.48	0.05	0.34	0.42	0.42	0.03
DEU	0.33	0.65	0.55	0.05	0.34	0.91	0.54	0.03
DNK	0.22	0.87	0.54	0.07	0.15	0.81	0.40	0.02
ESP	0.20	0.48	0.37	0.03	0.27	0.71	0.42	0.02
FIN	0.21	0.29	0.29	0.04	0.11	0.41	0.27	0.03
FRA	0.31	0.81	0.57	0.06	0.25	0.70	0.44	0.03
GBR	0.29	0.63	0.49	0.09	0.35	0.64	0.58	0.02
ITA	0.24	0.89	0.56	0.04	0.30	1.09	0.62	0.02
JPN	0.15	0.81	0.34	0.04	0.16	0.80	0.33	0.04
NLD	0.30	0.42	0.46	0.08	0.27	0.60	0.44	0.03
NOR	0.32	0.90	0.63	0.04	0.43	1.18	0.72	0.02
PRT	0.32	0.83	0.58	0.04	0.35	1.16	0.63	0.01
SWE	0.29	0.49	0.47	0.03	0.30	0.65	0.52	0.02
USA	0.38	0.54	0.52	0.03	0.35	1.00	0.57	0.02
World (pooled)	0.28	0.62	0.49	0.05	0.28	0.74	0.48	0.02

Notes: This table shows the Sharpe Ratios $SR = E(R - R_f) / \sigma(R - R_f)$ for excess returns for housing, equities, and total wealth, over bills only, for 16 countries and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). Also shown are standard deviations of real per capita consumption growth. The moments are computed for raw returns, and the units are raw, not percent. See text.

Table A.4: *Correlations of returns and consumption growth*

This table shows the correlation of returns and consumption growth $\rho(R, g)$ for housing, equities, and total wealth, for 16 countries and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). The moments are computed for raw returns. See text.

	(a) Full Sample			(b) Post-WW2 Balanced		
	$\rho(R, g)$			$\rho(R, g)$		
	Equities	Housing	Total	Equities	Housing	Total
AUS	0.23	0.02	0.18	0.25	0.28	0.31
BEL	0.03	0.18	0.10	0.19	0.40	0.29
CHE	0.09	0.19	0.15	0.26	0.33	0.33
DEU	0.25	0.24	0.30	0.37	0.49	0.45
DNK	-0.20	-0.01	-0.17	0.00	0.41	0.11
ESP	0.36	0.17	0.38	0.19	0.20	0.23
FIN	0.35	0.22	0.39	0.24	-0.06	0.22
FRA	0.22	0.39	0.33	0.22	0.48	0.34
GBR	-0.04	0.03	-0.03	0.21	0.53	0.41
ITA	0.07	-0.05	0.04	0.23	-0.03	0.21
JPN	0.37	0.23	0.39	0.39	0.36	0.43
NLD	0.00	0.14	0.06	0.31	0.53	0.44
NOR	0.10	0.04	0.11	0.40	0.37	0.48
PRT	0.25	0.38	0.35	0.21	0.04	0.22
SWE	0.28	0.18	0.32	0.27	0.30	0.38
USA	0.60	0.37	0.62	0.47	0.33	0.50
World (pooled)	0.11	0.13	0.14	0.26	0.30	0.33

Notes: This table shows the correlation of returns and consumption growth $\rho(R, g)$ for housing, equities, and total wealth, for 16 countries and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). The moments are computed for raw returns. See text.

Table A.5: Consumption betas

This table shows the consumption beta $\beta = \text{Cov}(R - R_f, g) / \text{Var}(g)$ for housing, equities, and total wealth, for 16 countries and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). The moments are computed for raw returns. See text.

	(a) Full Sample			(b) Post-WW2 Balanced		
	$\beta = \frac{\text{Cov}(R - R_f, g)}{\text{Var}(g)}$			$\beta = \frac{\text{Cov}(R - R_f, g)}{\text{Var}(g)}$		
	Equities	Housing	Total	Equities	Housing	Total
AUS	0.77	0.09	0.43	3.78	1.25	2.52
BEL	-0.07	0.22	0.07	2.66	1.27	1.97
CHE	0.15	0.13	0.14	2.46	1.14	1.80
DEU	1.95	1.09	1.52	4.54	1.56	3.05
DNK	-0.46	0.24	-0.11	0.29	1.86	1.07
ESP	2.52	0.69	1.60	2.23	0.47	1.35
FIN	3.50	1.27	2.39	2.57	-0.24	1.17
FRA	0.19	0.05	0.12	1.39	0.94	1.16
GBR	-0.06	0.07	0.01	2.25	2.52	2.39
ITA	0.40	-0.09	0.15	2.98	0.10	1.54
JPN	3.14	0.25	1.70	3.45	0.45	1.95
NLD	-0.01	0.13	0.06	3.03	1.81	2.42
NOR	0.61	0.21	0.41	5.00	1.17	3.09
PRT	0.90	0.32	0.61	3.39	0.06	1.72
SWE	1.71	0.35	1.03	2.45	0.85	1.65
USA	3.17	0.64	1.91	4.18	0.46	2.32
World (pooled)	0.43	0.23	0.33	2.78	0.93	1.85

Notes: This table shows the consumption beta $\beta = \text{Cov}(R - R_f, g) / \text{Var}(g)$ for housing, equities, and total wealth, for 16 countries and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). The moments are computed for raw returns. See text.

Table A.6: *Implied risk aversion parameters*

	(a) Full Sample			(b) Post-WW2 Balanced		
	γ			γ		
	Equities	Housing	Total	Equities	Housing	Total
AUS	33	215	53	69	243	111
BEL	147	26	48	99	185	137
CHE	42	71	56	45	55	50
DEU	28	52	37	28	66	38
DNK	-11	-1593	-33	712	146	210
ESP	21	88	35	63	236	94
FIN	24	52	29	16	-402	47
FRA	18	26	21	36	53	44
GBR	563	847	-1241	68	56	64
ITA	80	-591	290	51	-5084	124
JPN	11	86	23	10	53	18
NLD	-594	42	124	27	41	33
NOR	74	710	152	54	157	77
PRT	32	52	39	107	1868	191
SWE	35	93	49	41	93	60
USA	18	48	24	40	203	62
World (pooled)	46	80	60	41	109	60

Notes: This table shows the implied risk aversion parameter $\gamma = [\ln E(R) - \ln E(R_f)] / \text{Cov}(\ln R, \ln g)$ for housing, equities, and total wealth, for 16 countries and pooled World samples, for both the full period (1870–2015 including wars) and for the post-WW2 balanced sample (1963–2015). See text.