DISCUSSION PAPER SERIES

No. 10958

CLIMATE CHANGE AND LONG-RUN DISCOUNT RATES: EVIDENCE FROM REAL ESTATE

Stefano Giglio, Matteo Maggiori,
Johannes Ströbel and Andreas Weber

FINANCIAL ECONOMICS, MACROECONOMICS AND GROWTH, MONETARY ECONOMICS AND FLUCTUATIONS and PUBLIC ECONOMICS

ISSN 0265-8003

CLIMATE CHANGE AND LONG-RUN DISCOUNT RATES: EVIDENCE FROM REAL ESTATE

Stefano Giglio, Matteo Maggiori, Johannes Ströbel and Andreas Weber

Discussion Paper No. 10958
November 2015
Submitted 22 November 2015

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: (44 20) 7183 8801
www.cepr.org

CLIMATE CHANGE AND LONG-RUN DISCOUNT RATES: EVIDENCE FROM REAL ESTATE

Abstract

The optimal investment to mitigate climate change crucially depends on the discount rate used to evaluate the investment's uncertain future benefits. The appropriate discount rate is a function of the horizon over which these benefits accrue and the riskiness of the investment. In this paper, we estimate the term structure of discount rates for an important risky asset class, real estate, up to the very long horizons relevant for investments in climate change abatement. We show that this term structure is steeply downward-sloping, reaching 2.6% at horizons beyond 100 years. We explore the implications of these new data within both a general asset pricing framework that decomposes risks and returns by horizon and a structural model calibrated to match a variety of asset classes. Our analysis demonstrates that applying average rates of return that are observed for traded assets to investments in climate change abatement is misleading. We also show that the discount rates for investments in climate change abatement that reduce aggregate risk, as in disaster-risk models, are bounded above by our estimated term structure for risky housing, and should be below 2.6% for long-run benefits. This upper bound rules out many discount rates suggested in the literature and used by policymakers. Our framework also distinguishes between the various mechanisms the environmental literature has proposed for generating downward-sloping discount rates.

JEL Classification: G11, G12 and R30
Keywords: asset pricing, climate change, cost-benefit analysis, declining discount rates, environmental economics and real estate

Stefano Giglio stefano.giglio@chicagobooth.edu
Booth School of Business, University of Chicago, NBER and CEPR

Matteo Maggiori maggiori@fas.harvard.edu
Harvard University, NBER and CEPR

Johannes Ströbel jstroebe@stern.nyu.edu
Stern School of Business, New York University, NBER and CEPR

Andreas Weber aweber@stern.nyu.edu
Stern School of Business, New York University

Climate Change and Long-Run Discount Rates: Evidence from Real Estate

Stefano Giglio* Matteo Maggiori† Johannes Stroebel‡ Andreas Weber§

November 2015

Abstract

The optimal investment to mitigate climate change crucially depends on the discount rate used to evaluate the investment's uncertain future benefits. The appropriate discount rate is a function of the horizon over which these benefits accrue and the riskiness of the investment. In this paper, we estimate the term structure of discount rates for an important risky asset class, real estate, up to the very long horizons relevant for investments in climate change abatement. We show that this term structure is steeply downward-sloping, reaching 2.6% at horizons beyond 100 years. We explore the implications of these new data within both a general asset pricing framework that decomposes risks and returns by horizon and a structural model calibrated to match a variety of asset classes. Our analysis demonstrates that applying average rates of return that are observed for traded assets to investments in climate change abatement is misleading. We also show that the discount rates for investments in climate change abatement that reduce aggregate risk, as in disaster-risk models, are bounded above by our estimated term structure for risky housing, and should be below 2.6% for long-run benefits. This upper bound rules out many discount rates suggested in the literature and used by policymakers. Our framework also distinguishes between the various mechanisms the environmental literature has proposed for generating downward-sloping discount rates.

JEL Codes: G11, G12, R30.
Keywords: Environmental Economics, Declining Discount Rates, Climate Change, Real Estate, Cost-Benefit Analysis, Asset Pricing

---

*Booth School of Business, University of Chicago, NBER: stefano.giglio@chicagobooth.edu.
†Harvard University, NBER, CEPR: maggiori@fas.harvard.edu.
‡Stern School of Business, New York University, NBER, CEPR: johannes.stroebel@stern.nyu.edu.
§Stern School of Business, New York University: aweber@stern.nyu.edu.
We thank Arthur van Benthem, Robert Barro, John Campbell, Thom Covert, Robert Engle, Xavier Gabaix, Kenneth Gillingham, Christian Gollier, Bob Hall, Lars Hansen, Derek Lemoine, Martin Lettau, William Nordhaus, Monika Piazzesi, Nick Stern, Christian Traeger, Martin Weitzman as well as seminar participants at the Energy Policy Institute (EPIC) and the Becker Applied Economics Workshop at the University of Chicago, American Economic Association, World Bank, EUI, NYU, Harvard, and LMU Munich for helpful discussions and comments. We are particularly grateful to Michael Greenstone for his encouragement and guidance in writing this paper. We thankfully acknowledge the generous research support from the Harvard Weatherhead Center for International Affairs, the NYU Stern Center for the Global Economy and

Any consideration of the costs of meeting climate objectives requires confronting one of the thorniest issues in all climate-change economics: how should we compare present and future costs and benefits? [...] A full appreciation of the economics of climate change cannot proceed without dealing with discounting. **(Nordhaus, 2013)**

1 Introduction

Much of the economics literature on the optimal policy responses to climate change focuses on the trade-off between the immediate costs and the potentially uncertain long-run benefits of investments to reduce carbon emissions. Discount rates play a central role in this debate, since even small changes in discount rates can dramatically alter the present value of investments with very long horizons. As an example, assume that an investment to reduce carbon emissions costs $3 billion, and is expected to avoid environmental damages worth $100 billion in 100 years. At a discount rate of 3%, the present value of those damages is $5.2 billion, and the project should be implemented. At an only slightly higher discount rate, 5% for instance, the present value of the investment drops to $760 million, an order of magnitude smaller, and the investment no longer appears attractive.

Greenstone, Kopits and Wolverton (2013) remark that "the choice of a discount rate to be used over very long periods of time raises highly contested and exceedingly challenging scientific, economic, philosophical, and legal issues. As a result, there is no widespread agreement in the literature concerning the discount rates that should be used in an intergenerational context." In this paper, we make progress on this question along two dimensions. First, we provide new empirical evidence on the term structure of discount rates for an important asset class, real estate, up to the extremely long horizons that are relevant for analyzing climate change (hundreds of years). Second, we combine these new empirical facts with insights from asset pricing theory to discipline the debate on the appropriate choice of discount rates for investments in climate change abatement.

So far, the debate on appropriate discount rates has either relied on theoretical arguments, or has tried to infer discount rates from the realized returns of traded assets such as private capital, equity, bonds, and real estate.[1] For example, in the context of the canonical dynamic integrated climate-economy (DICE) model proposed by Nordhaus and Boyer (2000) and Nordhaus (2008), Nordhaus (2013) chooses a discount rate of 4% to reflect his

Business as well as from the Fama-Miller Center and the Initiative on Global Markets at the University of Chicago Booth School of Business. We thank iProperty and Rightmove for sharing part of their data, and Andreas Schaab for excellent research assistance.

[1]See also, for example, Kaplow, Moyer and Weisbach (2010), Schneider, Traeger and Winkler (2012), and Weisbach and Sunstein (2009) on a discussion of normative and descriptive approaches to discounting.

preferred estimate of the average rate of return to capital. We show that this common practice of valuing investments in climate change abatement by discounting cash flows using the average rate of return to some traded asset ignores important considerations regarding the *maturity* and *risk properties* of such investments.

As we review in Section 2, per-period discount rates could vary by horizon either because the risk profile of cash flows differs, or because households' sensitivity to risks that materialize over different horizons varies. In particular, asset pricing theory shows that the rate at which a particular expected cash flow should be discounted depends on the state of the world in which the cash flow is realized; cash flows that materialize in bad states are more desirable, and hence less risky for the investor. They should therefore be discounted at a lower rate. In addition, most traded assets are claims to cash flows at different horizons, each of which might require a different per-period discount rate. The average rate of return to an asset only captures the value-weighted average discount rate applied to all its cash flows. Therefore, this average rate is generally not informative for determining the appropriate discount rate for another asset such as an investment in climate change abatement, whose cash flows tend to occur at much longer horizons and which might have very different risk properties.

While the absence of assets with cash flows that exactly replicate those of investments in climate change abatement is an important limitation, it does not preclude the inference of essential information for applicable discount rates from the observed returns of traded assets. In this paper, we make progress along a number of the important empirical and theoretical dimensions in determining the appropriate discount rate for such investments.

Our first empirical contribution in Section 3 is to provide estimates of the term structure of discount rates for an important asset class, real estate, over a horizon of hundreds of years. This represents the first data-driven characterization of a term structure of discount rates for any asset over the horizons relevant for guiding the choice of discount rates for the distant benefits of climate change investments. Using a variety of empirical approaches, we estimate the average return to real estate to be above 6%. At the same time, recent estimates from Giglio, Maggiori and Stroebel (2015*b*) show that the discount rate for real estate cash flows 100 or more years in the future is about 2.6%. This combination of high average returns and low long-run discount rates implies a downward-sloping term structure of discount rates for real estate.

These findings reinforce how problematic it is to use the average rate of return to traded assets to discount investments in climate change abatement. Even if we assumed that climate change and real estate had similar risk properties at all horizons, using an average rate of return would suggest that climate change investments should be discounted at 6%. Instead, the appropriate discount rate for the very long-run benefits of these investments should be much lower, and their present value much higher.

After documenting a downward slope of discount rates for an important asset class, we next consider whether real estate might have different risk properties from climate change investments. We first document that real estate is indeed a risky asset: its returns are positively correlated with consumption growth, and it performs badly during consumption disasters, financial crises, and wars.[2]

To interpret the downward-sloping term structure of risky real estate, we build on insights from asset pricing theory and discuss the mechanisms that can produce such a slope. We also elaborate on the implications of these mechanisms for determining the appropriate discount rate for investments in climate change abatement.

Section 4 begins by providing a general decomposition of the term structure of risk and returns of the leading asset pricing models, based on Dew-Becker and Giglio (2013). This decomposition highlights that the shape of the term structure of discount rates for any asset can be attributed to the interaction of three main forces: 1) how much individuals care about long-term news relative to short-term news about consumption growth (i.e., the term structure of horizon-specific risk prices); 2) how much news about future consumption growth there is in the economy (e.g., whether the economy is subject to persistent or i.i.d. shocks); and 3) how exposed an asset is to news at different horizons. We operationalize this general decomposition by building on the canonical asset pricing model of Lettau and Wachter (2007). We parameterize this model to match an extensive set of observable discount rates for traded assets (equity and real bonds), as well as the term structure for real estate estimated in this paper. Since we find that the term structure of risk-free discount rates for the U.K. is approximately flat at 1%, the declining discount rates for risky assets, common to both equities and real estate, must be driven by a declining term structure of risk premia (see also van Binsbergen, Brandt and Koijen, 2012). The Lettau and Wachter (2007) model rationalizes this important feature of the data by inducing partial mean reversion in aggregate cash flows. In a similar spirit, Nakamura et al. (2013) motivate their downward-sloping term structure of discount rates for risky assets by documenting that consumption disasters tend to partially mean-revert in a large cross section of countries for the last century. Within the environmental literature, our results are therefore consistent with an interesting and empirically-motivated modification of models of climate change as a disaster risk (e.g., Weitzman, 2012; Barro, 2013), that would allow for a partial recovery of economic activity following a climate disaster.

[2]This is consistent with the average return to real estate of more than 6%, which is above the real risk-free rate of 1%, and thus includes a risk premium to compensate investors for bearing risk.

In the context of climate change as a possible source of disaster risk, we show that our results of a downward-sloping term structure and a very long-run discount rate of 2.6% for risky real estate provide an upper bound on discount rates for investments in climate change abatement. Since such investments reduce consumption risk, their discount rate has to be below that for real estate, which is a risky asset. We find that this upper bound is a simple yet powerful result that challenges a wide range of estimates previously used in the literature. For example, this bound is substantially below the 4% rate suggested by Nordhaus (2013) and the 4.6% suggested by Gollier (2013). Quantitatively, it is more in line with long-run discount rates that are close to the risk-free rate, as suggested by Weitzman (2012), or the 1.4% suggested by Stern (2006), or results by Barro (2013).[3] It is also close to the average recommended long-term social discount rate of 2.25% elicited by Drupp et al. (2015) in a survey of 197 experts, and falls within the range of 1% to 3% that more than 90% of these experts are comfortable with.

Our framework also helps to determine which models generate term structures, for risky or hedging assets, that are compatible with equilibrium outcomes and with the new empirical evidence presented in this paper. In particular, declining discount rate (DDR) functions have become increasingly common in environmental economics (Arrow et al., 2013*b*; Cropper et al., 2014; Farmer et al., 2015; Traeger, 2014*b*) and, in some cases, have been adopted by governmental institutions for the conduct of cost-benefit analyses.[4] While superficially our results are consistent with a DDR schedule, we find that the mechanisms that generate the downward slope differ dramatically across papers in the climate change literature, and are in many cases at odds with or orthogonal to our empirical findings. For example, the hyperbolic discounting of Laibson (1997) or the gamma discounting of Weitzman (2001) cannot rationalize our data. The hyperbolic subjective discount rate, while popular in the micro-behavior literature, has been shown by Barro (1999) and Luttmer and Mariotti (2003) to generate a flat term structure of risk-free discount rates in general equilibrium and to have no direct implications for risk premia. Gamma discounting is largely complementary to our analysis since it does not focus on risk premia, and our evidence on real estate does not rely on the polling of heterogeneous expert opinions that is at the core of Weitzman (2001)'s mechanism. Similarly, we find that introducing Epstein-Zin preferences makes it harder to rational-

[3]For models of climate change where costs materialize mostly in good states of the world, and in which climate change abatement investments behave as risky investments at all horizons, our findings suggest that climate change abatement investments should also be discounted at a rate that declines over maturity.

[4]France and the U.K. use discount rate schedules in which the discount rate applied today to benefits and costs occurring in the future declines over time (see HM Treasury, 2003; Lebègue, 2005). Relatedly, Metcalf and Stock (2015), Shapiro et al. (2010), and Stern (2014*a*,*b*) provide interesting discussions of the policy dimensions of economic research related to climate change.

ize the data because the pricing of long-run risks tends to induce an upward-sloping term structure of discount rates for risky assets and a downward-sloping one for risk-free assets. Along these lines, in an interesting contribution, Gollier (2013) notes that Epstein-Zin preferences and uncertainty about climate change, a form of long-run risk, counterfactually lead to an upward-sloping term structure of discount rates for risky assets, which also include climate change abatement investments in his model.

2 Discounting: The Role of Risk and Horizon

How should policymakers decide whether a particular investment in climate change abatement is worth pursuing? A common approach is to conduct a cost-benefit analysis to determine the societal net present value (NPV) of an investment project that is costly *today* and provides a stream of potentially uncertain *future* benefits (cash flows), with positive NPVs indicating socially beneficial projects. As highlighted in the introduction, discount rates play a central role in determining NPVs, since even small changes in discount rates can dramatically alter the NPV of investments with long horizons (see e.g., Arrow et al., 2013*a*; Dreze and Stern, 1987; Moyer et al., 2014; van Benthem, 2015; Weisbach, 2014).

In this section, we review the basic theoretical concepts and introduce the framework for our empirical and structural analysis in Sections 4 and 5. Section 2.1 describes how the appropriate rate for discounting a particular cash flow depends on both the riskiness and the maturity of that cash flow. Section 2.2 highlights what this implies for learning about the appropriate discount rates for climate change policies from observable assets that pay cash flows with different riskiness and maturity. The rest of the paper uses insights from the term structure of discount rates for one particular asset, real estate, to guide the choice of appropriate discount rates for investments in climate change abatement.

To introduce our basic notation, let us represent an investment at time t as a claim to a stream of future benefits (cash flows), D_{t+k}, $k = 1, 2..., n$, where n is the final maturity of the cash flows. For example, an investment to avoid one ton of CO_2 emissions today provides benefits in terms of mitigated climate change in each future period for hundreds of years. Each of these benefits, D_{t+k}, is stochastic and depends on the state of the world at time $t + k$. For example, the future benefits of reducing CO_2 emissions today could depend on how much the economy grows in the future. We denote the state of the world at time $t + k$ as $\omega_{t+k} \in \Omega_{t+k}$ and stress the dependence of benefits on its stochastic realization with the notation $D_{t+k}(\omega_{t+k})$. The set Ω_{t+k} includes all possible states of the world at time $t + k$, which can differ along many dimensions, including the health of the aggregate economy and the degree of environmental damage. In what follows, we will sometimes refer to general assets with maturity n that could pay cash flows such as dividends or rents at any point in time up to their maturity; these will simply be referred to with superscript n. A subset of these assets is the set of claims to a single cash flow at a specific point in time, maturity n; we will refer to these with superscript (n).

2.1 The Value of a Single-Cash-Flow Investment

We begin our analysis by studying the value of an investment that pays only one cash flow, at a specific point in time: $t+n$. This cash flow is not predetermined: it might be different in different states of the world, $\omega_{t+n} \in \Omega_{t+n}$. We denote the present value of the claim to this benefit as $P_t^{(n)}$. A classic tenet of asset pricing is that, under the relatively mild assumptions of no arbitrage and the law of one price, $P_t^{(n)}$ can be expressed as the weighted expected value of that cash flow across scenarios ω_{t+n}, where a benefit paid in each scenario is weighted by the importance investors assign to benefits in that state (see Hansen and Richard (1987), and Cochrane (2005) for a textbook treatment). Let $M_{t,t+n}(\omega_{t+n}) > 0$ denote the value that investors attach at time t to benefits in state ω_{t+n}. An asset is considered more *risky* if it pays off primarily in states of the world in which investors value that payoff less. If investors value benefits paid out earlier more than benefits paid out later, the weighting $M_{t,t+n}$ will also adjust for this time discounting. We can then write the value of an investment that yields D_{t+n} as:

$$P_t^{(n)} = \sum_{\omega_{t+n} \in \Omega_{t+n}} M_{t,t+n}(\omega_{t+n}) D_{t+n}(\omega_{t+n}) \pi_{t,t+n}(\omega_{t+n}) = E_t \left[M_{t,t+n} D_{t+n} \right], \qquad (1)$$

where $\pi_{t,t+n}(\omega_{t+n})$ is the conditional probability of state ω_{t+n}. The object $M_{t,t+n}$ is called the *stochastic discount factor* (SDF). In economic terms, the SDF reflects the marginal utility of a payoff in different states of the world. The value of the asset thus reflects both the physical properties of the asset (when and how much it pays in each state ω_{t+n}) and the preferences of investors (how much they value payoffs in each scenario ω_{t+n}).

An equivalent representation of $P_t^{(n)}$, which is more prevalent in policy analysis, is in terms of *discount rates*. The time and risk adjustments are then expressed using a per-period discount rate $\bar{r}_t^n$:

$$P_t^{(n)} = E_t \left[M_{t,t+n} D_{t+n} \right] = \frac{E_t \left[D_{t+n} \right]}{(1 + \bar{r}_t^n)^n}. \qquad (2)$$

Put differently, we can also think of prices as the expected value of the cash flow, not weighted by the SDF, $M_{t,t+n}$, but discounted at a per-period discount rate $\bar{r}_t^n$. The appropriate discount rate will differ across investments depending on which states of the world an investment pays benefits in, and the relative valuation of benefits across states of the world: more risky investments are valued less, and thus discounted at higher per-period discount rates.

2.2 The Importance of Horizon-Specific Risk Adjustments of Discount Rates

We now consider a multi-period-payoff investment project that pays stochastic benefits at different points in time up to maturity n. Any such asset can be thought of as the combination of many single cash-flow assets, each paying at specific points in time, $t + 1, t + 2, ..., t + n$. Therefore, the value of a multi-period-payoff investment project is the sum of the values of the individual single-period-payoff projects:

$$P_t^n = P_t^{(1)} + P_t^{(2)} + ... + P_t^{(n)}$$

Since the two representations discussed above for the one-period case also apply to the multi-period case, the value P_t^n can be written as:

$$\begin{aligned} P_t^n &= E_t\left[M_{t,t+1}D_{t+1} + M_{t,t+2}D_{t+2} + ... + M_{t,t+n}D_{t+n}\right] && (3) \\ &= \frac{E_t\left[D_{t+1}\right]}{1+\bar{r}_t^1} + \frac{E_t\left[D_{t+2}\right]}{(1+\bar{r}_t^2)^2} + ... + \frac{E_t\left[D_{t+n}\right]}{(1+\bar{r}_t^n)^n}. && (4) \end{aligned}$$

These two representations differ markedly from the valuation formula that is typically applied in cost-benefit analyses, which discounts each cash flow at the *same* per-period discount rate $\bar{r}_t$:

$$P_t^n = \frac{E_t\left[D_{t+1}\right]}{1+\bar{r}_t} + \frac{E_t\left[D_{t+2}\right]}{(1+\bar{r}_t)^2} + ... + \frac{E_t\left[D_{t+n}\right]}{(1+\bar{r}_t)^n}. \quad (5)$$

Representations 3 and 4 are always correct and equivalent; the last one is *only* correct if the discount rate $\bar{r}_t$ is chosen to match the risk and maturity of a particular asset. Therefore, $\bar{r}_t$ can only be applied to value the benefits of a project with exactly the same risk characteristics and exactly the same maturity as the asset from which $\bar{r}_t$ was derived in the first place. For the purpose of discounting the benefits of a project with different characteristics, the full term structure of discount rates $\bar{r}_t^1, \bar{r}_t^2, ..., \bar{r}_t^n$ needs to be known and appropriately adjusted for differences in risk characteristics. We highlight the importance of this by considering the valuation of three different investment projects below: A project with the same risk and payoff horizon as those of an observed traded asset (whose average per-period discount rate is $\bar{r}_t$); a project with the same risk properties but a different payoff horizon; and a project with different risk properties but the same payoff horizon.

Case 1: Same Risk, Same Horizon. Consider first an observable asset with maturity n and stochastic cash flows D_{t+1}, D_{t+2},..., D_{t+n} (if the asset has infinite maturity as in the case of the stock market, then $n = \infty$). Imagine we are able to observe the average discount rate of this asset, $\bar{r}_t$. Put differently, given an asset with maturity n and some risk profile, $\bar{r}_t$ is defined as the constant discount rate consistent with the asset's price. Now consider the case in which an investment in climate change abatement pays cash flows $\tilde{D}_{t+1}, ..., \tilde{D}_{t+n}$ that are different from the cash flows of the observed asset, but have the same risk characteristics (i.e., the same dependence on the state of the world ω_{t+n}). This is the only case in which cash flows from climate change abatement can be discounted at the *same* average rate as those from the observable asset. The value of the climate change investment, C_t^n, will be:

$$C_t^n = \frac{E_t\left[\tilde{D}_{t+1}\right]}{1+\bar{r}_t} + \frac{E_t\left[\tilde{D}_{t+2}\right]}{(1+\bar{r}_t)^2} + ... + \frac{E_t\left[\tilde{D}_{t+n}\right]}{(1+\bar{r}_t)^n}.$$

Case 2: Same Risk, Different Horizon. Since the risk preferences captured by $M_{t,t+k}$ potentially depend on the horizon, using average discount rates from one asset to discount cash flows from another investment is no longer valid if those cash flows materialize over different horizons. Take our example from above and assume that the asset's cash flows have the same riskiness as the cash flows from the investment in climate change abatement at each horizon. Assume further that the observable asset yields benefits in every period between time t and time $t+n$, while the investment in climate change abatement only yields benefits after maturity $\underline{n} > 1$. Since the riskiness of the cash flows of both investments is the same, one may be tempted to use the observed average discount rate $\bar{r}_t$ from the observable asset to discount climate change project cash flows. This turns out to be incorrect, however. The correct price is obtained as below:

$$C_t^{(\underline{n},n)} = \frac{E_t\left[\tilde{D}_{t+\underline{n}}\right]}{(1+\bar{r}_t^{\underline{n}})^{\underline{n}}} + \frac{E_t\left[\tilde{D}_{t+\underline{n}+1}\right]}{(1+\bar{r}_t^{\underline{n}+1})^{\underline{n}+1}} + ... + \frac{E_t\left[\tilde{D}_{t+n}\right]}{(1+\bar{r}_t^{n})^{n}},$$

where each dividend is discounted at the horizon-specific discount rate, $\bar{r}_t^{\underline{n}}, \bar{r}_t^{\underline{n}+1}, ..., \bar{r}_t^{n}$. Since $\bar{r}_t$ was obtained as the discount rate that applies to the observable asset, it reflects an average of *all* the horizon-specific discount rates $\bar{r}_t^1, \bar{r}_t^2, ..., \bar{r}_t^n$, including the ones for maturities up to $\underline{n}-1$. Since the climate change project does not accrue benefits at those horizons, its value should not depend on the discount rates between $t+1$ and $\underline{n}-1$.

To see this more clearly, suppose that investors are only worried about the states of the world in which the relatively near cash flows are being paid out (horizons 1 to $\underline{n}-1$), while they are not worried about risks for horizons higher than $\underline{n}$: for long maturities, investors only care about the expected payout from the asset, not the state of the world, in which it is paid out. They will discount the short-term cash flows at high rates, $\bar{r}_t^1, \bar{r}_t^2, ..., \bar{r}_t^{\underline{n}-1}$, but the longer-maturity cash flows at lower rates, $\bar{r}_t^{\underline{n}}, \bar{r}_t^{\underline{n}+1}, ..., \bar{r}_t^{n}$, reflecting their risk-neutrality at those horizons. The term structure of discount rates for this particular asset is thus downward-sloping. The claim to all cash flows may have a relatively high implied average discount rate, in particular if many of the cash flows accrue before $\underline{n}$. At the same time, if the benefits from a climate change investment had the same risk properties, but only accrued after $\underline{n}$, the correct present value for such an investment should *only* depend on the low discount rates $\bar{r}_t^{\underline{n}}, \bar{r}_t^{\underline{n}+1}, ..., \bar{r}_t^{n}$. It would thus be higher than under the relatively high *average* discount rate $\bar{r}_t$.

Case 3: Different Risk, Same Horizon. Beyond the timing of cash flows, a second potentially important difference between an observed asset's discount rates and those that apply to some investment project is the relative riskiness of the payoffs *at the same horizon*. As outlined before, riskiness here refers to whether an asset mostly pays in states of the world ω_{t+k} where payments are least valuable for the investor. Consider our example from above again. Assume that the asset as well as the climate change investment project only pay a single cash flow in period $t+n$. Further assume that the observed asset's cash flow is riskier than the investment's cash flow: for example, equities generally pay off in states of the world where the economy is doing well, while investments that mitigate the impact of climate disasters would pay off in states of the world where the economy is not doing well. The discount rate implied by the observable price of the asset will then be different from the appropriate discount rate for the investment project.

For concreteness, assume that there are only two equally likely states of the world – a good one (ω_{t+n}^G) and a bad one (ω_{t+n}^B). Assume that marginal utility in the good state of the world is lower than marginal utility in the bad state of the world, and assume that the observed asset pays out in the good state of the world only, while the investment project only pays out in the bad state of the world; both pay out the same amount if they pay out. This implies that $E_t\left[M_{t,t+n}D_{t+n}\right] < E_t\left[M_{t,t+n}\tilde{D}_{t+n}\right]$. It then follows from equation 2 that the investment project should be discounted at a lower rate than the asset.

3 The Term Structure of Real Estate Discount Rates

To illustrate the importance of allowing discount rates to vary with the maturity of the cash flow, we next analyze the term structure of discount rates in real estate markets. In recent work, Giglio, Maggiori and Stroebel (2015*b*) use unique data from the U.K. and Singapore to estimate how much value households attach to future real estate cash flows accruing over a horizon of hundreds of years. In these real estate markets, residential properties trade either as freeholds, which are permanent ownership contracts, or as leaseholds, which are pre-paid and tradable ownership contracts with finite maturity. The initial maturity of leasehold contracts generally varies between 99 years and 1,000 years. By comparing the relative prices of leasehold and freehold contracts for otherwise identical properties, the authors estimate the present value of owning a freehold after the expiration of the leasehold contract. They show how this present value is informative about the discount rate attached to real estate cash flows that occur in the very long run.

Figure I reports the estimates from Giglio, Maggiori and Stroebel (2015*b*). It shows the price discount of leaseholds with varying maturities compared to freeholds for otherwise identical properties. For the U.K. estimates, for example, the bucket with leaseholds of remaining maturity between 100 and 124 years shows that households are willing to pay 11% less for a leasehold with that maturity than for a freehold. Interpreted differently, at least 11% of the value of a freehold property is due to cash flows that accrue more than 100 years into the future. In general, leasehold discounts are strongly associated with maturity, with shorter leaseholds trading at bigger discounts: 17.6% for leaseholds with remaining maturity of 80-99 years, 11% for remaining maturities of 100-124 years, 8.9% for remaining maturities of 125-149 years, and 3.3% for remaining maturities of 150-300 years. Leaseholds with more than 700 years remaining maturity trade at the same price as freeholds. Pricing patterns are similar for properties in Singapore. The authors provide a detailed investigation of the institutional setup of leasehold and freehold contracts, and examine a number of possible explanations for the observed leasehold discounts. They conclude that leasehold price discounts are tightly connected to the contracts' maturity and that discount rates of around 2.6% for cash flows more than 100 years in the future are necessary to match their data.

In this section, we provide new evidence on the expected return and riskiness of real estate to study an important and previously unexplored dimension of real estate data: the term structure of real estate discount rates. Section 3.1 introduces our analysis of expected real estate returns, which we find to be relatively high, between 6.4% and 8.0%. Section 3.2 presents our analysis of the riskiness of real estate, which allows us to rationalize this relatively high average rate of return. This section shows that real estate returns are particularly low during consumption disasters and wars. In Section 3.3, we combine these new data with the estimates of Giglio, Maggiori and Stroebel (2015*b*) to provide evidence for the slope of the term structure of real estate discount rates. Our analysis suggests that this term structure is strongly downward-sloping, and thus cautions against the use of another asset's average rate of return to infer discount rates for very long-run benefits associated with investments in climate change abatement. In subsequent sections, we will use insights from asset pricing theory to inform what can be learned from the downward-sloping term structure of risky real estate cash flows about the optimal discount rate for investments in climate change abatement.

3.1 Average Rate of Return to Housing and Rental Growth Rate

We employ two complementary approaches to estimate the average return to real estate. The first approach, which we call the balance sheet approach, is based on the total value of the residential real estate and housing stock, and the total value of real estate and housing services consumed (the 'dividend' from the real estate and housing stock). We obtain this information from countries' national accounts.[5] We control for the growth of the real estate and housing stock over time to construct the return series for a representative property. The second approach, which we label the price-rent approach, starts from the price-rent ratio estimated in a baseline year and constructs a time series of returns by combining a house price index and a rental price index. This approach focuses on a representative portfolio of houses and, therefore, does not need to correct for changes in the real estate and housing stock. After adjusting for inflation, both methods provide estimates of the gross real returns to real estate, $E[R^G]$. To compute net returns, we subtract maintenance costs and depreciation (δ) and any tax-related decreases in returns (τ). We estimate net returns as $r = E[R] = E[R^G] - \delta - \tau$. See Appendix A.1 for details on these procedures and the underlying data sources used.

The top panel of Table I presents the estimated average real estate returns for the U.S., the U.K., and Singapore. Our estimates for real estate returns in the U.S. follow Favilukis, Ludvigson and van Nieuwerburgh (2010). While U.S. real estate returns are not the focus of this paper, since they cannot be compared with the long-run real estate discount rates in the U.K. and Singapore from Giglio, Maggiori and Stroebel (2015*b*), they provide a useful benchmark as they have been the subject of an extensive literature (e.g., Flavin

[5] To determine the total consumption of real estate services, these measures impute the value of the owner-occupied equivalent rents, the real estate services consumed by individuals from living in their own house. See Mayerhauser and Reinsdorf (2006) and McCarthy and Peach (2010) for details.

and Yamashita, 2002; Gyourko and Keim, 1992; Lustig and van Nieuwerburgh, 2005; Piazzesi, Schneider and Tuzel, 2007). The balance-sheet and the price-rent approaches provide similar estimates for the average annual real gross return ($E[R^G]$) in the U.S.: 10.3% and 9.6% respectively. We calibrate the impact of maintenance and depreciation at 2.5% and the property tax impact at 0.67%.[6] We conclude that average real net returns in the U.S. real estate market are between 6.4% and 7.1%. This is similar to the estimates in Flavin and Yamashita (2002), who find a real return to real estate of 6.6%, and Favilukis, Ludvigson and van Nieuwerburgh (2010), who find a real return of 9-10% before netting out depreciation and property taxes.

Columns 3 and 4 in Table I report our estimates for the Singapore real estate market. The balance-sheet and price-rent approaches provide similar estimates for the average annual real gross return ($E[R^G]$): 10.4% and 10.3%, respectively. We assume the cost of maintenance and depreciation to be 2.5%, in line with the estimates for the U.S., and the property tax impact to be 0.5%.[7] Our estimate of the real net returns in the Singapore real estate market is therefore between 7.3% and 7.4%.

Columns 5 and 6 of Table I report the estimates for the real estate market in the U.K. The balance-sheet and the price-rent approaches provide similar estimates for the average annual real gross return ($E[R^G]$) again: 10.5% and 9.8%, respectively. We maintain the calibration for the cost of maintenance and depreciation at 2.5%. There are no property taxes to be considered in the U.K. Average real net returns in the U.K. real estate market are therefore approximately between 7.3% and 8.0%.

Overall, these estimates show that real expected returns for real estate are at least 6.4% for the countries we consider. These estimates are in line with the existing literature, and robust to the different methodologies we use.[8] Our estimates are also consistent with the

[6]In a recent study, Harding, Rosenthal and Sirmans (2007) find that between 1983 and 2001, real estate (housing) depreciated at roughly 2.5% per year gross of maintenance. Malpezzi, Ozanne and Thibodeau (1987) provide an overview of the earlier literature on depreciation. For example, Leigh (1980) estimates the annual depreciation rate of housing units in the U.S. to be between 0.36% and 1.36%. Depreciation is also a key calibration parameter for much of the recent literature in macroeconomics that considers households' portfolio and consumption decisions with housing as an additional asset. Cocco (2005) chooses an annual depreciation rate equal to 1%; Díaz and Luengo-Prado (2008) use an annual depreciation rate of 1.5%. Property taxes in the U.S. are levied at the state level and, while there is variation across states, are generally around 1% of house prices. Property taxes, however, are deductible from federal income tax. We assume that the deductibility reflects a marginal U.S. federal income tax rate of 33%. The net impact is therefore $(1 - 0.33) * 0.01 = 0.67\%$.

[7]Singapore levies a 10% annual tax on the estimated rental income of the property. A lower tax rate applies to owner-occupied properties (6%), but we use the more conservative (higher) rate for rental properties. The tax impact on returns is the tax rate times the average rent-price ratio, estimated at 5%. Hence, $\tau = 0.1 * 0.05 = 0.5\%$.

[8]Also note that most movements in rent-price ratios are driven by movements in house prices and not by movements in rents (see Shiller, 2007); therefore, our estimates of returns are relatively unaffected by the

notion that average house price growth over extended periods of time is relatively low, as argued by Shiller (2006), with high rental yields being the key driver of real returns to real estate and housing. In particular, our estimated average capital gains are positive but relatively small for all three countries, despite focusing on samples and countries that are often regarded as having experienced major growth in house prices.

Our estimates of average returns to real estate imply a positive real estate risk premium. Intuitively, houses are risky because they have low payoffs during bad states of the world such as wars, financial crises, natural disasters, and epidemics. We formalize this intuition in the next section by analyzing how house prices react during such events, as well as by estimating their average correlation with consumption and personal disposable income.

Finally, we estimate the average real growth rate of rental income, which we denote by g. The estimated real growth rate of rents is low for the U.K., Singapore, and the U.S., which we add to compare our estimates with those in the existing literature again. We obtain an estimate of $g = 0.2\%$ for Singapore and a somewhat higher estimate of $g = 0.7\%$ for the U.K.[9] For the U.S., we estimate $g = 0.5\%$, an estimate in line with that of Campbell et al. (2009), who obtain a median growth rate of 0.4% per year.[10] This is consistent with Ambrose, Eichholtz and Lindenthal (2013), who find very low real rental growth in a long time series of rents for Amsterdam, and with Shiller (2006), who estimates long-run real house price growth rates to be very low, often below 1%.[11]

To recap, our estimates indicate an average real net rate of return to real estate of at least 6.4% and an average real growth rate of rents below 0.7% in both the U.K. and Singapore.

time periods chosen. Take the U.S as an example: by 2013 rent-price ratios in the U.S. had declined to their 2000 levels approximately. Ending our sample in 2005 would have produced a slightly lower average rent-price ratio. However, at the same time, this would exclude the house price crash after 2005, thus leading to higher estimated average capital gains. In the overall estimates of expected returns, the higher estimated capital gains would be offset by a lower estimated rent-price ratio.

[9]For the U.K., we use the CPI component "Actual rents for housing" (series D7CE) from the Office of National Statistics as a rental index, and the CPI series from the same source (series D7BT). Our sample covers the years 1996-2013. For Singapore, we obtain a time series of rental indices for the whole island from the Urban Redevelopment Authority (the official real estate arm of the government: ura.gov.sg), and use the CPI series from the National Statistical Office. Our sample covers the years 1990-2012.

[10]For the U.S., we follow Favilukis, Ludvigson and van Nieuwerburgh (2010) and compute rent growth using the BLS shelter index (the component of CPI related to shelter, item CU.S.R0000SAH1 from the Federal Reserve Bank of St. Louis). We obtain the CPI series from the same source, item CPIAUCSL. Our sample covers the years 1953-2013.

[11]The equivalence of these two long-run growth rates is necessary for rental yields to be stationary. Eichholtz (1997), Eitrheim and Erlandsen (2005) and Ambrose, Eichholtz and Lindenthal (2013) also confirm Shiller's observation of negligible long-run real house price growth in different countries using different data. Ambrose, Eichholtz and Lindenthal (2013) find evidence of cointegration between house prices and rents using very long-run real estate data for Amsterdam.

3.2 The Riskiness of Housing

In this section, we assess the riskiness of real estate assets. Given the relatively high rate of return on real estate documented in the previous section, one would expect real estate to perform particularly poorly in bad states of the world; finding evidence for the riskiness of real estate would therefore validate the relatively high estimated average rate of return. To see whether this is the case, we analyze the behavior of real house prices during financial crises and periods of rare disasters; we also estimate the average correlation between consumption as well as personal disposable income and house prices.

Panel A of Figure II shows the average reaction of real house prices to financial (banking) crises. The analysis is based on dates of financial crises in Schularick and Taylor (2012), Reinhart and Rogoff (2009) and Bordo et al. (2001) for 20 countries for the period 1870-2013, and on our own dataset of historical house price indices for these countries.[12] The beginning of a crisis is normalized to be time zero. The house price level is normalized to be one at the onset of the crisis. House prices rise on average in the three years prior to a crisis, achieve their highest level just before the crisis, and fall by as much as 7% in the three years following the onset of the crisis. This fall in house prices during crisis periods contributes to the riskiness of real estate as an asset. Panel (B) of Figure II shows the average behavior of house prices during the rare disasters based on Barro (2006) and Barro and Ursua (2008). The consumption disaster dates for the 20 countries included in our historical house price index dataset are those defined by Barro and Ursua (2008). The dotted line tracks the level of consumption: consumption falls for three years before reaching its trough (normalized to be time zero) and recovers in the subsequent three years. The solid line tracks the house price level: house prices fall together with consumption over the first three years of the disaster but fail to recover over the subsequent three years. The fall in house prices during these rare disasters also contributes to the riskiness of real estate as an asset.

Figure III shows the time series of house prices and marks crisis years for the U.K. and Singapore with shadowed bands.[13] The pattern of house price movements during crises in these two countries is similar to the average pattern described above. House prices peak and then fall during major crises in the sample: the 1974-76 and 1991 banking

[12] Appendix A.2.1 provides details of the crisis dates and the house price series. The raw data are available on the authors' websites. Reinhart and Rogoff (2008) also analyze real estate prices for 16 countries and 18 crises occurring in the period 1974-2008. We analyze real estate prices in 20 countries for 43 crises and 17 rare disasters occurring in the period 1870-2013.

[13] All crisis dates are from Reinhart and Rogoff (2009) except the periods 1997-98 and 2007-08 for Singapore. The latter dates have been added by the authors and correspond to the Asian financial crisis of 1997-98 and the global financial crisis of 2007-08.

crises in the U.K. as well as the 1982-83 banking crisis and the 1997 Asian financial crisis in Singapore.[14] Similarly, both countries experience a drop in house prices during the 2007-08 global financial crisis.

Figure IV shows the performance of house prices during major wars, namely World War I and II (WWI and WWII). In both cases, time zero is defined to be the start date of the war period, 1913 and 1939 for WWI and WWII, respectively. The dotted line tracks house prices of five countries for the duration of WWI (1913-1918).[15] House prices fell throughout the war with a total fall in real terms of around 30%. Similarly, the solid line tracks house prices of six countries for the duration of WWII (1939-1945).[16] House prices fell by 20% in real terms from 1939 to 1943 and then stabilized for the last two years of the war, 1944-45. Overall, we find wars to be periods of major declines in real house prices, which further contributes to the riskiness of real estate as an asset.

We also investigate the average correlation between consumption and house prices over the entire sample rather than just during crisis periods. Table II reports the correlation of house price changes with consumption changes over the entire sample and for each country. The correlation is positive for all 20 countries, except for France (-0.05), and often above 0.5. The estimated positive correlation between house prices and consumption reinforces the evidence that real estate is a risky asset: it has low payoffs in states of the world in which consumption falls and marginal utility is high. We also investigate the correlation between house price growth and alternative measures of economic activity by using data from Mack and Martínez-García (2011), and report the correlation between annual real house price growth and real personal disposable income growth in a panel of 20 developed and emerging countries (see Table III). The average correlation is 0.36, with a minimum of 0.05 for Luxembourg and a maximum of 0.63 for Spain. Overall, this evidence further corroborates the fact that real estate returns are risky.

Despite extensive efforts to collect an exhaustive database, our results are still limited by the relatively small number of crises for which house price data are available, and by the relatively low quality of house price time series before 1950. In addition, rental data is available only in extremely short time series. The lack of such rental data prevents us from performing a comprehensive study of the riskiness of the underlying cash flows of housing. For example, trying to detect the presence of long-run shocks for rent growth would even be challenging with time series that extend back for many more decades.

[14] The 1984 banking crisis in the U.K. is the sole exception with increasing house prices.

[15] Due to data availability for house price indices during this period, the countries included are Australia, France, Netherlands, Norway, and the United States.

[16] Due to data availability for house price indices during this period, the countries included are Australia, France, Netherlands, Norway, Switzerland, and the United States.

Nevertheless, our results suggest that real estate is an asset with risks broadly consistent with its estimated expected return.[17]

We summarize our results in the following stylized facts: (i) real estate is a risky asset that performs poorly in economic crises, (ii) correspondingly, it has an expected return of at least 6.4% per year; (iii) real rent growth rates are low, below 0.7% per year.

3.3 Matching Risk, Return and Leasehold Discounts in Reduced-Form

We started this section by presenting facts about the relative pricing of freeholds and leaseholds of different maturities. In general, leasehold discounts are strongly associated with maturity, with shorter leaseholds trading at bigger discounts: from 17.6% for leaseholds with remaining maturity of 80-99 years all the way down to 3.3% for remaining maturities of 150-300 years. The previous discussion also showed that, on average, real estate has a real expected rate of return of at least 6.4% per year. In this section, we focus on how economic theory might *jointly* match the significant leasehold discounts and the high average return to real estate. We highlight that only a model with declining discount rates for risky real estate cash flows is able to match the data, which also suggests a downward-sloping term structure of risk premia. In Section 4, we calibrate a leading reduced-form asset pricing model to rationalize these leasehold discounts and our estimates of the average rate of return to real estate, and propose one plausible explanation for the observed term structure of discount rates.

We start by considering the simple constant-discount-rate extension of the classic valuation model of Gordon (1982). We assume that rents (cash flows) arising in each future period are discounted at a constant rate r, so that the s-period discount function is e^{-rs}. We also assume that rents are expected to grow at a constant rate g, so that expected rents follow: $E_t[D_{t+s}] = D_t e^{gs}$.[18] In this model, a claim to the rents for n periods, the

[17] Note that our results are likely to underestimate the riskiness of real estate and housing due to three effects: index smoothing, declining rents during bad times, and the destruction of housing stock during wars and natural disasters. (1) House price indices are generally smoothed and therefore underestimate the true variation in house prices. (2) We only consider the behavior of house price changes (capital gains) and have not considered the behavior of rents (dividends). For the two countries for which long high-quality time series of rental indices are available (France for the period 1949-2010 and Australia for the period 1880-2013), we find rent growth to be positively correlated with consumption growth (0.36 and 0.15 respectively). (3) A substantial part of the housing stock tends to be destroyed during wars. Therefore, the return to a representative investment in real estate would be lower than the fall in index prices as it would incorporate the physical loss of part of the asset.

[18] Technically, g is the sum of the expected growth rate of rents and a Jensen inequality term. Given the low variance of rent growth and in the interest of intuitive results, we ignore the latter term and refer to g as the expected growth rate of rents.

n-maturity leasehold, is valued at:

$$P_t^n = \int_t^{t+n} e^{-r(s-t)} D_t e^{g(s-t)} ds = \frac{D_t}{r-g}(1 - e^{-(r-g)n}). \tag{6}$$

Correspondingly, the infinite maturity claim, the freehold, is valued at:

$$P_t = \lim_{n\to\infty} P_t^n = \frac{D_t}{r-g},$$

where we have imposed $r > g$, and the transversality condition $\lim_{n\to\infty} e^{-rn} P_{t+n} = 0$. In related work, Giglio, Maggiori and Stroebel (2015*a*) show that there is no evidence of a failure of this transversality condition in housing markets in the U.K. and Singapore, even during periods when a sizeable bubble was regularly thought to be present. The above valuation formula for infinite maturity claims is the classic formula by Gordon (1982). The price discount for an n-maturity leasehold with respect to the freehold is:

$$Disc_t^n \equiv \frac{P_t^n}{P_t} - 1 = -e^{-(r-g)n}. \tag{7}$$

For any given maturity, the price discount decreases (in absolute value) the higher the discount rate r and the lower the growth rate of rents g. The first effect occurs because a higher discount rate reduces the present value of rents accruing far in the future. The second effect occurs because a lower growth rate of rents lowers actual future rents.

At the estimated benchmark values of $r = 6.4\%$ and $g = 0.7\%$, the constant-discount-rate model implies a leasehold discount at 100 years of $Disc^{100} = -e^{-0.064*100} = -0.17\%$. In other words, the 100-year leasehold would be valued only 0.17% less than the freehold. The discount we find in the data is at least 11%, orders of magnitude higher. More generally, Figure V compares the logarithmic discounts obtained under our baseline calibration for different leasehold maturities (white bars) with those observed in the data for the U.K. and Singapore (black bars). The 700+ year leaseholds are valued essentially at a 0% discount to freeholds both in the data and in the model. However, the model cannot match the discounts observed for leaseholds with maturities of 300 years or less. For example, for U.K. leaseholds with 80-99 years remaining, we observe a log discount of 17.6% in the data. The log discount in the model for leaseholds with 80 to 99 years remaining is a mere 0.76%. Intuitively, a model of exponential discounting assigns essentially zero present value to cash flows occurring 100 or more years into the future when discounting at an effective rate $r - g$ of close to 6% or more.[19]

[19] This intuition is robust to even more conservative calibrations of r and g. We also evaluate a "high

While the long-run discounts could be matched by a calibration with a constant net discount rate of $r - g = 1.9\%$, such a calibration would not be consistent with new evidence of a high average return to real estate and a low growth rate of rents in this paper. This simple constant-discount-rate model highlights the challenge for economic theory posed by our results: to *jointly* rationalize both a high expected return to real estate as a risky asset and the low long-run discount rates necessary to match the observed discounts for leaseholds relative to freeholds.

In reduced form, we find our estimates to be consistent with a downward-sloping term structure of total discount rates for real estate cash flows. Discount rates have to be sufficiently high in the short to medium term to contribute to the high average expected returns on real estate estimated in Section 3.1, but also sufficiently low in the long term to match the observed discounts estimated for long-run cash flows in Giglio, Maggiori and Stroebel (2015*b*). Section 4 introduces a reduced-form asset pricing model that is calibrated to match key moments of multiple asset classes including real estate. Our analysis shows that downward-sloping term structures of discount rates are a common feature of risky assets, and highlights which mechanisms may help to explain those patterns. Section 5 will discuss the resulting implications for valuing investments in climate change abatement.

4 Price and Quantity of Risk Across the Term Structure

The previous section showed empirically that the term structure of discount rates for real estate – a risky asset – is steeply downward-sloping. In this section, we apply asset pricing theory to discuss a decomposition of this term structure into its building blocks: risk and return across the term structure. This decomposition will help us understand the forces that drive discount rates at different horizons, and provides a link between discount rates observed on tradable assets and investments in climate change abatement.

rent growth rate" scenario by setting $g = 2\%$, and a "low expected returns" scenario with $r = 5.4\%$ per year, significantly less than our lowest estimate. Figure V also shows the discounts obtained in the high-rent-growth and low-expected-return scenarios. Both robustness exercises increase the model's implied discounts only slightly. Even a calibration that allows for both low returns and high rent growth cannot match the data, especially at longer horizons.

4.1 Per-Period Discount Rates as Weighted Averages of Expected One-Period Returns

Most of the insights of asset pricing theory are clearest to interpret when thinking about one-period expected returns, independent of the maturity of the asset. Since any asset can be bought, held for one period, and then sold at the end of that period (before its maturity), looking at the one-period return is one way to reduce assets of different maturities to a common horizon. This allows us to compare their risk and return properties.

We start by introducing our main notation and by linking together the concepts of returns to maturity and one-period returns. In what follows, we will sometimes refer to general assets with maturity n that could pay cash flows such as dividends or rents at any point in time up to maturity; these will simply be denoted with superscript n. A subset of these assets is the set of claims to single cash flows at a specific point in time, maturity n; we will denote these with superscript (n).

Define the one-period (gross) return per dollar spent on any security with maturity n as the total amount obtained from buying the security and liquidating it after one period:

$$R^n_{t,t+1} \equiv \frac{P^{n-1}_{t+1} + D_{t+1}}{P^n_t},$$

Note that at the time the asset is sold, its maturity has shortened to $n-1$. The return to holding the security over multiple periods (and reinvesting all intermediate cash flows) can be found by compounding the one-period returns. For example, the return to maturity of any investment with maturity n is

$$R^n_{t,t+n} = \prod_{k=0}^{n-1} R^{n-k}_{t+k,t+1+k}.$$

Of particular interest to us are the one-period returns and discount rates for claims to a single cash flow D_{t+n}. In this case, the one-period returns in all but the last period are entirely driven by price movements (since D_{t+k} is zero for all k, except for $k=n$, the last cash flow at maturity). What makes the return to this security of particular interest is its intimate link to our per-period discount rates for horizon-specific cash flows, $\overline{r}^n_t$. To see this, we can rewrite the return to maturity of such a claim as:

$$R^{(n)}_{t,t+n} = \frac{D_{t+n}}{P^{(n)}_t};$$

we can do this since there are no dividends to be reinvested over the life of this security. Taking expectations on both sides, and then rearranging, we obtain:

$$P_t^{(n)} = \frac{E_t[D_{t+n}]}{E_t[R_{t,t+n}^{(n)}]}.$$

Comparing this equation with equation 2 in Section 2, we immediately see that the n-period expected return to maturity of a claim to a single dividend at $t+n$ is exactly the compounded discount rate to be applied to that security: $E_t[R_{t,t+n}^{(n)}] = (1+\bar{r}_t^n)^n$.

Next, we want to link these quantities to the one-period expected return, for which we are able to provide a very intuitive risk-return decomposition. The focus of this paper is on the *average* shape of the term structure of discount rates. Time-variation in discount rates, while important in the asset pricing literature, plays a second-order role in thinking about climate change investments. We therefore derive the link between one-period expected returns $E_t[R_{t,t+1}^{(n)}]$ and per-period discount rates $\bar{r}_t^n$ under the assumption that per-period discount rates for a cash flow with a particular maturity are constant over time; relaxing this assumption would complicate the intuition without adding any economically relevant elements to the analysis. If expected returns are constant over time (though they may be different across maturities, such that the term structure of discount rates is not necessarily flat at each point in time), we have

$$E_t[R_{t,t+n}^{(n)}] = \prod_{k=0}^{n-1} E_t[R_{t,t+1}^{(n-k)}],$$

where all the returns are for claims to single cash flows, D_{t+k}, at different horizons k. The formula shows that in this case, not only the realized but also the expected returns are linked through compounding. Since $E_t[R_{t,t+n}^{(n)}]$ is directly linked to $\bar{r}_t^n$ as shown above, we can then easily substitute and take logs (and recall that $\ln(1+x) \simeq x$), to obtain:

$$\bar{r}_t^n = \frac{1}{n} \sum_{k=1}^{n} \ln(E_t[R_{t,t+1}^{(k)}]). \tag{8}$$

Therefore, the discount rate for a particular horizon n is simply the average of the one-period expected returns for claims to cash flows at each horizon. A flat term structure of discount rates must then imply a flat term structure of expected one-period returns across maturities; in fact, expected one-period returns and discount rates across maturities would all be equal. In Section 4.2, we will build on this decomposition to further elaborate on the forces that shape the term structure of discount rates.

4.2 Decomposing the Term Structure of Expected One-Period Returns

Now that we have clarified the link between one-period returns and per-period discount rates, we focus on the one-period returns of securities with maturity n. We start by using the fundamental asset pricing equation introduced in Section 2.1 to decompose the expected one-period returns $R^{(n)}_{t,t+1}$ into a component that reflects time discounting and a component that reflects the riskiness of the underlying cash flow. It follows from:[20]

$$1 = E_t[M_{t,t+1}R^{(n)}_{t,t+1}]$$

that:

$$E_t[R^{(n)}_{t,t+1}] = R^f_{t,t+1} - Cov_t[R^{(n)}_{t,t+1}, M_{t,t+1}]R^f_{t,t+1},$$

where the first component $R^f_{t,t+1} = E_t[M_{t,t+1}]^{-1}$ is the one-period risk-free rate that reflects time discounting and the second component reflects an additional discount compensating the investor for bearing risk (the covariance with the SDF; it reflects whether this asset primarily pays off in good states of the world that have a low marginal utility of consumption). The risk premium has the opposite sign of the covariance between the stochastic discount factor (SDF) and the one-period return, $Cov_t[M_{t,t+1}, R^{(n)}_{t,t+1}]$. This reflects the fact that a claim with a higher return in states of the world in which extra resources are less valuable (i.e., when marginal utility $M_{t,t+1}$ is low), is less valuable to the investors, and thus has a positive risk premium. Finally, to highlight the fact that only *innovations* in the SDF matter for the purpose of understanding risk premia (rather than its mean, which instead pins down the risk-free rate), we can rewrite excess returns as:

$$E_t[R^{(n)}_{t,t+1}] - R^f_{t,t+1} = -Cov_t[R^{(n)}_{t,t+1}, M_{t,t+1} - E_t[M_{t,t+1}]]R^f_{t,t+1}.$$

As is common in asset pricing theory, we make the additional assumption that log returns $r^{(n)}_{t,t+1} \equiv \ln(R^{(n)}_{t,t+1})$ as well as the log stochastic discount factor $m_{t,t+1} \equiv \ln(M_{t,t+1})$ are jointly normally distributed, which allows us to simplify our expression for the risk premium to the covariance term alone (see, for example, Campbell and Vuolteenaho, 2004):

$$E_t[R^{(n)}_{t,t+1}] - R^f_{t,t+1} \simeq -Cov_t[r^{(n)}_{t,t+1}, m_{t,t+1} - E_t[m_{t,t+1}]], \quad (9)$$

[20] The fundamental asset pricing equation introduced in Section 2.1 (equation 1) can be restated as $P^{(n)}_t = E_t\left[M_{t,t+1}P^{(n-1)}_{t+1}\right]$, which implies $1 = E_t\left[M_{t,t+1}\frac{P^{(n-1)}_{t+1}}{P^{(n)}_t}\right] = E_t\left[M_{t,t+1}R^{(n)}_{t,t+1}\right]$.

The above notation highlights again that only *innovations* in the (log-)SDF, $m_{t,t+1} - E_t[m_{t,t+1}]$, matter for expected returns.

To highlight the forces that shape the term structure of expected one-period returns, we will focus on analyzing the set of linear and log-linear consumption-based asset pricing models, in which the stochastic discount factor is a function of consumption growth. This class of models encompasses the vast majority of modern asset pricing models, in particular those employed in climate change analysis, such as power utility models (Lucas, 1978), long-run risk models with Epstein–Zin preferences (Bansal and Yaron, 2004), and rare disaster models (Barro, 2006; Gabaix, 2012). As noted in Dew-Becker and Giglio (2013), these asset pricing models can be nested in the following general representation for the SDF innovations:[21]

$$m_{t,t+1} - E_t[m_{t,t+1}] = -\sum_{k=0}^{\infty} z_k \cdot (E_{t+1} - E_t)\Delta c_{t+1+k}, \tag{10}$$

where $\Delta c_{t+1} - E_t \Delta c_{t+1}$ (the first term of the sum, i.e. for $k = 0$) is the shock to current consumption growth, while $(E_{t+1} - E_t)\Delta c_{t+1+k}$ with $k > 0$ is *news* about future consumption growth at horizon k, received during the holding period (between t and $t + 1$).

The terms z_k depend only on the parameters of the utility function (not on the consumption growth process), and represent risk aversion regarding news about consumption growth at a particular horizon. They can be thought of as horizon-specific risk prices. Substituting equation 10 into equation 9, we can write the expected return of any asset by decomposing it across horizons:

$$\begin{aligned} E_t[R^{(n)}_{t,t+1}] - R^f_{t,t+1} &= z_0 Cov_t[R^{(n)}_{t,t+1}, \Delta c_{t+1}] + \\ &+ z_1 Cov_t[R^{(n)}_{t,t+1}, (E_{t+1} - E_t)\Delta c_{t+2}) + \\ &+ z_2 Cov_t[R^{(n)}_{t,t+1}, (E_{t+1} - E_t)\Delta c_{t+3}] + \\ &+ .... \end{aligned} \tag{11}$$

The above decomposition holds for *any* asset, and therefore holds for all claims to cash flows at one particular point in the future, D_{t+n}, which jointly characterize the term structure of discount rates as discussed above. In particular, it highlights that the shape of the term structure of expected one-period returns (and thus ultimately of horizon-specific discount rates) can be attributed to the interaction of three forces:

[21]More precisely, all of these models produce this representation of the SDF depending only on consumption growth news as long as the variance of consumption growth and higher moments are constant; the decomposition easily generalizes to cases with arbitrary affine processes.

1. The term structure of horizon-specific *risk prices* z_k, i.e., how much agents care about long-term news relative to short-term news. The higher z_k is for long maturities, the more worried agents are about the long-run risks in the economy. For example, Dew-Becker and Giglio (2013) show that in a power utility model, only z_0 is different from zero (and equal to risk aversion γ): the agent is *only* worried about the one-period innovation in consumption. On the other hand, an Epstein–Zin investor has $z_0 = \gamma$ like in the power utility model, but also $z_k = (\gamma - \frac{1}{\psi})\theta^k$ for $k > 0$, where ψ is the elasticity of intertemporal substitution and θ a parameter close to 1 related to the time discount factor. Epstein–Zin parameterizations with $\gamma > \frac{1}{\psi}$, as in standard calibrations of the long-run risk model, imply that agents are worried *both* about immediate consumption growth and news regarding future consumption.

And remembering that the *covariance* of two random variables can always be written as a product of the correlation and the standard deviations of those two random variables:

2. How much news about future consumption growth there is in the economy (which will affect the covariance terms for $k > 0$): if there is no news about future consumption growth, for example because it is unpredictable, all of the standard deviations that involve news about future consumption growth and hence all of the respective covariances will be equal to zero. This force reflects horizon-specific *risk quantities*.

3. How exposed each of the assets studied is to news at different maturities (which will be reflected in the correlations and thus the covariance terms). For example, the one-period return on a claim to very long-run consumption will be exposed to news about this very long-run consumption, whereas a claim to short-run consumption will not. We can think of this as horizon-specific *risk exposures*.

In principle, different combinations of these three forces can generate a downward-sloping term structure of expected returns and discount rates as observed in the data. However, they can also help us determine which combinations are *not* compatible with observed data. Before we discuss this point in more detail, we now turn to a specific model of the asset pricing literature, based on Lettau and Wachter (2007), that is able to match both the term structure of stocks and real bonds and, most importantly, the term structure of discount rates we estimate for real estate.

4.3 A Reduced-Form Asset Pricing Model for Stocks and Real Estate

4.3.1 Setup and Key Results

Uncertainty in this model is characterized by two shocks: a shock to *current* cash flow growth and a shock to *expected* future cash flow growth. Let ϵ_{t+1} be a 2×1-vector of independent standard normal shocks. Denote by D_t^i with $i \in \{s, h\}$ aggregate stock market dividends and aggregate real estate and housing rents in the economy, respectively, and let $d_t^i \equiv \ln(D_t^i)$. Aggregate cash flows evolve according to

$$\Delta d_{t+1}^i \equiv d_{t+1}^i - d_t^i = g_i + \kappa_i z_t + \sigma_i \epsilon_{t+1},$$

where z_t is a scalar that follows the $AR(1)$ process

$$z_{t+1} = \phi_z z_t + \sigma_z \epsilon_{t+1},$$

with persistence $0 \leq \phi_z < 1$, and where σ_i and σ_z are 1×2 row vectors. σ_i determines which combination of the two shocks in ϵ_{t+1} drives the transitory shock to dividend or rent growth. z_t is a persistent long-run shock to the growth rate of cash flows that affects both dividends and rents. To keep the model parsimonious, we assume that only *one* shock drives the long-term fluctuations of dividend and rent growth, although the two can load on it differently (through different κ_i). The one-period risk-free rate, denoted by $r^f \equiv \ln(R^f)$, is assumed to be constant, as in Lettau and Wachter (2007). In the spirit of an endowment economy in which endowments take the form of dividends, such that consumption equals dividends as markets clear, the stochastic discount factor is a function of dividend growth instead of consumption growth.[22] Moreover, following Lettau and Wachter (2007), only the current shock to dividend growth is priced, so that the stochastic discount factor in this economy can be directly specified as:

$$M_{t,t+1} = \exp\{m_{t,t+1}\} = \exp\{-r^f - \frac{1}{2}\overline{x}^2 - \overline{x}\epsilon_{d,t+1}\},$$

with

$$\epsilon_{d,t+1} = \frac{\sigma_d}{||\sigma_d||}\epsilon_{t+1},$$

where $\overline{x}$ is the price of current cash flow risk in the model. The risk premium of expected one-period returns on claims to cash flows at individual maturities can then be shown to

[22] While this keeps the model parsimonious, we show in Appendix A.3 how to re-cast the model with a consumption-based, rather than a dividend-based SDF.

follow[23]

$$E_t\left[R^{i,(n)}_{t,t+1} - R^f\right] \simeq (\sigma_i + B^i_z(n-1)\sigma_z)\frac{\sigma_d^{'}}{||\sigma_d||}\overline{x}, \tag{12}$$

with

$$B^i_z(n) = \kappa_i \frac{1-\phi_z^n}{1-\phi_z}. \tag{13}$$

Risk premia depend on the loadings on the different risk sources as well as the price of each source of risk. All claims load equally, irrespective of their maturity, on shocks to current cash flow growth Δd^i_{t+1} and $||\sigma_d||^{-1}\sigma_i\sigma_d^{'}\overline{x}$ is the associated price of risk. $B^i_z(n-1)$ is the loading on shocks to the growth rate of cash-flows z_t and $||\sigma_d||^{-1}\sigma_z\sigma_d^{'}\overline{x}$ is the associated price of risk. Notice that the term-structure variation is generated by the loading $B^i_z(n-1)$, which is a strictly increasing function of horizon n, starting at 0 and converging to $\frac{\kappa_i}{1-\phi_z}$ as $n \to \infty$.

Note that in this setup with a constant price of risk $\overline{x}$, the simple relationship between the per-period discount rate for maturity n and the expected one-period holding return of single-claim assets with maturities up to n applies: $\overline{r}^{i,n}_t = \frac{1}{n}\sum_{k=1}^{n}\ln(E_t[R^{i,(k)}_{t,t+1}])$. Therefore, the model will imply a downward slope in the term structure of per-period discount rates if it generates a downward slope in one-period expected returns over maturities.

4.3.2 Calibration and Asset Pricing Moments

Lettau and Wachter (2007) and Lettau and Wachter (2011) show that this type of model can successfully match a large number of asset pricing moments, including the term structure of bonds and equity.[24] To discipline our analysis, we maintained the calibration of Lettau and Wachter (2007) whenever possible.[25] We made one simplification by keeping the price of risk x constant, since for the purpose of the analysis in this paper time variation in risk prices is not of first-order importance (in Appendix A.3, we calibrate the full model with time-varying risk prices, and verify that our conclusions are robust). Since real estate was not present in Lettau and Wachter (2007), we add it to the model. We calibrate it by setting $g_h = 0.7\%$ to reflect the average rent growth estimated in Section 3.1, and choose volatility $\|\sigma_h\|$ as well as κ_h to target the expected return on real estate in the data (6.4%)

[23]See Appendix A.3 for the derivations of all results stated here.

[24]See also the exciting work of van Binsbergen, Brandt and Koijen (2012) and van Binsbergen et al. (2013) for the model implications for claims to individual stock market dividends at horizons of 1 to 10 years.

[25]In the interest of space, the calibration is discussed in Appendix A.3.

and the empirical term structure of real estate discount rates, in particular the estimate of the very long-run discount rate of 2.6%. We set the risk-free rate at 1%, based on the U.K. gilts real yield curve between 1998 and 2013 reported in Figure VI. These data show that the U.K. real yield curve is approximately flat on average, with a real yield of 1.4% for maturities between 1 and 25 years, and that there is some evidence for a mild downward slope at longer maturities.[26] Our calibration is summarized in Table IV.

Table V shows that our benchmark calibration, just as in Lettau and Wachter (2007), allows us to match the key moments for equity. The average annual return on the stock market is 9.1%. As a consequence of assuming away time-variation in the price of risk, the volatility of the stock market is significantly lower than in the data, but as we show in Appendix A.3, this is easily fixed once we allow for time-variation in x. Expected dividend growth is above 3% with a standard deviation of 14.5%. Importantly, the term structure of discount rates for equity is downward-sloping as shown in Panel A of Figure VII. It falls from 18.2% for claims to one year ahead dividends, to 6% for claims to one hundred year ahead dividends, all the way to 4.6% for one thousand year ahead dividends. This pattern of falling discount rates is consistent with empirical findings of van Binsbergen, Brandt and Koijen (2012), who decompose the S&P 500 index into a portfolio of claims to short-term dividends and a portfolio of claims to long-term dividends and show that expected returns on the portfolio of short-term claims are higher than the returns on long-term claims.

The expected return on real estate is 5.9%, just below our empirical net return estimates between 6.4% and 8.0%. Expected rent growth is 1.6%. Our calibration implies that returns as well as rent growth are about as volatile as equity, with 8.2% and 14.2%, respectively. In line with our empirical findings, rental yields contribute more than two thirds to total returns. The term structure of discount rates for real estate is downward-sloping as shown in Panel B of Figure VII. It falls from 17.8% for claims to one year ahead rents, to 3.9% for claims to rents one hundred years ahead, all the way to 2.3% for one thousand year ahead rents.

These results show that the extended Lettau and Wachter (2007) model can broadly

[26]The real yield curve is computed by the Bank of England and is available at `http://www.bankofengland.co.uk/statistics/Pages/yieldcurve/archive.aspx`. Figure VI plots the average shape of the real U.K. gilts curve for the period 1998-2013, as well as for two sub-periods: 1998-2007 and 2007-2013. The level of the yield curve shifted down during this latter period and the yield curve became hump-shaped. The U.K. government debt also includes some perpetual bonds: the War Loan and the Annuities. These bonds comprise a negligible part of the outstanding U.K. government debt (£2.6bn out of £1.5trn of debt outstanding), and are classified as small and illiquid issuances by the U.K. Debt and Management Office. They are excluded from our analysis, not only because they are nominal and we only use data on U.K. real gilts, but also because their negligible size, scarce liquidity, and callability make it hard to interpret their prices in terms of discount rates.

match the estimated term structure of discount rates for real estate, while at the same time maintaining the original calibration of the model that matched stock and bond discount rates. In the next subsection, we discuss in more detail the forces that allow the model to match the downward-sloping term structures of risky assets like real estate and equity, and highlight which modifications of the model would make the data harder or easier to match.

4.3.3 The Term Structure of Risk Prices, Quantities and Exposures

The assumption that the SDF depends exclusively on the dividend shock ϵ_d is crucial to understand the term structure implications of the model. It implies that agents, similar to the consumption-based power utility model of Lucas (1978), care about next period's dividend growth, but do not directly care about *news* concerning future cash flows ($z_0 = \overline{x} > 0$, but $z_k = 0$ for $k > 0$). Therefore, a claim to long-term cash flows commands a positive risk premium only if news about long-term cash flows are correlated with the current dividend shock. Furthermore, this risk premium decreases with the maturity of the claim if the cash flow process is mean-reverting, as is the case in the above calibration of the model. Intuitively, after a negative shock ϵ_d hits the economy, the economy is expected to recover and grow faster. This makes claims to short-run cash flows that do not benefit from this mean reversion riskier than long-run cash flows that do.

To formalize these two intuitions, we can look at the model through the lens of the price-quantity-exposure decomposition in equation 10. The innovations in the log stochastic discount factor $m_{t,t+1}$ depend *only* on the immediate cash flow shock Δd_{t+1}:

$$m_{t,t+1} - E_t[m_{t,t+1}] = -\sum_{k=0}^{\infty} z_k \cdot (E_{t+1} - E_t)\Delta d_{t+1+k} = -\overline{x} \cdot (E_{t+1} - E_t)\Delta d_{t+1},$$

where the second equality follows from $z_0 = \overline{x} > 0$ and $z_k = 0$ for all $k > 0$. We can now see immediately that agents care about next period's dividend growth, but do not directly care about *news* concerning future cash flows, and that the term structure of *risk prices* is equivalent to that of a power utility model, where $\overline{x}$ determines the relative risk aversion of the agent.[27] Similarly, from equation 11, we obtain:

$$E_t[R^{(n)}_{t,t+1}] - R^f = \overline{x} \cdot Cov_t[R^{(n)}_{t,t+1}, \Delta d_{t+1}] = (\sigma_i + B^i_z(n-1)\sigma_z)\frac{\sigma_d'}{||\sigma_d||}\overline{x},$$

where the last equality follows from equation 12. Since cash flows are mean-reverting,

[27] In the notation of Lucas (1978), $\overline{x} = \gamma$, and γ determines the agent's relative risk aversion.

the term structure of expected one-period holding returns is downward-sloping ($\sigma_z \sigma_d'$ is negative and $B_z^i(n)$ is monotonically increasing in maturity n). Put differently, it is the *quantity of risk* and the *exposure* to it across horizons that creates the downward slope in the term structure of risky assets.

In the following, we will go back to our general price-quantity-exposure decomposition in equation 11 to discuss which alternative asset pricing mechanisms proposed by the literature and previously applied to discounting climate change projects (see Section 5) might work towards or against matching our data on the term structure.

Long-Run Risks and Epstein–Zin Preferences. As discussed above, log-linear models with Epstein–Zin preferences feature $z_k > 0$ for $k > 0$, so that pure news about expected growth rates of cash flows directly affect the SDF. Since claims to long-run cash flows D_{t+n} are naturally exposed to *all* dividend growth shocks from t to $t+n$ ($D_{t+n} = D_t \exp[\Delta d_{t+1} + ... + \Delta d_{t+n}]$), claims become more exposed to long-run shocks as their maturity increases. Accordingly, their risk premium grows with maturity as more and more of the covariance terms in equation 11 are added up with positive weights. Therefore, introducing Epstein–Zin preferences would make it more difficult to match the data on downward-sloping term structures of discount rates for risky assets.

Rare disasters. In a model with permanent disasters such as Barro (2006) and Gabaix (2012), the covariance between returns and current dividend growth (the first term in equation 11) increases.[28] This would help to match expected returns of equities and real estate, even with low prices of risk (z_0) and a relatively low volatility of cash flow and consumption growth. Section 3.2 provided evidence that real estate is indeed exposed to disaster risk. A model of permanent disasters, however, generates a flat term structure of discount rates since all assets load equally on the disaster risk, irrespective of maturity.[29] Nakamura et al. (2013) show that this shortcoming is remedied if output growth in the economy is modelled to partially recover following a disaster. Importantly, the authors document that such an output process does a better job matching the data on consumption and output growth during disaster periods across multiple countries over the past century. In other words, the most comprehensive study on consumption behavior in disaster periods shows precisely the kind of mean reversion that is captured by our

[28] Note that while the Lettau and Wachter (2007) model solved in the previous subsection features log-normal shocks, the term-structure decomposition in equation 11 does not depend on assumptions about the distribution of the shocks and is compatible with rare disasters.

[29] In other words, the covariance term $Cov_t[R_{t,t+1}^{(n)}, \Delta d_{t+1}]$ would be the same for all n.

simple extension of the Lettau and Wachter (2007) model, and it is key in both setups to generate a downward-sloping term structure of discount rates.

To sum up, Section 3 has demonstrated that the term structure of discount rates of risky assets is downward-sloping for important and large asset classes. Our discussion of alternative asset-pricing mechanisms shows that no generally accepted modifications of simple power-utility preferences generate such a downward slope (Epstein–Zin preferences even counteract it!). Mean reversion in cash flows of risky assets emanates as a key mechanism to generate a downward-sloping term structure of discount rates. The next section will discuss what these findings imply for valuing climate change investments.

5 Implications for Environmental Policy

Our new empirical evidence on the term structure of real estate discount rates, combined with our structural analysis from asset pricing theory, provides a number of insights for evaluating climate change policy. In this section, we highlight the implications of our findings on the level and slope of the term structure of real estate discount rates for valuing investments in climate change abatement. We also discuss which of the existing frameworks for thinking about discounting the future benefits from climate change abatement are consistent with the observed downward-sloping term structure of discount rates (DDR) for risky assets.

I. Applying observed *average* returns of traded assets to discount cash flows from investments in climate change abatement can be misleading

Our evidence cautions against using the observed average rate of return of traded assets to discount investments in climate change abatement, even if the traded asset has the *same riskiness* as the investment in climate change abatement at all horizons. Many benefits of climate change abatement will materialize over several centuries, while average returns to most traded assets are largely driven by benefits over much shorter horizons. This creates a *maturity mismatch* between the cash flows of observed assets and those of investments in climate change abatement.

Consider for example the discount rate choices of Nordhaus and Boyer (2000) and Nordhaus (2008), who develop a classic workhorse model of climate change, the dynamic integrated climate-economy (DICE) model. Nordhaus (2013) suggests a constant rate of discounting for investments in climate change abatement of 4% based on an "opportunity cost approach." In particular, he considers the returns to alternative investments that the government could undertake instead of investing in climate change abatement; a discount rate of 4% reflects the author's preferred estimate of the average return to capital (see also Piketty, 2014). But the capital stock of the economy is an asset that pays cash flows at all horizons, including the short term. In thinking about an investment with primarily long-term payoffs, such as climate change abatement, the appropriate discount rate should be the one that applies to the long-term cash flows only.

To emphasize the quantitative importance of this distinction, remember that our estimates of average real estate returns are between 6.4% and 8.0% (an estimate that is within the range used by Nordhaus (2013) for real estate returns), while the returns to a claim to rents 100 or more years into the future are estimated to be below 2.6%. The benefits of climate change abatement that have a similar riskiness as real estate should be discounted at rates much closer to the (low) long-term discount rate rather than at the average rate of return, which reflects some value-weighted average of (high) short-term and (low) long-term discount rates. A choice of discount rates based on the average return to real estate as opposed to the appropriate long-run discount rate would select a rate that is more than two times too high, with dramatic effects on the present value of future benefits.

II. The appropriate discount rate for long-run climate-hedging investments is below 2.6%

Our empirical evidence in Section 3 shows that real estate is a risky asset with returns that covary positively with consumption growth and which performs poorly during disasters. Since its cash flows are risky, its discount rates are above the risk-free discount rate at all horizons.[30] This observation implies that in any model in which long-run investments in climate change abatement pay off in bad states of the world (and therefore act as a hedge), the appropriate discount rate needs to be below the long-run discount rate for real estate that we estimated at 2.6%. To formally derive this bound, we can express any discount rate as a combination of risk-free discounting and a risk premium correction:

$$2.6\% = \bar{r}^{h,n} = \bar{r}^{f,n} + RP^{h,n}.$$

Since real estate is risky, it has a positive risk-premium: $RP^{h,n} > 0$. Conversely, a climate change abatement investment, denoted cl, that reduces (hedges) risk at the same horizon

[30] To be precise, we provide evidence that claims to residential real estate cash flows at all horizons (i.e., a house including its grounds) are risky, which is supportive for the riskiness of its long-run cash flows.

will have a negative risk premium: $RP^{cl,n} < 0$. Therefore:

$$2.6\% = \bar{r}^{h,n} = \bar{r}^{f,n} + RP^{h,n} > \bar{r}^{f,n} + RP^{cl,n} = \bar{r}^{cl,n}$$

This is a simple yet powerful result. First, climate change abatement investments are indeed often considered to be hedges in the literature (Barro, 2013; Lemoine, 2015; Wagner and Weitzman, 2015; Weitzman, 2012). Second, this upper bound is lower than numerous estimates used in the existing literature and by policymakers for discounting investments in climate change abatement. For example, it is substantially below the 4% suggested by Nordhaus (2013) and the 4.6% suggested by Gollier (2013). Quantitatively, it is more in line with long-run discount rates that are close to the risk-free rate, as suggested by Weitzman (2012), or the 1.4% suggested by Stern (2006)[31], or results by Barro (2013). It is also close to the average recommended long-term social discount rate of 2.25% elicited by Drupp et al. (2015) in a survey of 197 experts, and falls within the range of 1% to 3% that more than 90% of these experts are comfortable with. At the same time, while some authors have proposed low long-run discount rates for climate change abatement, the theoretical arguments underlying these proposals vary widely, and some are at odds with our empirical evidence (see point III below).

In light of the general disagreement in the literature regarding the appropriate discount rate, the interagency group tasked by the U.S. government to value reductions in CO_2 chose three certainty-equivalent constant discount rates: 2.5%, 3%, and 5% per year. Our estimates provide a tight bound that is only consistent with the lowest rate of 2.5% for investments providing a long-run hedge against climate disasters. Greenstone, Kopits and Wolverton (2013) report the cost of 1 metric-ton of CO_2 to be \$57 when using our suggested 2.5% discount rate, but only \$11 when using a 5% discount rate, illustrating the impact of this bound on climate-change-related welfare calculations (see also Pizer et al., 2014).

III. Which downward-sloping discount rates (DDR)?

The existing economics literature on valuing investments in climate change abatement has developed a variety of models with vastly different theoretical implications for discounting cash flows from these investments. Some of this literature has argued for a downward-sloping term structure of discount rates (DDR) for investments in climate change abatement, where benefits that occur in later periods are discounted at a lower

[31] See also Stern (2008) for comments on the discussion of the Stern Review's approach to discounting.

per-period discount rate (Arrow et al., 2013*b*; Cropper et al., 2014; Farmer et al., 2015; Traeger, 2014*b*).[32] As discussed in the introduction, these arguments have motivated policy changes in France and the U.K., which have adopted a downward-sloping term structure of discount rates for evaluating long-run investments.

At a superficial level, our empirical finding of a downward-sloping discount rate for real estate cash flows might appear consistent with much of this literature proposing downward-sloping discount rates to evaluate investments in climate change abatement. On a deeper level however, our data and structural analysis help to restrict the space of valid mechanisms that may generate a downward-sloping term structure of discount rates for such investments in climate change abatement. In particular, since these investments are likely to pay off conditional on the state of the economy, it is important to understand the forces that drive the term structure of risk premia.

IIIa. Forces that struggle to generate DDRs for risky assets

We start by reviewing a number of economic mechanisms that have been proposed to justify downward-sloping term structures for climate change policies, but which at the same time imply a flat or even increasing term structure of discount rates for *risky* assets in equilibrium. These models therefore cannot rationalize our empirical finding of a downward-sloping term structure of discount rates for risky assets such as real estate.

Hyperbolic Discounting. A first argument for justifying downward-sloping discount rates adopts the hyperbolic discounting framework of Laibson (1997), where investors with time-inconsistent preferences apply higher rate of time preferences between the present and the immediate future and lower rates when comparing investments across periods further in the future. While this might be a good description of the preferences of individual agents, and one might be tempted to conjecture that it leads to a downward-sloping term structure of discount rates, Barro (1999) and Luttmer and Mariotti (2003) have pointed out that a hyperbolic subjective rate of time preferences actually generates a flat term structure of risk-free discount rates in general equilibrium models and has no direct implications for the term structure of risk premia.[33] It is then evident that these

[32] See, for example, Traeger (2013) for a discussion of the Weitzman-Gollier puzzle in this context.

[33] More precisely, in the presence of commitment there would be a one-off downward-sloping risk-free yield curve on the convergence path to the steady state. However, this downward slope would not affect risk premia and would disappear in the steady state. In the absence of commitment, the term structure is always flat. Cropper and Laibson (1998) point out that in this class of models there might be under-saving and a too-high risk-free rate, and it is therefore optimal for the government to increase aggregate investment. They correctly point out that this might lead to more climate change abatement investments, but that the argument equally holds for other investments.

kind of preferences by themselves are insufficient to rationalize our data.

Ramsey Rule. A second approach to motivating downward-sloping discount rates has extended the Ramsey Rule of deterministic models with a horizon-dependent precautionary savings term to a stochastic environment (see, for example, Arrow et al., 2013*b*; Dasgupta, 2008; Gollier, 2008, 2011, 2012; Groom et al., 2005; Karp and Traeger, 2013; Kopp et al., 2012). In particular, when preferences are time-separable with a flow utility of consumption u, the representative agent's first-order condition with respect to a real risk-free bond with maturity n is:

$$u'(C_t) = \delta^n R^f_{t,t+n} E_t[u'(C_{t+n})]$$

with δ the per-period time discount factor. Assuming standard power utility, $u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}$, with γ the coefficient of relative risk aversion, log-normal consumption growth as in Arrow et al. (2013*b*), and re-writing the discount factor in terms of the rate of time preference ρ: $\delta = \frac{1}{1+\rho}$, we can derive the following simplified expression for the risk-free discount rate at horizon n:[34]

$$\bar{r}^{f,n}_t = \ln \rho + \gamma \frac{1}{n} E_t[\Delta c_{t,t+n}] - \frac{1}{n}\frac{\gamma^2}{2} Var_t[\Delta c_{t,t+n}]$$

where $\Delta c_{t,t+n}$ is the logarithmic consumption growth between t and $t+n$, and $\bar{r}^{f,n}_t \equiv \frac{1}{n}\ln(R^f_{t,t+n})$ is the per-period risk-free discount rate up to horizon n.

There are only two ways to generate a downward-sloping term structure for per-period discount rates $\bar{r}^{f,n}_t$ within this framework. The first is to assume that consumption growth will slow down at long horizons, such that $\frac{1}{n}E_t[\Delta c_{t,t+n}]$ is declining with maturity. This mechanism is unlikely to generate a quantitatively large downward slope, nor can it generate a downward slope in steady state.

The second is to assume that uncertainty about consumption growth is strongly increasing with the horizon, so that the term $\frac{1}{n}\frac{\gamma^2}{2} Var_t[\Delta c_{t,t+n}]$ increases with maturity n. This would push down the long-run risk-free rate. Intuitively, the presence of substantial uncertainty about long-horizon growth in consumption gives rise to a precautionary motive that induces agents to save more for these horizons, ultimately pushing down equilibrium discount rates at long horizons (see Arrow et al., 2013*b*). Despite the intuitive

[34]The result is derived as follows: Parameterize the utility function and rearrange to obtain $1 = R^f_{t,t+n} E_t[\exp\{-n \cdot \ln \rho - \gamma \Delta c_{t,t+n}\}]$. Using the log-normal result that $E[\exp\{X\}] = \exp\{E[X] + \frac{1}{2}Var[X]\}$, we can rewrite: $1 = R^f_{t,t+n} \exp\{-n \cdot \ln \rho - \gamma E_t[\Delta c_{t,t+n}] + \frac{1}{2}\gamma^2 Var_t[\Delta c_{t,t+n}]\}]$. Taking logs, we obtain the result.

appeal of this mechanism, it has several theoretical and empirical limitations that our framework sheds light upon. First, it is only informative about the term structure of *risk-free* discount rates $\overline{r}_t^{f,n}$. Since investments in climate change abatement are unlikely to be risk-free, one also needs to understand the term structure of risk premia to determine appropriate discount rates.[35] Second, downward-sloping risk-free discount rates are inconsistent with the essentially flat, and certainly not strongly downward-sloping, real yield curve for U.K. government bonds reported in Figure VI. We conclude that the Ramsey framework provides only limited guidance for climate change abatement investment.

Gamma Discounting. A third approach, pioneered by Weitzman (2001), generates DDRs in a reduced-form environment where a number of experts, who are assumed to have a constant but heterogeneous discount rate, are polled regarding the appropriate level of the discount rate for environmental policy. The author then shows that if the disagreement among experts takes the form of a particular (Gamma) distribution, the resulting expert-average discount rate is hyperbolic and thus declining in the horizon of the investment. This analysis is theoretically interesting from a policy perspective and largely complementary to our own analysis. We have shown that downward-sloping term structures of discount rates are a pervasive phenomenon in the data and can be rationalized in both models and real life situations where the polling of experts plays no role.

Epstein–Zin Preferences. Finally, another strand of the climate change literature produces interesting shapes for the term structure of discount rates by introducing preferences that reflect a distaste for long-run risks (in particular Epstein–Zin preferences with a risk aversion, γ, that is higher than the inverse of the elasticity of intertemporal substitution, ψ). As discussed in Section 4, standard calibrations of Epstein–Zin preferences imply that risky assets that are positively exposed to risk in long-run consumption growth (such as long-maturity claims to aggregate cash flows) will have an *increasing* term structure of discount rates, inconsistent with our data for real estate (see also van Binsbergen, Brandt and Koijen, 2012). Standard calibrations of Epstein–Zin preferences in a long-run risk economy also generate a downward slope in the term structure of risk-free

[35]For example, in an exciting debate Gollier (2013) points out that in the DICE model climate change costs, being a by-product of production and consumption, tend to occur the most in good states of the economy, thus behaving like a hedge asset. Weitzman (2012), on the contrary, argues that climate costs behave like a risky asset because they occur the most in bad states of the economy via tipping points and feedback loops that induce low output and consumption as a result of a climatic disaster.

discount rates, a feature that is inconsistent with the yield curve data in Figure VI (see also Beeler and Campbell, 2012). For example, Gollier (2013) evaluates climate change policy in a long-run risk model that combines parametric uncertainty about the growth rate of the economy with uncertainty about the probability of rare disasters. He acknowledges that the model produces a counterfactual downward-sloping term structure of risk-free discount rates and an upward-sloping term structure of risky discount rates in its benchmark parametrization with Epstein–Zin preferences.[36]

IIIb. Forces that work towards generating DDRs for risky assets

As we show in Section 4, traditional preference specifications cannot, by themselves, produce a downward-sloping term structure of discount rates for risky assets, with some even working against such a downward slope. This suggests that the observed downward-sloping term structure of discount rates for risky assets must be largely driven by its cash flow dynamics.

In particular, our empirical finding of a downward-sloping term structure of discount rates for risky assets such as real estate, paired with our structural analysis in Section 4, suggests that important shocks in the economy exhibit partial mean reversion in cash flows. Intuitively, after a negative shock hits the economy, the economy is expected to recover and grow faster when cash flows are (partially) mean-reverting. This makes claims to short-run cash flows that do not benefit from this mean reversion riskier than long-run cash flows that do, and creates a downward-sloping term structure of discount rates. In this spirit, Nakamura et al. (2013) motivate their downward-sloping term structure of discount rates for risky assets by documenting that consumption disasters as in Barro (2006) and Gabaix (2012) tend to partially mean-revert in a large cross section of countries for the last century. Such mean reversion also rationalizes the downward-sloping risk premia in our parameterization of the Lettau and Wachter (2007) model in Section 4.3.

Our empirical estimates in Section 3.2 show that real estate is exposed to many economic disasters, including consumption disasters and wars. Specifically, due to its tight connection with immovable land, it is an asset class that is particularly exposed to risks that are focused geographically. As such, real estate cash flows might be informative about the response of the economy to climate disasters (e.g., real estate is likely to lose substantial value on average if sea levels rise significantly). The evidence on mean reversion after disasters in Nakamura et al. (2013) and our structural analysis suggest that the

[36] The long-run risk model of climate change by Bansal, Kiku and Ochoa (2013) behaves similarly. Cai, Judd and Lontzek (2013) also study climate change policy implications in a DSGE extension of the DICE model that includes beliefs about uncertain climate tipping points along with Epstein–Zin preferences.

economy as a whole might also exhibit features of mean reversion in cash flows following climate change disasters. Within the environmental literature, our results are therefore consistent with an interesting and empirically-motivated modification of models of climate change as a disaster risk (e.g., Barro, 2006; Weitzman, 2012), that would allow for a partial recovery of economic activity following a climate disaster. It remains an exciting open question for future research to investigate how such partial mean reversion can be rationalized in models that mix physical elements of climate change (tipping points, increasing ocean levels, etc.) with the likely response of economic activity (technological innovation, geographic relocation of production, etc.). Exciting new work in this broad direction is being undertaken by Crost and Traeger (2014), Desmet and Rossi-Hansberg (2015), Lemoine (2015), Lemoine and Traeger (2014), and Traeger (2014*a*, 2015).

IV. Social and private discounting

A well-known problem when thinking about climate change is the potential difference between private and social discount rates. If individuals are not fully internalizing all future benefits of a certain investment (for example in climate change abatement), observing private discount rates related to that investment will not be informative about the social discount rate a policymaker should apply in a normative analysis, i.e., when internalizing all these benefits.

When estimating discount rates from different assets and applying them to climate change policies, it is therefore crucial to consider assets for which individuals internalize all costs and benefits, just like a policymaker would.[37] For those assets, private and social discount rates are more closely aligned. The empirical evidence in this paper is based on one particular asset class, real estate, that is typically viewed as a very long-duration asset that passes from generation to generation. Based on a similar reasoning, a large literature has used hedonic regressions of house prices to estimate the values that policymakers should assign to environmental amenities and dis-amenities. For example, Chay and Greenstone (2005) estimate the capitalization of total suspended particulates air pollution into real estate values. Greenstone and Gallagher (2008) look at changes in house prices to estimate households' valuation of cleanups of hazardous waste cites. In agreement with this literature, we conclude that the *private* discount rates in real estate are highly informative and provide a tight upper bound for *social* discount rates in real

[37]See also Schneider, Traeger and Winkler (2012) for an interesting discussion of the challenges in using observable data based on individual decision-making in infinite-time-horizon models, which result from the discrepancy between individual decision horizons and time horizons relevant for discounting climate change investments.

estate; and, with appropriate risk and maturity adjustments, they are informative for the social discount rates of other investments such as climate change projects.

V. Time consistency and optimal intertemporal policy

The environmental literature on DDRs has emphasized the problem of time-inconsistency supposedly immanent in such discounting schedules. These worries mainly stem from the motivation of DDRs based on hyperbolic or gamma discounting. However, our discussion in Section 4 shows that mean-reverting cash flows in general, and mean-reverting disasters paired with standard power utility preferences in particular, are able to rationalize a downward-sloping term structure of discount rates for risky assets. Note that such a setup is void of any time-inconsistency problem. In particular, the specific calibration of our reduced-form model can be recovered from agents with time-separable utility and preferences over one-period consumption growth, $U = E_t[\sum_{k=0}^{\infty} \delta^k u(c_{t+k})]$, with u of the power-utility form, and consumption equal to the aggregate dividends paid by the stock market.

6 Conclusion

We estimate the term structure of discount rates for real estate over the very long horizons (100 or more years) crucial for evaluating investments in climate change abatement. We find these discount rates to be downward-sloping over maturity, with an average rate of return to real estate of at least 6.4% and a long-run discount rate of 2.6%. In light of this evidence, we provide a structural analysis based on asset pricing theory to show how these declining discount rates, which are also found in other risky asset classes such as equity, imply a declining term structure of risk premia, which we link to partial mean reversion in aggregate cash flows.

Based on our term structure estimates and our structural analysis, we derive several implications for the appropriate discount rate schedule for investments in climate change abatement. First, we show that applying observed average returns of traded assets to discount cash flows from investments in climate change abatement is misleading when discount rates vary substantially across horizons. Second, we show that as long as climate change is a form of aggregate risk, the long-run benefits of investments in its abatement should be discounted at a rate below 2.6%. Finally, we discuss which of the mechanisms used in the current literature to generate a downward-sloping term structure of discount rates are consistent with our empirical findings.

References

Ambrose, Brent W., Piet Eichholtz, and Thies Lindenthal. 2013. "House Prices and Fundamentals: 355 Years of Evidence." *Journal of Money, Credit and Banking*, 45(2-3): 477–491.

Arrow, Kenneth, Maureen Cropper, Christian Gollier, Ben Groom, Geoffrey Heal, Richard Newell, William Nordhaus, Robert Pindyck, William Pizer, Paul Portney, Thomas Sterner, Richard S. J. Tol, and Martin Weitzman. 2013*a*. "Determining Benefits and Costs for Future Generations." *Science*, 341(6144): 349–350.

Arrow, Kenneth, Maureen Cropper, Christian Gollier, Ben Groom, Geoffrey Heal, Richard Newell, William Nordhaus, Robert Pindyck, William Pizer, Paul Portney, et al. 2013*b*. "How Should Benefits and Costs Be Discounted in an Intergenerational Context? The Views of an Expert Panel." *Resources for the Future Discussion Paper No. 12-53*.

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.

Bansal, Ravi, Dana Kiku, and Marcelo Ochoa. 2013. "Climate Change, Growth, and Risk."

Barro, Robert J. 1999. "Ramsey Meets Laibson in the Neoclassical Growth Model." *Quarterly Journal of Economics*, 114(4): 1125–1152.

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121(3): 823–866.

Barro, Robert J. 2013. "Environmental Protection, Rare Disasters, and Discount Rates." National Bureau of Economic Research Working Paper 19258.

Barro, Robert J., and Jose F. Ursua. 2008. "Macroeconomic Crises since 1870." *Brookings Papers on Economic Activity*, 39(1): 255–350.

Beeler, Jason, and John Y. Campbell. 2012. "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment." *Critical Finance Review*, 1: 141–182.

Bordo, Michael, Barry Eichengreen, Daniela Klingebiel, and Maria Soledad Martinez-Peria. 2001. "Is the Crisis Problem Growing More Severe?" *Economic Policy*, 16(32): 51–82.

Cai, Yongyang, Kenneth L. Judd, and Thomas S. Lontzek. 2013. "The Social Cost of Stochastic and Irreversible Climate Change." *NBER Working Paper*, 18704.

Campbell, John Y, and Tuomo Vuolteenaho. 2004. "Bad Beta, Good Beta." *American Economic Review*, 94(5): 1249–1275.

Campbell, Sean D, Morris A Davis, Joshua Gallin, and Robert F Martin. 2009. "What Moves Housing Markets: A Variance Decomposition of the Rent–Price Ratio." *Journal of Urban Economics*, 66(2): 90–102.

Chay, Kenneth Y., and Michael Greenstone. 2005. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy*, 113(2).

Cocco, Joao F. 2005. "Portfolio Choice in the Presence of Housing." *Review of Financial Studies*, 18(2): 535–567.

Cochrane, John H. 2005. *Asset Pricing.* Princeton, NJ:Princeton University Press.

Cropper, Maureen, and David Laibson. 1998. *The Implications of Hyperbolic Discounting for Project Evaluation.* The World Bank.

Cropper, Maureen L., Mark C. Freeman, Ben Groom, and William A. Pizer. 2014. "Declining Discount Rates." *American Economic Review*, 104(5): 538–43.

Crost, Benjamin, and Christian P. Traeger. 2014. "Optimal CO2 Mitigation Under Damage Risk Valuation." *Nature Climate Change*.

Dasgupta, Partha. 2008. "Discounting Climate Change." *Journal of Risk and Uncertainty*, 37(2): 141–169.

Desmet, Klaus, and Esteban Rossi-Hansberg. 2015. "On the Spatial Economic Impact of Global Warming." *Journal of Urban Economics*, 88: 16 – 37.

Dew-Becker, Ian, and Stefano Giglio. 2013. "Asset Pricing in the Frequency Domain: Theory and Empirics." National Bureau of Economic Research Working Paper 19416.

Díaz, Antonia, and María José Luengo-Prado. 2008. "On the User Cost and Homeownership." *Review of Economic Dynamics*, 11(3): 584–613.

Dreze, Jean, and Nicholas Stern. 1987. "The Theory of Cost-Benefit Analysis." In *Handbook of Public Economics*. Vol. 2, , ed. A. J. Auerbach and M. Feldstein, Chapter 14, 909–989. Elsevier.

Drupp, Moritz A., Mark Freeman, Ben Groom, and Frikk Nesje. 2015. "Discounting Disentangled: An Expert Survey on the Determinants of the Long-Term Social Discount Rate." *Centre for Climate Change Economics and Policy Working Paper No. 195*.

Eichholtz, Piet. 1997. "A Long Run House Price Index: The Herengracht Index, 1628–1973." *Real Estate Economics*, 25(2): 175–192.

Eitrheim, Øyvind, and Solveig K Erlandsen. 2005. "House Prices in Norway, 1819–1989." *Scandinavian Economic History Review*, 53(3): 7–33.

Farmer, J. Doyne, John Geanakoplos, Jaume Masoliver, Miquel Montero, and Josep Perelló. 2015. "Value of the Future: Discounting in Random Environments." *Physical Review E*, 91: 052816.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn van Nieuwerburgh. 2010. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium." National Bureau of Economic Research Working Paper 15988.

Flavin, Marjorie, and Takashi Yamashita. 2002. "Owner-Occupied Housing and the Composition of the Household Portfolio." *American Economic Review*, 92(1): 345–362.

Gabaix, Xavier. 2012. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance." *Quarterly Journal of Economics*, 127(2): 645–700.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel. 2015*a*. "No-Bubble Condition: Model-Free Tests in Housing Markets." *Econometrica*, Forthcoming.

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

Gollier, Christian. 2008. "Discounting With Fat-Tailed Economic Growth." *Journal of Risk and Uncertainty*, 37(2/3): 171–186.

Gollier, Christian. 2011. "On the Understanding of the Precautionary Effect in Discounting." *Geneva Risk and Insurance Review*, 36: 95–111.

Gollier, Christian. 2012. *Pricing the Planet's Future: The Economics of Discounting in an Uncertain World*. Princeton, NJ:Princeton University Press.

Gollier, Christian. 2013. "Evaluation of Long-Dated Investments Under Uncertain Growth Trend, Volatility and Catastrophes." Toulouse School of Economics TSE Working Papers 12-361.

Gordon, Myron J. 1982. *The Investment, Financing, and Valuation of the Corporation*. Greenwood Press.

Greenstone, Michael, and Justin Gallagher. 2008. "Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program." *The Quarterly Journal of Economics*, 123(3): 951–1003.

Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton. 2013. "Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation." *Review of Environmental Economics and Policy*, 7(1): 23–46.

Groom, Ben, Cameron Hepburn, Phoebe Koundouri, and David Pearce. 2005. "Declining Discount Rates: The Long and the Short of it." *Environmental and Resource Economics*, 32(4): 445–493.

Gyourko, Joseph, and Donald B. Keim. 1992. "What Does the Stock Market Tell us About Real Estate Returns?" *Real Estate Economics*, 20(3): 457–485.
Hansen, Lars Peter, and Scott F Richard. 1987. "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models." *Econometrica*, 55: 587–614.
Harding, John P, Stuart S Rosenthal, and CF Sirmans. 2007. "Depreciation of Housing Capital, Maintenance, and House Price Inflation: Estimates From a Repeat Sales Model." *Journal of Urban Economics*, 61(2): 193–217.
HM Treasury. 2003. "Green Book."
Kaplow, Louis, Elisabeth Moyer, and David A. Weisbach. 2010. "The Social Evaluation of Intergenerational Policies and Its Application to Integrated Assessment Models of Climate Change." *The B.E. Journal of Economic Analysis & Policy*, 10(2).
Karp, Larry, and Christian P. Traeger. 2013. "Discounting." *Encyclopedia of Energy, Natural Resource, and Environmental Economics*, 286–292.
Kopp, Robert E., Alexander Golub, Nathaniel O. Keohane, and Chikara Onda. 2012. "The Influence of the Specification of Climate Change Damages on the Social Cost of Carbon." *Economics: The Open-Access, Open-Assessment E-Journal*, 6(2012-13): 1 – 40.
Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 112(2): 443–478.
Lebègue, Daniel. 2005. "Révision du taux d'actualisation des investissements publics. Rapport du Groupe d'Experts, Commisariat Gènèral du Plan."
Leigh, Wilhelmina A. 1980. "Economic Depreciation of the Residential Housing Stock of the United States, 1950-1970." *Review of Economics and Statistics*, 62(2): 200–206.
Lemoine, Derek. 2015. "The Climate Risk Premium." *Available at SSRN 2560031*.
Lemoine, Derek, and Christian P. Traeger. 2014. "Watch Your Step: Optimal Policy in a Tipping Climate." *American Economic Journal: Economic Policy*, 6(1): 137–166.
Lettau, Martin, and Jessica A. Wachter. 2007. "Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium." *Journal of Finance*, 62(1): 55–92.
Lettau, Martin, and Jessica A. Wachter. 2011. "The Term Structures of Equity and Interest Rates." *Journal of Financial Economics*, 101(1): 90–113.
Lucas, Robert E, Jr. 1978. "Asset Prices in an Exchange Economy." *Econometrica*, 46(6): 1429–45.
Lustig, Hanno N., and Stijn van Nieuwerburgh. 2005. "Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective." *Journal of Finance*, 60(3): 1167–1219.
Luttmer, Erzo G., and Thomas Mariotti. 2003. "Subjective Discounting in an Exchange Economy." *Journal of Political Economy*, 111(5): 959–989.
Mack, Adrienne, and Enrique Martínez-García. 2011. "A Cross-Country Quarterly Database of Real House Prices: A Methodological Note." Federal Reserve Bank of Dallas Globalization and Monetary Policy Working Paper 99.
Malpezzi, Stephen, Larry Ozanne, and Thomas G Thibodeau. 1987. "Microeconomic Estimates of Housing Depreciation." *Land Economics*, 63(4): 372–385.
Mayerhauser, Nicole, and Marshall Reinsdorf. 2006. "Housing Services in the National Economic Accounts." *US Bureau of Economic Analysis*.
McCarthy, Jonathan, and Richard Peach. 2010. "The Measurement of Rent Inflation." FRB of New York Staff Report 425.
Metcalf, Gilbert E., and James Stock. 2015. "The Role of Integrated Assessment Models in Climate Policy: A User's Guide and Assessment." *The Harvard Project on Climate Agreements Discussion Paper 15-68*.
Moyer, Elisabeth J., Mark D. Woolley, Nathan J. Matteson, Michael J. Glotter, and

David A. Weisbach. 2014. "Climate Impacts on Economic Growth as Drivers of Uncertainty in the Social Cost of Carbon." *The Journal of Legal Studies*, 43(2): 401 – 425.

Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa. 2013. "Crises and Recoveries in an Empirical Model of Consumption Disasters." *American Economic Journal: Macroeconomics*, 5(3): 35–74.

Nordhaus, William D. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies.* Yale University Press.

Nordhaus, William D. 2013. *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World.* Yale University Press.

Nordhaus, William D., and Joseph Boyer. 2000. *Warming the World: Economic Models of Global Warming.* MIT press.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2007. "Housing, Consumption and Asset Pricing." *Journal of Financial Economics*, 83(3): 531–569.

Piketty, Thomas. 2014. *Capital in the Twenty-first Century.* Harvard University Press.

Pizer, William, Matthew Adler, Joseph Aldy, David Anthoff, Maureen Cropper, Kenneth Gillingham, Michael Greenstone, Brian Murray, Richard Newell, Richard Richels, et al. 2014. "Using and Improving the Social Cost of Carbon." *Science*, 346(6214): 1189–1190.

Reinhart, Carmen M., and Kenneth S. Rogoff. 2008. "Is the 2007 US Sub-Prime Financial Crisis so Different? An International Historical Comparison." *American Economic Review: Papers and Proceedings*, 98(2): 339–344.

Reinhart, Carmen M., and Kenneth S. Rogoff. 2009. *This Time is Different: Eight Centuries of Financial Folly.* Princeton University Press.

Schneider, Maik T., Christian P. Traeger, and Ralph Winkler. 2012. "Trading Off Generations: Equity, Discounting, and Climate Change." *European Economic Review*, 56(8): 1621 – 1644.

Schularick, Moritz, and Alan M. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises." *American Economic Review*, 102(2): 1029–1061.

Shapiro, Harold T., Roseanne Diab, Carlos H. DeBrito Cruz, Maureen Cropper, Jingyun Fang, Louise O. Fresco, Syukuro Manabe, Goverdhan Mehta, Mario Molina, Peter Williams, Ernst-Ludwig Winacker, and Abdul H. Zakri. 2010. "Climate Change Assessments – Review of the Processes and Procedures of the IPCC." *Amsterdam, InterAcademy*.

Shiller, Robert J. 2006. "Long-Term Perspectives on the Current Boom in Home Prices." *The Economists' Voice*, 3(4): 1–11.

Shiller, Robert J. 2007. "Understanding Recent Trends in House Prices and Home Ownership." National Bureau of Economic Research Working Paper 15988.

Stern, Nicholas. 2006. *Stern Review on the Economics of Climate Change.* London, UK:Her Majesty's Treasury.

Stern, Nicholas. 2008. "The Economics of Climate Change." *American Economic Review*, 98(2): 1–37.

Stern, Nicholas. 2014*a*. "Ethics, Equity and the Economics of Climate Change Paper 1: Science and Philosophy." *Economics and Philosophy*, 30: 397–444.

Stern, Nicholas. 2014*b*. "Ethics, Equity and the Economics of Climate Change Paper 2: Economics and Politics." *Economics and Philosophy*, 30: 445–501.

Traeger, Christian P. 2013. "Discounting Under Uncertainty: Disentangling the Weitzman and the Gollier Effect." *Journal of Environmental Economics and Management*, 66(3): 573 – 582.

Traeger, Christian P. 2014*a*. "A 4-Stated DICE: Quantitatively Addressing Uncertainty Effects in Climate Change." *Environmental and Resource Economics*, 59(1): 1–37.

Traeger, Christian P. 2014*b*. "Why Uncertainty Matters: Discounting Under Intertemporal Risk Aversion and Ambiguity." *Economic Theory*, 56(3): 627–664.
Traeger, Christian P. 2015. "Analytic Integrated Assessment and Uncertainty."
van Benthem, Arthur. 2015. "What is the Optimal Speed Limit on Freeways?" *Journal of Public Economics*, 124: 44 – 62.
van Binsbergen, Jules, Michael Brandt, and Ralph Koijen. 2012. "On the Timing and Pricing of Dividends." *American Economic Review*, 102(4): 1596–1618.
van Binsbergen, Jules, Wouter Hueskes, Ralph Koijen, and Evert Vrugt. 2013. "Equity Yields." *Journal of Financial Economics*, 110(3): 503 – 519.
Wagner, Gernot, and Martin L. Weitzman. 2015. *Climate Shock: The Economic Consequences of a Hotter Planet.* Princeton, NJ:Princeton University Press.
Weisbach, David A. 2014. "Distributionally Weighted Cost-Benefit Analysis: Welfare Economics Meets Organizational Design." *Journal of Legal Analysis*, 7(1): 151 – 182.
Weisbach, David, and Cass R. Sunstein. 2009. "Climate Change and Discounting the Future: A Guide for the Perplexed." *Yale Law & Policy Review*, 27(2): 433 – 457.
Weitzman, Martin L. 2001. "Gamma Discounting." *American Economic Review*, 91(1): 260–271.
Weitzman, Martin L. 2012. "Rare Disasters, Tail-Hedged Investments, and Risk-Adjusted Discount Rates." National Bureau of Economic Research Working Paper 18496.

Figure I: Leasehold Price Discounts by Remaining Lease Length

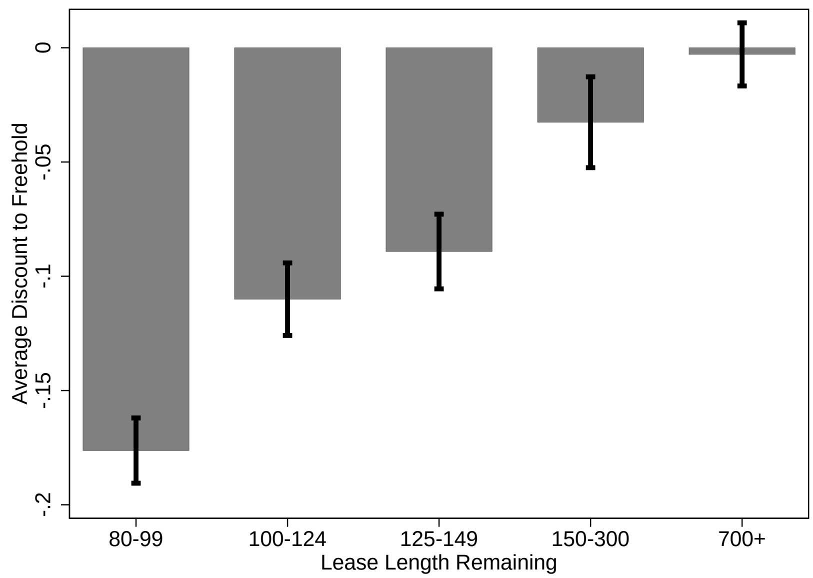

(A) Leasehold Discounts - U.K.

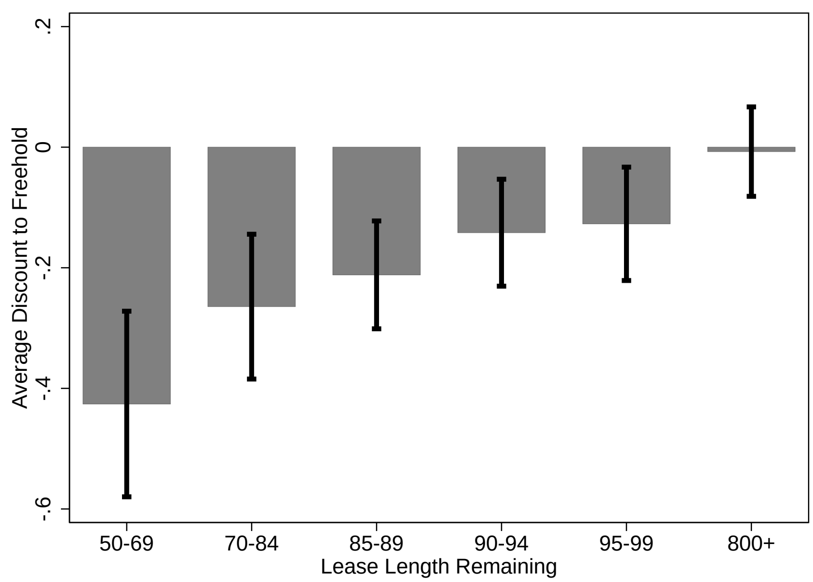

(B) Leasehold Discounts - Singapore

Note: Source: Giglio, Maggiori and Stroebel (2015*b*). The top panel plots log-price difference between leaseholds of varying remaining maturity and freeholds for flats sold in the U.K. between 2004 and 2013. The bottom panel plots log-price difference between leaseholds of varying remaining maturity and freeholds in Singapore between 1995 and 2013. The vertical axis is expressed in log-points, the horizontal axis shows leasehold discounts depending on the remaining term of the lease. The bars indicate the 95% confidence interval of the estimate. See original reference for further details on data and estimation.

Figure II: House Price Riskiness I

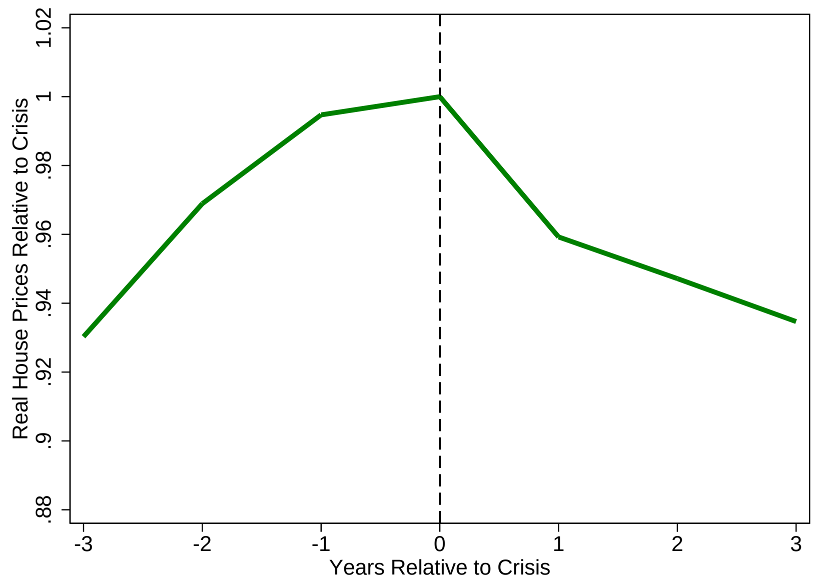

(A) Financial Crises

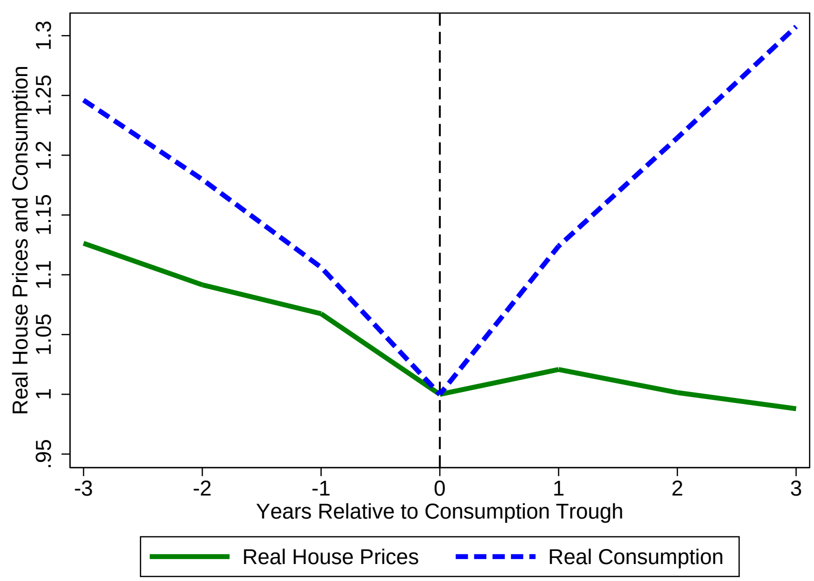

(B) Consumption Disasters

Note: The top panel shows average real house price movements relative to financial crises in Schularick and Taylor (2012), Bordo et al. (2001) and Reinhart and Rogoff (2009). The bottom panel shows average real house price movements and average real consumption relative to the trough of consumption disasters identified by Barro and Ursua (2008). House prices and consumption volumes during the reference year are normalized to 1. See Appendix A.2.1 and Appendix Table A.2 for a description of the countries included and the data series and crises considered here.

Figure III: House Price Riskiness II

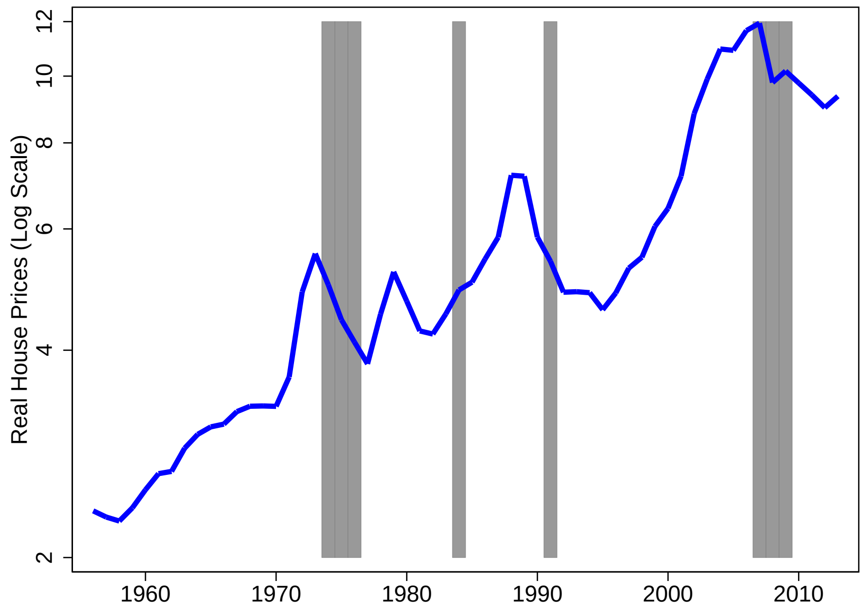

(A) U.K.

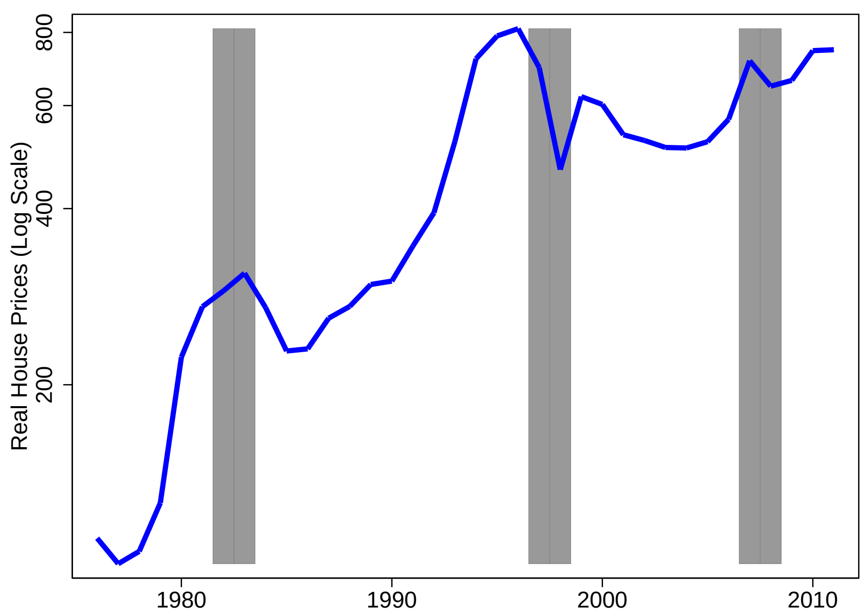

(B) Singapore

Note: The figures show the evolution of real house prices in the U.K. (top panel) and Singapore (bottom panel). Shaded regions for the U.K. are financial crises identified by Reinhart and Rogoff (2009): 1974-1976, 1984, 1991 and 2008-2009. Shaded regions for Singapore include the 1982-1983 financial crisis identified by Reinhart and Rogoff (2009), as well as the Asian financial crisis (1997-1998), and the 2007-2008 global financial crisis. See Appendix A.2.1 for a description of the data series studied here.

Figure IV: House Price Riskiness III

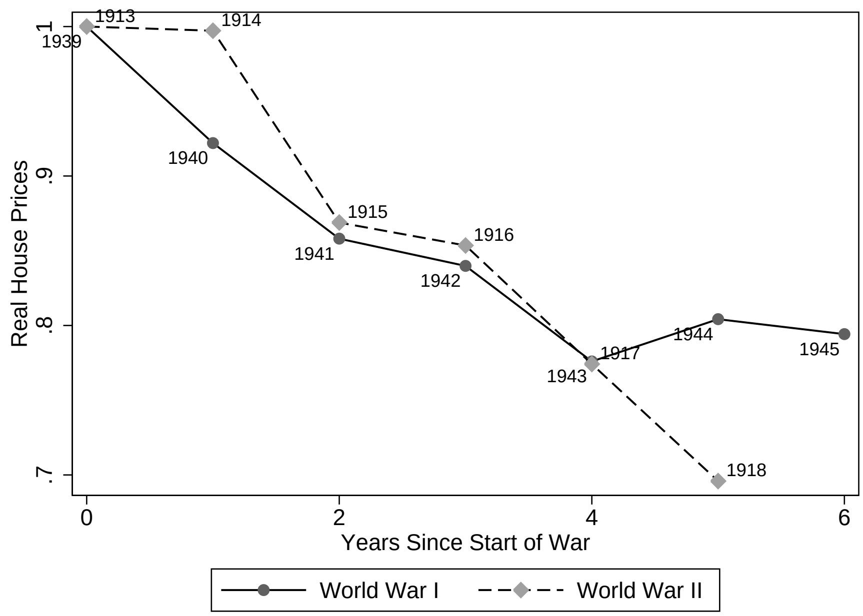

(A) Housing During World Wars

Note: Figure shows the evolution of real house prices for countries with available house-price time series during World War I (Australia, France, Netherlands, Norway, United States) and World War II (Australia, France, Netherlands, Norway, Switzerland, United States). See Appendix A.2.1 for a description of the data series studied here.

Figure V: Gordon Growth Model and Estimated Discounts

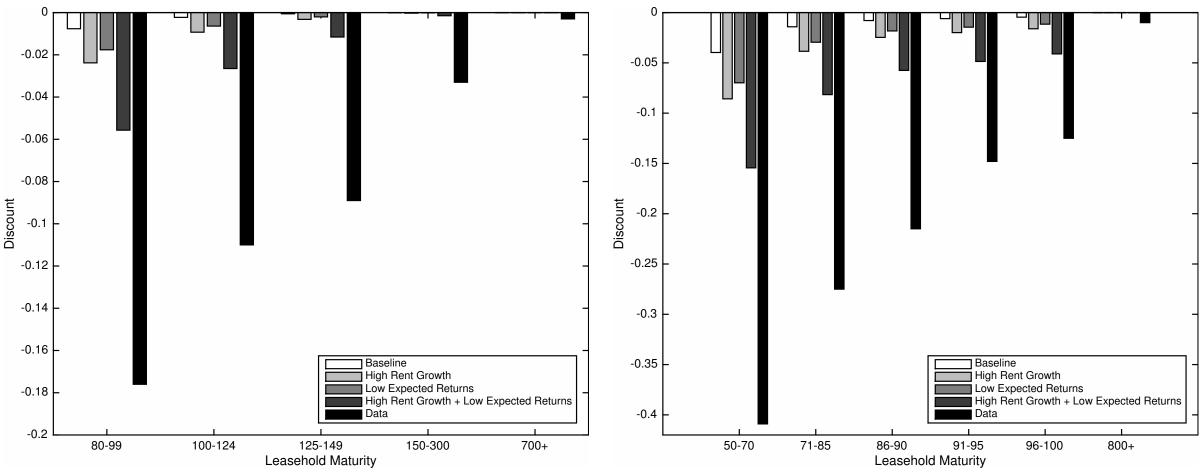

(A) U.K. (B) Singapore

Note: Figure shows the discounts for leaseholds observed in the U.K. (Panel A) and Singapore (Panel B) together with discounts predicted by a Gordon growth model with $r = 6.4\%$ and $g = 0.7\%$ (Baseline), $r = 6.4\%$ and $g = 2\%$ (High Rent Growth), $r = 5.4\%$ and $g = 0.7\%$ (Low Expected Returns), and $r = 5.4\%$ and $g = 2\%$ (High Rent Growth + Low Expected Returns).

Figure VI: U.K. Gilts Real Yield Curve

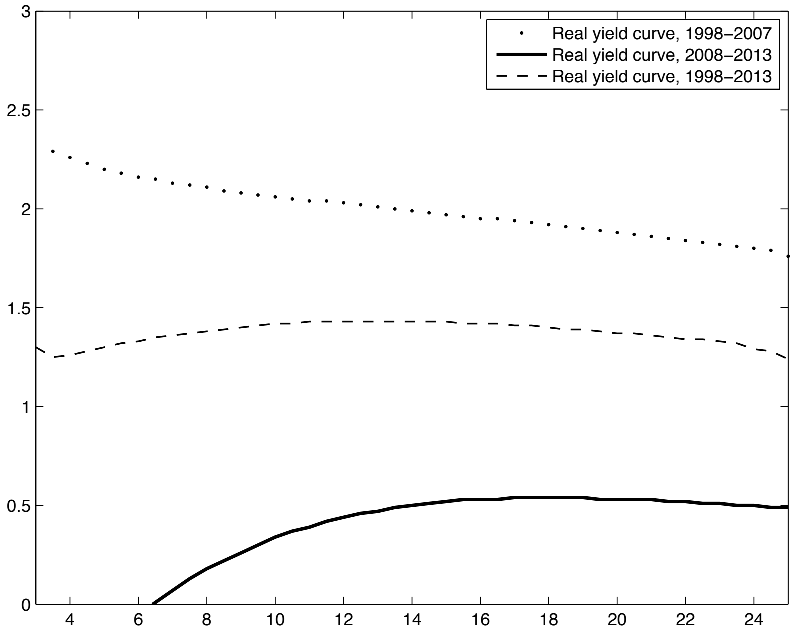

Note: The figure plots the real yield curve for U.K. gilts as computed by the Bank of England. It is available at http://www.bankofengland.co.uk/statistics/Pages/yieldcurve/archive.aspx. The figure plots the average shape of the real U.K. gilts curve for the period 1998-2013, as well as for two sub-periods: 1998-2007 and 2007-2013. The U.K. government debt also includes some perpetual bonds: the War Loan and the Annuities. These bonds comprise a negligible part of the outstanding U.K. government debt (£2.6bn out of £1.5trn of debt outstanding), and are classified as small and illiquid issuances by the U.K. Debt and Management Office. They are excluded from our analysis, not only because they are nominal and we only use data on U.K. real gilts, but also because their negligible size, scarce liquidity, and callability make it hard to interpret their prices in terms of discount rates.

Figure VII: Term Structure of Discount Rates

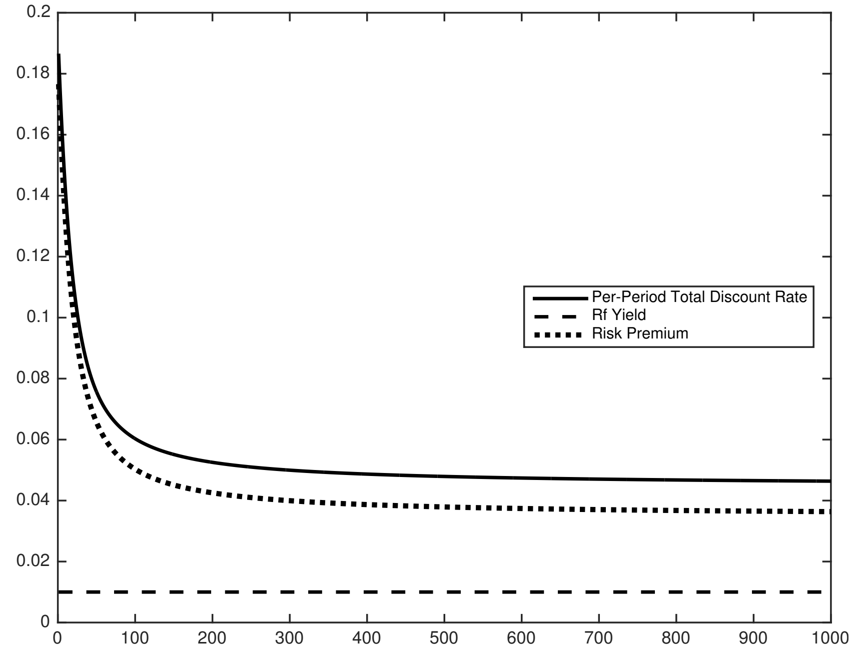

(A) Equity

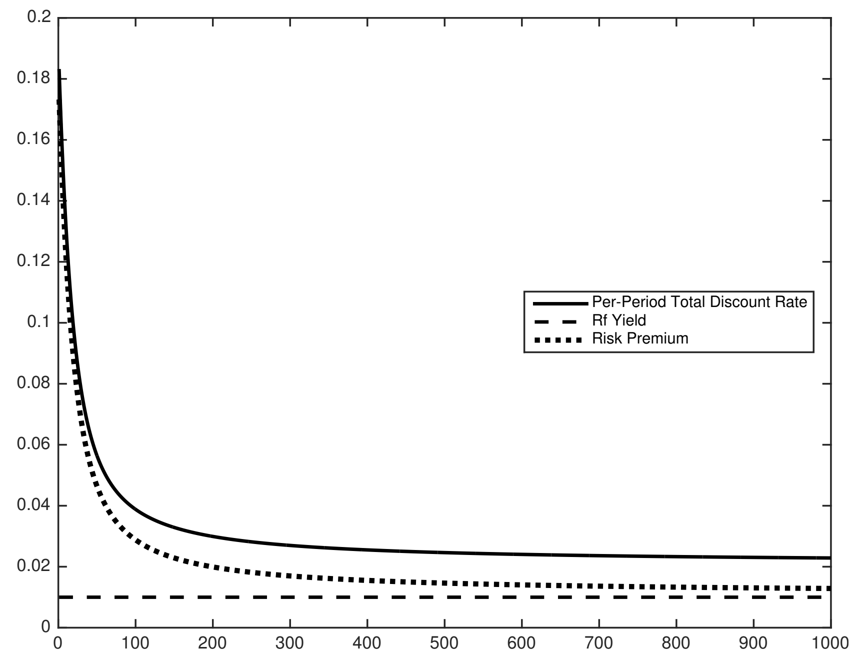

(B) Real Estate

Note: Panel A shows per-period discount rates, risk-free rates and risk premia of claims to single dividends with a maturity up to 1000 years implied by the reduced-form asset pricing model from Section 4.3. Panel B shows per-period discount rates, risk-free rates and risk premia of claims to single rents with a maturity up to 1000 years implied by the reduced-form asset pricing model from Section 4.3. The calibration details are discussed in Appendix A.3.2.1 and summarized in Table IV.

Table I: Expected Returns and Rental Growth

	United States		Singapore		United Kingdom	
	Balance Sheet	Price/Rent	Balance Sheet	Price/Rent	Balance Sheet	Price/Rent
Gross Return	10.3%	9.6%	10.4%	10.3%	10.5%	9.8%
Rental Yield	8.3%	9.2%	6.1%	6.0%	7.7%	5.8%
Capital Gain	2.0%	0.4%	4.3%	4.3%	2.8%	4%
Depreciation	2.5%	2.5%	2.5%	2.5%	2.5%	2.5%
Taxes	0.67%	0.67%	0.5%	0.5%	0%	0%
Net Return	**7.1%**	**6.4%**	**7.4%**	**7.3%**	**8.0%**	**7.3%**
Sample	1953-2012	1953-2012	1985-2012	1990-2012	1989-2012	1996-2012
Real Rent Growth		0.5%		0.2%		0.7%
Sample		1953-2012		1990-2012		1996-2012

Note: This table shows our estimates for net real returns to housing and real rent growth in the U.S., the U.K., and Singapore. See Appendix A.1 for details on the estimation procedures and the underlying data sources used.

Table II: Real House Price Growth and Real Consumption Growth

	Period	Real HP Growth		Real Cons. Growth		
		Mean	Std. Dev.	Mean	Std. Dev.	Correlation
Australia	1901 - 2009	2.51%	12.1%	1.51%	5.00%	0.102
Belgium	1975 - 2009	2.92%	6.06%	1.56%	1.49%	0.438
Canada	1975 - 2009	2.38%	7.69%	1.64%	1.71%	0.433
Denmark	1975 - 2009	1.99%	9.24%	1.03%	2.68%	0.538
Finland	1975 - 2009	2.17%	8.70%	2.09%	2.75%	0.710
France	1840 - 2009	2.06%	11.8%	1.53%	6.32%	-0.054
Germany	1975 - 2009	-0.45%	2.33%	1.71%	1.56%	0.494
Italy	1975 - 2009	1.28%	8.10%	1.69%	2.12%	0.165
Japan	1975 - 2009	0.02%	4.45%	2.00%	1.59%	0.502
Netherlands	1814 - 2009	2.79%	21.0%	1.57%	7.49%	0.078
New Zealand	1975 - 2009	2.46%	8.09%	0.90%	2.34%	0.578
Norway	1830 - 2009	1.77%	11.6%	1.78%	3.83%	0.243
Singapore	1975 - 2009	7.18%	19.5%	3.37%	2.98%	0.348
South Africa	1975 - 2009	1.13%	10.1%	0.90%	2.98%	0.707
South Korea	1975 - 2009	0.58%	7.93%	4.58%	4.43%	0.370
Spain	1975 - 2009	3.14%	8.07%	1.54%	2.57%	0.593
Sweden	1952 - 2009	1.55%	6.04%	1.66%	1.98%	0.536
Switzerland	1937 - 2009	0.47%	7.93%	1.55%	3.85%	0.187
U.K.	1952 - 2009	2.89%	9.55%	2.22%	2.12%	0.700
U.S	1890 - 2009	0.49%	7.36%	1.80%	3.41%	0.148

Note: The table shows time series properties of annual growth rates of real house prices (as described in Appendix A.2.1) and real consumption, as collected by Barro and Ursua (2008). Column (1) shows the sample considered. Columns (2) and (3) show the mean and standard deviation of real house price growth. Columns (4) and (5) show the mean and standard deviation of real consumption growth. Column (6) shows the correlation of real house price growth and real consumption growth.

Table III: Real House Price Growth and Personal Disposable Income Growth

	Real HP Growth		Real PDI Growth		
	Mean	Std. Dev.	Mean	Std. Dev.	Correlation
Australia	3.20%	6.89%	1.43%	2.77%	0.093
Belgium	2.80%	5.87%	1.17%	2.27%	0.436
Canada	2.51%	7.63%	1.37%	2.10%	0.489
Switzerland	0.94%	4.73%	1.12%	1.63%	0.445
Germany	-0.29%	2.31%	1.27%	1.70%	0.288
Denmark	1.57%	8.99%	1.09%	2.29%	0.211
Spain	2.05%	8.26%	0.83%	2.46%	0.631
Finland	2.04%	8.19%	2.07%	3.21%	0.482
France	2.52%	5.23%	1.22%	1.58%	0.358
U.K.	3.53%	8.54%	2.20%	2.74%	0.355
Italy	0.60%	8.28%	0.82%	2.44%	0.325
Japan	-0.24%	4.28%	1.55%	1.40%	0.587
S. Korea	0.59%	7.70%	3.95%	4.58%	0.235
Luxembourg	3.94%	6.68%	2.84%	3.75%	0.054
Netherlands	2.32%	9.43%	0.48%	3.25%	0.472
Norway	2.76%	7.23%	2.22%	2.52%	0.064
New Zealand	2.20%	7.73%	0.98%	3.45%	0.530
Sweden	1.50%	7.27%	1.34%	2.28%	0.431
U.S.	1.13%	3.89%	1.60%	1.56%	0.371
S. Africa	0.88%	9.65%	0.53%	3.05%	0.373

Note: This table shows time series properties of quarterly frequency annual growth rates of real house prices and personal disposable income between 1975 and Q2 2013, as collected by Mack and Martínez-García (2011). Columns (1) and (2) show the mean and standard deviation of real house price growth. Columns (3) and (4) the mean and standard deviation of real personal disposable income growth. Column (5) shows the correlation of real house price growth with real personal disposable income growth.

Table IV: Parameters of the Model

Calibrated Variables	Value
g_s	2.28%
g_h	0.7%
κ_s	1
κ_h	1.14
r^f	1%
$\overline{x}$	0.625
ϕ_z	0.91
$\|\sigma_s\|$	14.5%
$\|\sigma_h\|$	14.2%
$\|\sigma_z\|$	0.32%
Correlation of Δd^s and z shocks	-0.83

Implied Parameters	Value
σ_s	[0.0724,0]
σ_h	[0.071,0]
σ_z	[-0.00136,0.0009]

Note: The table shows the calibration of the reduced-form asset pricing model in Section 4.3. The unconditional mean of dividend and rent growth rates, the risk-free rate, the persistence variables, and the conditional standard deviations are in annual terms. The stock market including x is calibrated as in Lettau and Wachter (2007). The unconditional mean of the dividend growth rate and the conditional standard deviation of dividend growth are set to match their respective data counterparts. Under the assumption that the conditional mean of dividend growth can be identified with the log consumption-dividend ratio in the data, autocorrelation ϕ_z and the correlation between shocks to Δd^s and shocks to z are set to match their respective data counterparts. $\|\sigma_z\|$ is set to reflect the low predictability of dividend growth in the data (for details, see Lettau and Wachter, 2007). The unconditional mean of the rent growth rate is set to match the data. Volatility $\|\sigma_h\|$ and κ_h are calibrated to match the expected return on real estate in the data as well as the empirical term structure of real estate discount rates. The risk-free rate is set at 1%, based on the U.K. gilts real yield curve between 1998 and 2013 reported in Figure VI.

Table V: Simulated Moments for the Market and Cash Flow Growth

Moment	Equity	Housing
$E[R]$	9.1%	5.9%
$\sigma(R)$	9.0%	8.2%
Expected cash flow growth	3.4%	1.6%
$\sigma(\Delta d)$	14.5%	14.2%
Expected capital gains	2.7%	1.0%
Expected yields	6.5%	4.9%

Note: The table shows key moments from simulating the reduced-form asset pricing model in Section 4.3 for 50,000 quarters. Time series are aggregated to an annual frequency. R is the return on the market. Small letters indicate log variables. Possible differences between expected returns and the sum of expected capital gains and expected yields are due to rounding errors.

Appendix to "Climate Change and Long-Run Discount Rates: Evidence from Real Estate"

Stefano Giglio Matteo Maggiori Johannes Stroebel Andreas Weber

Not for publication

A.1 Details on the Estimation Procedures of the Average Return for Residential Real Estate

This section describes the methodology and data used to compute average real returns and rent growth for residential properties. We report the details of the calculations in an online worksheet.

The balance-sheet approach Following Favilukis, Ludvigson and van Nieuwerburgh (2010), this approach uses information about the value of the stock of residential real estate to estimate the value (price) of housing and total household expenditure on housing as a measure of the value of rents in each period. Since we are only interested in the return to a representative property, we need to control for changes in the total housing stock. We proxy for the change in the stock by population growth, assuming that at least over long periods the per capita stock of housing is constant. We derive the gross return to housing in each period as:

$$R^G_{t,t+1} = \frac{V_{t+1} + CE_{t+1}}{V_t} \frac{\pi_t}{\pi_{t+1}} \frac{L_t}{L_{t+1}},$$

where V is the value of the housing stock, CE is the household expenditure on housing, π is the CPI price level index, and L is population.

- For the U.S., we follow Favilukis, Ludvigson and van Nieuwerburgh (2010) and use data from the Flow of Funds (obtained from the Federal Reserve Board and the Federal Reserve Bank of St. Louis). For the value of the housing stock, we sum two series: owner-occupied real estate and tenant-occupied real estate (FL155035005, FL115035023) from the Flow of Funds. From the Federal Reserve Bank of St. Louis we obtain: (i) household expenditure on housing services (DHUTRC1A027NBEA); (ii) Population estimates (POP); and (iii) the Consumer Price Index (U.S.ACPIBLS).
- For the U.K., using the same procedure, we combine the value of the total stock of

housing (series CGRI) and the total expenditures on housing (series ADIZ) from the National Accounts (available from the Office of National Statistics). From the same source, we obtain the CPI (series D7BT). We adjust for the change in the stock of housing using the population growth series from ONS for the U.K.

- For Singapore, using a similar procedure, we obtain the value of the private residential stock of housing (series M013181.1.1.1 P017199) and the private consumption expenditure on housing and utilities (series M013131.1.4 P017135) from the National Accounts (available from the Department of Statistics Singapore). We obtain the CPI series from the same source. Finally, we obtain the series for the population growth (that proxies for the change in the stock of housing wealth) from the World Bank (series SP.POP.GROW).

The price-rent approach This approach constructs a time series of returns by combining information from a house price index, a rent index, and an estimate of the rent-to-price ratio in a baseline year. Without loss of generality, suppose we know the rent-to-price ratio at time $t = 0$. We can derive the time series of the rent-to-price ratio as:

$$\frac{P_t}{D_{t+1}} = \frac{P_t}{P_{t-1}} \frac{D_t}{D_{t+1}} \frac{P_{t-1}}{D_t}; \qquad \frac{P_0}{D_1} \quad \text{given},$$

where P is the price index and D the rental index. Notice that, given a baseline year $\frac{P_0}{D_1}$, only information about the growth rates in prices and rents are necessary for the calculations. We then compute real returns using the formula:

$$R^G_{t,t+1} = \left(\frac{D_{t+1}}{P_t} + \frac{P_{t+1}}{P_t} \right) \frac{\pi_t}{\pi_{t+1}},$$

where π is the CPI price level index again.

- For the U.S., we follow Favilukis, Ludvigson and van Nieuwerburgh (2010) and use the Case-Shiller 10-City House Price Index (series SPCS10RSA from the Federal Reserve Bank of St. Louis), and compute rent growth using the BLS shelter index (the component of CPI related to shelter, item CU.S.R0000SAH1 from the Federal Reserve Bank of St. Louis). However, unlike Favilukis, Ludvigson and van Nieuwerburgh (2010), we choose 2012 as a baseline year for the rent-price ratio, which we estimate to be 0.1; the choice of the baseline year is motivated by the availability of high quality data. In particular, we obtain two independent estimates for the rent-price ratio. The first estimate is the price-rent ratio implied by the balance-sheet approach. The second estimate is a direct estimate obtained from data by the real estate portal Trulia. Figure A.1 shows the distribution of rent-price ratios across the 100 largest MSAs provided by Trulia.[1] Both independent estimates imply a rent-price ratio of 10% in 2012. Figure A.2 suggests that these rent-price ratios are close to their long-run average.

- For the U.K., we use the house price index from the U.K. Land Registry to compute price appreciation and we use the CPI component "Actual rents for housing" (series D7CE) from the Office of National Statistics as a rental index. As a baseline, we used the 6% rent-price ratio obtained from the balance-sheet approach for 2012.

- For Singapore, we obtain a time series of price and rental indices for the whole island from the Urban Redevelopment Authority (the government's official housing arm government: ura.gov.sg). To estimate the baseline rent-price ratio, we use data from for-sale and for-rent listings provided by iProperty.com, Asia's largest online property listing portal. We observe approximately 105,000 unique listings from 2012, about 46% of which are for-rent listings. To estimate the rent-price ratio, we run the following regression, which pools both types of listings. The methodology is similar to the one used to construct rent-price ratios for the U.S. in Figure A.1:

$$\ln\left(ListingPrice\right)_{i,t} = \alpha + \beta_i ForRent_i + \gamma Controls_{i,t} + \epsilon_{i,t} \qquad \text{(A.1)}$$

  The dependent variable, *ListingPrice*, is equal to the list-price in "for-sale" listings, and equal to the annual rent in "for-rent" listings. $ForRent_i$ is an indicator variable that is equal to one if the listing is a for-rent listing. The results are reported in Table A.1. In column (1), we control for postal code by quarter fixed effects. The estimated coefficient on β_i suggests a rent-price ratio of $e^{\beta_i} = 4.5\%$. In columns (2) - (4), we also control for other characteristics of the property, such as the property type, the number of bedrooms, bathrooms as well as the size, age, and the floor of the building. In columns (3) and (4), we tighten fixed effects to the month by postal code level and the month by postal code by number of bedrooms level respectively. In all specifications, the estimated rent-price ratio for 2012 is 4.4% or 4.5%. Finally, we note that if we used the rent-price ratio obtained from the Balance Sheet approach as a baseline estimate in 2012 instead, we would obtain a higher total return (as the baseline in 2012 would be 6% rather than 4.5%). We conservatively choose 4.5%.

[1] We thank Jed Kolko and Trulia for providing these data. Trulia observes a large set of both for-sale and for-rent listings. The rent-price ratio is constructed using an MSA-level hedonic regression of $\ln(price)$ on property attributes, zip-code fixed effects, and a dummy for whether the unit is for sale or for rent. The rent-to-price ratio is constructed by taking the exponent of the coefficient on this dummy variable.

A.2 Details on the Riskiness of Housing

A.2.1 The Riskiness of Housing – Details on Main Analyses

This section provides the details underlying the analysis carried out in Section 3. Section A.2.2 will provide additional evidence for the riskiness of housing.

Table A.2 reports the availability of house price data and the associated financial crises and rare disasters. The first column in Table A.2 shows the time coverage of house price indices for each country. For some countries, we can go far back in time; for example, we sourced data as far back as 1819 for Norway, 1890 for the U.S., and 1840 for France. The second and third column report the dates of any banking crises or rare consumption disasters that occur in each country over the time period provided in the first column. Banking crises dates for all countries, except Singapore, Belgium, Finland, New Zealand, South Korea, and South Africa, are from Schularick and Taylor (2012). Banking crises dates for the countries not covered by Schularick and Taylor (2012) are from Reinhart and Rogoff (2009).[2] Rare disaster dates in the last column indicate the year of the trough in consumption during a consumption disaster as reported by Barro and Ursua (2008).

For each country, we obtained the longest continuous and high-quality time series of house price data available. To make the data comparable across countries and time periods, we focus on real house prices at an annual frequency. Finally, to increase historical comparability across time within each time series, we report each index for the unit of observation, for instance a city, for which the longest possible high quality time series is available. For example, since a house price index for France is only available since 1936, but a similar index is available for Paris since 1840, we focus on the Paris index for the entire history from 1840-2012. We stress, however, that for each index and country we have carried out an extensive comparison with alternative indices, in particular with indices available for the most recent time period, in order to ensure that we are observing consistent patterns in the data. In the following, we detail the sources for each of the 20 countries in our sample:

- **Australia**: Real annual house price indices are from Stapledon (2012). For our analysis, we use the arithmetic average of the indices (rebased such that 1880 = 100) for Melbourne and Sydney.

- **Belgium, Canada, Denmark, Finland, Germany, Japan, Italy, New Zealand, South Africa, South Korea, and Spain**: Real annual house price indices are from the

[2] For this second set of countries and dates, we have also consulted Bordo et al. (2001), who confirm all dates in Reinhart and Rogoff (2009), except 1985 for South Korea and 1989 for South Africa.

Federal Reserve Bank of Dallas.[3] The sources and methodology are described in Mack and Martínez-García (2011).

- **France**: Nominal annual house price index and CPI are available from the Conseil Général de l'Environnement et du Développement Durable (CGEDD).[4] We obtain the real house price index by deflating the nominal index using the CPI. For our analysis, we use the longer time series available for the Paris house price index.

- **Netherlands**: Nominal annual house price index for Amsterdam and CPI for the Netherlands are available from Eichholtz (1997) and Ambrose, Eichholtz and Lindenthal (2013).[5] We obtain the real house price index by deflating the nominal index using the CPI.

- **Norway**: Nominal annual house price index and CPI are from the Norges Bank.[6] We obtain the real house price index by deflating the nominal index using the CPI.

- **Singapore**: Nominal annual house price index for the whole island is from the Urban Redevelopment Authority (`http://www.ura.gov.sg`). CPI is from Statistics Singapore. We obtain the real house price index by deflating the nominal index using the CPI.

- **Sweden**: Nominal house price index for one-or-two-dwelling buildings and CPI are from Statistics Sweden. We obtain the real house price index by deflating the nominal index by CPI.

- **Switzerland**: Nominal house price index for Switzerland is available from Constantinescu and Francke (2013). Among the various indices the authors estimate, we focus on the local linear trend (LLT) index. The data are available for the period 1937-2007. We update the index for the period 2007-2012 by using the percentage growth of the house price index for Switzerland available from the Dallas Fed.[7] The CPI index for Switzerland is from the Office fédéral de la statistique (OFS). We obtain the real house price index by deflating the nominal index using the CPI.

[3] The data are available at: `http://www.dallasfed.org/institute/houseprice/`, last accessed February 2014.

[4] `http://www.cgedd.developpement-durable.gouv.fr/les-missions-du-cgedd-r206.html`, last accessed February 2014.

[5] Part of the data are available on Eicholtz' website at: `http://www.maastrichtuniversity.nl/web/Main/Sitewide/Content/EichholtzPiet.htm`, last accessed February 2014.

[6] `http://www.norges-bank.no/en/price-stability/historical-monetary-statistics/`, last accessed February 2014.

[7] This source is described in the second bullet point above.

- **U.K.:** Annual nominal house price data are from the Nationwide House Price Index. We divide the nominal index by the U.K. Office of National Statistics "long term indicator of prices of consumer goods and services" to obtain the real house price index. The Nationwide index has a missing value for the year 2005, for that year we impute the value based on the percentage change in value of the house price index produced by the England and Wales Land Registry.
- **U.S.**: Real annual house price data are originally from Shiller (2000). Updated data are available on the author's website.[8]

For all countries, the real annual consumption data are from Barro and Ursua (2008) and available on the authors' website.[9]

Figure II is produced by combining the time series of house prices and consumption described in the previous subsection with the dates for banking crises and rare disasters in Table A.2. When taking averages across countries in Panel A of Figure II for the 6 year windows around a banking crisis, the following countries have missing observations for the house price series: France data are unavailable for the year 2011 following the 2008 crisis, Netherlands data are unavailable for the years 2010 and 2011 following the 2008 crisis, and South Africa data are unavailable for the year 1974 before the 1977 crises. In these cases, the crises are still included in the sample but the average reported in the figure excludes these missing country-year observations.[10]

A.2.2 The Riskiness of Housing – Additional Evidence

In this section, we provide additional details and evidence for the riskiness of real estate to complement the analysis in Section 3.

Figure A.3 plots the growth rates of rents and personal consumption expenditures (PCE) in the U.S. since 1929. In periods of falling PCE, in particular the Great Depression, rents also fell noticeably. The bottom panel shows a (weak) positive relationship between the growth rates of rents and personal consumption expenditures. This suggests that housing rents tend to increase when consumption increases and marginal utility of consumption is low. Figure A.4 indicates that rents in London are positively correlated with house prices in London, but more volatile.

[8] Available at: `http://aida.wss.yale.edu/~shiller/data.htm`, last accesses February 2014.

[9] Available at: `http://rbarro.com/data-sets/`, last accessed February 2014.

[10] In unreported results, we have verified that the result is essentially unchanged if we exclude the 2008 crisis for France and the Netherlands and the 1975 crisis for South Africa from the data.

A.3 Details on the Reduced-Form Asset Pricing Model

This section presents the reduced-form asset pricing model, which extends Lettau and Wachter (2007) to include real estate and wealth as two additional asset classes and casts the stochastic discount factor as a function of consumption. Section A.3.1 reviews and discusses the assumptions on aggregate cash flows and the stochastic discount factor, and presents the key results for the term structure of discount rates. Section 4.3 in the main part of the paper features a simplified version of this model with a constant price of risk and a dividend-based stochastic discount factor. Section A.3.2 describes the calibration for this simplified version of the model as well as the calibration and consistent asset pricing results for the full model. Section A.3.3 derives the key results analytically.

A.3.1 Model and Key Results

A.3.1.1 Setup

Uncertainty in this model is characterized by three shocks: a shock to *current* cash flow growth, a shock to *expected* future cash flow growth, and a shock to the price of risk, which is directly specified in reduced-from. Let ϵ_{t+1} be a 3×1-vector of independent standard normal shocks. Denote by D_t^i with $i \in \{s, h, c\}$ aggregate stock market dividends, aggregate real estate and housing rents, and aggregate consumption in the economy, respectively, at time t, and let $d_t^i \equiv \ln(D_t^i)$. Aggregate cash flows evolve according to

$$\Delta d_{t+1}^i \equiv d_{t+1}^i - d_t^i = g_i + \kappa_i z_t + \sigma_i \epsilon_{t+1},$$

where z_t is a scalar that follows the $AR(1)$ process

$$z_{t+1} = \phi_z z_t + \sigma_z \epsilon_{t+1},$$

with persistence $0 \leq \phi_z < 1$, and where σ_i and σ_z are 1×3-row vectors. σ_i determines which combination of the three shocks in ϵ_{t+1} drives the transitory shock to dividend, rent, or consumption growth. z_t is a persistent long-run shock to the growth rate of cash flows that affects all three, dividends, rents, and consumption. To keep the model parsimonious, we assume that only *one* shock drives the long-term fluctuations of dividend, rent, and consumption growth, although all three can load on it differently (through different κ_i). The one-period risk-free rate, denoted by $r^f \equiv \ln(R^f)$, is assumed to be constant, as in Lettau and Wachter (2007). We let the stochastic discount factor be a function of consumption in the general setup. Only the current shock to cash flow growth is priced, so that the stochastic discount factor in this economy can be directly specified as:

$$M_{t,t+1} = \exp\{m_{t,t+1}\} = \exp\{-r^f - \frac{1}{2}x_t^2 - x_t\epsilon_{c,t+1}\},$$

with

$$\epsilon_{c,t+1} = \frac{\sigma_c}{||\sigma_c||}\epsilon_{t+1},$$

where x_t is the (time-varying) price of risk, which follows the $AR(1)$ process

$$x_{t+1} = (1-\phi_x)\overline{x} + \phi_x x_t + \sigma_x\epsilon_{t+1},$$

with $0 \leq \phi_x < 1$ and 1×3-row vector σ_x. Shocks to the price of risk are assumed to be uncorrelated with shocks to Δd^i_{t+1} and z_{t+1}.

A.3.1.2 Pricing Claims to Single-Period Cash Flows of Different Maturities

The price-dividend (price-rent, wealth-consumption) ratio of claims to single cash flows of maturity n is exponentially affine and follows

$$PD_t^{i,(n)} = \exp\{A^i(n) + B^i_x(n)x_t + B^i_z(n)z_t\},$$

where

$$A^i(n) = A^i(n-1) - r^f + g_i + B^i_x(n-1)(1-\phi_x)\overline{x} + \frac{1}{2}V^i(n-1)V^i(n-1)',$$

$$B^i_x(n) = B^i_x(n-1)(\phi_x - \sigma_x\frac{\sigma_c'}{||\sigma_c||}) - (\sigma_i + B^i_z(n-1)\sigma_z)\frac{\sigma_c'}{||\sigma_c||},$$

$$B^i_z(n) = \kappa_i\frac{1-\phi_z^n}{1-\phi_z},$$

$$V^i(n-1) = \sigma_i + B^i_z(n-1)\sigma_z + B^i_x(n-1)\sigma_x.$$

with $A^i(0) = B^i_x(0) = B^i_z(0) = 0$. Note that $\frac{\sigma_c}{||\sigma_c||}$ is simply $[1,0,0]$ given the baseline calibration. As discussed in Lettau and Wachter (2007), $A^i(n)$ is a constant that determines the level of price-dividend (price-rent, wealth-consumption) ratios. We have $B^i_z(n) > 0$ for all n. Intuitively, the higher z_t, the higher expected cash flow growth and the higher the price of an asset that pays a single cash flow in the future. Moreover, $B^i_z(n)$ increases in n and converges to $\frac{\kappa_i}{1-\phi_z}$ as n grows large, reflecting persistence in cash flow growth and the cumulation of shocks between t and $t+n$ in D^i_{t+n}. Since shocks to the price of risk and shocks to cash flow growth are uncorrelated by assumption, $B^i_x(n) < 0$

for all n and prices decrease (and risk premia increase) as the price of risk increases.

A.3.1.3 Returns on Claims to Single-Period Cash Flows

Denote one-period returns on claims to cash flows at maturity n of asset i by $R^{i,(n)}_{t,t+1}$ and corresponding log-returns by $r^{i,(n)}_{t,t+1}$. The risk premium of expected one-period returns on claims to cash flows at individual maturities then follows (see Lettau and Wachter, 2007)

$$\begin{aligned} E_t\left[R^{i,(n)}_{t,t+1} - R^f\right] &\simeq E_t\left[r^{i,(n)}_{t,t+1} - r^f\right] + \frac{1}{2}Var_t\left[r^{i,(n)}_{t,t+1}\right] \\ &= (\sigma_i + B^i_x(n-1)\sigma_x + B^i_z(n-1)\sigma_z)\frac{\sigma_c^{'}}{||\sigma_c||}x_t. \end{aligned}$$

Risk premia depend on the loadings on the different risk sources as well as the price of each source of risk. All claims load equally, irrespective of their maturity, on shocks to current cash flow growth Δd^i_{t+1} and $||\sigma_d||^{-1}\sigma_i\sigma_d^{'}x_t$ is the associated price of risk. $B^i_z(n-1)$ is the loading on shocks to the growth rate of cash-flows z_t and $||\sigma_d||^{-1}\sigma_z\sigma_d^{'}x_t$ is the associated price of risk. Notice that the loading on shocks to the price of risk is irrelevant when x_{t+1} and Δd^i_{t+1} are assumed to be uncorrelated. As a consequence, the term-structure variation is generated by the loading $B^i_z(n-1)$, which is a strictly increasing function of horizon n, starting at 0 and converging to $\frac{\kappa_i}{1-\phi_z}$ as $n \to \infty$.

A.3.1.4 Per-Period Discount Rates

The per-period discount rate $\bar{r}^{i,n}_t$ is implicitly given by

$$PD^{i,(n)}_t \equiv \frac{P^{i,(n)}_t}{D^i_t} = \frac{E_t[\frac{D^i_{t+n}}{D^i_t}]}{(1+\bar{r}^{i,n}_t)^n}.$$

Using $\bar{r}^{i,n}_t \simeq \ln(1+\bar{r}^{i,n}_t)$ and defining $pd^{i,(n)}_t \equiv \ln(PD^{i,(n)}_t)$, we can solve for the the per-period discount rate as

$$\begin{aligned} \bar{r}^{i,n}_t &= \frac{1}{n}\left(\ln(E_t[\frac{D^i_{t+n}}{D^i_t}]) - pd^{i,(n)}_t\right) \\ &= g_i + \frac{1}{n}\left(\kappa_i\frac{(1-\phi_z^n)}{1-\phi_z}z_t + 0.5V^{i,n}\right) - \frac{1}{n}pd^{i,(n)}_t, \end{aligned}$$

where $V^{i,n}$ captures a Jensen's inequality adjustment that arises from taking the expectation of a variable that is exposed to cumulative log-normal shocks; its functional form is detailed in Section A.3.3.2. For a more intuitive expression, we can take the unconditional expectation:

$$E[\overline{r}_t^{i,n}] = g_i + \frac{1}{2n} V^{i,n} - \frac{1}{n} E[pd_t^{i,(n)}].$$

It is increasing in expected cash flow growth and decreasing in the price-dividend ratio.

In the main body of the paper, which considers the case of a constant price of risk $x_t = \overline{x}$, the simple relationship between the per-period discount rate for maturity n and the expected one-period holding return of single-claim assets with maturities up to n applies:

$$\overline{r}_t^{i,n} = \frac{1}{n} \sum_{k=1}^{n} \ln(E_t[R_{t,t+1}^{i,(k)}]).$$

A.3.1.5 Pricing the Aggregate Market and Market Returns

The price-dividend (price-rent, wealth-consumption) ratio of the market is just the sum of the respective ratios of the claims to single-period cash flows:

$$PD_t^i = \sum_{n=1}^{\infty} PD_t^{i,(n)}.$$

The expected return of each market can be found as the weighted average of the expected return of the claims to single-period cash flows:

$$E_t[R_{t,t+1}^i] = \sum_{n=1}^{\infty} \left[\frac{PD_t^{i,(n)}}{PD_t^i} E_t[R_{t,t+1}^{i,(n)}] \right] = E_t \left[\frac{P_{t+1}^i + D_{t+1}^i}{P_t^i} \right].$$

A.3.1.6 Constant Price of Risk and a Dividend-Based Stochastic Discount Factor

In Section 4.3 of the main part of the paper, we discuss a simplified version of the above setup to keep the model parsimonious. In particular, we fix the price of risk at $x_t = \overline{x}$, since for the purpose of the analysis in this paper, time variation in risk prices is not of first-order importance. Here, we also calibrate the full model with time-varying risk prices, and verify that our conclusions are robust. In the main part of the paper, we also express the stochastic discount factor as a function of dividend growth instead of consumption growth for simplicity. Consequently, all expressions from the main text can be recovered by setting $x_t = \overline{x}$, $\sigma_x = 0$, and by replacing $\sigma_c / ||\sigma_c||$ with $\sigma_d / ||\sigma_d||$.

A.3.2 Calibration and Results

A.3.2.1 Constant Price of Risk and a Dividend-Based SDF

Lettau and Wachter (2007) and Lettau and Wachter (2011) show that this type of model can be calibrated to successfully match a large number of asset pricing moments, including the term structure of bonds and equity.[11] To discipline our analysis, we maintained the calibration of Lettau and Wachter (2007) whenever possible. We made one simplification by keeping the price of risk x constant, since for the purpose of the analysis in this paper, time variation in risk prices is not of first-order importance. Returns, prices, and dividends for equity are based on corresponding moments of the S&P 500 index. Average dividend growth is set to 2.28%. $\|\sigma_d\|$ is set to match the unconditional standard deviation of annual dividend growth at 14.5%. The predictable component of dividend growth is captured by z_t; $\|\sigma_z\|$ is set to 0.32% to reflect low predictability in the sample, and κ_s is set to 1. Importantly for the term structure results, the correlation between shocks to z_t and shocks to Δd_t^s is calibrated to -0.83.[12] This implies that the underlying fundamental cash flow process is mean-reverting: drops in cash flows are partially reversed through higher growth in the future. To calibrate the additional process for rents we set $g_h = 0.7\%$ to reflect the average rent growth estimated in Section 3. We choose a volatility $\|\sigma_h\|$ of 14.2% as well as a κ_h of 1.14 to target the expected return on real estate in the data (6.4%) and the empirical term structure of real estate discount rates, in particular the estimate of the very long-run discount rate of 2.6%. As in Lettau and Wachter (2007), all remaining correlations are assumed to be zero for simplicity. To summarize, we set quarterly $\sigma_s = [0.0724, 0]$, $\sigma_h = [0.071, 0]$, and $\sigma_z = [-0.00136, 0.0009]$.[13] We set the risk-free rate at 1%, based on the U.K. gilts real yield curve between 1998 and 2013 reported in Figure VI. These data show that the U.K. real yield curve is approximately flat on average, with a real yield of 1.4% for maturities between 1 and 25 years, and that there is some evidence for a mild downward slope at longer maturities.[14] Our calibration is

[11]See also the exciting work of van Binsbergen, Brandt and Koijen (2012) and van Binsbergen et al. (2013) for the model implications for claims to individual stock market dividends at horizons of 1 to 10 years.

[12]See Lettau and Ludvigson (2005) for further discussion on how the consumption-dividend ratio can be used as an empirical proxy for the calibration of z_t.

[13]Note that we have restricted ourselves to two shocks to keep the number of state variables low and the model parsimonious. As a consequence, $\|\sigma_h\|$ should capture the systematic risk of rents only.

[14]The real yield curve is computed by the Bank of England and is available at `http://www.bankofengland.co.uk/statistics/Pages/yieldcurve/archive.aspx`. Figure VI plots the average shape of the real U.K. gilts curve for the period 1998-2013, as well as for two sub-periods: 1998-2007 and 2007-2013. The level of the yield curve shifted down during this latter period and the yield curve became hump-shaped. The U.K. government debt also includes some perpetual bonds: the War Loan and the Annuities. These bonds comprise a negligible part of the outstanding U.K. government debt (£2.6bn out of £1.5trn of debt outstanding), and are classified as small and illiquid issuances by the U.K. Debt and

summarized in Table IV. Results are displayed in Table V and discussed in the main part of the paper.

A.3.2.2 Time-Varying Price of Risk and a Consumption-Based SDF

We also calibrate a general version of the model with a time-varying price of risk and a consumption-based stochastic discount factor and verify that our conclusions are robust. For all parameters that appear in both versions of the model, we keep the calibration described in Section A.3.2.1. The only exception is the volatility of rents; since we calibrate the process for rents numerically to target the expected return on real estate and the empirical term structure of real estate discount rates estimated in Section 3, we have to adjust $\|\sigma_h\|$ to 14.5%; we keep κ_h at 1.14. The process for x_t is calibrated to fit various moments of returns to equities, in particular the volatility of stock prices; we follow Lettau and Wachter (2007) and set ϕ_x to 0.87 and $\|\sigma_x\|$ to 24%. For the consumption process, we follow Bansal and Yaron (2004) and equalize consumption and dividend growth rates (at 2.28%), and set volatility ($\|\sigma_c\|$) to 4.83% and κ_c to 0.33, one third of the respective values for the dividend growth process. Table A.3 summaries our calibration of the extended model.

Table A.4 shows that our extended model matches the same key moments as our simplified model in the main part of the paper. The average annual return on the stock market is 7.3%. Due to time-variation in the price of risk, the volatility of returns is 17.0%, in line with the data. Expected dividend growth is above 3% with a standard deviation of 14.5% as before. Importantly, the term structure of discount rates for equity is downward-sloping as shown in Panel A of Figure A.5. It falls from 18.5% for claims to one year ahead dividends, to 5.0% for claims to one hundred year ahead dividends, all the way to 4.1% for one thousand year ahead dividends. The expected return on real estate is 5.7%, just below our empirical net return estimates between 6.4% and 8.0% again. Expected rent growth is 1.8%. Our calibration implies that returns as well as rent growth are about as volatile as equity, with 13.3% and 14.5% respectively. In line with our empirical findings, rental yields contribute more than two thirds to total returns. The term structure of discount rates for real estate is again downward-sloping, as shown in Panel B of Figure A.5. It falls from 18.5% for claims to one year ahead rents, to 3.7% for claims to one hundred year ahead rents, all the way to 2.5% for one thousand year ahead rents. The expected return on wealth is 2.4% and its volatility is 4.4%, about one third of the

Management Office. They are excluded from our analysis, not only because they are nominal and we only use data on U.K. real gilts, but also because their negligible size, scarce liquidity, and callability make it hard to interpret their prices in terms of discount rates.

respective values for equity. Expected consumption growth is 2.4%. Our calibration also implies that consumption growth is about one third as volatile as dividends are (4.8%). The term structure of discount rates for wealth is downward-sloping as well, as shown in Panel C of Figure A.5. It falls from 6.9% for claims to one year ahead rents, to 2.6% for claims to one hundred year ahead rents, all the way to 2.2% for one thousand year ahead rents.[15]

A.3.3 Solving the Model

A.3.3.1 Derivation of Prices of Claims to Single-Period Cash Flows

As in Lettau and Wachter (2007), conjecture

$$PD_t^{i,n} = \exp\{A^i(n) + B_x^i(n)x_t + B_z^i(n)z_t\} \equiv F^i(x_t, z_t, n).$$

From the fundamental asset pricing equation, we know

$$E_t\left[M_{t,t+1}R_{t,t+1}^{i,(n)}\right] = 1,$$

and

$$P_t^{i,(n)} = E_t\left[M_{t,t+1}P_{t+1}^{i,(n-1)}\right],$$

with boundary condition

$$P_t^{i,(0)} = D_t^i.$$

By the boundary condition,

$$A^i(0) = B_x^i(0) = B_z^i(0) = 0.$$

Substituting F^i, we get

$$E_t\left[M_{t,t+1}\frac{D_{t+1}^i}{D_t^i}F^i(x_{t+1}, z_{t+1}, n-1)\right] = F^i(x_t, z_t, n).$$

[15] Note that our calibration implies a very high wealth-consumption ratio; this is a moment that is not of any direct relevance for our analysis. Since wealth is not very risky, its return is close to the risk-free rate. Moreover, our high quality data on the risk-free rate guides us to choose a lower risk-free rate than in the original calibration of Lettau and Wachter (2007). At the same time, we choose a relatively high growth rate for consumption. In sum, this pushes $r - g$ close to zero for wealth, which creates a very high wealth-consumption ratio.

Substituting for F^i and $M_{t,t+1}$, we get

$$E_t[\exp\{-r_f - \frac{1}{2}x_t^2 - x_t\epsilon_{c,t+1} + \Delta d^i_{t+1} + A^i(n-1) + B^i_x(n-1)x_{t+1} + B^i_z(n-1)z_{t+1}\}]$$

$$= \exp\{A^i(n) + B^i_x(n)x_t + B^i_z(n)z_t\}.$$

Substituting functional forms for $t+1$-variables:

$$E_t[\exp\{-r_f - \frac{1}{2}x_t^2 - x_t\epsilon_{c,t+1} + [g_i + \kappa_i z_t + \sigma_i\epsilon_{t+1}] + A^i(n-1) + B^i_x(n-1)\left[(1-\phi_x)\overline{x} + \phi_x x_t + \sigma_x\epsilon_{t+1}\right]$$

$$+B^i_z(n-1)\left[\phi_z z_t + \sigma_z\epsilon_{t+1}\right]\}]$$

$$= \exp\{A^i(n) + B^i_x(n)x_t + B^i_z(n)z_t\}.$$

Solving the expectation:

$$\exp\{-r_f - \frac{1}{2}x_t^2 + [g_i + \kappa_i z_t] + A^i(n-1) + B^i_x(n-1)\left[(1-\phi_x)\overline{x} + \phi_x x_t\right]$$

$$+B^i_z(n-1)\left[\phi_z z_t\right] + \frac{1}{2}[x_t^2 - 2x_t\left(\sigma_i + B^i_x(n-1)\sigma_x + B^i_z(n-1)\sigma_z\right)\frac{\sigma_c^{'}}{||\sigma_c||} + V^i(n-1)V^i(n-1)']\}$$

$$= \exp\{A^i(n) + B^i_x(n)x_t + B^i_z(n)z_t\},$$

where:

$$V^i(n-1) = \sigma_i + B^i_x(n-1)\sigma_x + B^i_z(n-1)\sigma_z.$$

Collecting terms:

$$\exp\{[-r_f + g_i + A^i(n-1) + B^i_x(n-1)(1-\phi_x)\overline{x} + \frac{1}{2}V^i(n-1)V^i(n-1)']$$

$$+[B^i_x(n-1)\phi_x - \left(\sigma_i + B^i_x(n-1)\sigma_x + B^i_z(n-1)\sigma_z\right)\frac{\sigma_c^{'}}{||\sigma_c||}]x_t + [\kappa_i + B^i_z(n-1)\phi_z]z_t\}$$

$$= \exp\{A^i(n) + B^i_x(n)x_t + B^i_z(n)z_t\}.$$

Matching terms:

$$A^i(n) = A^i(n-1) - r^f + g_i + B^i_x(n-1)(1-\phi_x)\overline{x} + \frac{1}{2}V^i(n-1)V^i(n-1)',$$

$$B^i_x(n) = B^i_x(n-1)(\phi_x - \sigma_x\frac{\sigma_c'}{||\sigma_c||}) - (\sigma_i + B^i_z(n-1)\sigma_z)\frac{\sigma_c^{'}}{||\sigma_c||},$$

$$B_z^i(n) = \kappa_i \frac{1-\phi_z^n}{1-\phi_z}.$$

A.3.3.2 Derivation of Per-Period Discount Rates

To compute the per-period discount, we need to compute the expectation of cash flow growth, since $\bar{r}_t^{i,n}$ is that discount rate such that:

$$PD_t^{i,(n)} = \frac{E_t[\frac{D_{t+n}^i}{D_t^i}]}{(1+\bar{r}_t^{i,n})^n}$$

Now,

$$E_t[\frac{D_{t+n}^i}{D_t^i}] = E_t[\exp\{\sum_{k=1}^{n} \Delta d_{t+k}^i\}]$$

which is equal to

$$E_t[\exp\{(g_i + \kappa_i z_t + \sigma_i \epsilon_{t+1})$$

$$+(g_i + \kappa_i z_{t+1} + \sigma_i \epsilon_{t+2})$$

$$+(g_i + \kappa_i z_{t+2} + \sigma_i \epsilon_{t+3})$$

$$\vdots$$

$$+(g_i + \kappa_i z_{t+n-1} + \sigma_i \epsilon_{t+n})\}]$$

$$=$$

$$E_t[\exp\{(g_i + \kappa_i z_t + \sigma_i \epsilon_{t+1})$$

$$+(g_i + \kappa_i(\phi_z z_t + \sigma_z \epsilon_{t+1}) + \sigma_i \epsilon_{t+2})$$

$$+(g_i + \kappa_i(\phi_z z_{t+1} + \sigma_z \epsilon_{t+2}) + \sigma_i \epsilon_{t+3})$$

$$\vdots$$

$$=$$

$$E_t[\exp\{(g_i + \kappa_i z_t + \sigma_i \epsilon_{t+1})$$

$$+(g_i + \kappa_i(\phi_z z_t + \sigma_z \epsilon_{t+1}) + \sigma_i \epsilon_{t+2})$$

$$+(g_i + \kappa_i(\phi_z(\phi_z z_t + \sigma_z \epsilon_{t+1}) + \sigma_z \epsilon_{t+2}) + \sigma_i \epsilon_{t+3})$$

$$\vdots$$

Collecting terms, this is equal to

$$\exp\{g_i n\} E_t[\exp\{(\sigma_i + \kappa_i\sigma_z + \kappa_i\phi_z\sigma_z + ... + \kappa_i\phi_z^{n-2}\sigma_z)\epsilon_{t+1}$$

$$+(\sigma_i + \kappa_i\sigma_z + \kappa_i\phi_z\sigma_z + ... + \kappa_i\phi_z^{n-3}\sigma_z)\epsilon_{t+2}$$

$$\vdots$$

$$+(\sigma_i + \kappa_i\sigma_z)\epsilon_{t+n-1}$$

$$+(\sigma_i)\epsilon_{t+n}$$

$$+\kappa_i(1 + \phi_z + \phi_z^2 + ... + \phi_z^{n-1})z_t\}]$$

Since these are geometric sums, we can rewrite

$$(\sigma_i + \kappa_i\sigma_z + \kappa_i\phi_z\sigma_z + ... + \kappa_i\phi_z^{n-2}\sigma_z) = (\sigma_i + \kappa_i\sigma_z(1 + \phi_z + ... + \phi_z^{n-2})) = (\sigma_i + \kappa_i\sigma_z\frac{(1-\phi_z^{n-1})}{1-\phi_z}),$$

$$(\sigma_i + \kappa_i\sigma_z + \kappa_i\phi_z\sigma_z + ... + \kappa_i\phi_z^{n-3}\sigma_z) = (\sigma_i + \kappa_i\sigma_z\frac{(1-\phi_z^{n-2})}{1-\phi_z}),$$

$$\vdots$$

to get

$$= \exp\{g_i n\} E_t[\exp\{(\sigma_i + \kappa_i\sigma_z\frac{(1-\phi_z^{n-1})}{1-\phi_z})\epsilon_{t+1}$$

$$+(\sigma_i + \kappa_i\sigma_z\frac{(1-\phi_z^{n-2})}{1-\phi_z})\epsilon_{t+2}$$

$$\vdots$$

$$+(\sigma_i + \kappa_i\sigma_z)\epsilon_{t+n-1}$$

$$+\sigma_i\epsilon_{t+n}$$

$$+\kappa_i\frac{(1-\phi_z^n)}{1-\phi_z}z_t\}].$$

We can then compute the variance term as

$$V_t^{i,n} = Var_t[(\sigma_i + \kappa_i\sigma_z\frac{(1-\phi_z^{n-1})}{1-\phi_z})\epsilon_{t+1} + (\sigma_i + \kappa_i\sigma_z\frac{(1-\phi_z^{n-2})}{1-\phi_z})\epsilon_{t+2}$$

$$\ldots + (\sigma_i + \kappa_i\sigma_z)\epsilon_{t+n-1} + (\sigma_i)\epsilon_{t+n}].$$

Since all ϵ are independent, the variance reduces to

$$V_t^{i,n} = V^{i,n} = \sum_{k=1}^{n} \tilde{V}^{i,k},$$

where:

$$\tilde{V}^{i,1} = \sigma_i \sigma_i',$$

$$\tilde{V}^{i,2} = (\sigma_i + \kappa_i \sigma_z)(\sigma_i + \kappa_i \sigma_z)',$$

$$\tilde{V}^{i,3} = (\sigma_i + \kappa_i \sigma_z \frac{(1-\phi_z^2)}{1-\phi_z})(\sigma_i + \kappa_i \sigma_z \frac{(1-\phi_z^2)}{1-\phi_z})',$$

$$\vdots$$

$$\tilde{V}^{i,n} = (\sigma_i + \kappa_i \sigma_z \frac{(1-\phi_z^{n-1})}{1-\phi_z})(\sigma_i + \kappa_i \sigma_z \frac{(1-\phi_z^{n-1})}{1-\phi_z})'.$$

Therefore,

$$E_t[\frac{D_{t+n}^i}{D_t^i}] = E_t[\exp\{\sum_{k=1}^{n} \Delta d_{t+k}^i\}]$$

$$= \exp\{g_i n + \kappa_i \frac{(1-\phi_z^n)}{1-\phi_z} z_t + 0.5 V^{i,n}\},$$

and thus:

$$PD_t^{i,(n)} = \frac{\exp\{g_i n + \kappa_i \frac{(1-\phi_z^n)}{1-\phi_z} z_t + 0.5 V^{i,n}\}}{(1+r_n^i)^n}.$$

Using $\bar{r}^{i,n} \simeq \ln(1+\bar{r}^{i,n})$ and defining $pd_t^{i,(n)} = \ln(PD_t^{i,(n)})$, we get

$$\bar{r}_t^{i,n} = g_i + \frac{1}{n}\left(\kappa_i \frac{(1-\phi_z^n)}{1-\phi_z} z_t + 0.5 V^{i,n}\right) - \frac{1}{n} pd_t^{i,(n)}.$$

Appendix References

Ambrose, Brent W., Piet Eichholtz, and Thies Lindenthal. 2013. "House Prices and Fundamentals: 355 Years of Evidence." *Journal of Money, Credit and Banking*, 45(2-3): 477–491.

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.

Barro, Robert J., and Jose F. Ursua. 2008. "Macroeconomic Crises since 1870." *Brookings Papers on Economic Activity*, 39(1): 255–350.

Bordo, Michael, Barry Eichengreen, Daniela Klingebiel, and Maria Soledad Martinez-Peria. 2001. "Is the Crisis Problem Growing More Severe?" *Economic Policy*, 16(32): 51–82.

Constantinescu, Mihnea, and Marc Francke. 2013. "The Historical Development of the Swiss Rental Market - A New Price Index." *Journal of Housing Economics*, 22(2): 135 – 145.

Eichholtz, Piet. 1997. "A Long Run House Price Index: The Herengracht Index, 1628–1973." *Real Estate Economics*, 25(2): 175–192.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn van Nieuwerburgh. 2010. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium." National Bureau of Economic Research Working Paper 15988.

Lettau, Martin, and Jessica A. Wachter. 2007. "Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium." *Journal of Finance*, 62(1): 55–92.

Lettau, Martin, and Jessica A. Wachter. 2011. "The Term Structures of Equity and Interest Rates." *Journal of Financial Economics*, 101(1): 90–113.

Lettau, Martin, and Sydney C. Ludvigson. 2005. "Expected Returns and Expected Dividend Growth." *Journal of Financial Economics*, 76(1): 583–626.

Mack, Adrienne, and Enrique Martínez-García. 2011. "A Cross-Country Quarterly Database of Real House Prices: A Methodological Note." Federal Reserve Bank of Dallas Globalization and Monetary Policy Working Paper 99.

Reinhart, Carmen M., and Kenneth S. Rogoff. 2009. *This Time is Different: Eight Centuries of Financial Folly.* Princeton University Press.

Schularick, Moritz, and Alan M. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises." *American Economic Review*, 102(2): 1029–1061.

Shiller, Robert. 2000. *Irrational Exuberance.* Princeton University Press, Princeton.

Stapledon, Nigel. 2012. "Trends and Cycles in Sydney and Melbourne House Prices from 1880 to 2011." *Australian Economic History Review*, 52(3): 293–317.

van Binsbergen, Jules, Michael Brandt, and Ralph Koijen. 2012. "On the Timing and Pricing of Dividends." *American Economic Review*, 102(4): 1596–1618.

van Binsbergen, Jules, Wouter Hueskes, Ralph Koijen, and Evert Vrugt. 2013. "Equity Yields." *Journal of Financial Economics*, 110(3): 503 – 519.

Appendix Figures

Figure A.1: Cross-Sectional Distribution of the Rent-to-Price Ratio in the U.S.

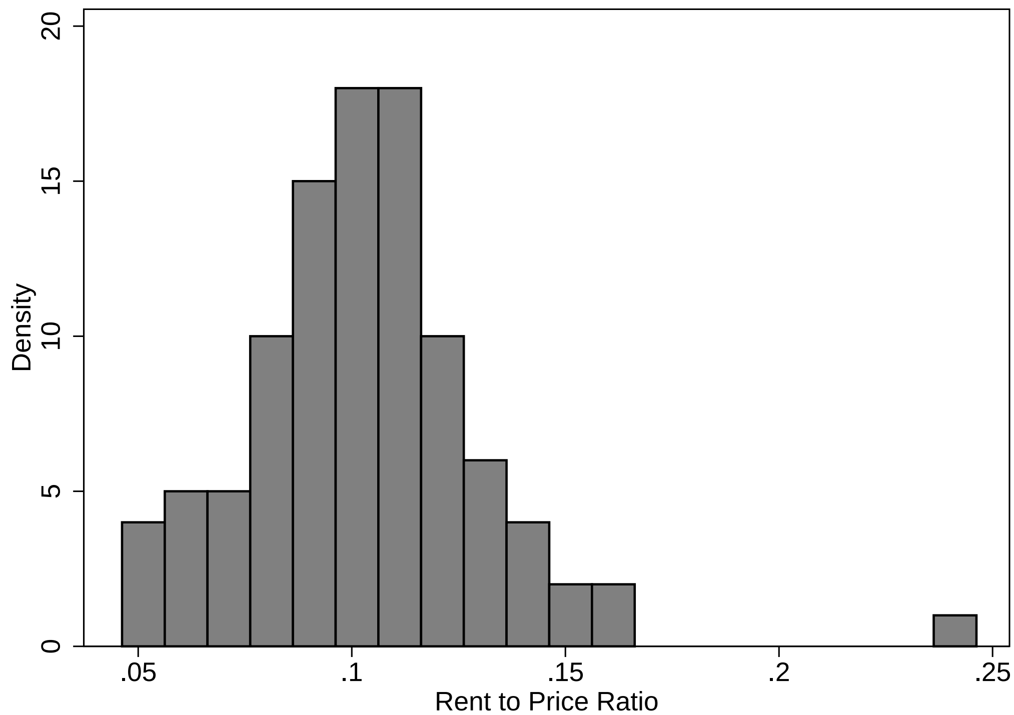

Note: The figure shows the distribution of the rent-to-price ratio for the 100 largest MSAs in the U.S. in 2012 as constructed by Trulia, which observes a large set of both for-sale and for-rent listings. It is constructed using a metro-level hedonic regression of $\ln(price)$ on property attributes, zip-code fixed effects, and a dummy for whether the unit is for rent. The rent-to-price ratio is constructed by taking the exponent of the coefficient on this dummy variable.

Figure A.2: Price-to-Rent Ratio Time-Series in the U.S.

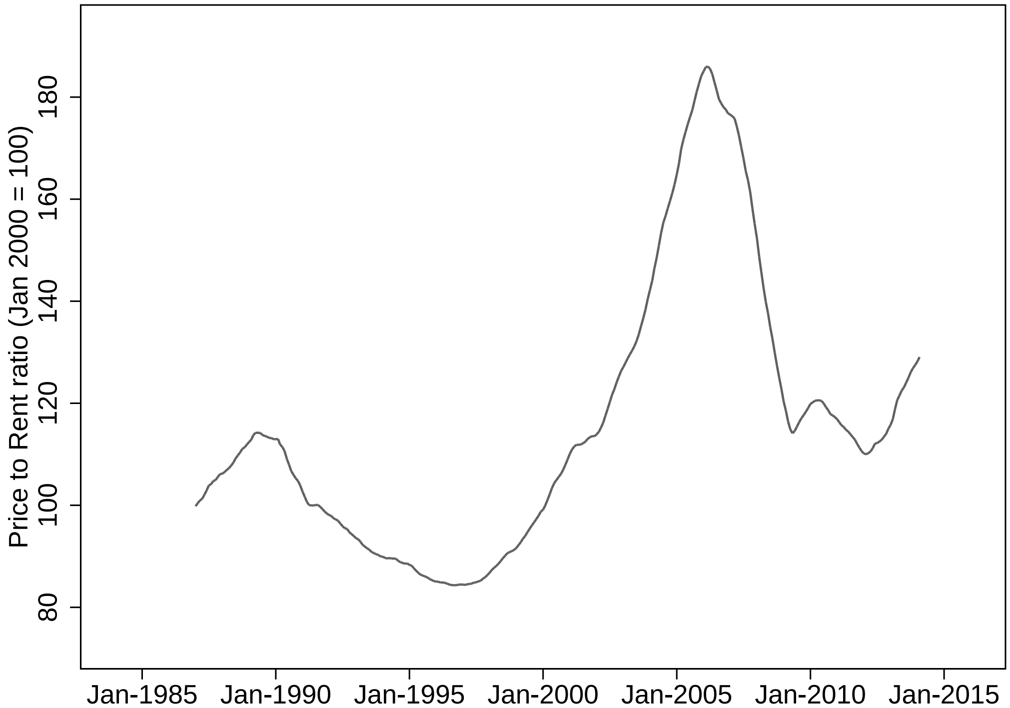

Note: The figure shows the time series of the price-rent ratio in the U.S., constructed as the ratio of the 10-City Case-Shiller House Price Index (Fred Series: SPCS10RSA) and the Bureau of Labor Statistic's Consumer Price Index for All Urban Consumers: Owners' Equivalent Rent of Residences (Fred Series: CUSR0000SEHC). The index ratio is normalized to 100 in January 2000.

Figure A.3: Rent Growth vs. PCE Growth in the U.S.

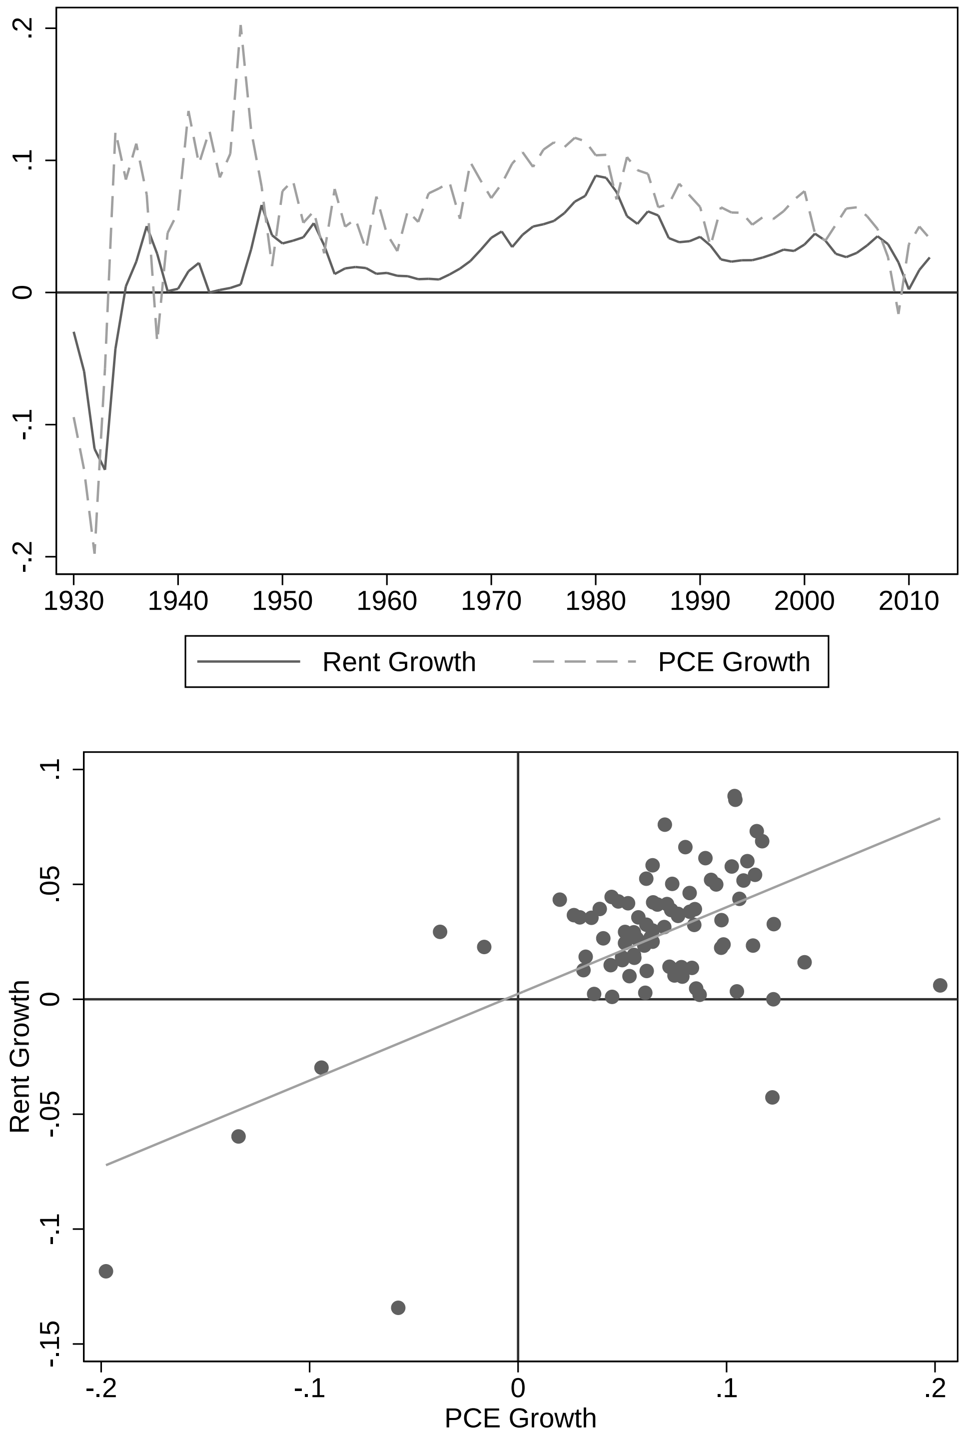

Note: The figure shows the annual growth rates of the "Consumer Price Index for All Urban Consumers: Rent of Primary Residence" (FRED ID: CUUR0000SEHA) and "Personal Consumption Expenditure" (FRED ID: PCDGA) since 1929.

Figure A.4: House Prices and Rents in Prime Central London Areas during the 2007-09 Financial Crisis

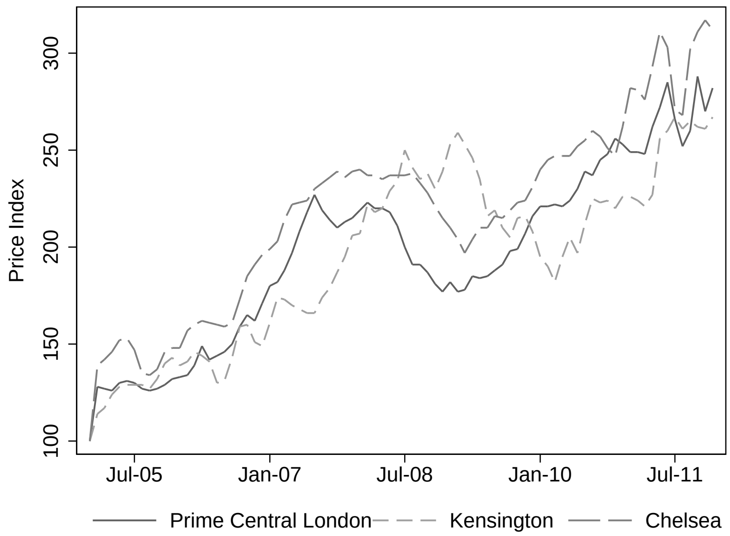

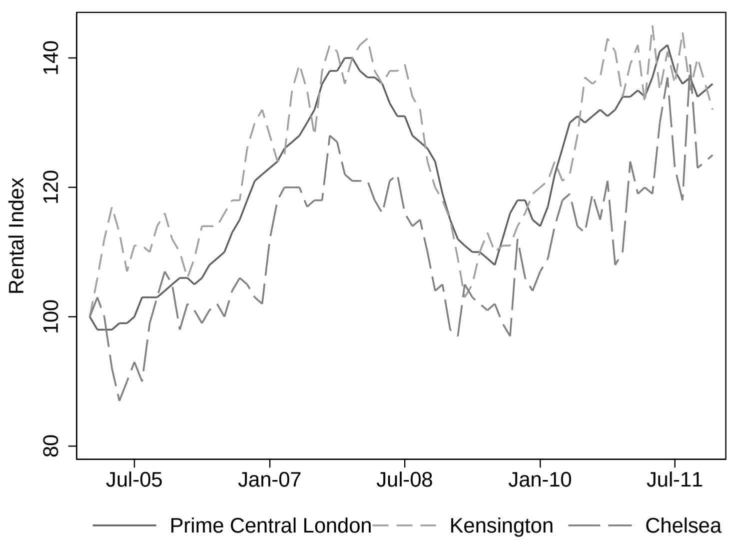

Note: The figure shows the time series of house prices and rents for Prime Central London, Kensington, and Chelsea for the period January 2005 to January 2012. The series are monthly and available from John D Wood & Co. at http://www.johndwood.co.uk/content/indices/london-property-prices/

Figure A.5: Term Structure of Discount Rates, Full Model

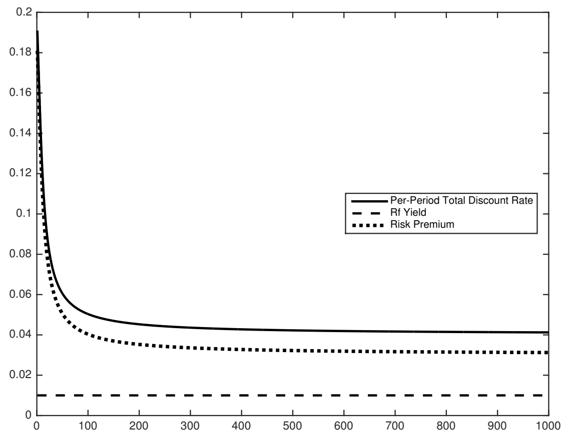

(A) Equity

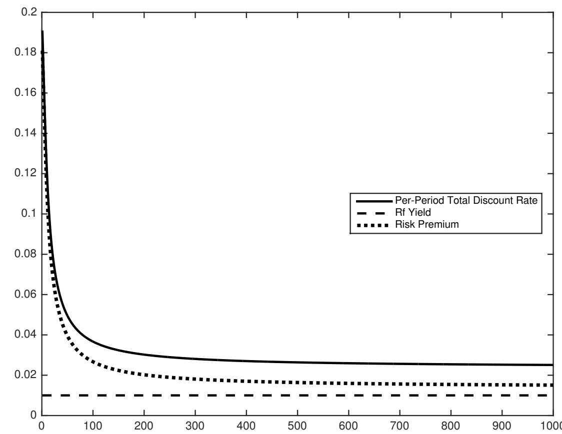

(B) Real Estate

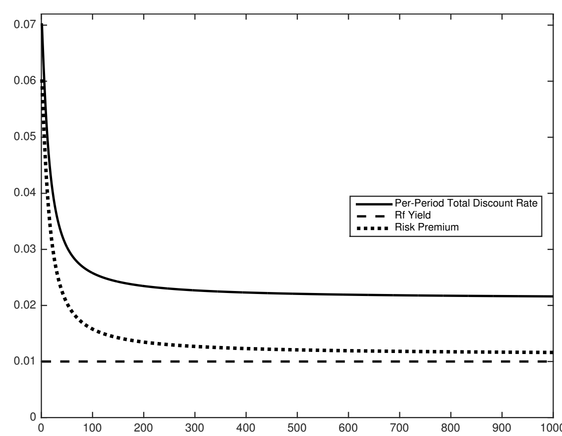

(C) Consumption

Note: Panel A shows per-period discount rates, risk-free rates and risk premia of claims to single dividends with a maturity up to 1000 years implied by the reduced-form asset pricing model from Section A.3. Panel B shows per-period discount rates, risk-free rates and risk premia of claims to single rents with a maturity up to 1000 years implied by the reduced-form asset pricing model from Section A.3. Panel C shows per-period discount rates, risk-free rates and risk premia of claims to single consumption strips with a maturity up to 1000 years implied by the reduced-form asset pricing model from Section A.3. The calibration details are discussed in Appendix A.3.2.2 and summarized in Table A.3.

Appendix Tables

Table A.1: Rent-to-Price Ratio Singapore - 2012

	(1)	(2)	(3)	(4)
For-Rent Dummy	-3.095***	-3.131***	-3.123***	-3.107***
	(0.044)	(0.019)	(0.014)	(0.025)
Fixed Effects	Quarter × Postal Code	Quarter × Postal Code	Month × Postal Code	Month × Postal Code × Bedrooms
Controls	·	✓	✓	✓
Implied Rent-to-Price Ratio	4.5%	4.4%	4.4%	4.5%
R-squared	0.804	0.873	0.872	0.872
N	106,145	105,189	105,189	105,189

Note: This table shows results from regression (A.1). The rent-to-price ratio is constructed by taking the exponent of the coefficient on this dummy variable. The dependent variable is the price (for-sale price or annualized for-rent price) for properties listed on iProperty.com in Singapore in 2012. Fixed effects are included as indicated. In columns (2) to (4), we also control for characteristics of the property: we include dummy variables for the type of the property (condo, house, etc.), indicators for the number of bedrooms and bathrooms, property age, property size (by adding dummy variables for 50 equal-sized buckets), information on the kitchen (ceramic, granite, etc.), which floor the property is on, and the tenure type for leaseholds. Standard errors are clustered at the level of the fixed effect. Significance levels are as follows: * ($p<0.10$), ** ($p<0.05$), *** ($p<0.01$).

Table A.2: House Prices, Banking Crises, Rare Disasters

	House Price Index Time Period	Banking Crises	Rare Disasters
Australia	1880 - 2013	1893, 1989	1918, 1932, 1944
Belgium	1975 - 2012	2008	
Canada	1975 - 2012		
Denmark	1975 - 2012	1987	
Finland	1975 - 2012	1991	1993
France	1840 - 2010	1882, 1889, 1907, 1930, 2008	1871, 1915, 1943
Germany	1975 - 2012	2008	
Italy	1975 - 2012	1990, 2008	
Japan	1975 - 2012	1992	
Netherlands	1649 - 2009	1893, 1907, 1921, 1939, 2008	1893, 1918, 1944
New Zealand	1975 - 2012	1987	
Norway	1819 - 2013	1899, 1922, 1931, 1988	1918, 1921, 1944
Singapore	1975 - 2012	1982	
South Africa	1975 - 2012	1977, 1989	
South Korea	1975 - 2012	1985, 1997	1998
Spain	1975 - 2012	1978, 2008	
Sweden	1952 - 2013	1991, 2008	
Switzerland	1937 - 2012	2008	1945
U.K.	1952 - 2013	1974, 1984, 1991, 2007	
U.S.	1890 - 2012	1893, 1907, 1929, 1984, 2007	1921, 1933

Note: The table shows the time series availability of house price indices in the first column. The second and third column report dates of banking crises or rare consumption disasters for each country in the time period provided in the first column. Banking crisis dates for all countries, except Singapore, Belgium, Finland, New Zealand, South Korea, and South Africa, are from Schularick and Taylor (2012). Banking crisis dates for the countries not covered by Schularick and Taylor (2012) are from Reinhart and Rogoff (2009). Rare disaster dates indicate the year of the trough in consumption during a consumption disaster as reported by Barro and Ursua (2008).

Table A.3: Parameters of the Full Model

Calibrated Variables	Value
g_s	2.28%
g_h	0.7%
g_c	2.28%
κ_s	1
κ_h	1.14
κ_c	0.33
r^f	1%
$\overline{x}$	0.625
ϕ_z	0.91
ϕ_x	0.87
$\lVert\sigma_s\rVert$	14.5%
$\lVert\sigma_h\rVert$	14.5%
$\lVert\sigma_c\rVert$	4.83%
$\lVert\sigma_z\rVert$	0.32%
$\lVert\sigma_x\rVert$	24.0%
Correlation of Δd^s and z shocks	-0.83
Correlation of x with all other shocks	0

Implied Parameters	Value
σ_s	[0.0724,0,0]
σ_h	[0.0724,0,0]
σ_c	[0.0241,0,0]
σ_z	[-0.00136,0.0009,0]
σ_x	[0,0,0.12]

Note: The table shows the calibration of the reduced-form asset pricing model in Section A.3. The unconditional mean of dividend, rent, and consumption growth rates, the risk-free rate, the persistence variables, and the conditional standard deviations are in annual terms. The stock market is calibrated as in Lettau and Wachter (2007). The unconditional mean of the dividend growth rate and the conditional standard deviation of dividend growth are set to match their respective data counterparts. Under the assumption that the conditional mean of dividend growth can be identified with the log consumption-dividend ratio in the data, autocorrelation ϕ_z and the correlation between shocks to Δd^s and shocks to z are set to match their respective data counterparts. $\lVert\sigma_z\rVert$ is set to reflect the low predictability of dividend growth in the data. The unconditional mean of the rent growth rate is set to match the data. Volatility $\lVert\sigma_h\rVert$ as well as κ_h are calibrated to match the expected return on housing in the data as well as the empirical term structure of housing discount rates. The process of consumption is calibrated to the unconditional mean of dividend growth and to one third of the volatility of dividend growth. x is calibrated to fit various moments of returns to equities as in Lettau and Wachter (2007), in particular the volatility of stock prices. The risk-free rate is set at 1%, based on the U.K. gilts real yield curve between 1998 and 2013 reported in Figure VI.

Table A.4: Simulated Moments for the Full Model

Moment	Equity	Housing	Consumption
$E[P/D]$	32.0	27.5	17,700
$\sigma(p-d)$	38.5%	36.5%	12.9%
AC of $p-d$	0.90	0.91	0.92
$E[R]$	7.3%	5.7%	2.4%
$\sigma(R)$	17.0%	13.3%	4.4%
AC of R	-0.05	-0.04	-0.04
Sharpe ratio	0.37	0.35	0.32
Expected cash flow growth	3.4%	1.8%	2.4%
$\sigma(\Delta d)$	14.5%	14.5%	4.8%
AC of Δd	-0.04	-0.03	-0.03
Expected capital gains	3.6%	1.5%	2.4%
Expected yields	3.7%	4.1%	0%

Note: The table shows key moments from simulating the reduced-form asset pricing model in Section A.3 for 50,000 quarters. Time series are aggregated to an annual frequency. R is the return on the market. Small letters indicate log variables. Possible differences between expected returns and the sum of expected capital gains and expected yields are due to rounding errors.