The Review of Financial Studies

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

Stefano Giglio
Yale University, NBER, and CEPR

Matteo Maggiori
Stanford University, NBER, and CEPR

Krishna Rao
Zillow

Johannes Stroebel
Stern School of Business, New York University, NBER, and CEPR

Andreas Weber
Stern School of Business, New York University

We show that housing markets provide information about the appropriate discount rates for valuing investments in climate change abatement. Real estate is exposed to both consumption and climate risk and its term structure of discount rates is downward sloping, reaching 2.6% for payoffs beyond 100 years. We use a tractable asset pricing model that incorporates features of climate change to show that the term structure of discount rates for climate-hedging investments is thus upward sloping but bounded above by the risk-free rate. At horizons at which risk-free rates are unavailable, the estimated housing discount rates provide an upper bound. (*JEL* G11, G12, R30)

Received June 9, 2020; editorial decision January 26, 2021 by Editor Stijn Van Nieuwerburgh. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

We thank several referees, Arthur van Benthem, Robert Barro, John Campbell, Thom Covert, Simon Dietz, Robert Engle, Xavier Gabaix, María Pilar Martínez-García, Kenneth Gillingham, Christian Gollier, Michael Greenstone, Ben Groom, Bob Hall, Lars Hansen, Geoffrey Heal, Larry Karp, Derek Lemoine, Martin Lettau, Antony Millner, Stijn van Nieuwerburgh, William Nordhaus, Monika Piazzesi, Nick Stern, Thomas Sterner, Christian Traeger, Martin Weitzman, and numerous conference and seminar participants. We thankfully acknowledge generous research support from the Harvard Weatherhead Center for International Affairs, the Center for the Global Economy and Business, the Fama-Miller Center, the Initiative on Global Markets, the Norwegian Finance Initiative, NBIM, and the Global Risk Institute. We thank Zillow, iProperty, and Rightmove for sharing part of their data, and Konhee Chang, Hongbum Lee, Sung Lee, Gil Nogueira, and Andreas Schaab for excellent research assistance. Supplementary data can be found on *The Review of Financial Studies* web site. Send correspondence to Johannes Stroebel, johannes.stroebel@nyu.edu.

The Review of Financial Studies 34 (2021) 3527–3571

doi:10.1093/rfs/hhab032 Advance Access publication March 25, 2021

> 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.
> — William Nordhaus[1]

Much of the economics literature on the optimal policy response to climate change focuses on the trade-off between the immediate costs and the potentially uncertain long-run benefits of reducing 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 that pay off over long horizons. For 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 a slightly higher discount rate, such as 5%, the present value of the investment drops to $760 million and the investment is no longer attractive. However, despite the importance of these discount rates for optimal policy design, economists and policy makers do not agree on what discount rates should be used to value investments in climate change mitigation.

In this paper, we make progress on this question by exploring the information that private market discount rates contain about how to appropriately value investments in climate change abatement. First, we provide new 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 facts with insights from asset pricing theory to discipline the debate on the appropriate choice of discount rates for an investment in climate change abatement, which involves similar horizons as the housing asset but has a different risk profile.

Much of the prior debate on the appropriate discount rates for climate change investments has either relied on theoretical arguments or tried to infer discount rates from the realized returns of traded assets, such as private capital, equity, bonds, and real estate. For example, in the context of the dynamic integrated climate-economy (DICE) model, Nordhaus (2013) chooses a discount rate of 4% to reflect his preferred estimate of the average rate of return to capital.[2] 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 often ignores important considerations regarding the *maturity* and *risk properties* of such investments.

---

[1] Quotation comes from Nordhaus (2013).

[2] See also Kaplow, Moyer, and Weisbach (2010), Schneider, Traeger, and Winkler (2012), and Weisbach and Sunstein (2009) for discussions of normative and descriptive approaches to discounting.

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, different assets pay off their cash flows at different maturities. Because risk in the economy is different for different horizons and preferences for risks can vary with the horizon as well, horizon-specific discount rates must be used when evaluating investments with different maturity profiles. The average rate of return to a particular asset, for example, capital, only reflects the discount rate appropriate for that particular stream of cash flows. It is thus generally not informative for determining the appropriate discount rate for another asset, such as an investment in climate change abatement, which has benefits that tend to be delayed until much longer horizons and which have very different risk properties.

In theory, then, to understand the appropriate discount rate for investments in climate change abatement we would want to look at traded assets with similar *riskiness* and *horizons*. While this is difficult in practice, we show that researchers can still extract relevant information from the observed private market returns of assets, such as real estate. This information can then be used together with asset pricing models to adjust for the maturity and riskiness of cash flows of investments in climate change abatement.

Our first empirical contribution 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 investments in climate change abatement.[3] Using a variety of approaches, we estimate the average return to real estate to be around 6%. This contributes to a recent research effort to better document the return properties of residential real estate as an asset class (e.g., Favilukis, Ludvigson, and Van Nieuwerburgh 2017; Jordà et al. 2019; Chambers, Spaenjers, and Steiner 2019; Eichholtz et al. 2020). At the same time, recent estimates from Giglio, Maggiori, and Stroebel (2015) 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 (expected) returns and low long-run discount rates implies a downward-sloping term structure of discount rates for real estate. Intuitively, since real estate assets are claims to cash flows (rents) at all horizons, their expected rate of return is an average of the discount rates on short-run and long-run cash flows. If average returns are higher than long-run discount rates, then short-run discount rates must be higher than long-run discount rates (and higher than average returns).

[3] van Binsbergen, Brandt, and Koijen (2012) provide evidence of a downward-sloping term structure of discount rates for equities over a 1- to 10-year horizon. van Binsbergen and Koijen (2017) review related evidence across a number of asset classes.

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

Of course, this simple comparison ignores potential differences in risk properties of investments in climate change abatement and real estate. We thus also document the risk properties of real estate. We first show that real estate is indeed a risky asset: its returns are positively correlated with consumption growth, and therefore with the marginal utility of consumption, and it performs badly during consumption disasters, financial crises, and wars. This is consistent with the average return to real estate of about 6%, which is above the real risk-free rate, and thus includes a risk premium to compensate investors for bearing risk.

We then document that real estate is exposed specifically to climate change risk, and that this risk is reflected in house prices. This is an important step in helping us link the discount rates applicable to real estate and the discount rates for investments in climate change abatement. For this analysis, we work with a proprietary data set of housing transaction prices as well as for-sale and for-rent listings for properties located in the coastal states of Florida, New Jersey, North Carolina, and South Carolina. Properties in these states are exposed to climate change risk due to both rising sea levels and hurricanes. To obtain a measure of each property's physical exposure to climate risk, we geo-code the addresses of all properties to identify those properties that will be flooded with a 6-feet increase in the sea level, as measured by NOAA.

Since physical exposure to climate risk is correlated with unobserved property amenities, such as beach access, we cannot simply compare the prices across properties that are differentially exposed to such risk in order to estimate the price impact of climate risk. Instead, we test whether the prices of properties that are more exposed to climate change decline in relative terms when the perception of climate risk increases. We measure perception of climate risk in the housing market by performing a systematic textual analysis of the for-sale listings to measure the frequency with which climate-related text (e.g., mentions of hurricanes or flood zones) appears in the written description of the listed properties. The fraction of listings that include such texts is the basis for a "Climate Attention index" that we construct at both the ZIP-code-quarter and ZIP-code-year levels. Our interpretation of this index is that it reflects households' perceptions of the risk of future climate change on the cash flows from real estate in those locations.

We use data on the universe of property transactions from these states to conduct hedonic regressions that explore how the transaction prices of properties in the flood zone vary differentially when the "Climate Attention index" changes, controlling for property characteristics and various fixed effects. Our analysis shows that when the fraction of property listings that mention climate change doubles, there is a 2% to 3% relative decrease in the prices of properties that are in the flood zone compared to otherwise comparable properties in the same ZIP code that are not in a flood zone. This result survives in a specification with property fixed effects, which only identifies the pricing of climate risks from multiple transactions of the same property in periods with differential perceptions of these risks. Furthermore, we show that annual rents of exposed and non-exposed properties do not vary differentially with movements in our "Climate Attention index." This confirms that our estimates of differential price movements are not driven by differential changes in the flow utilities, but instead result from a differential change in the risks associated with future cash flows.

Based on these findings, we conclude that real estate prices directly reflect climate risk, making it a particularly interesting asset to study the valuation of investments to mitigate such risks. These findings are consistent with a quickly growing literature in finance that has documented the exposure of real estate to physical climate risk factors, such as rising sea levels and wildfires (e.g., Hallstrom and Smith 2005; McKenzie and Levendis 2010; Atreya and Ferreira 2015; Bakkensen and Barrage 2017; Gibson, Mullins, and Hill 2017; Eichholtz, Steiner, and Yönder 2019; McCoy and Walsh 2018; Ortega and Taspinar 2018; Bernstein, Gustafson, and Lewis 2019; Garnache and Guilfoos 2019; Baldauf, Garlappi, and Yannelis 2020).[4] Relative to much of this literature, our use of time- and space-varying measures of climate risk attention and our focus on rents in addition to home sales allow us to address a number of alternative interpretations of the observed relative price differences between properties that are differentially exposed to climate risk.

To explore the implications of the downward-sloping term structure of risky real estate for valuing investments in climate change abatement, we build a tractable asset pricing model that incorporates crucial features of climate change and its related risks. Our aim is not to provide an entirely new asset pricing model, nor is it to fully incorporate the micro foundations of physical models of climate change. Rather, we aim to provide a transparent and portable framework to show how the insights of modern asset pricing theory can be used together with inputs from a physical model of climate change to inform the appropriate discount rates for investments in climate change abatement.[5]

[4] Other research has explored the extent to which other asset classes, such as equities and fixed-income assets, are exposed to climate risk (Engle et al. 2020; Huynh and Xia 2020; Painter 2020). See Giglio, Kelly, and Stroebel (2020) for a review of this literature.

[5] This modeling approach relates to exciting new work that mixes 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.) as undertaken by Crost and Traeger (2014), Lemoine (2021), Lemoine and Traeger (2014), and others.

Our baseline model builds on the view that climate change is a form of disaster risk (see, Weitzman 2012; Barro 2015 for prominent articulations of this view): it is a rare event with potentially devastating consequences for the economy. We embed this view in a general equilibrium model with a representative agent and complete markets based on the endowment economy studied by Lucas (1978). We further modify this classic setup to reflect two important messages of the climate change literature.

First, we incorporate feedback loops between the state of the economy and the time-varying probability of a climate disaster. In particular, we allow the probability of a disaster to increase endogenously over time when the economy grows at a faster rate. Intuitively, this feature captures the notion that faster growth accumulates more environmental damages, such as greenhouse gas emissions and pollution, thereby increasing the probability of adverse climatic events, akin to tipping points (see, Alley et al. 2003, Lemoine and Traeger 2014). These damages in turn might feed on themselves, for example, because rising temperatures lead to even more carbon emissions for the same level of production. Our model captures these vicious cycles by allowing the probability of a further disaster to increase after a disaster occurs (see, Cox et al. 2000).

Second, we allow for economic growth to pick up temporarily after a disaster. This feature captures the potential adaptation of the economy following a disaster, and reflects a variety of adaptation measures, including relocating production to less affected areas, investments to prevent further damages (e.g., sea walls), and investments, such as air conditioning, that allow for productive work despite adverse climate conditions (see the discussions in, Brohé and Greenstone 2007; Desmet and Rossi-Hansberg 2015; Burke and Emerick 2016; Barreca et al. 2016). While we only capture these forces in reduced form, we show that they play a crucial role in capturing a more realistic evolution of the economy in response to climate change. In addition, this mean reversion of cash flows allows the model to match our data on the term structure of risky real estate. For assets exposed to the disaster risk, the partial mean reversion of the economy after a disaster implies that short-term cash flows are riskier than long-term cash flows, which only occur after the economy has partially recovered. This mechanism is central to generating downward-sloping term structures of discount rates: the riskier short-term cash flows are discounted at higher rates than the safer long-term cash flows.

Since climate change is a form of disaster risk, investments in the mitigation of this risk are hedges: similar to insurance policies, they pay off primarily in bad states of the world, and are thus particularly valuable. This has a number of implications for the discount rates used to value their cash flows. The first implication is that the *shape* of the term structure of discount rates for investments to abate climate change is the opposite of what we estimate for the term structure of housing, a risky asset. In fact, the term structure for abatement investments should be *upward sloping*: hedging against effects of the disaster on short-term cash flows is more valuable than hedging the effects on long-term cash flows, since these long-term cash flows are affected less due to adaptation.

Importantly, however, this upward-sloping term structure does not imply that the *level* of discount rates for investments in climate change abatement is high at any horizon. In fact, it should be *below the risk-free rate at all horizons*, reflecting the investment's hedge characteristics. For shorter horizons, we can observe the real risk-free rate (given by real bond yields) directly in the data, providing us with a tight upper bound (1% – 2%) on the discount rate for short-term cash flows from investments in climate change abatement. For longer horizons, there are no reliable estimates of the level of the risk-free interest rate. However, our model suggests that the very long-run discount rate of 2.6% for risky real estate provides an upper bound on the risk-free rate, and therefore also on the discount rates for long-term cash flows from investments in climate change abatement. This simple upper bound is a 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). 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). 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.

Note that our finding that the appropriate term structure to discount cash flows from climate change abatement is *low but upward-sloping* contrasts with a number of papers that have argued for using declining discount rates for valuing investments in climate change abatement (Arrow et al. 2013; Cropper et al. 2014; Farmer et al. 2015; Traeger 2014). These arguments have motivated policy changes in France and the United Kingdom, which have adopted a downward-sloping term structure of discount rates for evaluating long-run investments, including those in climate change abatement. While these differences do not have a substantial effect on the actual discount rates used to value the long-run cash flows from such investments (they are relatively low, at approximately 2%, both under the term structures used in those countries and under our upward-sloping term structure), the two have substantially different implications for the economic mechanisms to create these low long-run discount rates.[6] In addition, they have substantially different implications for evaluating the payoffs from climate abatement investments that may accrue at shorter

[6] The literature in climate change economics has sometimes motivated a downward slope in the discount rates for investments in climate change abatement with an extension of the Ramsey rule to include uncertainty about consumption growth that increases with the horizon. This would have the effect of pushing down the long-run risk-free rate due to a precautionary savings motive that increases in the horizon (see, Arrow et al. 2013). However, the predictions of this framework are inconsistent with the relatively flat term structure of real interest rates observed in the data. Moreover, the Ramsey framework does not consider the riskiness of cash flows and therefore has no predictions on the term structure of risk premiums. Consistent with this, the guidance on discount rates provided by governments recommending declining discount rates for cost-benefit analysis, usually does not indicate that the discount rate should vary with the risk properties of the investments.

horizons. The calibration of our model suggests that climate disasters cause the most damage immediately after they hit, making it most valuable to hedge the immediate costs. As a result, the correct discount rates for investments that yield shorter-term protection against climate change disasters should be substantially below the risk-free rate of 1%-2%. In contrast, the downward-sloping term structures used in France and United Kingdom suggest discount rates of 4% and 3.5%, respectively, for the first 30 years of a project's cash flows.

Finally, in addition to exploring the discount rates appropriate for climate change mitigation within our disaster-risk view of climate change risk, we can use our model to understand discounting of climate investments in alternative models of climate change risk. In particular, our specification for the economy and climate change dynamics is general enough to also nest, under a different parametrization, an important alternative view of climate change: that of the DICE models of Nordhaus and Boyer (2000) and Nordhaus (2008), in which (a) climate change acts as a tax on output and climate damages are higher when the economy is doing well, and (b) uncertainty about the path of the economy is the main driver of uncertainty about climate change. Under this parameterization, climate change mitigation investments pay off mostly in *good* states of the world (when the economy is expanding). The appropriate discount rates for these *risky* investments are thus *above* the risk-free rate. In this class of models, the climate "tax rate" can be increasing with the level of economic activity, so that the damages are disproportionally higher during booming economies. In our framework, such a feature implies discount rates for investments in climate change abatement that are *high and increasing with the horizon*. Intuitively, this occurs because a bad shock to the economy lowers both climate damages and the growth rate of damages over time. Our framework explains why the "disaster" view and the "tax" view of climate change have diametrically opposed predictions for the appropriate discount rates for investments in climate change abatement.

1. Risk and Return Properties of Real Estate

As described in the introduction, private market discount rates have the potential to inform the valuation of investments in climate change abatement. In this section, we discuss a number of reasons why real estate discount rates are particularly valuable from this perspective. First, we show that real estate is both risky in general (i.e., it pays off more in good states of the world) and exposed to climate risk in particular. Second, we show that, for real estate, private markets reveal information about the term structure of discount rates for horizons of up to hundreds of years. This feature of real estate is particularly beneficial to learn about the valuation of investments in climate change abatement, for which the potential benefits can stretch over very long time periods.

1.1 The riskiness of housing: Exposure to climate risk

We first provide direct evidence that climate risk is priced in real estate markets, with increased climate risk leading to relatively lower prices for more exposed properties. Our analysis has to overcome a number of empirical challenges. First, when comparing prices of properties that are differentially exposed to climate risk, it is difficult to control for all amenities that might be correlated with exposure to climate risk. For example, beachfront properties are more exposed to climate risk than properties further inland—they are more likely to be flooded when sea levels rise—but they might still sell at a premium because of the value of the beach access. Controlling for such difficult-to-measure amenities in hedonic regressions is challenging, which introduces concerns about omitted variable bias.

To overcome this challenge, we therefore investigate how the prices of properties that are differentially exposed to climate risk change in response to a change in that climate risk. As long as the amenity value of beach access does not change when climate risk changes, this analysis is informative about the pricing of climate risk in housing markets. However, such a "differences-in-differences" analysis presents a second challenge: true climate risk is a relatively slow-moving object that does not provide much of the time-series variation required to identify how it is priced. Our approach is to instead exploit the much more substantial time-series variation in the *attention* paid to climate risk in the housing market. Indeed, even though true climate risk might not change much from year to year, we show that the extent to which homebuyers focus on these risks changes much more frequently, and we would thus expect the pricing implications of climate risk to be particularly strong when households pay more attention to these risks.

1.1.1 Data construction. Our empirical analysis builds on a number of data sets. Our baseline data contain the universe of for-sale and for-rent property listings from Zillow, a major online real estate data provider. We obtained listings from four coastal states with properties that are potentially exposed to climate risk through rising sea levels: Florida, New Jersey, North Carolina, and South Carolina. For each listing, we observe the textual description of the property provided by the real estate agents, in addition to the listing date and listing price. The For-rent listings cover the period between the first quarter of 2011 and the second quarter of 2017. The For-sale listings extend back to the first quarter of 2008.

Our second data set contains the universe of public record assessor and transaction deeds data for the same states since the start of 2008. These data include detailed property characteristics, such as information on the property size and the number of bathrooms and bedrooms, as well as transaction prices and dates for all property sales.

To measure different properties' exposures to climate risk, we geo-code their addresses and map them to geographic shapefiles provided by the National

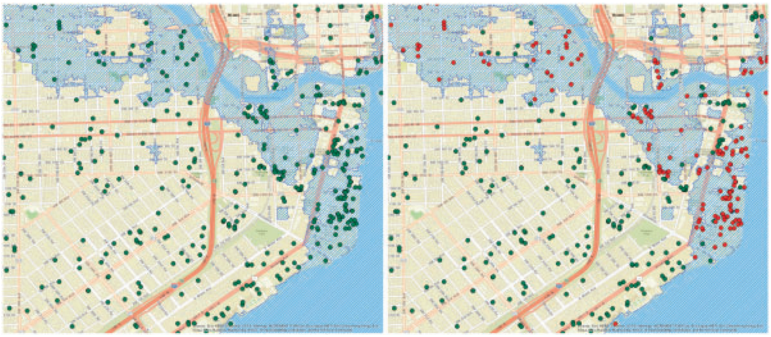

Figure 1
Illustration of identifying properties in the flood zone
This figure illustrates how we identify properties in the flood zone of downtown Miami, Florida. On the left, we plot each property as a green dot and overlay the NOAA's flood map. Then, on the right, we geo-code to identify the properties that fall under the flood zone and represent them as red dots. As seen above, properties closer to the coastal line are more likely to be situated in the flood zone.

Oceanic and Atmospheric Administration (NOAA) that indicate which regions will be flooded should sea levels rise by 6 feet or more. While flooding risk is only one of a number of climate risk factors, it is an important and easily measurable risk for properties in the coastal regions of the states analyzed in our study. Properties that are more exposed to climate risk on this measure tend to be closer to the waterfront, but there is substantial variation in exposure to climate risk across properties in the same narrow geography (see Figure 1, which shows the variation in our measure of climate risk exposure for downtown Miami).

We use our property listings data to build a novel measure of attention to climate risk. We construct this "Climate Attention index" by calculating the proportion of for-sale listings with property descriptions that contain climate change-related words and phrases such as "hurricanes," "FEMA," "floodplain," and "flood risk." Most of the flagged listings include descriptions that highlight that a specific property is *less* exposed to climate risk (e.g., "Not in a flood zone, it's high and dry!"). We believe that this is sensible: if you are selling a property with particular exposure to climate risk, for example, because it sits in a flood zone, you would not want to highlight this negative feature in a property listing. However, if you are selling a house that is *not* exposed to climate risk, this is something worth highlighting in a property listing, in particular in areas and at times when potential buyers pay more attention to these risks. Internet Appendix A.2 provides more details on the construction of the Climate Attention index, which we make publicly available to other researchers in the replication package associated with this paper.

There is substantial spatial and time-series variation in this measure of climate risk attention. The top panel of Figure 2 provides a heatmap of the Climate Attention index for Florida, pooling across all listings in our sample at the ZIP code level (the Internet Appendix includes corrsponding maps for

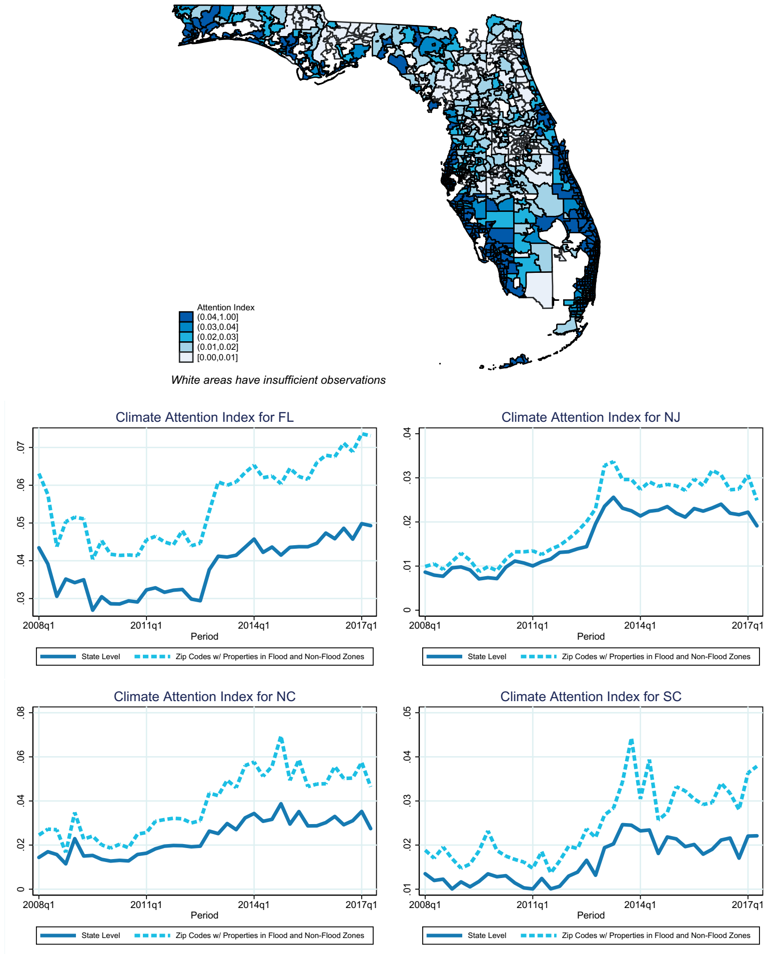

Figure 2
Climate attention index in the cross-section and time series
The top panel visualizes a heatmap of our Climate Attention index in Florida at the ZIP code level. The Climate Attention index is defined as the fraction of for-sale listings whose description includes climate-related text for the period from 2008Q1 to 2017Q2. The other panels illustrate the quarterly time series of the Climate Attention index aggregated at the state level as well as for ZIP codes that include at least some properties in the flood zone.

the three other states). Properties near the coast are more susceptible to climate risk. Consistent with this, the Climate Attention index is substantially higher for these properties in the cross-section. The other panels of Figure 2 illustrate the time series of the Climate Attention index for each of the four states in our sample, both for the whole state (black solid line) and only for ZIP codes that include at least some properties in a flood zone (blue dashed line). Consistent with the heatmap, the attention paid to climate risk is substantially higher in ZIP codes that are located in parts of the country where properties will be flooded if sea levels rise substantially. There is also sizable time-series variation in the Climate Attention index within geographies. For example, in New Jersey, the Climate Attention index nearly tripled between 2011 and 2013, around the time of Hurricane Sandy, which rendered more than 20,000 homes in the state uninhabitable.

1.1.2 Empirical analysis. Next, we estimate how climate risk is priced in real estate markets. Our baseline hedonic regression is given by Equation (1):

$$log(Price)_{i,h,g,t} = \alpha + \beta log(Index_{g,t}) \times FloodZ_h + \gamma FloodZ_h + \delta X_h + \phi_g \times \psi_t + \epsilon_{i,h,g,t}. \quad (1)$$

The unit of observation is a transaction i, of property h, in ZIP code g, at time t. The dependent variable is the log of the transaction price. We flexibly control for various property characteristics in X_h. We also include ZIP code-quarter fixed effects, $\phi_g \times \psi_t$, to capture differential house price movements across ZIP codes and time. We interact the log of the Climate Attention index, $log(Index_{g,t})$, with the Flood Zone indicator, $FloodZ_h$.[7] This allows us to estimate the effects of changing climate attention for properties that are differentially exposed to physical climate risks. We also include the Flood Zone indicator directly, allowing us to control for the unconditional price effect of being located in a flood zone as well as of any unobserved property amenities that are correlated with this measure of exposure to climate risk.

Column 1 of Table 1, panel A, shows estimates from this regression when we measure the Climate Attention index at the ZIP-code-year level. All else equal, properties that lie in the flood zone trade at a (statistically insignificant) premium to properties that are not in the flood zone, consistent with those properties also having more attractive amenities, such as proximity to the beach. More importantly, we estimate a statistically significant negative β-coefficient. A doubling in the Climate Attention index is associated with a relative 2.4% decline in the transaction prices of properties in the flood zone. The direct effect of increasing climate attention on all properties is absorbed by the ZIP-code-quarter fixed effects.[8] Column 2 measures the Climate Attention index at the ZIP-code-quarter level, and presents similar estimates. In columns 3 and 4 of Table 1, panel A, we include property fixed effects in the regressions from columns 1 and 2. In these specifications, the estimates of β are identified off

---

[7] To deal with the (small) number of ZIP-code-years with no listing mentioning climate change, we add a small constant (0.01) to the Climate Attention index before taking logs. Our results are robust to variation in the constant added and to the linear (instead of log-linear) inclusion of the Climate Attention index.

[8] While the coefficients for the control variables are not of primary interest in this study, Internet Appendix A.2.2 shows that they are consistent with estimates from the literature (e.g., Kurlat and Stroebel 2015; Stroebel 2016): for example, larger and more recently upgraded homes trade at a premium.

Table 1
Transaction prices and rent prices: Hedonic analysis

	A. Transaction prices *Dependent variable:* log(transaction prices)					
	(1)	(2)	(3)	(4)	(5)	(6)
Flood Zone	0.004	0.014			0.085***	
	(0.015)	(0.013)			(0.006)	
log(Index by ZIP-Year)	−0.024***		−0.029**			
× Flood Zone	(0.005)		(0.010)			
log(Index by ZIP-Quarter)		−0.020***		−0.021**		
× Flood Zone		(0.004)		(0.007)		
Index by ZIP-Quarter					−0.210**	−0.367***
× Flood Zone					(0.071)	(0.091)
Property controls	✓	✓	.	.	✓	.
ZIP × Quarter FE	✓	✓	✓	✓	✓	✓
Property FE	.	.	✓	✓	.	✓
R-squared	.585	.585	.721	.721	.585	.721
N	7,287,000	7,233,113	3,485,238	3,443,265	7,233,113	3,443,265

	B. Rent prices *Dependent variable:* log(rent prices)					
	(1)	(2)	(3)	(4)	(5)	(6)
Flood Zone	0.041***	0.033**			−0.034***	
	(0.012)	(0.011)			(0.006)	
log(Index by ZIP-Year)	0.018***		0.005			
× Flood Zone	(0.004)		(0.005)			
log(Index by ZIP-Quarter)		0.015***		0.003		
× Flood Zone		(0.004)		(0.003)		
Index by ZIP-Quarter					0.415***	0.016
× Flood Zone					(0.072)	(0.042)
Property controls	✓	✓	.	.	✓	.
ZIP × Quarter FE	✓	✓	✓	✓	✓	✓
Property FE	.	.	✓	✓	.	✓
R-squared	.728	.728	.942	.942	.728	0.942
N	2,142,433	2,142,240	1,191,657	1,191,642	2,142,240	1,191,642

This table shows results from Regression 1. The dependent variable is the log of the transaction price in panel A and the log of the rental listing price in panel B. In columns 1, 2, and 5, we control for various property characteristics, such as the property size, property age, and the number of bedrooms. In columns 3, 4, and 6, we include property fixed effects. The Flood Zone indicator and the property controls are naturally dropped in these regressions due to perfect multicollinearity. Index by ZIP-Year and Index by ZIP-Quarter represent the fraction of listings whose description includes climate-related texts at the ZIP-code-year level and the ZIP-code-quarter level, respectively. Standard errors are clustered at the ZIP-code-quarter level and in parentheses. * $p < .05$; ** $p < .01$; *** $p < .001$.

properties that we observe transacting more than once. The estimates are nearly identical, suggesting that our baseline findings are not driven by unobserved property characteristics. In columns 5 and 6, we include the raw Climate Attention index rather than the log of the index. Interpreting the magnitudes suggests that a 1 percentage point increase in the number of listings that suggest particular attention to climate risk is associated with a 0.2% – 0.4% decrease in the transaction price.

One concern with the estimates presented above is that they might not just capture the pricing of future climate change risk, but that our estimates also might be picking up changes in the flow utility of climate risk-exposed properties that could be correlated with climate risk attention. For example, it could be that climate risk attention rises after damaging storms that have a particularly strong direct effect on the utility of living in properties located in flood zones. To show that such a confounding story is not driving our results, panel B of Table 1 runs regressions similar to Equation (1), but now uses the log of the rental listing price as the dependent variable. In contrast to the transaction price regression, rental prices of properties exposed to climate risk increase during periods of increasing attention paid to climate risk, though the effect declines and is not statistically significant when we include property fixed effects. This is reassuring, because it suggests that our findings for transaction prices are not the result of a decline in the flow utility of these properties when climate risk increases. Instead, the decline in transaction prices most likely results from the increased present discounted cost of climate risk.[9]

1.1.3 Key takeaways. The evidence provided above shows that real estate has substantial exposure to climate risk, and thus fulfills an important criterion for us to use housing discount rates to learn about how to value investments in climate change abatement.

1.2 The riskiness of housing: Exposure to consumption risk

Next, we show that in addition to being exposed to climate risk, real estate is exposed to consumption risk: its returns are higher in states of the world where the marginal utility of consumption is lower. To show this, we analyze the behavior of real house prices during financial crises and periods of rare consumption disasters; we also estimate the correlation between house prices and consumption as well as personal disposable income.

Panel A of Figure 3 shows the average reaction of real house prices during 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 data set of historical house price indexes for these countries. Internet Appendix A.3.1 provides the details of the crisis dates and the house price series. 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, which are usually characterized by high marginal utilities of consumption, contributes to the riskiness of real estate as an asset.

Panel B of Figure 3 shows the average behavior of house prices during the rare consumption disasters as defined by Barro (2006). The consumption disaster

[9] The positive effect on rents that we observe in some of these specifications could, for example, be the result of general equilibrium effects in the housing market. Increased attention to climate risk makes individuals who are interested in living near the coast less likely to want to buy a house. Individuals who choose instead to shift into the rental market could be driving up rents.

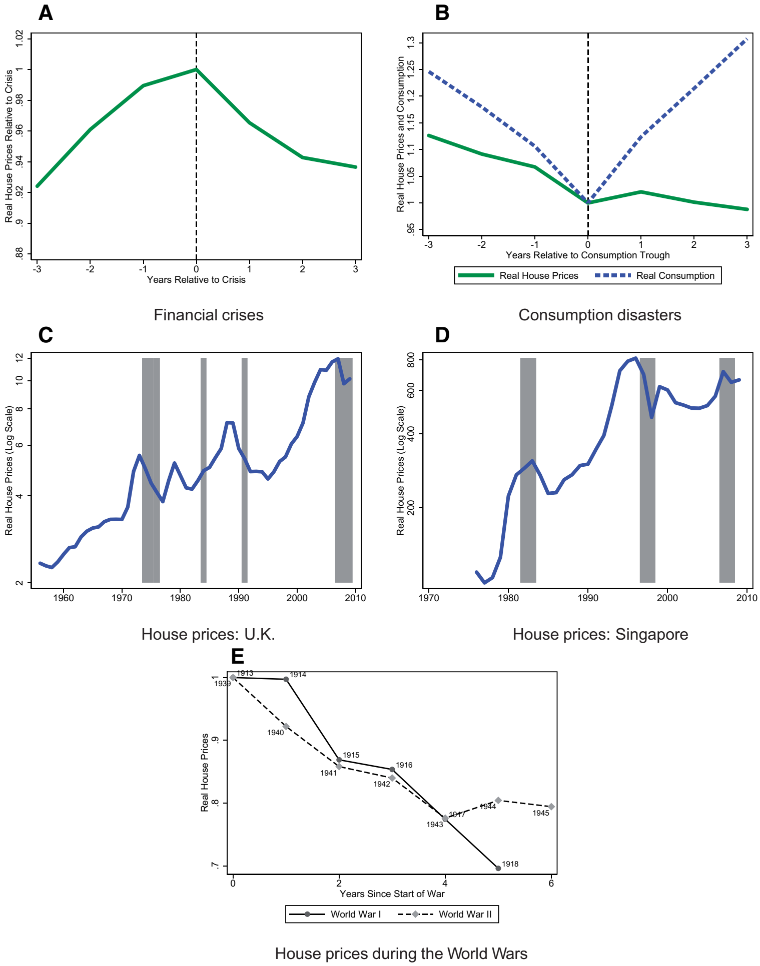

Figure 3
House price riskiness
Panel A shows average real house price movements relative to those during financial crises. Panel B shows average real house price movements and average real consumption relative to the trough of consumption disasters. Panels C and D show the evolution of real house prices in the United Kingdom and Singapore, respectively. Shaded regions represent financial crises. Panel E shows the evolution of real house prices for countries with available house price time series during World War I and World War II. See Internet Appendix A.3.1 for a description of the data series.

dates for the 20 countries included in our historical house price index data set are those defined by Barro and Ursua (2008). The dotted line tracks the level of consumption: following the start of a disaster, 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 class.

Panels C and D of Figure 3 show the time series of house prices and crisis years for the United Kingdom and Singapore—the countries with the best data on the term structure of housing discount rates (see below).[10] 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 such as the 2007-2008 global financial crisis. The 1984 banking crisis in the United Kingdom is the sole exception with increasing house prices.

Panel E of Figure 3 shows the performance of house prices during World War I and World War 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 with data availability for the duration of WWI (1913-1918): Australia, France, Netherlands, Norway, and the United States. House prices fell throughout the war with a total decline in real terms of around 30%. Similarly, the solid line tracks house prices of six countries—now also including Switzerland—for the duration of WWII (1939-1945). House prices fell by 20% in real terms from 1939 to 1943 and then stabilized for the last two years of the war, 1944-1945. 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.[11]

We also investigate the average correlation between consumption and house prices over the entire sample rather than just during crisis periods. Table 2 reports the correlation of house price changes with consumption changes as well as consumption betas 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. Accordingly, consumption betas are also positive except for France (−0.10) and often above 1.0. The estimated positive correlation between house prices and consumption and the positive consumption betas reinforce the evidence that real estate is a risky asset: it has low payoffs in states of

[10] All crisis dates are from Reinhart and Rogoff (2009), except for the periods 1997-1998 and 2007-2008 for Singapore. The latter dates have been added by the authors and correspond to the Asian financial crisis of 1997-1998 and the global financial crisis of 2007-2008.

[11] Despite extensive efforts to collect an exhaustive database. We 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 series before 1950. In addition, rental data are generally unavailable, preventing us from performing a comprehensive study of the riskiness of the underlying cash flows of housing. Nevertheless, our results suggest that real estate is an asset that has relatively lower payoffs during economic crises. Note that our results are likely to underestimate the riskiness of real estate and housing because of three effects. (1) House price indexes 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 indexes 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 sizable part of the housing stock is often destroyed during wars. Thus, 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.

Table 2
Real house price growth and real consumption growth

		Real HP growth		Real cons. growth			
	Period	Mean	SD	Mean	SD	Correlation	Cons. beta
Australia	1901-2009	2.51%	12.1%	1.51%	4.99%	.102	0.248
Belgium	1975-2009	2.92%	6.06%	1.59%	1.51%	.439	1.761
Canada	1975-2009	2.38%	7.69%	1.61%	1.73%	.433	1.929
Denmark	1975-2009	1.99%	9.24%	0.98%	2.71%	.538	1.838
Finland	1975-2009	2.17%	8.70%	2.07%	2.79%	.710	2.214
France	1840-2009	2.06%	11.8%	1.49%	6.32%	−.054	−0.101
Germany	1975-2009	−0.45%	2.33%	1.64%	1.52%	.494	0.755
Italy	1975-2009	1.28%	8.10%	1.75%	2.18%	.165	0.614
Japan	1975-2009	0.02%	4.45%	1.97%	1.60%	.503	1.394
Netherlands	1814-2009	2.79%	20.8%	1.57%	7.49%	.078	0.215
New Zealand	1975-2009	2.46%	8.09%	1.00%	2.30%	.580	2.044
Norway	1830-2009	1.77%	11.6%	1.78%	3.83%	.243	0.737
Singapore	1975-2009	7.18%	19.5%	3.43%	4.03%	.348	1.685
South Africa	1975-2009	1.13%	10.1%	0.92%	3.02%	.707	2.365
South Korea	1975-2009	0.58%	7.93%	4.62%	4.49%	.370	0.652
Spain	1975-2009	3.14%	8.07%	1.56%	2.60%	.593	1.837
Sweden	1952-2009	1.55%	6.04%	1.63%	1.99%	.536	1.627
Switzerland	1937-2009	0.47%	7.17%	1.48%	3.82%	.187	0.350
U.K.	1952-2009	2.89%	9.55%	2.26%	2.11%	.700	3.169
U.S.	1890-2009	0.49%	7.36%	1.84%	3.41%	.148	0.320

The table shows the time-series properties of annual growth rates of real house prices (as described in Internet Appendix A.3.1) 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. Column 7 shows the consumption beta of house prices.

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 23 developed and emerging countries (see Table 3). The average correlation is 0.33 and positive for all 23 countries, except for Croatia (−0.35), otherwise with a minimum of 0.04 for Norway and a maximum of 0.62 for Japan. The "personal disposable income beta" is positive for all countries, except Croatia (−0.16), and often above 1.0 again. Overall, this evidence further corroborates the fact that real estate returns are risky.

1.3 The term structure of real estate discount rates

Next, we provide evidence on an important and previously unexplored dimension of real estate data: the term structure of housing discount rates. We first present our analysis of expected real estate returns, which we find to be relatively high, between 5.5% and 7.4%. We then combine these new data with the estimates of Giglio, Maggiori, and Stroebel (2015) to provide evidence for the slope of the term structure of real estate discount rates. Our analysis suggests that this term structure is downward sloping, and thus cautions against using real estate's *average* rate of return to infer discount rates for very *long-run* benefits associated with investments in climate change abatement. In

Table 3
Real house price growth and personal disposable income growth

	Real HP growth		Real PDI growth			
	Mean	SD	Mean	SD	Correlation	PDI beta
Australia	3.54%	6.67%	1.37%	2.10%	.156	0.495
Belgium	2.53%	5.50%	0.92%	2.30%	.431	1.031
Canada	2.91%	7.49%	1.35%	2.18%	.466	1.604
Switzerland	1.12%	4.58%	1.17%	1.53%	.425	1.275
Germany	0.07%	2.52%	1.28%	1.64%	.237	0.365
Denmark	1.73%	8.58%	1.13%	2.29%	.224	0.839
Spain	−0.09%	10.6%	0.81%	2.27%	.409	1.909
Finland	1.90%	7.71%	1.92%	2.97%	.470	1.219
France	2.28%	5.13%	1.12%	1.61%	.332	1.056
U.K.	3.47%	8.49%	2.05%	2.26%	.420	1.575
Ireland	3.36%	9.44%	1.89%	3.33%	.574	1.627
Italy	0.33%	8.15%	0.89%	2.48%	.363	1.195
Japan	−0.39%	4.24%	1.49%	1.44%	.622	1.835
South Korea	0.64%	7.36%	3.97%	4.38%	.245	0.412
Luxembourg	4.16%	6.40%	2.76%	3.63%	.067	0.117
Netherlands	2.31%	9.12%	0.74%	3.01%	.467	1.414
Norway	2.65%	6.92%	2.22%	2.05%	.037	0.126
New Zealand	2.90%	7.73%	1.13%	3.41%	.486	1.103
Sweden	2.00%	7.01%	1.40%	2.39%	.467	1.371
U.S.	1.36%	3.88%	1.59%	1.54%	.322	0.812
South Africa	0.49%	9.13%	0.34%	2.37%	.474	1.824
Croatia	1.16%	12.3%	8.79%	27.0%	−.345	−0.158
Israel	3.05%	8.83%	2.74%	7.37%	.129	0.155

This table shows the time-series properties of quarterly frequency annual growth rates of real house prices and personal disposable income between 1975 and 2016, 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 show 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. Column 6 shows the personal disposable income beta of house prices.

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.

1.3.1 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 *price-rent approach*, starts from a 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: Without loss of generality, suppose we know the price-rent ratio at time $t=0$. We can then derive the time series of the price-rent ratio as

$$\frac{P_t}{D_{t+1}} = \frac{P_t}{P_{t-1}} \frac{D_t}{D_{t+1}} \frac{P_{t-1}}{D_t}; \qquad \frac{D_1}{P_0} \quad \text{given,} \tag{2}$$

where P is the price index and D the rental index. Note that, given a baseline price-rent ratio, only information about the growth rates in prices and rents is necessary for these calculations. Gross real housing returns are then

$$R^G_{t,t+1} = \left(\frac{D_{t+1}}{P_t} + \frac{P_{t+1}}{P_t} \right) \frac{\pi_t}{\pi_{t+1}}, \tag{3}$$

where π is a price level index to adjust for inflation. To compute expected net returns $E[R]$, we subtract maintenance costs and depreciation (δ) and any tax-related decreases in returns (τ):

$$E[R]=E[R^G]-\delta-\tau. \tag{4}$$

The second approach, which we label the *balance-sheet approach*, follows Favilukis, Ludvigson, and Van Nieuwerburgh (2017) and Piketty and Zucman (2014): We obtain data on the value of the residential housing stock from countries' national accounts to estimate the value of the housing stock (i.e., its price), and data on the net capital income earned on the housing stock (i.e., the "dividend" earned on the housing stock). Since we are only interested in the return to a representative property, we need to control for changes in the total housing stock to derive the net return to housing in each period as

$$R_{t,t+1}=\frac{V_{t+1}+NI_{t+1}}{V_t}\frac{\pi_t}{\pi_{t+1}}\frac{S_t}{S_{t+1}}, \tag{5}$$

where V is the value of the housing stock, NI is net capital income on housing, π is a price level index that adjusts for inflation, and S is the stock of housing.

To adjust for the quality and quantity of the housing stock, we use several complementary approaches. In our first approach, we proxy for the change in the housing stock by population growth. In alternative specifications, we control for the change in the housing stock with the growth in residential housing units or the growth in residential floor space. In our most conservative approach, we rely on (constant-quality) quantity indexes, which allows us to directly control for quality as well as "pure" quantity changes in the housing stock at the same time. For the United States, we can also draw on holding gains from the national revaluation accounts, which directly hold the aggregate stock of housing constant. Finally, even though our main interest lies in net returns to housing, the national accounts also allow us to estimate maintenance costs and depreciation (δ) and tax-related decreases in returns (τ), and hence gross returns to housing $E[R^G]$, which we compare to our results from the price-rent approach.

Table 4 presents estimates of the return to housing for three countries. We explore data from the United Kingdom and Singapore, since we are able to measure very long-run discount rates for these countries (see below); we also provide estimates for the United States for comparison, since they have been the subject of an extensive literature (e.g., Flavin and Yamashita 2002; Lustig and van Nieuwerburgh 2005; Piazzesi, Schneider, and Tuzel 2007). Internet Appendix A.4 provides the details of our approach and the underlying data sources.

United States. For the United States, our preferred estimates using the price-rent approach are based on a price-rent ratio from Trulia that includes a utilities correction (column 2); our preferred results using the balance-sheet

Table 4
Expected returns and rental growth

	United States										United Kingdom					Singapore
	Price/rent					Balance sheet					Price/rent		Balance sheet			Price/rent
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Gross return	10.2%	9.7%	8.9%	9.4%	10.0%	9.7%	9.3%	9.1%	8.9%	8.7%	9.5%	9.6%	10.2%	9.9%	9.7%	9.9%
Rental yield	9.5%	8.9%	8.1%	8.9%	9.0%	7.1%	7.1%	7.0%	7.0%	7.0%	6.8%	6.9%	7.0%	6.9%	6.9%	5.6%
Capital gain	0.8%	0.8%	0.8%	0.5%	1.0%	2.7%	2.3%	2.1%	1.8%	1.6%	2.7%	2.7%	3.3%	3.0%	2.8%	4.2%
Depreciation	2.5%	2.5%	2.5%	2.5%	2.5%	2.3%	2.3%	2.3%	2.3%	2.3%	2.5%	2.5%	2.4%	2.4%	2.4%	2.5%
Taxes	0.67%	0.67%	0.67%	0.67%	0.67%	1.1%	1.1%	1.1%	1.1%	1.1%	0%	0%	0%	0%	0%	0.6%
Net return	**7.1%**	**6.5%**	**5.7%**	**6.3%**	**6.8%**	**6.3%**	**5.9%**	**5.7%**	**5.5%**	**5.2%**	**7.0%**	**7.1%**	**7.9%**	**7.6%**	**7.4%**	**6.8%**
Rent growth	0.7%	0.7%	0.7%	0.7%	0.9%	-	-	-	-	-	1.4%	1.4%	-	-	-	−0.4%
Baseline-P/R	Trulia	Trulia⁻	Bal	Trulia⁻	Trulia⁻	-	-	-	-	-	Bracke	Bal	-	-	-	iProp
Price index	CS	CS	CS	FHFA	CS	-	-	-	-	-	LR	LR	-	-	-	URA
Rent index	CPI-S	CPI-S	CPI-S	CPI-S	PCE-H	-	-	-	-	-	CPIH	CPIH	-	-	-	URA
Stock adj.	-	-	-	-	-	Pop	Units	Floor	Reval	QI	-	-	Pop	Units	QI	-
Sample	1953-2016										1988-2016					1990-2016

This table shows our estimates for net real returns to housing and real rent growth in the United States, the United Kingdom, and Singapore based on the price-rent approach and the balance-sheet approach. The *price-rent approach* starts from a 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. Baseline-P/R is the source of the baseline price-rent ratio, either a direct estimate or based on the balance-sheet approach (Bal). In the United States, Trulia^{-} includes an adjustment for utilities possibly included in Trulia's gross rents. CS is the Case-Shiller house price index; FHFA is the FHFA house price index; CPI-S is the shelter component of the CPI; and PCE-H is the housing component of the PCE price index. In the United Kingdom, LR stands for Land Registry, and CPIH is the housing component of the consumer price index including housing. In Singapore, iProp stands for iProperty.com, and URA stands for Urban Redevelopment Authority. In the *balance-sheet approach*, the total value of the residential housing stock is used to estimate the value of housing, and net capital income earned on the housing stock is used to estimate net rents. To estimate the return on a representative property, we control for changes in the total housing stock by the growth in population (Pop), housing units (Units), housing floor space (Floor), or quality-adjusted quantity indexes (QI). The U.S. Financial Accounts also publish aggregate holding gains for each sector in the economy in the revaluation accounts (Reval), which directly hold the aggregate stock of housing constant. See Internet Appendix A.4 for further details on the estimation procedures and the underlying data sources used. Numbers may not add up due to rounding.

approach use direct holding gains from the revaluation accounts (column 9). We also provide robustness checks that use alternative price and rental indexes as well as alternative price-rent ratios for the price-rent approach, and various corrections for the growth in the housing stock for the balance-sheet approach. Both approaches provide similar estimates for the average annual real gross return ($E[R^G]$): 9.7% based on the preferred estimate from the price-rent approach and 8.9% based on the preferred estimate from the balance-sheet approach. We estimate a maintenance and depreciation impact of 2.3% using the balance-sheet approach and calibrate the impact of maintenance and depreciation at 2.5% for the price-rent approach based on prior results from Harding, Rosenthal, and Sirmans (2007). Our balance-sheet estimates imply a tax impact of 1.1%, and we assume a property tax impact of 0.67% for a representative household for the price-rent approach. This results in average real net returns of between 5.5% and 6.5% for the United States housing market. These estimates are similar to the estimates in Flavin and Yamashita (2002), who find a real return to real estate of 6.6%, and the estimates in Favilukis, Ludvigson, and Van Nieuwerburgh (2017), who find a real return of 9%-10% before netting out depreciation and property taxes.

United Kingdom. Columns 11 to 15 of Table 4 report the estimates for the real estate market in the United Kingdom The price-rent and the balance-sheet approaches produce similar estimates for the average annual real gross return ($E[R^G]$): 9.5% for the price-rent approach and 9.7% for the balance-sheet approach. We estimate a maintenance and depreciation cost of 2.4% using the balance-sheet approach and maintain a calibration of 2.5% for the price-rent approach. There are no property taxes to be considered in the United Kingdom Average real net returns in the United Kingdom real estate market are therefore between 7.0% and 7.4%.[12]

Singapore. Column 16 in Table 4 reports our price-rent approach estimate for the Singapore real estate market at 9.9%. We assume the cost of maintenance and depreciation to be 2.5%, in line with our estimates for the United States, and the property tax impact to be 0.6%. Our estimate of the real net return in the Singapore real estate market is thus 6.8%. We do not calculate complementary balance-sheet approach estimates for Singapore for two reasons: First, more than three quarters of residential dwellings are not in the private housing market but publicly governed and developed by the Housing and Development Board (HDB). Unfortunately, the national accounts data do not allow us to separate these out with sufficient accuracy. Second, the national accounts data do not allow us to determine the total consumption of real estate services excluding relevant costs, that is, net rents, with sufficient accuracy.

12 Numbers for the balance sheet approach may not add up due to rounding when moving from gross to net returns.

Average rate of return: summary. Overall, these estimates show that expected real returns for real estate are around 6% or higher for the countries we consider. These estimates are robust to the different methodologies we use. They are also in line with contemporaneous work from Jordà et al. (2017) that finds average returns to housing of around 7% before taxes across a number of countries. Our estimates are also consistent with 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 fact, 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. Consistent with our results from Section 1.2, our estimates of average returns to real estate imply a positive real estate risk premium.

Growth rate of rental income. Finally, we estimate the average real growth rate of rental income using the same data sources, which we denote by g. For all three countries, the estimated real growth rate of rents is low. For the United States, we estimate $g=0.7\%$, an estimate in line with that of Campbell et al. (2009), who obtain a median growth rate of 0.4% per year. We obtain a slightly higher estimate of $g=1.4\%$ for the United Kingdom and a slightly lower estimate of $g=-0.4\%$ for Singapore (largely driven by a few deflationary periods). These results are 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% (the equivalence of these two long-run growth rates is necessary for rental yields to be stationary).

1.3.2 Long-Run housing discount rates. In recent work, Giglio, Maggiori, and Stroebel (2015) use unique data from the United Kingdom and Singapore to estimate how much value households attach to future real estate cash flows accruing over a horizon of hundreds of years (see also Giglio, Maggiori, and Stroebel 2016). In these real estate markets, residential properties trade either as freeholds, which are permanent ownership contracts, or as leaseholds, which are prepaid 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.

The red bars in Figure 8 represent the estimates from Giglio, Maggiori, and Stroebel (2015). They show the price discount of leaseholds with varying maturities compared to freeholds for otherwise identical properties. For the United Kingdom 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, 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: between 17.6% for leaseholds with remaining maturity of 80-99 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 the data from both countries.

1.3.3 Takeaway: The term structure of discount rates in the housing market. In this section, we show that (a) real estate has a real expected rate of return of above 6% per year, and (b) the relative pricing of freeholds and leaseholds implies discount rates of around 2.6% for rents 100 years or more in the future. Since the average return on real estate is simply a weighted average of the average returns of all of its cash flows (at all maturities), these two facts together are informative about the shape of the term structure of discount rates for the housing asset. It needs to be low at the long end, in order to match the 2.6% discount rate applied to the long-term housing claims. But it needs to be high enough at the short end to imply an *average* discount rate of 6%. In other words, the term structure of discount rates for the housing asset needs to be downward sloping in order to explain the data. In the next section, we introduce an asset pricing model of real estate that is able to match these moments, and discuss its implications for valuing investments in climate change abatement.

2. Valuing Investments in Climate Change Abatement in a World with Declining Discount Rates for Risky Assets

The previous section provided evidence that the term structure of discount rates for real estate, a risky asset, is downward sloping, and that real estate is an asset class that is directly exposed to climate change risk. In this section, we introduce a general equilibrium model to study the link between climate change risk, the term structure of discount rates for real estate, and consumption. Our model has two objectives: First, it provides a quantitative framework, calibrated to asset markets and our new evidence on the term structure of housing discount rates, from which one can extract appropriate discount rates for climate-change-abatement investments at any horizon. Second, it allows us to nest, in reduced form, a number of different approaches to modeling the economic impact of climate change, ranging from the "tax view" in the spirit of Nordhaus and Boyer (2000) and Nordhaus (2008), to the "disaster view" in the spirit of Weitzman (2012, 2014). This allows us to understand these views' different predictions for the discount rates of investments in climate change abatement.

2.1 A general equilibrium model with climate change risks

Our model builds on a modified version of the Lucas (1978) representative-agent economy with power utility preferences. To provide a simple analytical framework for climate change, we introduce a production sector that, while exogenous, allows for important feedback effects between the growth rate of the economy and the probability of rare and adverse climate shocks that destroy parts of the output. The setup is rich enough to allow for key climate-related dynamics in the economy, including an endogenous relationship between consumption and climate risk, while also being stylized enough to be solved in closed form up to simple recursive expressions.

Model setup. We assume that aggregate consumption follows

$$\Delta c_{t+1} = \mu + x_t - J_{t+1}, \tag{6}$$

$$x_{t+1} = \mu_x + \rho x_t + \phi J_{t+1}, \tag{7}$$

where c_t is the log of aggregate consumption; since the economy is closed and does not feature investment, c_t also corresponds to aggregate output in equilibrium.[13] J_t is a jump process that takes value $\xi \in (0,1)$ with probability λ_t in each period, and value 0 otherwise. We interpret J as climate risk: a rare but possibly large negative shock to output. The climate disaster probability λ_t depends endogenously on the dynamics of the economy (see below).[14] The process x_t captures persistent changes in the growth rate of consumption and plays a key role in determining the term structure of discount rates.

As is standard in financial economics, we allow for a separate cash flow process, d_t, for risky assets—which in our model corresponds to the rents of real estate—to capture the idea that asset markets only reflect a subset of total economic activity. The process for these rent cash flows is similar to the one for aggregate consumption:

$$\Delta d_{t+1} = \mu_d + y_t - \eta J_{t+1}, \tag{8}$$

$$y_{t+1} = \mu_y + \omega y_t + \psi J_{t+1}. \tag{9}$$

[13] Note that we assume complete markets, so all risk is shared perfectly across households. The equilibrium effects of incorporating heterogeneity and incomplete risk-sharing in asset pricing models has been studied in a long literature (see, e.g., Constantinides and Duffie 1996). We leave the exploration of the specific implications of climate change risks to future research.

[14] Our model is designed to help researchers and policy makers understand how to value investments in climate change abatement. We thus remove any risk sources not related to climate risk. Other shocks could be added without changing the qualitative implications of the model.

The main difference between real estate rents and consumption is the larger exposure of rents to the underlying economic shocks, represented by climate risk J. This is captured by the multiplier $\eta > 1$. In our case, η reflects the empirical observation that housing has an above-average exposure to climate risk, primarily due to the immovability of land. Analogous to x_t, the process y_t captures persistent changes in the growth rate of rents. Having different processes for x_t and y_t allows for flexibility in the calibration of this model to different specific settings.

Our setup allows for partial mean reversion in the growth rate of consumption and rents after a disaster. Formally, after a disaster strikes, the growth rate of the economy increases ($\psi > 0, \phi > 0$) and this increase is persistent ($\rho > 0, \omega > 0$). As we show below, this partial mean reversion plays a crucial role in explaining the term structure of discount rates for risky assets (see also Gourio 2008; Lettau and Wachter 2011; Nakamura et al. 2013; Belo, Collin-Dufresne, and Goldstein 2015; Hasler and Marfe 2016). In the context of climate risk modeling, the partial mean reversion captures the notion that economic activity picks up after a climate disaster as the economy adapts to new climatic circumstances. Numerous papers have highlighted the importance of such adaptation processes, including Brohé and Greenstone (2007), Deschênes and Greenstone (2011), Desmet and Rossi-Hansberg (2015), Burke and Emerick (2016), and Barreca et al. (2016). Yet, since there have not been many global climate disasters (especially in modern data), such feature remains a possibility rather than an empirical regularity.

The last component of our model is an endogenous climate disaster probability, λ_t:

$$\lambda_{t+1} = \mu_\lambda + \alpha\lambda_t + \nu x_t + \chi J_{t+1}. \tag{10}$$

In designing this process, we aim to capture some of the main features of physical models of climate change, while at the same time maintaining a tractable solution to the asset pricing model. Two features of this process make it particularly useful for bringing climate risk into an asset pricing framework:

1. The disaster probability λ_t is an endogenous function of the growth rate of the economy. Since x_t—which captures the expected deviation of the growth rate of the economy from the trend—enters additively and positively ($\nu > 0$) in Equation (10), the probability of a disaster increases over time when the economy grows at a faster rate. Intuitively, this feature captures the notion that faster growth accumulates more environmental damages, such as greenhouse gas emissions and pollution, thereby increasing the probability of adverse climatic events, akin to tipping points (see, Alley et al. 2003; Lenton et al. 2008; Overpeck and Cole 2006; Lemoine and Traeger 2014; Franklin and Pindyck 2018).

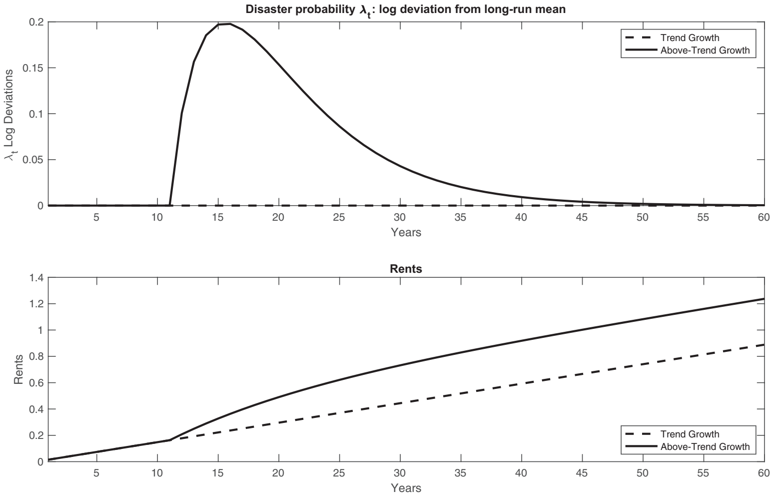

Figure 4
Sample paths: Trend growth and above-trend growth
The figure shows two sample paths of the economy under our baseline calibration. The top panel reports the log deviation of the climate disaster probability, λ_t, from its mean. The bottom panel reports the path of log rents, d_t. The dotted line represents the baseline path along which the economy grows at its deterministic trend. The solid line represents a temporary deviation from the trend in which growth accelerates.

2. The climate disaster probability λ_t increases following the occurrence of a disaster ($\chi > 0$), thus allowing climate shocks to induce a self-reinforcing cycle in which each shock increases the probability of the next shock (see, e.g., Cox et al. 2000; Melillo et al. 2017). Note that, in contrast to the mean reversion in cash flows described above, this is a force that pushes toward making long-run cash flows more risky.

To illustrate the richness of these patterns, Figures 4 and 5 plot two sample paths of the economy. Figure 4 shows a path in which no disaster occurs, but the economy grows above trend for a sustained period of time, starting in year 10.[15] The top panel shows log deviations of the disaster probability λ_t from its steady-state value. We set the steady state value to 3% to reflect the Barro (2006) estimate of the average probability of a consumption disaster. The bottom panel shows the path of log rents, d_t, over time: rents (and the economy overall) increase at a decreasing rate, reaching a permanently higher level as a result of the growth spurt. This sustained economic expansion induces a progressive increase in the probability of a climate disaster until the economy returns to its steady-state growth rate. The lags of the effect of greenhouse gas emissions and pollution on the disaster probability can be substantial: The disaster probability

15 As we will discuss in Section 2.2, trend growth is calibrated at 2%. We assume that in period 10, growth of both consumption and rents increases to 5% and then slowly reverts to long-run trend growth. Since growth is a persistent process in the model, growth is above trend for approximately 20 years in this sample path.

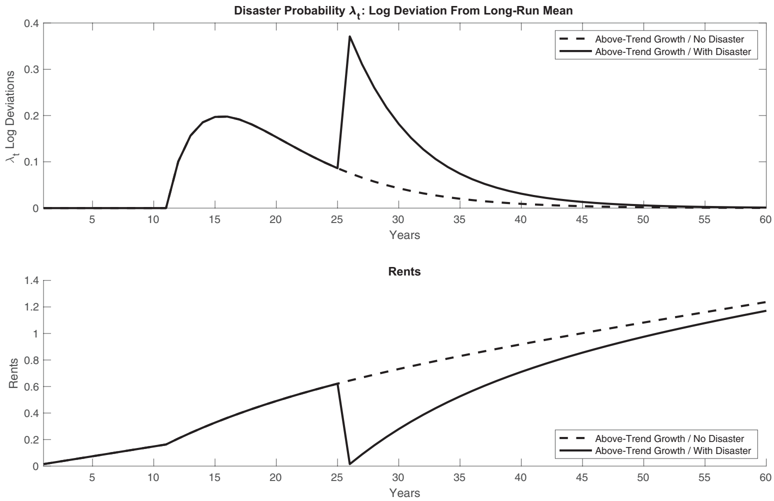

Figure 5
Sample paths: Above-trend growth, with and without a disaster
The figure shows two sample paths of the economy under our baseline calibration. The top panel reports the log deviation of the climate disaster probability, λ_t, from its mean. The bottom panel reports the path of log rents, d_t. The dotted line represents a path along which the economy grows at its deterministic trend, then experiences an increase in growth (the same as the solid line in Figure 4). The solid line represents an alternative path along which the increased probability of disasters due to the temporary acceleration in the economy leads to the occurrence of a disaster after year 25.

reaches its maximum approximately 7 years after the growth spurt has started. Since the model is stationary, the disaster probability ultimately reverts to its mean, but the half-life of the shock is extremely long at 14 years.

Figure 5 instead shows a path in which the economy expands above trend, starting in year 10 (exactly as before), but with a climate disaster occurring after year 25. This disaster induces a large drop in consumption and rents. The dynamics of climate risks are particularly interesting. As before, the disaster probability increases as the economy accelerates. Once the disaster hits, the probability of a future disaster increases further. It takes almost 40 years (in a sample path chosen to have no further disasters) from the original growth spurt shock for the probability of a disaster to revert to its long-run mean. The bottom panel of Figure 5 also illustrates the mean reversion in the growth rate of the economy. After a disaster strikes, the growth rate of the economy increases ($\psi > 0$) and this increase is persistent ($\omega > 0$).

The term structure of discount rates for risky assets. Despite the richness of the underlying dynamics of the economy, we are able to solve quasi-analytically for the term structure of housing risk premiums and the risk-free rate. We derive these objects assuming the existence of a representative agent who maximizes lifetime utility and faces a complete set of financial instruments. In our baseline model, the period utility function features constant relative-risk aversion (γ) as in Lucas (1978):

$$U(C_t)=\delta^t \frac{C_t^{1-\gamma}}{1-\gamma}, \tag{11}$$

where δ is the rate of time preference. In Internet Appendix A.6, we derive the prices of claims to consumption and rents at different horizons. Here, we focus on the crucial forces determining the term structure of housing discount rates. Formally, we are interested in the per-period discount rate of maturity n, denoted $\overline{r}_t^n$, that equates the price of a single cash flow $E_t[D_{t+n}]$ of maturity n, denoted $P_t^{(n)}$, with its present discounted value:[16]

$$P_t^{(n)}=\frac{E_t[D_{t+n}]}{(1+\overline{r}_t^n)^n}. \tag{12}$$

As we show in Internet Appendix A.5, the term structure of discount rates $\overline{r}_t^n$ is closely linked to the term structure of one-period expected returns. Intuitively, the appropriate discount rate for cash flows of horizon n is simply the average across one-period expected returns $E_t[R_{t,t+1}^{(n)}]$ for claims to cash flows at each horizon up to n, where the holding-period returns over the next period are given by $R_{t,t+1}^{(n)}=P_{t+1}^{(n-1)}/P_t^{(n)}$. More formally, this can be written as[17]

$$\overline{r}_t^n \simeq \frac{1}{n}\sum_{k=1}^{n}\ln(E_t[R_{t,t+1}^{(k)}]). \tag{13}$$

The expected one-period return for the n-maturity claim, $E_t[R_{t,t+1}^{(n)}]$, can in turn be thought of as the sum of the one-period risk-free rate, $R_{t,t+1}^f$, which is independent of the maturity of the claim, and a risk premium, $E_t[R_{t,t+1}^{(n)}]-R_{t,t+1}^f$, which varies with the horizon n. While Internet Appendix A.6 provides the full solution of the model, and while the calibrated results presented below use this full solution, we next focus on a simple approximate solution that captures the main forces that shape the term structure of one-period excess returns in our model here:

$$\begin{aligned} E_t[R_{t,t+1}^{(n)}]-R_{t,t+1}^f &\simeq \gamma\ Cov_t[r_{t,t+1}^{(n)},\Delta c_{t+1}] \\ &=\gamma\left[\eta-\psi e_{d,n-1}-\phi b_{d,n-1}-\chi f_{d,n-1}\right]\xi^2\lambda_t(1-\lambda_t). \end{aligned} \tag{14}$$

The first equality above uses the fact that the log stochastic discount factor under power utility preferences is $m_{t,t+1}=\log\delta-\gamma\Delta c_{t+1}$, and the second equality

[16] We label objects that relate to single cash flows at a specific maturity n by superscript (n). The set of claims to a single cash flow at maturity n that we are interested in is a subset of a more general class of assets with maturity n that could pay cash flows, such as rents at any point in time up to that maturity. We denote prices and returns of claims to more general classes of assets with maturity n with superscript n.

[17] See Internet Appendix A.5.1 for a derivation. The result holds exactly when the term structure of discount rates is constant over time (though it can have any shape over maturities n). For example, a flat term structure of discount rates implies a flat term structure of expected one-period returns across maturities, and vice versa.

represents the solution to the model, where b_d, e_d and f_d solve recursive equations derived analytically and reported in Equations (A.14c) to (A.14e) in Internet Appendix A.6.

Equation (14) highlights the components that drive the downward-sloping term structure of discount rates for risky assets (i.e., housing) in our framework. The *level* of the term structure is determined by the aggregate amount of risk in the economy, $\xi^2\lambda_t(1-\lambda_t)$, and by the agent's risk aversion, γ. Neither has a differential effect on the risk premiums of cash flows with different maturities (there are no n-subscripts).

The *shape* of the term structure is determined by the terms inside the bracket. The term b_d arises from the term structure of risk-free assets (term premia) and is essentially constant in realistic calibrations of the model that match the flat term structure of risk-free rates in the data. The term f_d is a quantitatively small adjustment for the risk that arises from changes in the disaster probability λ_t. The component that quantitatively dominates the shape of the term structure of housing risk premiums is $\psi e_{d,n-1}$, which captures the term structure of exposures of claims of different maturity to the climate disaster (and the ensuing recovery). The model is parsimonious enough to admit an analytical solution for $e_{d,n}$ (see Equation (A.14e) in the Internet Appendix):

$$e_{d,n} = \frac{1-\omega^n}{1-\omega}. \tag{15}$$

Since this term enters negatively in Equation (14), positive values for both ψ and ω imply a declining term structure of risk premiums for rents. Recall that ψ determines the degree of mean reversion of the growth rate of the economy after a climate disaster, and ω captures the persistence of this growth rate increase. When $\psi > 0$, as in our baseline calibration below, rents (partially) mean-revert after a climate disaster. This mean reversion in cash flows implies that the occurrence of a disaster is worse for short-term claims than it is for long-term claims because immediate short-term cash flows drop by more than cash flows that are further in the future.

Remark on preferences. The previous discussion highlights that, in our setting, the observed downward-sloping term structure of risk premiums for housing is generated by the dynamics of the cash flows (*risk quantities*) rather than by the term structure of *risk prices*, which are flat. One might wonder whether more sophisticated preferences, such as Epstein-Zin preferences that are popular in both the asset pricing and climate change literatures, could also generate this downward slope. We discuss in Internet Appendix A.5.3 that this is not the case. In fact, introducing Epstein-Zin preferences would push the slope of the term structure of discount rates for risky assets upward. To match the data on a downward-sloping term structures of discount rates for risky real estate, we would thus require an even stronger mean reversion in cash flows.[18] More generally, we are not aware of a standard representation of preferences that would push toward a downward-sloping term structure of discount rates for risky assets, such as real estate. As a result, capturing the observed downward slope through the dynamics of risk quantities, as we do in our model, seems like the natural approach to us, in particular given that the required dynamics are highly consistent with empirical research on the adaptation to climate change.

2.2 Calibration

In this section, we turn to calibrating our model. The objective of our calibration is not to quantitatively match all conceivable moments of real and financial variables; since our model is only driven by a single climate disaster shock, J, we would certainly fail in such an exercise along many dimensions. Furthermore, history and scientific evidence only provide incomplete and uncertain guidance on many key parameters related to climate events. What we strive for instead is a reasonable calibration that can match core moments of the data as they relate to the discounting of climate-change-abatement investments and, in particular, match our new evidence on the risk and return properties of real estate, including the term structure of discount rates from Section 1.

2.2.1 Baseline calibration. Whenever possible, we calibrate parameters following the existing asset pricing literature. The remaining parameters are calibrated to match some of our new moments estimated in Section 1. For example, we follow the asset pricing literature and set risk aversion $\gamma = 10$, the drop in consumption following a disaster $\xi = 21\%$, and the exposure of risky cash flows to the climate shock $\eta = 3$ (see Bansal and Yaron 2004; Barro 2006; Barro and Jin 2011). Average consumption growth in the absence of a disaster is set to $\mu = 2\%$. The remaining parameters of the consumption process are chosen to generate a recovery in consumption growth after disasters ($\phi > 0$), and persistent growth rates ($\rho > 0$). The magnitude of these parameters ($\phi = 0.025$, $\rho = 0.85$) targets a term structure of real interest rates that is slightly upward sloping with a level of around 1%, matching our empirical estimates based on the U.K. gilts real yield curve between 1998 and 2016 reported in Figure 6. These data show that the U.K. real yield curve is approximately flat on average, with a real yield close to 1% for maturities between 1 and 25 years.[19]

In our calibration, rents are not only correlated with consumption but also share many dynamics with consumption, including recovery after disasters

[18] The long-run risk model of climate change by Bansal, Kiku, and Ochoa (2013) and the model of Cai, Judd, and Lontzek (2013) both use Epstein-Zin preferences.

[19] Figure 6 plots the average shape of the real U.K. gilts curve for the period 1998-2016, as well as for two subperiods: 1998-2007 and 2008-2016. The level of the yield curve shifted down during this latter period and the yield curve became hump shaped. More recently, as more and more U.K. government bonds with longer maturities have been issued, reliable prices for such longer maturities have also become available. In 2016, the Bank of England therefore started to extend the real yield curve up to maturities of 40 years. For the short time period, when data on such long maturities are available, the yield curve is essentially flat for these longer maturities as well.

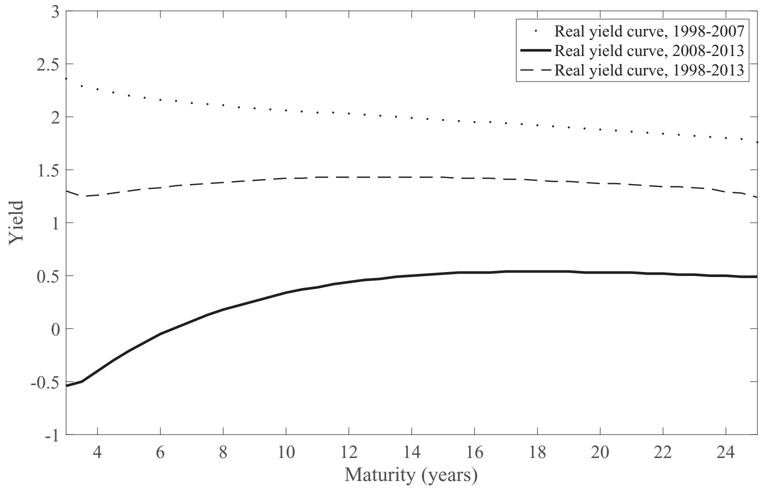

Figure 6
U.K. gilts real yield curve
The figure plots the real yield curve for U.K. gilts as computed by the Bank of England for the period 1998-2016, as well as for two subperiods: 1998-2007 and 2008-2016. It is available at http://www.bankofengland.co.uk/statistics/Pages/yieldcurve/archive.aspx, last accessed July 2017. Until 2015, the U.K. government debt also included some perpetual bonds: war loans and annuities. These bonds composed a negligible part of the outstanding U.K. government debt (£2.6bn out of £1.5trn of debt outstanding) and were classified as small and illiquid issuances by the U.K. Debt and Management Office. In 2015, following the passage of the Finance Act, all outstanding perpetuities were called in by the British government. 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 difficult to interpret their prices in terms of discount rates.

($\psi > 0$) and persistent rent growth ($\omega > 0$). The magnitudes of these parameters ($\psi = 0.24$, $\omega = 0.915$) are chosen to match the shape and the level of the observed term structure of discount rates in the housing market as described in Section 1. Finally, we set the steady-state conditional probability of disasters, $\overline{\lambda}$, to 3% per year, following the estimates in Barro (2006). The remaining parameters for the λ_t-process are chosen to obtain economically reasonable interactions between the real economy and the disaster probability, while at the same time matching the term structure of the risk-free rate. The risk-free rate is directly affected by the disaster probability dynamics through the precautionary savings channel; an increase in the disaster probability decreases the rate by increasing precautionary savings. In particular, the disaster probability is persistent ($\alpha = 0.75$), increases after a jump ($\chi = 0.05$), and increases when expected consumption growth x_t is above its trend ($\nu = 0.1$). Finally, we impose that rents and consumption have the same long-run growth rate and that x_t and y_t have mean zero.[20]

[20] The resultant parameter restrictions are $\mu_d = \mu + (\eta - 1)\bar{\lambda}\xi$, $\mu_x = -\bar{\lambda}\phi\xi$, and $\mu_y = -\bar{\lambda}\psi\xi$. See Internet Appendix A.6 for details.

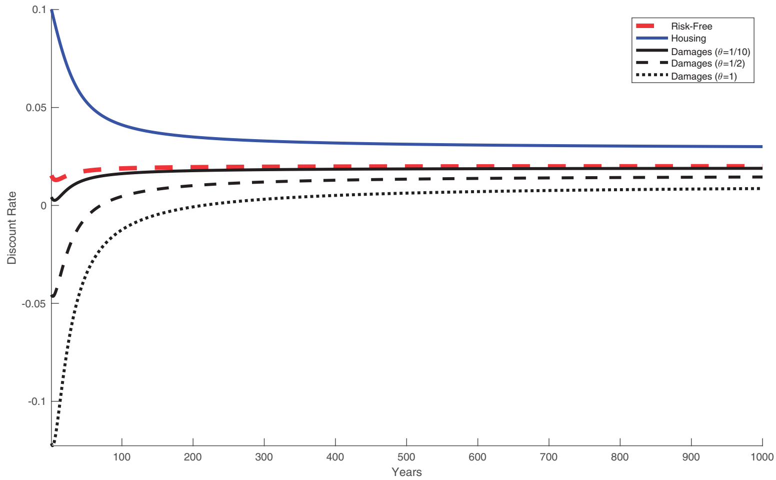

Figure 7
Discount rates for risk-free bonds, housing, damages
The figure shows the per-period discount rate corresponding to different assets for maturities 1 to 1,000 years, in our baseline model calibration. The top line represents the term structure of discount rates for the risky housing asset. The dashed line below it represents the real risk-free asset, that is, the real yield curve. The three black lines in the bottom represent different calibrations of the damage process with $\theta \in \{1,\ 0.5,\ 0.1\}$.

2.2.2 Calibration-implied housing term structure and climate risk elasticities. Figure 7 plots the term structure of discount rates that the calibrated model implies for risk-free and risky assets. The model is able to match the approximately flat term structure of risk-free rates observed in the data with an average level of around 1.0%. The model also produces a strongly declining term structure of discount rates for real estate, starting around 10% per year at short horizons and decreasing to around 3% per year at long horizons, matching the declining term structure of discount rates that we estimated for housing. To further assess how well we match our estimates of the real estate data, the two panels of Figure 8 report the leasehold price discounts estimated for housing in the United Kingdom and Singapore, together with the ones implied by the calibrated model, to highlight the close fit between the model and the data. The model also matches the average rate of return on housing (at around 5.5%) that we have independently estimated in the data.

Our model also helps rationalize the cross-sectional regularity that houses that are differentially exposed to climate risk have different price elasticities with respect to *news* about climate change. This is qualitatively consistent with the evidence reported in Section 1.1. That section focused on increase perception of climate risk, that is, future climate affecting future rents, but not current rents. Of course, current prices react immediately since they correspond to the present value of future rents. In our calibration, a house with a 1% higher exposure to climate change (i.e., higher η), responds to a one percentage point increase in the probability of a climate disaster (λ_t) with a price decline

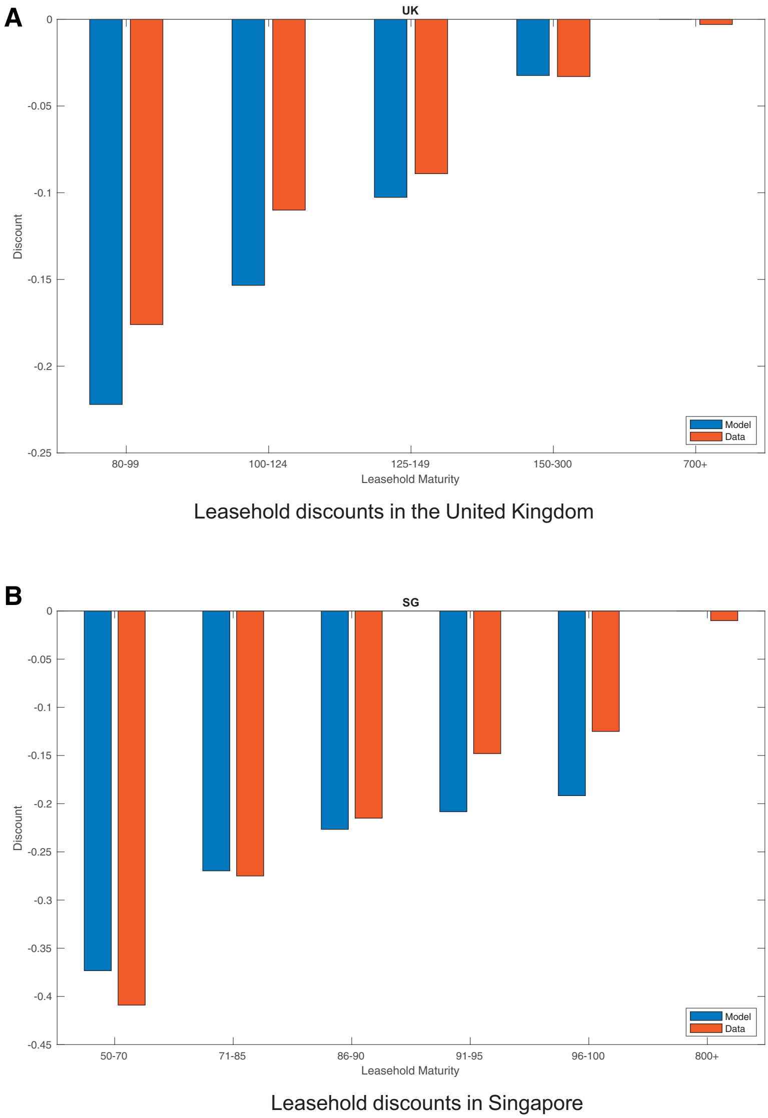

Figure 8
Leasehold discounts
The figure shows the discount rates for housing assets predicted by the model (left bars) and in the data (right bars), for the United Kingdom (upper panel) and Singapore (lower panel). The discounts are estimated from a hedonic regression and reported in log points.

that is 0.4 percentage points larger relative to a house with lower exposure. Unfortunately, these magnitudes are not directly comparable to the estimates from Section 1.1, since such a comparison would require us to map changes in our Climate Attention index to (perceived) changes in λ_t.

2.3 Valuing investments in climate change abatement

We start this section by using our calibrated model to derive appropriate discount rates for various types of investments in climate change abatement. Since our model nests the key ideas of a variety of standard models in climate change economics, we then proceed to show how our results compare to and improve upon the implications of two key views in climate change economics: the "tax" view, pioneered by Nordhaus and Boyer (2000) and Nordhaus (2008), and the "disaster view," pioneered by Weitzman (2012, 2014).

2.3.1 Valuing investments in the benchmark model. To derive the appropriate discount rates for investments in climate change abatement, we need to model their cash flows and their relation with the climate shocks. We model climate change investments as assets that compensate the investor for future damages to production and rents due to climate change, akin to insurance policies on climate change. Climate change mitigation investments in our framework are not large enough to affect equilibrium consumption; they are infinitesimal investments that are therefore informative about marginal valuations. We denote the process of damages to rents due to climate change by Q_t and model its (log) dynamics as

$$\Delta q_{t+1} = \mu_q - y_t + \eta J_{t+1}. \tag{16}$$

Intuitively, the occurrence of a climate disaster in our model induces an immediate destruction of rents ($\eta\xi$), but these damages revert over time as the economy adapts (captured by y_t).[21] Investments in climate change abatement provide a payoff that at least partially offsets the damages. Specifically, we assume that an insurance contract insures a fixed proportion θ of the growth rate of damages: $\theta \Delta q_{t+1}$. Values of θ range from 1 (full insurance) to close to zero (no insurance). Of course, it is possible to specify alternative types of climate change mitigation investments, for example, some that mitigate the long-run damages more strongly than in this specification. One advantage of our fully specified model is that it allows researchers to explore different types of climate interventions.

Figure 7 reports the appropriate discount rates for investments in climate change abatement of different maturity for three values of θ: 1, 0.5, and 0.1. Higher values of θ correspond to lower black lines.[22] The figure highlights a number of crucial results from our model:

1. **Appropriate discount rates for investments in climate change abatement are always below the risk-free rate.** This feature comes

[21] We set $\mu_q = \mu_d - 2\bar{\lambda}\eta\xi$ so that damages and rents have the same long-run growth rate.

[22] In some calibrations, the appropriate discount rates are negative, especially at shorter maturities. This is not surprising given the insurance contract nature of the investments; investors are simply willing to pay a price today that is above the expected payoff of the project.

from these investments being a hedge: they pay off following a climate disaster, and therefore in states of the world with high marginal utility. For relatively short horizons, we have estimated a real risk-free rate of about 1%, providing us with a tight upper-bound on the appropriate discount rates.

2. **At long horizons, the term structure of housing discount rates provides an upper bound on the appropriate discount rate.** While the risk-free rate provides a theoretically tight upper bound at all horizons, no reliable data exist on risk-free rates beyond horizons of about 30 years. This makes direct measurements of the risk-free-rate upper bound at horizons relevant for investments in climate change abatement infeasible. However, Section 1 described observed discount rates on risky housing for such long horizons, allowing us to bound the long-run discount rates for assets that are safer than real estate, including investments in climate change abatement, to be below 2.6%.

   Importantly, a discount rate below 2.6% (and even more so 1%) is lower than many 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). 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 (2015). 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. Moreover, 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 the 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.

3. **The term structure of discount rates for investments in climate change abatement is upward sloping, making the housing discount rates a tighter bound for longer horizons.** Appropriate discount rates for investments in climate change abatement increase with the horizon, which disproportionally tightens our upper bound as the horizon increases. This feature is driven by the same mean reversion in cash flows that generates the downward slope in the term structure of risky assets (such as real estate): since cash flows that are further in the future are

reduced less by a climate disaster, the benefits of reducing its effects are smaller, too.

Note that the implied *low but upward-sloping* term structure of discount rates for investments in climate change abatement contrasts with a number of papers that have argued for using declining discount rates for valuing investments in climate change abatement (Arrow et al. 2013; Cropper et al. 2014; Farmer et al. 2015; Traeger 2014). These arguments have motivated policy changes in France and the United Kingdom, which have adopted a downward-sloping term structure of discount rates for evaluating long-run investments, including those in climate change abatement. While this disagreement about the term structures does not have a substantial effect on the actual level of discount rates to value the *long-run cash flows*—they are relatively low, at approximately 2%, under both the term structures used in those countries and under our estimated upward-sloping term structure—the two rely on different economic mechanisms. Importantly, they also make substantially different predictions for evaluating the payoffs from abatement investments that may accrue at shorter horizons. The calibration of our model suggests that climate disasters cause the most damage immediately after they hit, making it most valuable to hedge the immediate costs. As a result, the discount rates are substantially below the risk-free rate of 1%-2%. In contrast, the downward-sloping term structures used in France and the United Kingdom suggest discount rates of 4% and 3.5%, respectively, for the first 30 years of a project's cash flows.

2.3.2 Alternative models: The "tax" view versus the "disaster" view of climate change. Modeling climate change risk and its effects on the economy is a daunting task, both because the physical processes driving climate change are not fully understood and because of the sparsity of historical data to predict how climate change will affect the economy. It is unsurprising, therefore, that the literature has approached the modeling of climate change and its effects on the economy in many different ways.

One view, pioneered by Nordhaus and Boyer (2000) and Nordhaus (2008), thinks of climate change akin to a tax on output. When output is high, pollution and the costs of climate change are also high. In this view, the main source of uncertainty about the future of climate is the future path of the economy. If the economy does well, pollution and climate change damages will be high; if the economy deteriorates, pollution and damages will be low. Investments in climate change abatement are thus risky, as they pay off in states of the world in which the economy is already doing well.

The alternative view follows Weitzman (2012, 2014): climate change is a disaster-type risk that, if it materializes, causes output to drop (see also, Barro 2015; Lemoine 2021; Wagner and Weitzman 2015). In this "disaster" interpretation, climate change itself represents the main source of uncertainty, and is itself a source of aggregate risk for the economy. Alternatively, this “disaster" view of climate change can also represent the case in which uncertainty about the future path of the economy (and not uncertainty about the climate per se) is the dominant source of uncertainty, but nonlinearities in the feedback from the economy to climate change are so pronounced that sufficiently high economic expansion might ultimately lead to a disaster (if a tipping point is reached). In these cases, investments in climate change abatement are then hedges that reduce aggregate risk, because they pay off when consumption is low (after a climate disaster materializes).

Our own calibrated model is a special case of this “disaster” view of climate change. However, our framework is general enough to nest both of these views and to shed light on the very different implications they have for the appropriate discount rates for investments in climate change abatement. To highlight this, Equation (17) presents a generalized version of the dynamics of climate damages (Equation (16) in our calibrated baseline specification):

$$\Delta q_{t+1} = \mu_q - \pi_q y_t + \eta_q J_{t+1}. \tag{17}$$

Different parameters of the model primitives under either the “tax view” or the “disaster view” map onto different values of μ_q, π_q, and η_q in this general specification. For example, by setting $\pi_q = 1$ and $\eta_q = \eta$, we recover our benchmark specification in Equation (16). To illustrate the discounting implications of the two different views, we compute the implied term structure of discount rates for a benchmark investment in climate change abatement that provides partial insurance. We assume the investment’s payoff process to follow $\Delta q_{t+1}/10$, thus hedging 10% of the innovation in climate change damages.

The basic “disaster” view of climate change. Our framework can be made consistent with the core of Weitzman’s original argument if we set the probability of a climate disaster to be constant ($\lambda_t = \bar{\lambda}$), remove the mean reversion in the economy ($x_t = y_t = 0$), and set $\pi_q = 0$ and $\eta_q = \eta$ in Equation (17). The climate-change-damages process then follows:

$$\Delta q_{t+1} = \mu_q + \eta J_{t+1}. \tag{18}$$

Figure 9 reports the term structure of discount rates for the climate-change-abatement investment described above in this Weitzman-type model (lowest solid line). We can see that the original Weitzman logic implies discount rates that are low, indeed lower than the risk-free rate, but also invariant across horizons. This invariance across horizons clearly conflicts with our evidence on horizon-dependent term structures of discount rates for assets exposed to climate risk (such as housing).

Relative to this original Weitzman view, our model adds two features that allow us to capture richer dynamics in climate change damages and to match the empirically-observed horizon-dependent term structure of discount rates: mean reversion (adaptation) in the economy, and climate risk that depends

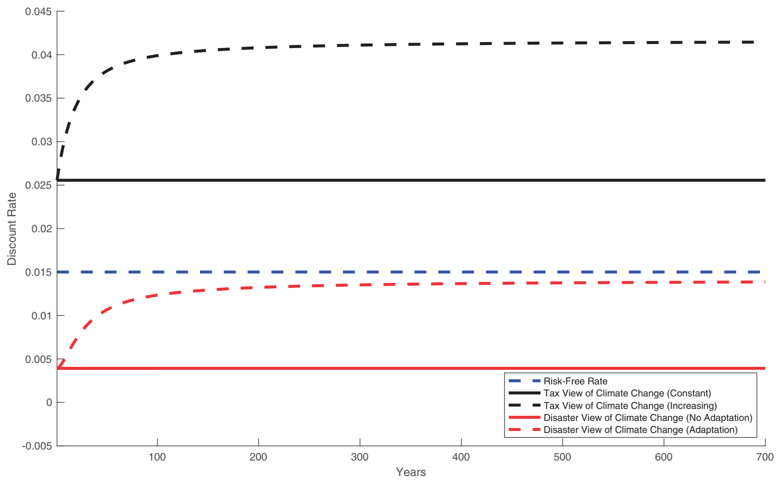

Figure 9
Leading models of climate change: Predictions for discount rates
The figure shows the per-period discount rate appropriate for climate-change-abatement investments under different models of climate change damages. In all these models, climate damages follow the process $\Delta q = \mu_q - \pi_q y_t + \eta_q J_{t+1}$; the discount rates in the figure correspond to those applied to an investment whose payoff is $\Delta q/10$. The "constant tax view" of climate change views damages as a constant fraction of output, so that $\mu_q = \mu_d$, $\eta_q = -\eta$, and $\pi_q = 0$. The "increasing tax view" views damages as a fraction of output that increases in good times, so that $\mu_q = \mu_d$, $\eta_q = -\eta$, and $\pi_q = k$. The "disaster view" with no mean-reversion views climate change damages as inducing a drop in output, so that $\mu_q = 0$, $\eta_q = \eta$, and $\pi_q = 0$. Finally, the "disaster view" with mean reversion corresponds to our baseline case, with $\mu_q = 0$, $\eta_q = \eta$, and $\pi_q = 1$.

endogenously on the growth rate of the economy as well as the occurrence of climate shocks. To illustrate, if we reintroduce mean reversion as in our benchmark calibration into this Weitzman economy, the climate-change-abatement investment starts to pay off whenever a climate disaster occurs (captured by the term ηJ_{t+1}), and continues to pay off at a declining rate in future periods (captured by the term $-\pi_q y_t$), reflecting higher economic growth due to adaptation. The lowest dashed line in Figure 9 indeed replicates our baseline results from Figure 7 and confirms that the discount rates for this investment are below the risk-free rate at all horizons, but increase with the horizon.

The basic "tax" view of climate change. We can also use our framework to explore the "tax" view of climate change elaborated on most prominently by Nordhaus. For exposition, we start by considering a simplified environment in which the tax rate that climate change imposes on the economy is constant, and the fundamental source of uncertainty stems from shocks to the economy. Such a setup corresponds to a linear damage function in the DICE model. The payoff to an investment in climate change abatement is then equivalent to the tax revenue from the climate tax, which can be captured by setting $Q_t = \tau D_t$, where τ is the climate tax rate. We keep all other processes in our economy unchanged, but remove the mean reversion ($x_t = y_t = 0$) to stay within the neoclassical-growth-model spirit of Nordhaus' DICE model.[23] Since the tax is constant, the payoff to climate-change-abatement investments behaves exactly like output. In particular, we can parameterize Equation (17) by setting $\mu_q = \mu_d$, $\pi_q = 0$, and $\eta_q = -\eta$. The process for damages from climate change then becomes

$$\Delta q_{t+1} = \Delta d_{t+1} = \mu_d - \eta J_{t+1}. \tag{19}$$

It follows immediately that investments in climate-change abatement in this setting are *risky*, since their payoff is positively correlated with consumption (see also Gollier 2013). This is reflected by the negative loading on J_{t+1} in the above equation ($\eta_q = -\eta$). Note that this loading is positive in the corresponding equation for the "disaster" view (Equation (18)); these different loadings are at the core of the starkly different predictions that these two views offer for discounting investments in climate change abatement. Indeed, Figure 9 shows that in the "tax" view (in which shocks to the economy are the fundamental source of uncertainty and the relationship between production and climate change is not very nonlinear), discount rates are high, above the risk-free rate, and invariant across horizons (solid black line). The first implication, high discount rates, is a key characteristic of the "tax" view of climate change. The second implication, a flat term structure, derives from our assumption of a constant tax rate.

A richer model in the spirit of Nordhaus (2008) allows for the tax rate to increase with economic activity, such that damages are disproportionally higher when the output of the economy is higher. Yet, the nonlinearities are not sufficiently strong to actually imply *lower* consumption in paths of high economic growth compared with paths of low economic growth (as in the tipping point literature captured by the "disaster" view). We capture the essence of this argument by assuming that tax proceeds follow $Q_t = \tau_{t-1} D_{g,t}$, where τ_{t-1} increases when output is high as specified below. We obtain that

$$D_t = D_{g,t}(1 - \tau_{t-1}), \tag{20}$$

where $D_{g,t}$ are rents in the absence of climate damages, which we refer to as gross rents, and D_t are net rents. We assume that gross rents follow $\Delta d_{g,t+1} = \mu_d - \eta J_{t+1}$ as before. Net of the climate tax, rents then follow

$$\Delta d_{t+1} = \Delta d_{g,t+1} + \left[\ln(1-\tau_t) - \ln(1-\tau_{t-1})\right] = \Delta d_{g,t+1} + y_t, \tag{21}$$

where $y_t = \left[\ln(1-\tau_t) - \ln(1-\tau_{t-1})\right]$ follows the same process as specified in Equation (9).[24] As we can see, this richer Nordhaus view still implies an output

[23] In Weitzman's work, and in our benchmark model, the shocks J_{t+1} are a direct manifestation of climate change disasters, and we parametrized them accordingly. In Nordhaus' work, climate change is a tax on the economy, and the shocks J_{t+1} are to be interpreted as not directly related to climate change (e.g., they may capture shocks to productivity instead). We focus on showing how the views of Nordhaus and Weitzman can be mapped into our model and highlight their starkly different predictions for discount rates here. Since the difference in predictions is stark in a qualitative sense already (i.e., different signs of the covariance of climate risk with consumption), we thought it best not to recalibrate shocks when analyzing the implications of Nordhaus' view in our framework.

[24] That is, we are implicitly defining the tax rate to follow $[\ln(1-\tau_{t+1}) - \ln(1-\tau_t)] - [\ln(1-\tau_t) - \ln(1-\tau_{t-1})] = \mu_y + \omega\{[\ln(1-\tau_{t1}) - \ln(1-\tau_{t-1})] - [\ln(1-\tau_{t-1}) - \ln(1-\tau_{t-2})]\} + \psi J_{t+1}$.

process that can (approximately) be nested in Equation (8) of our baseline model; the only difference lies in the interpretation of some of these processes. Note that this setup now also generates mean reversion in cash flows. However, while mean reversion in cash flows comes from adaptation to climate events in our baseline model, mean reversion is mechanically induced by the increasing schedule of the climate-change tax rate (τ) with respect to the level of economic activity in the present setup. As we will see, this leads to starkly different implications for discount rates.

Climate-change damages in this setup are given by $\Delta q_{t+1} = \Delta d_{g,t+1} + \left[\ln \tau_t - \ln \tau_{t-1}\right]$. These damages are similar to those in the simpler Nordhaus setup in Equation (19), but now include an extra term, $\left[\ln \tau_t - \ln \tau_{t-1}\right]$, that derives from time variation in the climate tax rate. To preserve the linearity and tractability of our model, we capture these damages in approximate form:[25]

$$\Delta q_{t+1} = \Delta d_{g,t+1} + \left[\ln \tau_t - \ln \tau_{t-1}\right] \approx \Delta d_{g,t+1} - k y_t = \mu_d - k y_t - \eta J_{t+1}. \qquad (22)$$

The above process is now a special case of Equation (17), in which $\mu_q = \mu_d$, $\pi_q = k$, and $\eta_q = -\eta$. As in the simpler constant-tax version of the Nordhaus view discussed above, investments in climate change abatement are risky (their payoffs are still positively correlated with output, $\eta_q = -\eta$). However, as shown in Figure 9, the increasing tax rate (captured by $-k y_t$) now induces the discount rates for climate-change-abatement investments not only to be high (above the risk-free rate) but also to be increasing with the horizon (dashed black line). Intuitively, when the economy does badly, expected climate damages are persistently low, thus making long-term investments in climate-change abatement even riskier than short-term investments.

2.4 Key takeaways

The evidence in Section 1 uncovered a downward-sloping term structure of discount rates for real estate, an asset that has substantial exposure to both consumption risk and climate risk. The general equilibrium model developed in this section is able to match this downward-sloping term structure of discount rates by leveraging a simple mechanism: mean reversion in cash flows as the economy adapts after a climate disaster. The implication of this mean reversion is declining risk exposures of higher-maturity cash flows, since a climate disaster that strikes today has larger effects on immediate cash flows than on distant cash flows.

Our modeling exercise has allowed us to establish a number of simple yet powerful results on appropriate discount rates for investments in climate change abatement that hedge disaster-type climate risks: (a) that these discount rates

25 We use the approximation $\ln \tau_t - \ln \tau_{t-1} \approx -k y_t$ and choose $k = \frac{1-\bar{\tau}}{\bar{\tau}}$. Recall that $y_t = \ln(1-\tau_t) - \ln(1-\tau_{t-1}) \approx \frac{\tau_{t-1} - \tau_t}{1-\tau_{t-1}}$, so that $\ln \tau_t - \ln \tau_{t-1} \approx \frac{\tau_t - \tau_{t-1}}{\tau_{t-1}} = -y_t \frac{1-\tau_{t-1}}{\tau_{t-1}}$. Since we do not want the loading on y_t to be time varying, we set $k = \frac{1-\bar{\tau}}{\bar{\tau}}$, where $\bar{\tau}$ is the steady-state tax rate.

are *bounded above* by the risk-free rate, (b) that for horizons at which we do not observe estimates of the risk-free rate, the estimated long-run discount rates for housing provide a relatively tight bound, and (c) that the upward-sloping nature of discount rates for investments in climate change abatement means that this bound gets tighter as the horizon lengthens. In addition, our calibrated model—which generates a term structure of discount rates at all horizons—can be used to value actual climate change mitigation investments.

We also show how our model implications on discount rates for climate-change-abatement investments improve upon and differ from two key views of climate change economics: the "tax" view of Nordhaus and Boyer (2000) and Nordhaus (2008), and the "disaster view" of Weitzman (2012, 2014). In particular, while appropriate discount rates in our model are *low and bounded above* by the risk-free rate as in Weitzman-type settings, they are also *upward sloping* and increasing toward the risk-free rate upper bound with horizon. This mirrors our empirical findings on the downward-sloping term structure of discount rates for risky real estate and reflects the investment's nature as a hedge. In Nordhaus-type settings by contrast, discount rates are *above* the risk-free rate, and on top of that also *increasing away* from the risk-free rate as the horizon increases when damages are a convex function of output.

3. Conclusion

In this paper, we showed how discount rates estimated from private markets, such as the housing market, can be informative about appropriate discount rates for investments in climate change abatement. While much is still unknown about the dynamics of climate change and its impacts on the economy, the seminal work by Nordhaus, Weitzman, Gollier, and others has substantially advanced our understanding of these issues. Our empirical and structural analysis contributes to this line of work, furthers our understanding of existing models, and provides new challenges for the next generation of models hoping to capture the interaction of climate change, asset markets, and the economy.

References

Alley, R. B., J. Marotzke, W. D. Nordhaus, J. T. Overpeck, D. M. Peteet, R. A. Pielke, R. T. Pierrehumbert, P. B. Rhines, T. F. Stocker, L. D. Talley, et al. 2003. Abrupt climate change. *Science* 299:2005–10.

Ambrose, B. W., P. Eichholtz, and T. Lindenthal. 2013. House prices and fundamentals: 355 years of evidence. *Journal of Money, Credit and Banking* 45:477–91.

Arrow, K., M. Cropper, C. Gollier, B. Groom, G. Heal, R. Newell, W. Nordhaus, R. Pindyck, W. Pizer, P. Portney, et al. 2013. How should benefits and costs be discounted in an Intergenerational context? The views of an expert panel. Discussion Paper, Stanford University.

Atreya, A., and S. Ferreira. 2015. Seeing is believing? Evidence from property prices in inundated areas. *Risk Analysis* 35:828–48.

Bakkensen, L. A., and L. Barrage. 2017. Flood risk belief heterogeneity and coastal home price dynamics: Going under water? Working Paper, University of Arizona.

Baldauf, M., L. Garlappi, and C. Yannelis. 2020. Does climate change affect real estate prices? Only if you believe in it. *Review of Financial Studies* 33:1256–95.

Bansal, R., and A. Yaron. 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59:1481–509.

Bansal, R., D. Kiku, and M. Ochoa. 2013. Climate change, growth, and risk. Working Paper, Duke University.

Barreca, A., K. Clay, O. Deschenes, M. Greenstone, and J. S. Shapiro. 2016. Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century. *Journal of Political Economy* 124:105–9.

Barro, R. J. 2006. Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics* 121:823–66.

———. 2015. Environmental protection, rare disasters, and discount rates. 82:1–23.

Barro, R. J., and J. F. Ursua. 2008. Macroeconomic crises since 1870. *Brookings Papers on Economic Activity* 39:255–350.

Barro, R., J. and T. Jin. 2011. On the size distribution of macroeconomic disasters. *Econometrica* 79:1567–89.

Belo, F., P. Collin-Dufresne, and R. S. Goldstein. 2015. Dividend dynamics and the term structure of dividend strips. *Journal of Finance* 70:1115–60.

Bernstein, Asaf, Matthew T Gustafson, and Ryan Lewis. 2019. Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134:253–72.

Bordo, M., B. Eichengreen, D. Klingebiel, and M. Soledad Martinez-Peria. 2001. Is the Crisis Problem Growing More Severe? *Economic Policy* 16:51–82.

Brohé, A., and M. Greenstone. 2007. The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather. *American Economic Review* 97:354–85.

Bureau of Economic Analysis. 2017. Concepts and methods of the U.S. national income and product accounts. Report, Washington, DC.

Burke, M., and K. Emerick. 2016. Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy* 8:106–40.

Cai, Y., K. L. Judd, and T. S. Lontzek. 2013. The social cost of stochastic and irreversible climate change. Working Paper, Ohio State University.

Campbell, S. D, M. A. Davis, J. Gallin, and R. F. Martin. 2009. What moves housing markets: A variance decomposition of the rent-price ratio. *Journal of Urban Economics* 66:90–102.

Chambers, D., C. Spaenjers, and E. Steiner. 2021. The rate of return on real estate: Long-run micro-level evidence. *Review of Financial Studies* 34:3572–607.

Constantinides, G. M., and D. Duffie. 1996. Asset pricing with heterogeneous consumers. *Journal of Political Economy* 104:219–40.

Cox, P. M., R. A. Betts, C. D. Jones, S. A. Spall, and I. J. Totterdell. 2000. Acceleration of global warming due to carbon-cycle feedbacks in a coupled climate model. *Nature* 408:184–7.

Cropper, M. L., M. C. Freeman, B. Groom, and W. A. Pizer. 2014. Declining discount rates. *American Economic Review* 104:538–43.

Crost, B., and C. P. Traeger. 2014. Optimal CO2 mitigation under damage risk valuation. *Nature Climate Change* 4:631–6.

Deschênes, O. and M. Greenstone. 2011. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics* 3:152–85.

Desmet, K., and E. Rossi-Hansberg. 2015. On the spatial economic impact of global warming. *Journal of Urban Economics* 88:16–37.

Drupp, M. A., M. Freeman, B. Groom, and F. Nesje. 2015. Discounting disentangled: An expert survey on the determinants of the long-term social discount rate. Working Paper, Centre for Climate Change Economics and Policy.

Eichholtz, P., E. Steiner, and E. Yönder. 2019. Where, when, and how do sophisticated investors respond to flood risk? Working Paper, Cornell University.

Eichholtz, P., Ma. Korevaar, T. Lindenthal, and R. Tallec. 2020. The total return and risk to residential real estate. Working Paper, Cornell University.

Engle, R. F., S. Giglio, B. Kelly, H. Lee, and J. Stroebel. 2020. Hedging climate change news. *Review of Financial Studies* 33:1184–216.

Farmer, J. D., J. Geanakoplos, J. Masoliver, M. Montero, and J. Perelló. 2015. Value of the future: Discounting in random environments. *Physical Review E* 91:052816.

Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh. 2017. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy* 125:140–223.

Flavin, M., and T. Yamashita. 2002. Owner-occupied housing and the composition of the household portfolio. *American Economic Review* 92:345–62.

Franklin, S. L., Jr., and R. S. Pindyck. 2018. Tropical forests, tipping points, and the social cost of deforestation. *Ecological Economics* 153:161–71.

Garnache, C. and T. Guilfoos. 2019. A city on fire? Effect of salience on risk perceptions. Working Paper, University of Oslo.

Gibson, M., J. T. Mullins, and A. Hill. 2017. Climate change, flood risk, and property values: Evidence from New York City. Working Paper, Williams College.

Giglio, S., B. T. Kelly, and J. Stroebel. 2020. Climate finance. Working Paper, Yale School of Management.

Giglio, S., M. Maggiori, and J. Stroebel. 2015. Very long-run discount rates. *Quarterly Journal of Economics* 130:1–53.

———. 2016. No-bubble condition: Model-free tests in housing markets. *Econometrica* 84:1047–91.

Gollier, C. 2013. Evaluation of long-dated investments under uncertain growth trend, volatility and catastrophes. Working Papers, Toulouse School of Economics.

Gourio, F. 2008. Disasters and recoveries. *American Economic Review* 98:68–73.

Greenstone, M., E. Kopits, and A. Wolverton. 2013. Developing a social cost of carbon for US regulatory analysis: A methodology and interpretation. *Review of Environmental Economics and Policy* 7:23–46.

Hallstrom, D. G., and V. K. Smith. 2005. Market responses to hurricanes. *Journal of Environmental Economics and Management* 50:541–61.

Harding, J. P., S. S. Rosenthal, and C. F. Sirmans. 2007. Depreciation of housing capital, maintenance, and house price inflation: Estimates from a repeat sales model. *Journal of Urban Economics* 61:193–217.

Hasler, M., and R. Marfe. 2016. Disaster recovery and the term structure of dividend strips. *Journal of Financial Economics* 122:116–34.

Huynh, T., and Y. Xia. 2020. Climate change news Risk and corporate bond returns. *Journal of Financial and Quantitative Analysis.* Advance Access published September 16, 2020, 10.1017/S0022109020000757.

Jordà, Ò., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor. 2017. The rate of return on everything, 1870-2015. Working Paper, Federal Reserve Bank of San Francisco.

———. 2019. The rate of return on everything, 1870–2015. *Quarterly Journal of Economics* 134:1225–98.

Kaplow, L., E. Moyer, and D. A. Weisbach. 2010. The social evaluation of intergenerational policies and its application to integrated assessment models of climate change. *B.E. Journal of Economic Analysis & Policy* 10:1–32.

Kurlat, P., and J. Stroebel. 2015. Testing for information asymmetries in real estate markets. *Review of Financial Studies* 28:2429–61.

Lemoine, D. 2021. The climate risk premium: how uncertainty affects the social cost of carbon. *Journal of the Association of Environmental and Resource Economists*, 8:27–57.

Lemoine, D., and C. P. Traeger. 2014. Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy* 6:137–66.

Lenton, T. M, H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf, and H. J. Schellnhuber. 2008. Tipping elements in the Earth's climate system. *Proceedings of the National Academy of Sciences* 105:1786–93.

Lettau, M., and J. A. Wachter. 2011. The term structures of equity and interest rates. *Journal of Financial Economics* 101:90–113.

Lucas, R. E, Jr. 1978. Asset prices in an exchange economy. *Econometrica* 46:1429–45.

Lustig, H. N. and S. van Nieuwerburgh. 2005. Housing collateral, consumption Insurance, and risk premia: An empirical perspective. *Journal of Finance* 60:1167–219.

Mack, A., and E. Martínez-García. 2011. A cross-country quarterly database of real house prices: A methodological note. Working Paper, Federal Reserve Bank of Dallas Globalization and Monetary Policy.

McCoy, S. J., and R. P. Walsh. 2018. Wildfire risk, salience & housing demand. *Journal of Environmental Economics and Management* 91:203–28.

McKenzie, R., and J. Levendis. 2010. Flood hazards and urban housing markets: The effects of Katrina on New Orleans. *Journal of Real Estate Finance and Economics* 40:62–76.

Melillo, J. M., S. D. Frey, K. M. DeAngelis, W. J. Werner, M. J. Bernard, F. P. Bowles, G. Pold, M. A. Knorr, and A. S. Grandy. 2017. Long-term pattern and magnitude of soil carbon feedback to the climate system in a warming world. *Science* 358:101–5.

Nakamura, E., J. Steinsson, R. Barro, and J. Ursúa. 2013. Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics* 5:35–74.

Nordhaus, W. D. 2008. *A question of balance: Weighing the options on global warming policies.* New Haven, CT: Yale University Press.

———. 2013. The climate casino: Risk, uncertainty, and economics for a warming world. New Haven, CT: Yale University Press.

Nordhaus, W. D., and J. Boyer. 2000. *Warming the world: Economic models of global warming.* Cambridge: MIT Press.

Ortega, F., and S. Taspinar. 2018. Rising sea levels and sinking property values: Hurricane Sandy and New York's housing market. *Journal of Urban Economics* 106:81–100.

Overpeck, J. T., and J. E. Cole. 2006. Abrupt change in Earth's climate system. *Annual Review of Environmental Resources* 31:1–31.

Painter, M. 2020. An inconvenient cost: The effects of climate change on municipal bonds. *Journal of Financial Economics* 135:468–82.

Piazzesi, M., M. Schneider, and S. Tuzel. 2007. Housing, consumption and asset pricing. *Journal of Financial Economics* 83:531–69.

Piketty, T., and G. Zucman. 2014. Capital is back: Wealth-income ratios in rich countries 1700-2010. *Quarterly Journal of Economics* 129:1255–310.

Reinhart, C. M., and K. S. Rogoff. 2009. *This time is different: Eight centuries of financial folly.* Princeton, NJ: Princeton University Press.

Schneider, M. T., C. P. Traeger, and R. Winkler. 2012. Trading off generations: Equity, discounting, and climate change. *European Economic Review* 56:1621–44.

Schularick, M., and A. M. Taylor. 2012. Credit booms gone bust: Monetary policy, leverage cycles, and financial crises. *American Economic Review* 102:1029–61.

Shiller, R. J. 2006. Long-term perspectives on the current boom in home prices. *Economists*' Voice 3:1–11.

Stern, N. 2006. Stern review on the economics of climate change. London, UK: Her Majesty's Treasury.

Stroebel, J. 2016. Asymmetric information about collateral values. *Journal of Finance* 71:1071–112.

Traeger, C. P. 2014. Why uncertainty matters: Discounting under intertemporal risk aversion and ambiguity. *Economic Theory* 56:627–64.

van Binsbergen, J., and R. S. J. Koijen. 2017. The term structure of returns: Facts and theory. *Journal of Financial Economics* 124:1–21.

van Binsbergen, J., M. Brandt, and R. Koijen. 2012. On the timing and pricing of dividends. *American Economic Review* 102:1596–618.

Wagner, G., and M. L. Weitzman. 2015. *Climate shock: The economic consequences of a hotter planet.* Princeton, NJ: Princeton University Press.

Weisbach, D., and C. R. Sunstein. 2009. Climate change and discounting the future: A guide for the perplexed. *Yale Law & Policy Review* 27:433–57.

Weitzman, M. L. 2012. Rare disasters, tail-hedged investments, and risk-adjusted discount rates. Working Paper, Harvard University.

———. 2014. Fat tails and the social cost of carbon. *American Economic Review* 104:544–46.