

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

Stefano Giglio* Matteo Maggiori† Krishna Rao‡ Johannes Stroebe§ Andreas Weber¶

Abstract

We explore what private market data can tell us about the appropriate discount rates for valuing investments in climate change abatement. We estimate the term structure of discount rates for real estate up to the very long horizons relevant for investments in climate change abatement. The housing term structure is downward-sloping, reaching 2.6% at horizons beyond 100 years. We also show that real estate is exposed to both consumption risk and climate risk. We explore the implications of these new data using a tractable asset pricing model that incorporates important features of climate change. Climate change is modeled as a rare catastrophic event, the probability of which increases with economic growth. Economic activity partially mean reverts following a climate disaster, capturing the ability of the economy to adapt. As a result, short-run cash flows are more exposed to climate risk than long-run cash flows, allowing us to match the observed housing term structure. The model and data provide simple yet powerful guidance for appropriate discount rates for investments that hedge climate disaster risk. The term structure of these discount rates is upward-sloping but bounded above by the risk-free rate. For extremely far horizons at which we do not observe the risk-free rate, the estimated long-run discount rates for housing (a risky asset) provide an upper bound that becomes tighter with maturity. This suggests that the appropriate discount rates for investments in climate change abatement are low at all horizons, substantially below those conventionally used for valuing these investments and for determining the social cost of carbon.

JEL Codes: G11, G12, R30.

Keywords: Climate Change, Discounting, Cost-Benefit Analysis, Real Estate

*Yale University, NBER, CEPR: stefano.giglio@yale.edu.

†Harvard University, NBER, CEPR: maggiori@fas.harvard.edu.

‡Zillow: krishnar@zillow.com.

§Stern School of Business, New York University, NBER, CEPR: johannes.stroebe@stern.nyu.edu.

¶Stern School of Business, New York University: aweber@stern.nyu.edu.

This version: March 2018. We thank Arthur van Benthem, Robert Barro, John Campbell, Thom Covert, Simon Dietz, Robert Engle, Xavier Gabaix, María Pilar Martínez-García, Kenneth Gillingham, Christian Gollier, Ben Groom, Bob Hall, Lars Hansen, Geoffrey Heal, Larry Karp, Derek Lemoine, Martin Lettau, Antony Millner, William Nordhaus, Monika Piazzesi, Nick Stern, Thomas Sterner, Christian Traeger, Martin Weitzman, and numerous conference and seminar participants. We are particularly grateful to Michael Greenstone for his encouragement and guidance. 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, 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.

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

1 Introduction

Much of the economics literature on the optimal policy response to climate change focuses on the trade-off between the immediate costs and the potentially uncertain long-run benefits of investments to reduce carbon emissions. Discount rates play a central role in this debate, since even small changes in discount rates can dramatically alter the present value of investments that pay off over very 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, 5% for instance, the present value of the investment drops to \$760 million and the investment no longer appears attractive.

Greenstone, Kopits and Wolverton (2013) remark that “the choice of a discount rate to be used over very long periods of time raises highly contested and exceedingly challenging scientific, economic, philosophical, and legal issues. As a result, there is no widespread agreement in the literature concerning the discount rates that should be used in an intergenerational context.” In this paper, we make progress on this question by exploring the information that private market discount rates contain about how to appropriately value investments in climate change abatement. First, we provide new empirical evidence on the term structure of discount rates for an important asset class, real estate, up to the extremely long horizons that are relevant for analyzing climate change (hundreds of years). Second, we combine these new empirical facts with insights from asset pricing theory to discipline the debate on the appropriate choice of discount rates for 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 has 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.¹ We show that this common practice of valuing investments in climate change abatement by discounting cash flows using the average rate of return to some traded asset ignores important considerations regarding the *maturity* and *risk properties* of such investments.

In particular, discount rates might vary with the horizon of an investment's payoffs either because the risk profile of the cash flows varies, or because households' sensitivity to risks that materialize over different horizons varies. 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. For example, a short-maturity investment will pay off sooner than a long-maturity investment. 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, whose cash flows and benefits tend to be delayed until much longer horizons and which might therefore 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 *horizon*. This of course is difficult in practice. However, as we show in this paper, it does not preclude researchers from extracting relevant information from the observed private market returns (like housing or bonds) and using this information 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.² Using a variety of empirical approaches, we estimate the average return to real estate

¹See also Kaplow, Moyer and Weisbach (2010), Schneider, Traeger and Winkler (2012), and Weisbach and Sunstein (2009) for discussions of the normative and descriptive approaches to discounting.

²Binsbergen, Brandt and Kojen (2012) provide evidence of a downward sloping term structure of discount rates for equities over a 1-10 year horizon.

to be around 6%. At the same time, recent estimates from Giglio, Maggiori and Stroebe (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 higher average (expected) returns and lower 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, thus implying a downward-sloping term structure.

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 very 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.³

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. We use 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-foot 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

³This is consistent with the average return to real estate of about 6%, which is above the real risk-free rate of 1%, and thus includes a risk premium to compensate investors for bearing risk.

change exposure. Instead, we test whether the prices of properties that are more exposed to climate change fall more 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 level. Our interpretation of this index is that it reflects households’ perceptions of the risk of future climate change on the cash flows from owning the property.

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 other comparable properties in the same zip code. 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. 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.

In order to explore the implications of the downward-sloping term structure of risky real estate for valuing climate change abatement, we build a simple and 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, tractable, and portable framework to show how the insights of modern asset pricing 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.

Our baseline model builds on the view that climate change is a form of disaster risk (Weitzman, 2012; Barro, 2013): it is a rare event with potentially devastating consequences for the economy. We embed this view in a general equilibrium model with a represen-

tative 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.⁴

First, we allow for 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; Lenton et al., 2008; Overpeck and Cole, 2006; Lemoine and Traeger, 2014; Franklin and Pindyck, 2017). These damages in turn might feed on themselves, for example because rising temperatures lead to even more gas emission for the same level of production. Our model captures these vicious cycles in reduced form by allowing the probability of a further disaster to increase after a disaster occurs (see Cox et al., 2000; Melillo et al., 2017).

Second, we allow for economic growth to pick up temporarily after a disaster. This 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, and investments (like air-conditioning rooms as temperatures rise) that allow for productive work despite adverse climate conditions (see, for example, Brohé and Greenstone, 2007; Deschênes and Greenstone, 2011; 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, as we discuss in detail in the paper, 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. Intuitively, while short-term cash flows are exposed to the full extent of the disaster (should it occur), long-term cash flows are less exposed since they 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 act as hedges: similar to insurance policies, they pay off primarily in bad states of the

⁴This 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 (2015), Lemoine and Traeger (2014), and Traeger (2014a, 2015).

world, and are thus particularly valuable. This has a number of important 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 in climate change abatement is the opposite of what we estimate for the term structure of housing, a risky asset. In fact, the term structure should be *upward-sloping*. In other words, 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. 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 (real bond yield) 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 upper bound is a simple yet powerful result that challenges a wide range of estimates previously used in the literature. For example, this bound is substantially below the 4% rate suggested by Nordhaus (2013). Quantitatively, it is more in line with long-run discount rates that are close to the risk-free rate, as suggested by Weitzman (2012), or the 1.4% suggested by Stern (2006), or results by Barro (2013). 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.

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, 2014b). These arguments have motivated policy changes in France and the U.K., which have adopted a downward-sloping term structure of discount rates for evaluating long-run investments, 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.⁵ In addition, they have substantially different implications for evaluating the payoffs from climate 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 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 U.K. suggest discount rates of 4% and 3.5%, respectively, for the first 30 years of a project's cash flows.

Finally, in addition to using our model to explore the discount rates appropriate for climate change mitigation within our disaster-risk view of climate change risk, we can use our model to also explore discounting of climate investments in alternative models of climate change risk. In particular, our specification for the economy and for the climate change dynamics is general enough that it can 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 climate change acts as a tax on output: damages are higher when the economy is doing well. In this view, climate change mitigation investments pay off mostly in *good* states of the world and they are, therefore, *risky* investments. Under this parameterization, the appropriate discount rates are *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 disproportionately higher during booming economies. In our framework, such a feature implies appropriate 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.

⁵The 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 strongly 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 premia. This is consistent with the guidance on discount rates provided by governments recommending declining discount rates for cost-benefit analysis, which usually do not indicate that the discount rate should vary with the risk properties of the investments.

2 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 up to horizons of 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.

Real estate also minimizes another concern with using private market rates to learn about valuing investments in climate change abatement. If individuals are not fully internalizing all benefits or costs of a certain private market investment, observing private discount rates related to that investment is not informative about the social discount rate a policymaker should apply in a normative analysis. We argue that buyers of real estate internalize the vast majority of the costs and benefits associated with the investment. After an appropriate risk and maturity adjustment, housing discount rates are therefore informative about private discount rates for other investments, for example those in climate change mitigation. Since the latter investments are typically thought of as carrying positive externalities, social discount rates should be bounded above by the private discount rates. We therefore view our estimates of climate change mitigation discount rates as an upper bound to the social discount rates for such investments. It is worth noting that a similar reasoning has been used in the literature to link hedonic analyses of house prices to the valuation that policymakers should assign to environmental amenities and dis-amenities (e.g., Chay and Greenstone, 2005; Greenstone and Gallagher, 2008).

2.1 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 lower prices for more exposed properties.⁶ Our analysis has to overcome a number of empirical challenges. First, when comparing prices of

⁶This analysis contributes to a recent literature that has explored the pricing of flood risk in real estate prices (e.g., Hallstrom and Smith, 2005; McKenzie and Levendis, 2010; Atreya and Ferreira, 2015; Bakkensen and Barrage, 2017; Bernstein, Gustafson and Lewis, 2017; Gibson, Mullins and Hill, 2017; Ortega and Taspinar, 2017). We believe we are the first to directly measure and study households' perceptions of climate risk in extremely localized real estate markets. As we describe below, this allows us to pursue tests of the pricing of climate risk in real estate that are not confounded by the potential for omitted variables.

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, beach front 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. Completely controlling for such hard-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 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.

2.1.1 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.⁷ 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-sales listings extend back to the first quarter of 2008.

Our second dataset 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

⁷While flooding risk is only one of a number of climate risk factors, it is certainly a first-order risk for properties in the coastal regions of the states we analyze. Focusing on rising sea levels and the associated flooding risk as our climate risk factor has the advantage that exposure to this risk is precisely measurable and can vary within small geographies. This allows us to compare the pricing of properties that are differentially exposed to climate risk, but that are otherwise very similar.

and bedrooms, as well as transaction prices and dates for all property sales.

To measure different properties' exposure to climate risk, we geo-code their addresses and map them to geographic shapefiles provided by the National Oceanic and Atmospheric Administration (NOAA) that indicate which regions will be flooded should sea levels rise by six feet or more. 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 I).

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. Appendix A.2 provides more details on the construction of the Climate Attention Index.

There is substantial spatial and time-series variation in this measure of climate risk attention. The top panel of Figure II provides a heatmap of the Climate Attention Index for Florida, pooling across all listings in our sample at the zip code level. 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 II 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 might be flooded if sea levels rise substantially. There is also substantial time-series variation in the Climate Attention index within geographies. For example, in New Jersey, the Climate Attention Index nearly triples between 2011 and 2013, around the time of Hurricane Sandy, which rendered more than 20,000 homes in the state uninhabitable.

2.1.2 Empirical Analysis

We next estimate how climate risk is priced in real estate markets. Our baseline hedonic regression is given by equation 1:

$$\log(\text{Price})_{i,h,g,t} = \alpha + \beta \log(\text{Index}_{g,t}) \times \text{FloodZone}_h + \gamma \text{FloodZone}_h + \delta Z_h + \phi_g \times \psi_t + \epsilon_{i,h,g,t} \quad (1)$$

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 Z_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, $\log(\text{Index}_{g,t})$, with the Flood Zone indicator, FloodZone_h . 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 I, 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 I, Panel A, we include property fixed effects in the regressions. In these specifications, the estimates of β are only identified off 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.

⁸While the coefficients on the control variables are not of primary interest in this study, Appendix A.2.2 shows that they are consistent with estimates from the literature: for example, larger and more recently upgraded homes trade at a premium.

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 might also 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 I 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 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

Key Take-Aways. The evidence provided in this section shows that real estate as an asset class has substantial exposure to climate risk, and therefore fulfills the first important criterion for us to use information from housing discount rates to learn about how to value investments in climate change abatement.

2.2 The Riskiness of Housing – Exposure to Consumption Risk

We next 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 when 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 consumption as well as personal disposable income and house prices.

Panel A of Figure III 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 dataset of historical house price indices for these countries.¹⁰ The beginning of a crisis is normalized to be time zero. The house price

⁹The 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. If instead these individuals choose to move into the rental market, this could be driving up rents.

¹⁰Appendix A.3.1 provides details of the crisis dates and the house price series. Reinhart and Rogoff (2008) also analyze real estate prices for 16 countries and 18 crises occurring in the period 1974-2008. We analyze real estate prices in 20 countries for 43 crises and 17 rare disasters occurring in the period 1870-2013.

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 III shows the average behavior of house prices during the rare consumption disasters as defined by Barro (2006). The consumption disaster dates for the 20 countries included in our historical house price index dataset are those defined by Barro and Ursua (2008). The dotted line tracks the level of consumption: 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.¹¹

Figure IV shows the time series of house prices and crisis years for the U.K. and Singapore – the countries with the best data on the term structure of discount rates (see below).¹² The pattern of house price movements during crises in these two countries is similar to the average pattern described above. House prices peak and then fall during major crises in the sample: the 1974-76 and 1991 banking crises in the U.K. as well as the 1982-83 banking crisis and the 1997 Asian financial crisis in Singapore.¹³ Similarly, both countries experience a drop in house prices during the 2007-08 global financial crisis.

¹¹Despite extensive efforts to collect an exhaustive database, our results are still limited by the relatively small number of crises for which house price data are available, and by the relatively low quality of house price time series before 1950. In addition, rental data is available only in extremely short time series. The lack of such rental data prevents us from performing a comprehensive study of the riskiness of the underlying cash flows of housing. For example, trying to detect the presence of long-run shocks for rent growth would even be challenging with time series that extend back for many more decades. 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 due to three effects: index smoothing, declining rents during bad times, and the destruction of housing stock during wars and natural disasters. (1) House price indices are generally smoothed and therefore underestimate the true variation in house prices. (2) We only consider the behavior of house price changes (capital gains) and have not considered the behavior of rents (dividends). For the two countries for which long high-quality time series of rental indices are available (France for the period 1949-2010 and Australia for the period 1880-2013), we find rent growth to be positively correlated with consumption growth (0.36 and 0.15 respectively). (3) A substantial part of the housing stock tends to be destroyed during wars. Therefore, the return to a representative investment in real estate would be lower than the fall in index prices as it would incorporate the physical loss of part of the asset.

¹²All crisis dates are from Reinhart and Rogoff (2009) except the periods 1997-98 and 2007-08 for Singapore. The latter dates have been added by the authors and correspond to the Asian financial crisis of 1997-98 and the global financial crisis of 2007-08.

¹³The 1984 banking crisis in the U.K. is the sole exception with increasing house prices.

Figure V shows the performance of house prices during major wars, namely World War I and II (WWI and WWII). In both cases, time zero is defined to be the start date of the war period, 1913 and 1939 for WWI and WWII, respectively. The dotted line tracks house prices of five countries for the duration of WWI (1913-1918).¹⁴ House prices fell throughout the war with a total fall in real terms of around 30%. Similarly, the solid line tracks house prices of six countries for the duration of WWII (1939-1945).¹⁵ House prices fell by 20% in real terms from 1939 to 1943 and then stabilized for the last two years of the war, 1944-45. Overall, we find wars to be periods of major declines in real house prices, which further contributes to the riskiness of real estate as an asset.

We also investigate the average correlation between consumption and house prices over the entire sample rather than just during crisis periods. Table II reports the correlation of house price changes with consumption changes 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 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 III). 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.

2.3 The Term Structure of Real Estate Discount Rates

In this section, we provide new evidence on an important and previously unexplored dimension of real estate data: the term structure of housing discount rates. The first two sections present our analysis of expected real estate returns, which we find to be relatively high, between 5.5% and 7.3%. In Section 2.3.3, we combine these new data with

¹⁴Due to data availability for house price indices during this period, the countries included are Australia, France, Netherlands, Norway, and the United States.

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

the estimates of Giglio, Maggiori and Stroebe (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 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.

2.3.1 Average Rate of Return to Housing and Rental Growth Rate – Measurement

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:

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

where P is the price index and D the rental index.¹⁶ Gross real housing returns are then:

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

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

$$E[R] = E[R^G] - \delta - \tau.$$

The second approach, which we label the *balance-sheet approach*, follows Favilukis, Ludvigson and van Nieuwerburgh (2010) 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).¹⁷ Since

¹⁶Notice that given a baseline price-rent ratio, only information about the growth rates in prices and rents, but no information on levels is necessary for these calculations.

¹⁷To determine the total net income from the consumption of real estate services, these measures impute the value of owner-occupied equivalent rents, the real estate services consumed by individuals from living in their own house. See Mayerhauser and Reinsdorf (2006) and McCarthy and Peach (2010) for details.

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:¹⁸

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

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.¹⁹

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 indices, which allows us to directly control for quality as well as “pure” quantity changes in the housing stock at the same time.²⁰ For the U.S., 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.

2.3.2 Average Rate of Return to Housing and Rental Growth Rate – Estimates

We estimate the return to housing for three countries.²¹ We explore data from the United Kingdom and Singapore, since we are able to directly measure very long-run discount

¹⁸Since we are interested in the returns to housing as an unlevered asset, we do not have to lever capital gains or correct net capital income on housing for mortgage interest. Given the riskiness of housing we document in the following section, levered returns should be higher than unlevered returns.

¹⁹Note that discounting the rental yield with stock growth is rather conservative – it effectively assumes that all rental income is earned on the $t + 1$ -housing stock.

²⁰Quantity indices are a widely used concept in national accounts and aim to capture changes in the value of an asset that are not driven by (constant-quality) price changes. Quality changes are treated as changes in quantity in such a decomposition. See Bureau of Economic Analysis (2017) for more details on quantity indices, for example.

²¹For the balance-sheet approach, national accounts typically provide information on the housing stock split by sectors and we focus on housing returns to owner-occupied and non-corporate tenant-occupied properties. Owner-occupied and non-corporate tenant-occupied housing account for the vast majority of housing wealth in most countries. In the U.S., for example, these sectors accounted for more than 90% of housing wealth over our sample period. Excluded are the public and non-profit sectors, due to their non-commercial nature, as well as the corporate sector, due to the relative concentration of its housing portfolios in non-representative urban centers. See Appendix A.4.2 for more details.

rates for these countries (see Section 2.3.3); we also provide estimates for the U.S. for comparison, since they have been the subject of an extensive literature (e.g., Flavin and Yamashita, 2002; Gyourko and Keim, 1992; Lustig and van Nieuwerburgh, 2005; Piazzesi, Schneider and Tuzel, 2007). Most data are sourced from each country’s various national statistical agencies. The U.S. price index is from Shiller (2000). We obtain a benchmark price-rent ratio for the U.S. from Trulia, and estimate a benchmark price-rent ratio using data from iProperty.com (Asia’s largest online property listing portal) and a similar methodology for Singapore. For the U.K., we rely on estimates from Bracke (2015). All results are robust to using benchmark price-rent ratios implied by the balance-sheet approach instead of our direct estimates. Appendix A.4 discusses further details on these procedures and the underlying data sources.

United States. The top panel of Table IV presents the estimated average real estate returns. Besides our main specifications, we also provide robustness checks that use alternative price and rental indices 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. Our preferred results are based on the baseline price-rent ratio from Trulia including a utilities correction (column 2), and on the balance-sheet approach that relies on direct holding gains from the revaluation accounts (column 9).²²

Both approaches provide similar estimates for the average annual real gross return ($E[R^G]$): 9.7% based on the price-rent approach and 8.9% based on 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 Malpezzi, Ozanne and Thibodeau (1987).²³ Our balance-sheet estimates imply a tax impact of 1.1% and we assume a prop-

²²Our benchmark price-rent ratio from Trulia might include rental properties where some utilities are covered by the monthly rent. Using data from the balance-sheet approach, we adjust the Trulia price-rent ratio for that possibility in this specification (see Appendix A.4.1 for details). We prefer the balance-sheet approach that relies on direct holding gains from the revaluation accounts over the one that relies on a quantity index for the U.S., since the holding gains are derived from the same data sources as our measure of housing wealth, which guarantees an exact sectoral match and thus internal consistency (see Appendix A.4.2 for details).

²³Harding, Rosenthal and Sirmans (2007) find that between 1983 and 2001, real estate (housing) depreciated at roughly 2.5% per year gross of maintenance. Malpezzi, Ozanne and Thibodeau (1987) provide an overview of the earlier literature on depreciation. For example, Leigh (1980) estimates the annual depreciation rate of housing units in the U.S. to be between 0.36% and 1.36%. Depreciation is also a key calibration parameter for much of the recent literature in macroeconomics that considers households’ portfolio and consumption decisions with housing as an additional asset. Cocco (2005) chooses an annual depreciation rate equal to 1%; Díaz and Luengo-Prado (2008) use an annual depreciation rate of 1.5%.

erty tax impact of 0.67% for a representative household for the price-rent approach.²⁴ This results in average real net returns between 5.5% and 6.5% for the U.S. 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 Favilukis, Ludvigson and van Nieuwerburgh (2010), who find a real return of 9 – 10% before netting out depreciation and property taxes.

United Kingdom. Columns 11 to 15 of Table IV report the estimates for the real estate market in the U.K. The price-rent and the balance-sheet approaches provide similar estimates for the average annual real gross return ($E[R^G]$) again, 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 U.K. Average real net returns in the U.K. real estate market are therefore between 7.0% and 7.3%.

Singapore. Column 16 in Table IV 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 U.S., and the property tax impact to be 0.6%.²⁵ 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: Firstly, 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).²⁶ Unfortunately, the national accounts data do not allow us to separate these out with sufficient accuracy. Secondly, the national accounts data do not allow us

²⁴Property taxes in the U.S. are levied at the state level and, while there is variation across states, are generally around 1% of house prices. Property taxes, however, are deductible from federal income tax. We assume that the deductibility reflects a marginal U.S. federal income tax rate of 33%. The net impact is therefore $(1 - 0.33) * 0.01 = 0.67\%$.

²⁵Before 2003, Singapore levied a 10% annual tax on the estimated rental income of the property. A lower tax rate of 4% applied to owner-occupied properties. Starting in 2011 for owner-occupiers and in 2014 for landlords, Singapore has introduced increasingly progressive tax schemes that start at 0% and cap out at 16% for owner-occupiers, and start at 10% and cap out at 20% for landlords since 2015. Even though homeownership rates are around 90% during our sample period (numbers based on series *M810401 - Resident Households By Tenancy* as published by Statistics Singapore on its website <http://www.singstat.gov.sg>), we use the more conservative (higher) rate of 10% for rental properties as it has prevailed for most of our sample period. The tax impact on returns is the tax rate times the average rent-price ratio, estimated at around 6%. Hence, $\tau = 0.1 * 0.06 = 0.6\%$. For details, see <https://www.iras.gov.sg/irashome/Property/Property-owners/Working-out-your-taxes/Property-Tax-Rates-and-Sample-Calculations/>, last accessed February 2018.

²⁶Numbers are based on series *M400751 - Residential Dwellings* as published by Statistics Singapore on its website <http://www.singstat.gov.sg>, last accessed February 2018.

to determine the total consumption of real estate services excluding relevant costs, that is net rents, with sufficient accuracy.

Average Rate of Return: Summary. Overall, these estimates show that real expected returns for real estate are around 6% or higher for the countries we consider. These estimates are robust to the different methodologies we use.²⁷ They are also in line with more recent work from Jordà et al. (2017) that extends our work to 16 countries and finds average returns to housing of around 7% before taxes.²⁸ 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 2.2, 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 from the same data sources, which we denote by g . The estimated real growth rate of rents is low for the U.K., Singapore, and the U.S., which we add again to compare our estimates with those in the existing literature. For the U.S., 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 U.K. and a slightly lower estimate of $g = -0.4\%$ for Singapore.²⁹ Overall, 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

²⁷Also note that most movements in rent-price ratios are driven by movements in house prices and not by movements in rents (see Shiller, 2007); therefore, our estimates of returns are relatively less affected by the time periods chosen. Take the U.S. as an example: by 2013 rent-price ratios in the U.S. had declined to their 2000 levels approximately. For example, ending our sample in 2005 vs. 2012 would have produced a slightly lower average rent-price ratio. However, at the same time, this would exclude the house price crash after 2005, thus leading to higher estimated average capital gains. In the overall estimates of expected returns, the higher estimated capital gains would be partially offset by a lower estimated rent-price ratio.

²⁸Interestingly, Jordà et al. (2017) also show that these returns are only marginally lower than equity returns in the post-1950s period that we consider.

²⁹The relatively high rent growth rate in the U.K. is driven by high growth rates in the late 80s and early 90s. Since capital gains were also low during this period, excluding this time period would increase our estimates for U.K. housing returns. It should be noted that the U.K. changed its methodology for tracking prices, including rents, from the Retail Price Index (RPI) to the Consumer Price Index (CPI) in 1995 and that pre-1995 rental growth rates are measured with greater uncertainty as a consequence. The slightly negative rent growth rate in Singapore is driven by a deflationary period in the early 2000s and another deflationary period since 2012.

estimates long-run real house price growth rates to be very low, often below 1%.³⁰

Key Take-Aways. To recap, our estimates indicate an average real net rate of return to real estate of around 7% for the U.K. and Singapore. These results are consistent with complementary estimates of real estate returns between 5.5% and 6.5% for the United States. We also find an average real growth rate of rents slightly below zero for Singapore and around 1.4% for the U.K. Again, both results are consistent with complementary estimates of rental growth rates of 0.7% for the United States.

2.3.3 Long-Run Housing Discount Rates

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

Figure VI reports the estimates from Giglio, Maggiori and Stroebel (2015). It shows the price discount of leaseholds with varying maturities compared to freeholds for otherwise identical properties. For the U.K. estimates, for example, the bucket with leaseholds of remaining maturity between 100 and 124 years shows that households are willing to pay 11% less for a leasehold with that maturity than for a freehold. Interpreted differently, at least 11% of the value of a freehold property is due to cash flows that accrue more than 100 years into the future. In general, leasehold discounts are strongly associated with maturity, with shorter leaseholds trading at bigger discounts: 17.6% for leaseholds with remaining maturity of 80-99 years, 11% for remaining maturities of 100-124 years, 8.9% for remaining maturities of 125-149 years, and 3.3% for remaining maturities of 150-300 years. Leaseholds with more than 700 years remaining maturity trade at the same price as freeholds. Pricing patterns are similar for properties in Singapore. The authors provide

³⁰The equivalence of these two long-run growth rates is necessary for rental yields to be stationary. Eichholtz (1997), Eitrheim and Erlandsen (2005) and Ambrose, Eichholtz and Lindenthal (2013) also confirm Shiller's observation of negligible long-run real house price growth in different countries using different data. Ambrose, Eichholtz and Lindenthal (2013) find evidence of cointegration between house prices and rents using very long-run real estate data for Amsterdam.

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.

2.3.4 Take-away: The Term Structure of Discount Rates in the Housing Market

In this section we showed (i) that real estate has a real expected rate of return of above 6% per year, and (ii) 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.

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

The previous section provided empirical 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 empirical 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' qualitatively and quantitatively different

predictions for the discount rates of investments in climate change abatement.

3.1 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. In order 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. It also allows us to nest the key assumptions of the climate models by Nordhaus (2008) and Weitzman (2012, 2014), something we discuss in greater detail in Section 3.3.

3.1.1 Model Setup

We assume that aggregate consumption follows:

$$\Delta c_{t+1} = \mu + x_t - J_{t+1}, \quad (2)$$

$$x_{t+1} = \mu_x + \rho x_t + \phi J_{t+1}, \quad (3)$$

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. J_t is a jump process that takes value $\zeta \in (0, 1)$ 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. We allow the probability of a climate shock, λ_t , to be time-varying, and in particular to depend endogenously on the dynamics of the economy.³¹ The process x_t captures persistent changes in the growth rate of the economy; as we discuss below, x_t plays an important 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 (here: 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

³¹Our 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, but this would result in a less tractable benchmark model.

one for aggregate consumption:

$$\Delta d_{t+1} = \mu_d + y_t - \eta J_{t+1}, \quad (4)$$

$$y_{t+1} = \mu_y + \omega y_t + \psi J_{t+1}. \quad (5)$$

The main difference between real estate rents and consumption is the larger exposure of rents to the underlying economic shocks (here: climate change 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.

Our setup allows for partial mean reversion in the growth rate of consumption and rents after a disaster. As we show below, this partial mean reversion plays a crucial role in explaining the term structure of discount rates for risky assets. Nakamura et al. (2013) were the first to highlight this prominent feature of the data. It has become a standard ingredient of asset pricing models aiming to explain downward-sloping term structures of discount rates in various asset markets and found continued empirical support (Gourio, 2008; Lettau and Wachter, 2011; Belo, Collin-Dufresne and Goldstein, 2015; Hasler and Marfe, 2016). In the context of climate risk modeling, this feature captures the notion that economic activity picks up after a climate disaster as the economy adapts to new climatic circumstances. Numerous papers have highlighted 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). Formally, after a disaster strikes, the growth rate of the economy increases ($\psi > 0, \phi > 0$) and this increase is persistent ($\rho > 0, \omega > 0$).

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

$$\lambda_{t+1} = \mu_\lambda + \alpha \lambda_t + \nu x_t + \chi J_{t+1}. \quad (6)$$

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 a particularly useful way to bring climate risk into an asset pricing framework:

1. The disaster probability λ_t is an endogenous function of the growth rate of the economy: Remember that x_t is the expected deviation of the growth rate of the economy from trend. Since it enters additively and positively ($\nu > 0$) in equation 6, the probability of a disaster increases 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; Lenton et al., 2008; Overpeck and Cole, 2006; Lemoine and Traeger, 2014; Franklin and Pindyck, 2017).

2. The climate disaster probability λ_t increases following the occurrence of a disaster ($\chi > 0$), thus allowing for vicious spirals in climate change: climate shocks induce a self-reinforcing cycle in which each shock increases the probability of the next shock (see, for example, Cox et al., 2000; Melillo et al., 2017).

To illustrate the richness of these patterns, Figures VII and VIII plot two sample paths of the economy. Figure VII shows a path in which no disaster occurs, but the economy grows above trend for a sustained period of time, starting in year 10. 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.³² 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 spur. 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 reaches its maximum approximately 7 years after the growth spur 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 VIII 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 spur shock for the probability of a disaster to revert back to its long-run mean. The bottom panel of Figure VIII also illustrates the mean reversion in the growth rate of the economy. After a disaster strikes, the growth rate of the economy increases

³²As we discuss in greater detail in Section 3.2, 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.

($\psi > 0$) and this increase is persistent ($\omega > 0$).

3.1.2 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 entire term structure of housing risk premia 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) and Barro (2006), and is given by:

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

where δ is the rate of time preference. In 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 real estate discount rates. Formally, we are interested in the per-period discount rate of maturity n , denoted \bar{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:³³

$$P_t^{(n)} = \frac{E_t[D_{t+n}]}{(1 + \bar{r}_t^n)^n}. \quad (7)$$

As we show formally in Appendix A.5, the term structure of discount rates \bar{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 $R_{t,t+1}^{(n)} = \frac{P_{t+1}^{(n-1)}}{P_t^{(n)}}$. More formally:³⁴

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

For example, a flat term structure of discount rates implies a flat term structure of expected one-period returns across maturities, and vice versa.

³³Note that we label objects that relate to single cash flows at a specific maturity n with 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 dividends or 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 .

³⁴See Appendix A.5.1 for a derivation. The result holds exactly when the term structure of discount rates is constant over time t (though it can have any shape over maturities n).

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 Appendix A.6 provides the full solution of the model, and while the calibrated results presented below use this full solution, we 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:

$$\begin{aligned} E_t[R_{t,t+1}^{(n)}] - R_{t,t+1}^f &\simeq \gamma \text{Cov}_t[r_{t,t+1}^{(n)}, \Delta c_{t+1}] \\ &= \gamma [\eta - \psi e_{d,n-1} - \phi b_{d,n-1} - \chi f_{d,n-1}] \xi^2 \lambda_t (1 - \lambda_t). \end{aligned} \quad (9)$$

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 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 Appendix A.6.

Equation 9 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 have a differential effect on the risk premia of cash flows with different maturities (there are no n -subscripts).

The *shape* of the term structure is determined by the terms inside the square 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 premia is $\psi e_{d,n-1}$, which captures the term structure of exposures of claims of different maturity to a disaster (and the ensuing recovery). The model is parsimonious enough to admit an analytical solution for $e_{d,n}$ (see Appendix equation A.14d):

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

Since this term enters negatively in equation 9, positive values for both ψ and ω imply a declining term structure of risk premia 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 farther in the future. As discussed above, mean reversion in cash flows captures the ability of the economy to (partially) adapt to climate disasters over time in our context, a feature supported by existing empirical research.

Remark on Preferences. The previous discussion highlights that, in our setting, the observed downward-sloping term structure of risk premia 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 literature, could also generate this downward slope. We discuss in 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 upwards. 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.³⁵ More generally, we are not aware of a standard representation of preferences that would push towards 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.

3.2 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 an 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 empirical evidence on the risk and return properties of real estate, including the term structure of discount rates from Section 2.

³⁵The 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.

3.2.1 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 2. 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 IX. 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.³⁶ In our calibration, rents are not only correlated with consumption, but also share many of its dynamics, including recovery after disasters ($\psi > 0$) 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 2. Finally, we set the steady-state conditional probability of disasters, $\bar{\lambda}$, 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 (which is directly affected by the disaster probability dynamics through the precautionary savings channel). 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.³⁷

³⁶Figure IX plots the average shape of the real U.K. gilts curve for the period 1998-2016, as well as for two sub-periods: 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 where data on such long maturities are available, the yield curve is essentially flat for these longer maturities as well.

³⁷The resulting 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 Appendix A.6 for details.

3.2.2 Calibration-Implied Housing Term Structure and Climate Risk Elasticities

Figure X 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 XI report the leasehold price discounts estimated for housing in the U.K. 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 generates the cross-sectional prediction that houses that are differentially exposed to climate risk have different price elasticities with respect to *news* about climate change, again consistent with the new evidence reported in Section 2.1. 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 that is 3.9 percentage points larger relative to a house with lower exposure. In the next section, we use our calibrated model to derive appropriate discount rates for investments in climate change abatement.

3.3 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).

3.3.1 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:

$$\Delta q_{t+1} = \mu_q - y_t + \eta J_{t+1}. \quad (10)$$

Intuitively, the occurrence of a climate disaster in our model induces an immediate destruction of rents ($\eta\tilde{\zeta}$), but these damages revert over time as the economy adapts (captured by y_t).³⁸ 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).

Figure X 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.³⁹ The figure highlights a number of crucial results from our model:

1. **The discount rate is always below the risk-free rate.** Appropriate discount rates for investments in climate change abatement are always below the risk-free rate, across all horizons. This feature comes from the investment being a hedge – it pays 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 good 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, we do observe discount rates on risky housing for such long horizons, and our results from Section 2 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%.

³⁸We set $\mu_q = \mu_d - 2\bar{\lambda}\eta\tilde{\zeta}$ so that damages and rents have the same long-run growth rate.

³⁹In some calibrations the appropriate discount rates are negative, especially at shorter maturities. This is not surprising given the insurance-contract nature of the investments; it simply means that investors are willing to pay a price today that is above the expected payoff of the project.

Importantly, a discount rate below 2.6% (and even more so 1%) is lower than numerous estimates used in the existing literature and by policymakers for discounting investments in climate change abatement. For example, it is substantially below the 4% suggested by Nordhaus (2013). 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),⁴⁰ or results by Barro (2013). It is also close to the average recommended long-term social discount rate of 2.25% elicited by Drupp et al. (2015) in a survey of 197 experts, and falls within the range of 1% to 3% that more than 90% of these experts are comfortable with. 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 CO₂ 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 CO₂ 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 (see also Pizer et al., 2014).

3. **The term structure of the climate mitigation discount rate 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 farther 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, 2014b). These arguments have motivated policy changes in France and the U.K., which have adopted a downward-sloping term structure of discount rates for evaluating long-run investments, 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

⁴⁰See also Stern (2008) for comments on the discussion of the Stern Review’s approach to discounting.

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 climate-change-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 U.K. suggest discount rates of 4% and 3.5%, respectively, for the first 30 years of a project’s cash flows.

3.3.2 Alternative Models: The “Tax” View vs. 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 if large temperature changes were to occur. It is unsurprising, therefore, that the literature on the economics of climate change 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. In this view, when output is high, pollution and the costs of climate change are also high. Investments in climate change abatement are thus risky, as they pay off in states of the world when the economy is already doing well.

An alternative view, following Weitzman (2012, 2014), presents climate change as a disaster-type risk that, if it materializes, causes output to drop (see also Barro, 2013; Lemoine, 2015; Wagner and Weitzman, 2015). In this “disaster” interpretation, climate change itself is a source of aggregate risk for the economy. Investments in climate change abatement are thus hedges that reduce aggregate risk. Our own calibrated model is a special case of this view of climate change.

Our framework allows us 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 11 presents a generalized version of the dynamics of climate damages (equation 10 in our calibrated baseline specification):

$$\Delta q_{t+1} = \mu_q - \pi_q y_t + \eta_q J_{t+1}. \quad (11)$$

Different parameters of the model primitives under either the “tax view” or the “disaster

view” map into 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 10.

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 “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 11. The climate-change-damages process then follows:

$$\Delta q_{t+1} = \mu_q + \eta J_{t+1}. \quad (12)$$

Figure XII 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 endogenously on the growth rate of the economy as well as the occurrence of climate shocks. To illustrate, if we re-introduce 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 XII indeed replicates our baseline results from Figure X and confirms that the discount rates for this investment are below the risk-free rate at all horizons, but increase with the horizon.

The “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. 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.⁴¹ Since the tax is constant, the payoff to climate-change-abatement investments behaves exactly like output. In particular, we can parameterize equation 11 by setting $\mu_q = \mu_d$, $\pi_q = 0$, and $\eta_q = -\eta$. The process for damages from climate change then becomes:

$$\Delta q_{t+1} = \Delta d_{t+1} = \mu_d - \eta J_{t+1}. \quad (13)$$

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 12); 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 XII shows that in the “tax” view, 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 a tax rate that increases with the level of economic activity, such that damages are disproportionately higher when the output of the economy is higher. 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

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

where $D_{g,t}$ are rents in the absence of climate change damages, which we refer to as gross

⁴¹In 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.

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:

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

where $y_t = [\ln(1 - \tau_t) - \ln(1 - \tau_{t-1})]$ follows the same process as specified in equation 5.⁴² As we can see, this richer Nordhaus view still implies an output process that can (approximately) be nested in equation 4 of our baseline model; the only difference lies in the interpretation of some of these processes. Note that this setup now also implies 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 this Nordhaus setup. As we will see, this leads to starkly different implications for discount rates.

Climate-change damages in this setup are given by: $\Delta q_{t+1} = \Delta d_{g,t+1} + [\ln \tau_t - \ln \tau_{t-1}]$. These damages are similar to those in the simpler Nordhaus setup in equation 13, but now include an extra term, $[\ln \tau_t - \ln \tau_{t-1}]$, 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:⁴³

$$\Delta q_{t+1} = \Delta d_{g,t+1} + [\ln \tau_t - \ln \tau_{t-1}] \approx \Delta d_{g,t+1} - ky_t = \mu_d - ky_t - \eta J_{t+1}. \quad (14)$$

The above process is now a special case of equation 11, 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 XII, the increasing tax rate (captured by $-ky_t$) now induces 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.

⁴²That is, we are implicitly defining the tax rate to follow the process $[\ln(1 - \tau_{t+1}) - \ln(1 - \tau_t)] - [\ln(1 - \tau_t) - \ln(1 - \tau_{t-1})] = \mu_y + \omega\{[\ln(1 - \tau_{t1}) - \ln(1 - \tau_{t-1})] - [\ln(1 - \tau_{t-1}) - \ln(1 - \tau_{t-2})]\} + \psi J_{t+1}$.

⁴³We use the approximation $\ln \tau_t - \ln \tau_{t-1} \approx -ky_t$ 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.

3.4 Key Take-Aways

The evidence in Section 2 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 we 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 to new climatic circumstances 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.

This 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: (i) that these discount rates are *bounded above* by the risk-free rate, (ii) 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 (iii) 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 – that generates an entire term structure of discount rates at all horizons – can be used to value different types of climate change mitigation investments, that vary by the maturity of their cash flows and by the amount of mitigation they provide.

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 towards 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.

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 climate change economics models.

4 Conclusion

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. Our findings highlight that for such an analysis to be useful, one needs to consider the entire term structure of discount rates as well as the risk properties of the assets from which discount rates are inferred. We also showed how the implied discount rates for investments in climate change abatement depend crucially on their assumed risk properties in turn.

While our estimates of housing discount rates provide an important first look at the term structures of risky assets over extremely long horizons – and therefore at the appropriate term structure of discount rates for valuing investments in climate change abatement – we hope that future work on the risk and return properties of other asset classes over similar horizons will further discipline our thinking and modeling. In addition, as we obtain additional scientific evidence on the physical processes of climate change and their interaction with economic activity, the calibration of asset pricing models designed to inform cost-benefit analyses or estimates of the social cost of carbon will improve further.

References

- Alley, Richard B, Jochem Marotzke, William D Nordhaus, Jonathan T Overpeck, Dorothy M Peteet, Roger A Pielke, Raymond T Pierrehumbert, Peter B Rhines, Thomas F Stocker, Lynne D Talley, et al.** 2003. "Abrupt climate change." *Science*, 299(5615): 2005–2010.
- Ambrose, Brent W., Piet Eichholtz, and Thies Lindenthal.** 2013. "House Prices and Fundamentals: 355 Years of Evidence." *Journal of Money, Credit and Banking*, 45(2-3): 477–491.
- Arrow, Kenneth, Maureen Cropper, Christian Gollier, Ben Groom, Geoffrey Heal, Richard Newell, William Nordhaus, Robert Pindyck, William Pizer, Paul Portney, et al.** 2013. "How Should Benefits and Costs Be Discounted in an Intergenerational Context? The Views of an Expert Panel." *Resources for the Future Discussion Paper No.* 12-53.
- Atreya, Ajita, and Susana Ferreira.** 2015. "Seeing is believing? Evidence from property prices in inundated areas." *Risk Analysis*, 35(5): 828–848.
- Bakkensen, Laura A, and Lint Barrage.** 2017. "Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics: Going Under Water?" National Bureau of Economic Research.
- Bansal, Ravi, and Amir Yaron.** 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, 59(4): 1481–1509.
- Bansal, Ravi, Dana Kiku, and Marcelo Ochoa.** 2013. "Climate Change, Growth, and Risk."
- Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph 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(1): 105–159.
- Barro, Robert J.** 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121(3): 823–866.
- Barro, Robert J.** 2013. "Environmental Protection, Rare Disasters, and Discount Rates." National Bureau of Economic Research Working Paper 19258.
- Barro, Robert J., and Jose F. Ursua.** 2008. "Macroeconomic Crises since 1870." *Brookings Papers on Economic Activity*, 39(1): 255–350.
- Barro, Robert J, and Tao Jin.** 2011. "On the size distribution of macroeconomic disasters." *Econometrica*, 79(5): 1567–1589.
- Belo, Frederico, Pierre Collin-Dufresne, and Robert S Goldstein.** 2015. "Dividend dynamics and the term structure of dividend strips." *The Journal of Finance*, 70(3): 1115–1160.
- Bernstein, Asaf, Matthew Gustafson, and Ryan Lewis.** 2017. "Disaster on the Horizon: The Price Effect of Sea Level Rise."
- Binsbergen, Jules van, Michael Brandt, and Ralph Koijen.** 2012. "On the Timing and Pricing of Dividends." *American Economic Review*, 102(4): 1596–1618.
- Bordo, Michael, Barry Eichengreen, Daniela Klingebiel, and Maria Soledad Martinez-Peria.** 2001. "Is the Crisis Problem Growing More Severe?" *Economic Policy*, 16(32): 51–82.
- Bracke, Philippe.** 2015. "House Prices and Rents: Microevidence from a Matched Data Set in Central London." *Real Estate Economics*, 43(2): 403–431.
- Brohé, Arnaud, and Michael Greenstone.** 2007. "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather." *The American Economic Review*, 97(1): 354–385.
- Bureau of Economic Analysis.** 2017. "Concepts and Methods of the U.S. National Income and Product Accounts."

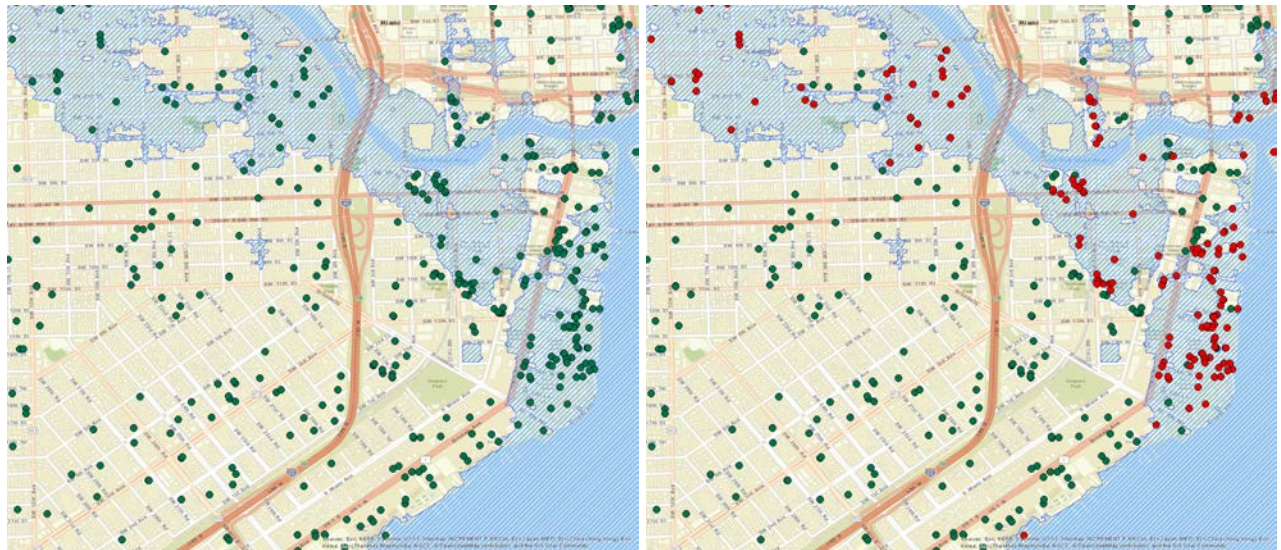
- Burke, Marshall, and Kyle Emerick.** 2016. "Adaptation to climate change: Evidence from US agriculture." *American Economic Journal: Economic Policy*, 8(3): 106–40.
- Cai, Yongyang, Kenneth L. Judd, and Thomas S. Lontzek.** 2013. "The Social Cost of Stochastic and Irreversible Climate Change." *NBER Working Paper*, 18704.
- Campbell, Sean D, Morris A Davis, Joshua Gallin, and Robert F Martin.** 2009. "What Moves Housing Markets: A Variance Decomposition of the Rent–Price Ratio." *Journal of Urban Economics*, 66(2): 90–102.
- Chay, Kenneth Y., and Michael Greenstone.** 2005. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy*, 113(2).
- Cocco, Joao F.** 2005. "Portfolio Choice in the Presence of Housing." *Review of Financial Studies*, 18(2): 535–567.
- Cox, Peter M, Richard A Betts, Chris D Jones, Steven A Spall, and Ian J Totterdell.** 2000. "Acceleration of global warming due to carbon-cycle feedbacks in a coupled climate model." *Nature*, 408(6809): 184–187.
- Cropper, Maureen L., Mark C. Freeman, Ben Groom, and William A. Pizer.** 2014. "Declining Discount Rates." *American Economic Review*, 104(5): 538–43.
- Crost, Benjamin, and Christian P. Traeger.** 2014. "Optimal CO2 Mitigation Under Damage Risk Valuation." *Nature Climate Change*.
- Deschênes, Olivier, and Michael Greenstone.** 2011. "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US." *American Economic Journal: Applied Economics*, 3(4): 152–185.
- Desmet, Klaus, and Esteban Rossi-Hansberg.** 2015. "On the Spatial Economic Impact of Global Warming." *Journal of Urban Economics*, 88: 16 – 37.
- Díaz, Antonia, and María José Luengo-Prado.** 2008. "On the User Cost and Homeownership." *Review of Economic Dynamics*, 11(3): 584–613.
- Drupp, Moritz A., Mark Freeman, Ben Groom, and Frikk Nesje.** 2015. "Discounting Disentangled: An Expert Survey on the Determinants of the Long-Term Social Discount Rate." *Centre for Climate Change Economics and Policy Working Paper No. 195*.
- Eichholtz, Piet.** 1997. "A Long Run House Price Index: The Herengracht Index, 1628–1973." *Real Estate Economics*, 25(2): 175–192.
- Eitrheim, Øyvind, and Solveig K Erlandsen.** 2005. "House Prices in Norway, 1819–1989." *Scandinavian Economic History Review*, 53(3): 7–33.
- Farmer, J. Doyne, John Geanakoplos, Jaume Masoliver, Miquel Montero, and Josep Perelló.** 2015. "Value of the Future: Discounting in Random Environments." *Physical Review E*, 91: 052816.
- Favilukis, Jack, Sydney C. Ludvigson, and Stijn van Nieuwerburgh.** 2010. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium." National Bureau of Economic Research Working Paper 15988.
- Flavin, Marjorie, and Takashi Yamashita.** 2002. "Owner-Occupied Housing and the Composition of the Household Portfolio." *American Economic Review*, 92(1): 345–362.
- Franklin, Sergio L, and Robert S Pindyck.** 2017. "Tropical Forests, Tipping Points, and the Social Cost of Deforestation." National Bureau of Economic Research.
- Gibson, Matthew, Jamie T Mullins, and Alison Hill.** 2017. "Climate change, flood risk, and property values: Evidence from New York City." Working Paper.
- Giglio, Stefano, Matteo Maggiori, and Johannes Stroebe.** 2015. "Very Long-Run Discount Rates." *The Quarterly Journal of Economics*, 130(1): 1–53.
- Gollier, Christian.** 2013. "Evaluation of Long-Dated Investments Under Uncertain Growth Trend, Volatility and Catastrophes." Toulouse School of Economics TSE Working Papers 12-361.
- Gourio, François.** 2008. "Disasters and recoveries." *The American Economic Review*,

- 98(2): 68–73.
- Greenstone, Michael, and Justin Gallagher.** 2008. “Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program.” *The Quarterly Journal of Economics*, 123(3): 951–1003.
- Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton.** 2013. “Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation.” *Review of Environmental Economics and Policy*, 7(1): 23–46.
- Gyourko, Joseph, and Donald B. Keim.** 1992. “What Does the Stock Market Tell us About Real Estate Returns?” *Real Estate Economics*, 20(3): 457–485.
- Hallstrom, Daniel G, and V Kerry Smith.** 2005. “Market responses to hurricanes.” *Journal of Environmental Economics and Management*, 50(3): 541–561.
- Harding, John P, Stuart S Rosenthal, and CF Sirmans.** 2007. “Depreciation of Housing Capital, Maintenance, and House Price Inflation: Estimates From a Repeat Sales Model.” *Journal of Urban Economics*, 61(2): 193–217.
- Hasler, Michael, and Roberto Marfe.** 2016. “Disaster recovery and the term structure of dividend strips.” *Journal of Financial Economics*, 122(1): 116–134.
- Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor.** 2017. “The Rate of Return on Everything, 1870–2015.” *Working Paper*.
- Kaplow, Louis, Elisabeth Moyer, and David A. Weisbach.** 2010. “The Social Evaluation of Intergenerational Policies and Its Application to Integrated Assessment Models of Climate Change.” *The B.E. Journal of Economic Analysis & Policy*, 10(2).
- Leigh, Wilhelmina A.** 1980. “Economic Depreciation of the Residential Housing Stock of the United States, 1950–1970.” *Review of Economics and Statistics*, 62(2): 200–206.
- Lemoine, Derek.** 2015. “The Climate Risk Premium.” *Available at SSRN 2560031*.
- Lemoine, Derek, and Christian P. Traeger.** 2014. “Watch Your Step: Optimal Policy in a Tipping Climate.” *American Economic Journal: Economic Policy*, 6(1): 137–166.
- Lenton, Timothy M, Hermann Held, Elmar Kriegler, Jim W Hall, Wolfgang Lucht, Stefan Rahmstorf, and Hans Joachim Schellnhuber.** 2008. “Tipping elements in the Earth’s climate system.” *Proceedings of the national Academy of Sciences*, 105(6): 1786–1793.
- Lettau, Martin, and Jessica A. Wachter.** 2011. “The Term Structures of Equity and Interest Rates.” *Journal of Financial Economics*, 101(1): 90–113.
- Lucas, Robert E, Jr.** 1978. “Asset Prices in an Exchange Economy.” *Econometrica*, 46(6): 1429–45.
- Lustig, Hanno N., and Stijn van Nieuwerburgh.** 2005. “Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective.” *Journal of Finance*, 60(3): 1167–1219.
- Mack, Adrienne, and Enrique Martínez-García.** 2011. “A Cross-Country Quarterly Database of Real House Prices: A Methodological Note.” Federal Reserve Bank of Dallas Globalization and Monetary Policy Working Paper 99.
- Malpezzi, Stephen, Larry Ozanne, and Thomas G Thibodeau.** 1987. “Microeconomic Estimates of Housing Depreciation.” *Land Economics*, 63(4): 372–385.
- Mayerhauser, Nicole, and Marshall Reinsdorf.** 2006. “Housing Services in the National Economic Accounts.” *US Bureau of Economic Analysis*.
- McCarthy, Jonathan, and Richard Peach.** 2010. “The Measurement of Rent Inflation.” FRB of New York Staff Report 425.
- McKenzie, Russell, and John Levendis.** 2010. “Flood hazards and urban housing markets: The effects of Katrina on New Orleans.” *The Journal of Real Estate Finance and Economics*, 40(1): 62–76.
- Melillo, Jerry M, Serita D Frey, Kristen M DeAngelis, William J Werner, Michael J**

- Bernard, FP Bowles, Grace Pold, MA Knorr, and AS Grandy.** 2017. "Long-term pattern and magnitude of soil carbon feedback to the climate system in a warming world." *Science*, 358(6359): 101–105.
- Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa.** 2013. "Crises and Recoveries in an Empirical Model of Consumption Disasters." *American Economic Journal: Macroeconomics*, 5(3): 35–74.
- Nordhaus, William D.** 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.
- Nordhaus, William D.** 2013. *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World*. Yale University Press.
- Nordhaus, William D., and Joseph Boyer.** 2000. *Warming the World: Economic Models of Global Warming*. MIT press.
- Ortega, Francesc, and Suleyman Taspinar.** 2017. "Rising Sea Levels and Sinking Property Values: The Effects of Hurricane Sandy on New York's Housing Market."
- Overpeck, Jonathan T, and Julia E Cole.** 2006. "Abrupt change in Earth's climate system." *Annu. Rev. Environ. Resour.*, 31: 1–31.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel.** 2007. "Housing, Consumption and Asset Pricing." *Journal of Financial Economics*, 83(3): 531–569.
- Piketty, Thomas, and Gabriel Zucman.** 2014. "Capital is Back: Wealth-Income Ratios in Rich Countries 1700–2010." *The Quarterly Journal of Economics*, 129(3): 1255–1310.
- Pizer, William, Matthew Adler, Joseph Aldy, David Anthoff, Maureen Cropper, Kenneth Gillingham, Michael Greenstone, Brian Murray, Richard Newell, Richard Richels, et al.** 2014. "Using and Improving the Social Cost of Carbon." *Science*, 346(6214): 1189–1190.
- Reinhart, Carmen M., and Kenneth S. Rogoff.** 2008. "Is the 2007 US Sub-Prime Financial Crisis so Different? An International Historical Comparison." *American Economic Review: Papers and Proceedings*, 98(2): 339–344.
- Reinhart, Carmen M., and Kenneth S. Rogoff.** 2009. *This Time is Different: Eight Centuries of Financial Folly*. Princeton University Press.
- Schneider, Maik T., Christian P. Traeger, and Ralph Winkler.** 2012. "Trading Off Generations: Equity, Discounting, and Climate Change." *European Economic Review*, 56(8): 1621 – 1644.
- Schularick, Moritz, and Alan M. Taylor.** 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises." *American Economic Review*, 102(2): 1029–1061.
- Shiller, Robert.** 2000. *Irrational Exuberance*. Princeton University Press, Princeton.
- Shiller, Robert J.** 2006. "Long-Term Perspectives on the Current Boom in Home Prices." *The Economists' Voice*, 3(4): 1–11.
- Shiller, Robert J.** 2007. "Understanding Recent Trends in House Prices and Home Ownership." National Bureau of Economic Research Working Paper 15988.
- Stern, Nicholas.** 2006. *Stern Review on the Economics of Climate Change*. London, UK: Her Majesty's Treasury.
- Stern, Nicholas.** 2008. "The Economics of Climate Change." *American Economic Review*, 98(2): 1–37.
- Traeger, Christian P.** 2014a. "A 4-States DICE: Quantitatively Addressing Uncertainty Effects in Climate Change." *Environmental and Resource Economics*, 59(1): 1–37.
- Traeger, Christian P.** 2014b. "Why Uncertainty Matters: Discounting Under Intertemporal Risk Aversion and Ambiguity." *Economic Theory*, 56(3): 627–664.
- Traeger, Christian P.** 2015. "Analytic Integrated Assessment and Uncertainty."
- Wagner, Gernot, and Martin L. Weitzman.** 2015. *Climate Shock: The Economic Consequences*

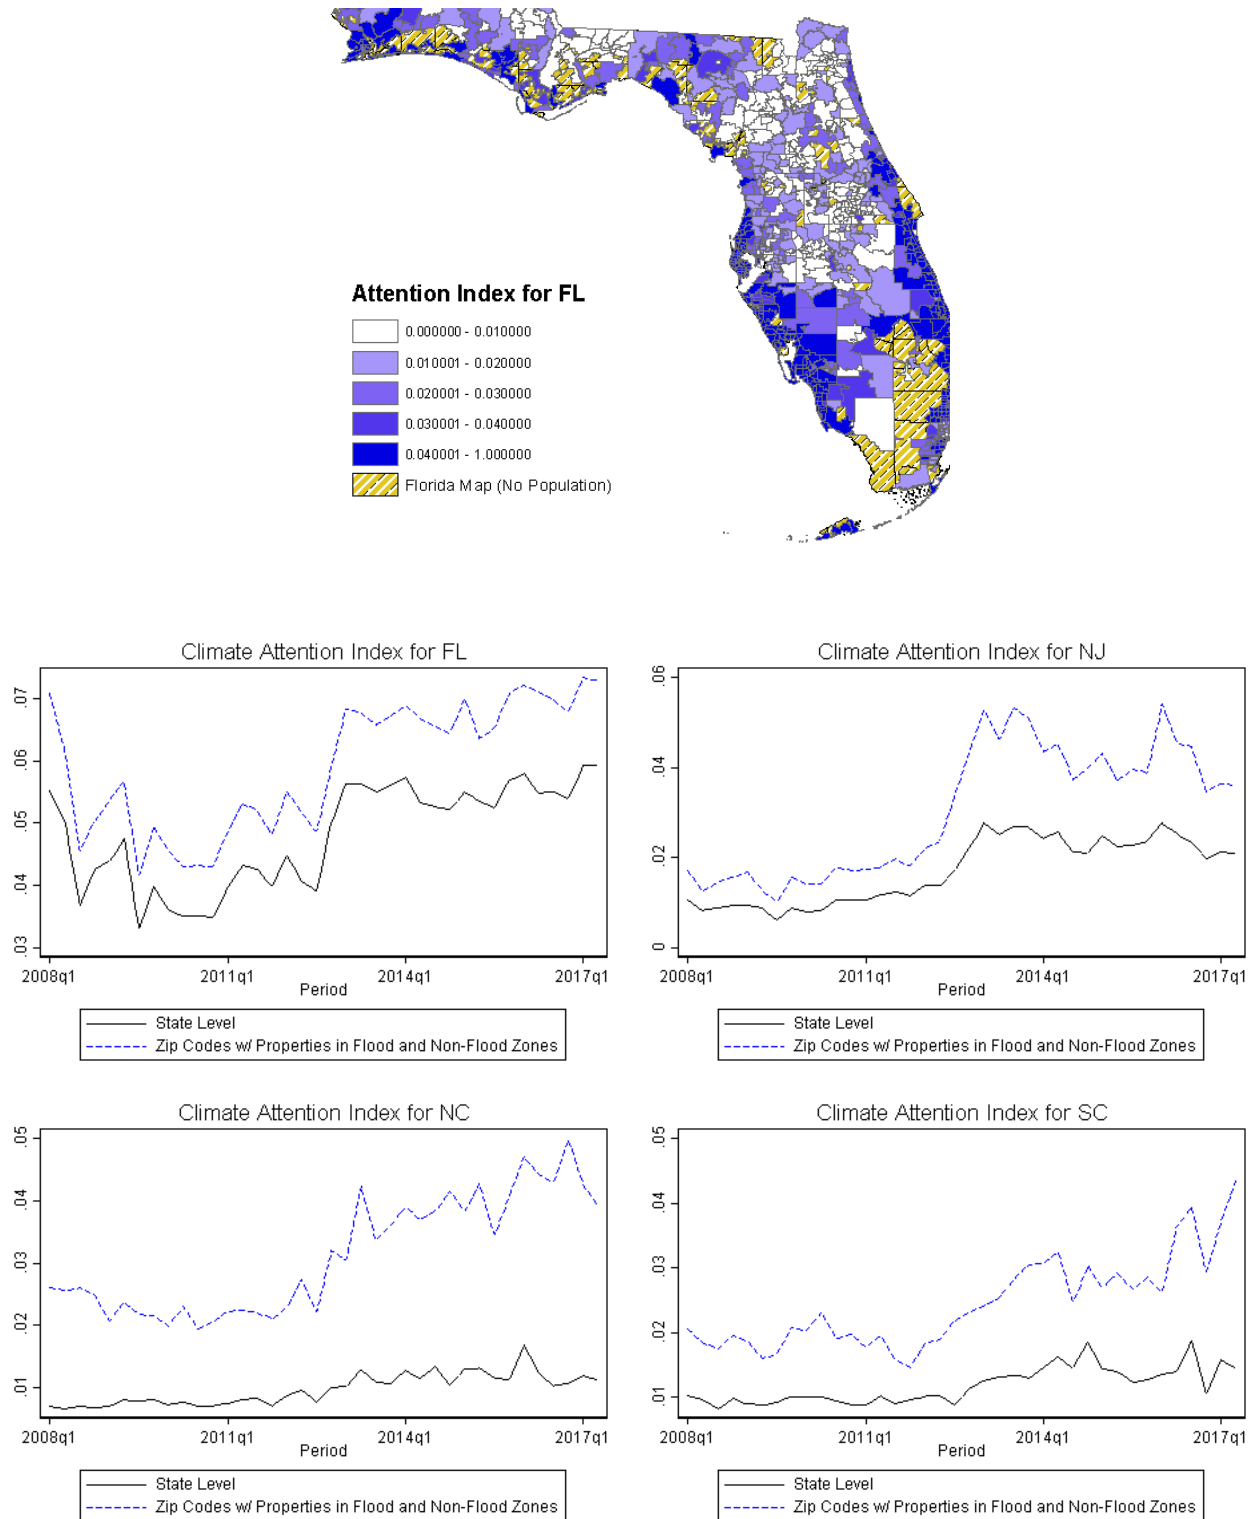
- of a Hotter Planet*. Princeton, NJ:Princeton University Press.
- Weisbach, David, and Cass R. Sunstein.** 2009. "Climate Change and Discounting the Future: A Guide for the Perplexed." *Yale Law & Policy Review*, 27(2): 433 – 457.
- Weitzman, Martin L.** 2012. "Rare Disasters, Tail-Hedged Investments, and Risk-Adjusted Discount Rates." National Bureau of Economic Research Working Paper 18496.
- Weitzman, Martin L.** 2014. "Fat tails and the social cost of carbon." *The American Economic Review*, 104(5): 544–546.

Figure I: Illustration of Identifying Properties in the Flood Zone




Note: 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 geocode to identify the properties that fall under the flood zone and color them as red dots. As seen above, as properties are closer to the coastal line, they are more likely to be in the flood zone.

Figure II: Climate Attention Index in the Cross-Section and Time-Series



Note: The top panel shows 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 darker the color is for a zip code, the more attention to climate risk people pay. The other panels illustrate the quarterly time series of the Climate Attention Index aggregated at the state level and the level of zip codes that include both properties in the Flood Zone and properties not in the Flood Zone. Not only does the index vary over time, but it also varies across different states. Most notably, as the higher index of the zip code level indicates, when people live within or near the Flood Zone, they pay more attention to climate risk.

Figure III: House Price Riskiness I




(A) Financial Crises




(B) Consumption Disasters

Note: The top panel shows average real house price movements relative to financial crises in Schularick and Taylor (2012), Bordo et al. (2001), and Reinhart and Rogoff (2009). The bottom panel shows average real house price movements and average real consumption relative to the trough of consumption disasters identified by Barro and Ursua (2008). House prices and consumption volumes during the reference year are normalized to 1. See Appendix A.3.1 and Appendix Table A.7 for a description of the countries included and the data series and crises considered here.

Figure IV: House Price Riskiness II




(A) U.K.



(B) Singapore

Note: The figures show the evolution of real house prices in the U.K. (top panel) and Singapore (bottom panel). Shaded regions for the U.K. are financial crises identified by Reinhart and Rogoff (2009): 1974-1976, 1982-1983, 1991 and 2008-2009. Shaded regions for Singapore include the 1982-1983 financial crisis identified by Reinhart and Rogoff (2009), as well as the Asian financial crisis (1997-1998), and the 2007-2008 global financial crisis. See Appendix A.3.1 for a description of the data series studied here.


Figure V: House Price Riskiness III




(A) Housing During World Wars

Note: Figure shows the evolution of real house prices for countries with available house-price time series during World War I (Australia, France, Netherlands, Norway, United States) and World War II (Australia, France, Netherlands, Norway, Switzerland, United States). See Appendix A.3.1 for a description of the data series studied here.

Figure VI: Leasehold Price Discounts by Remaining Lease Length




(A) Leasehold Discounts - U.K.



(B) Leasehold Discounts - Singapore


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

Figure VII: Sample Paths: Trend Growth and Above-Trend Growth




Note: 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 in which the economy grows at its deterministic trend. The solid line represents a temporary deviation from trend in which growth accelerates.

Figure VIII: Sample Paths: Above-Trend Growth, With and Without a Disaster




Note: 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 in which the economy grows at its deterministic trend, then experiences an increase in growth (the same as the solid line in Figure VII). The solid line represents an alternative path in which the increased probability of disasters due to the temporary acceleration in the economy leads to the occurrence of a disaster after year 25.

Figure IX: U.K. Gilts Real Yield Curve



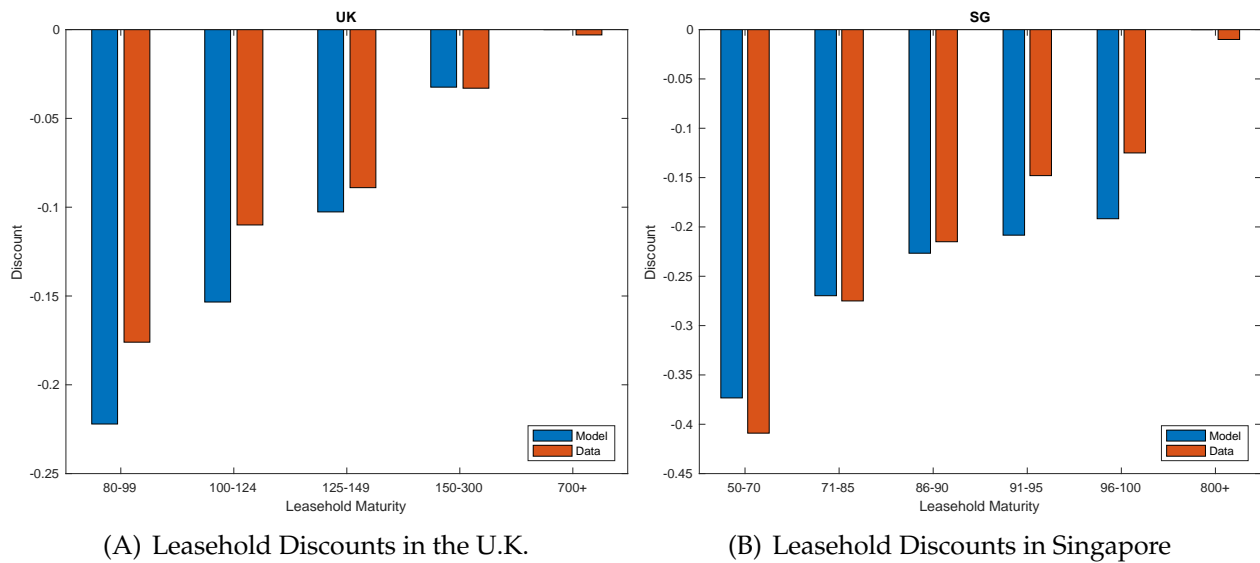
Note: The figure plots the real yield curve for U.K. gilts as computed by the Bank of England for the period 1998-2016, as well as for two sub-periods: 1998-2007 and 2008-2016. It is available at <http://www.bankofengland.co.uk/statistics/Pages/yieldcurve/archive.aspx>, last accessed July 2017. Until 2015, the U.K. government debt also included some perpetual bonds: the War Loans and the Annuities. These bonds comprised 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 hard to interpret their prices in terms of discount rates.

Figure X: Discount Rates for Risk-Free Bonds, Housing, Damages



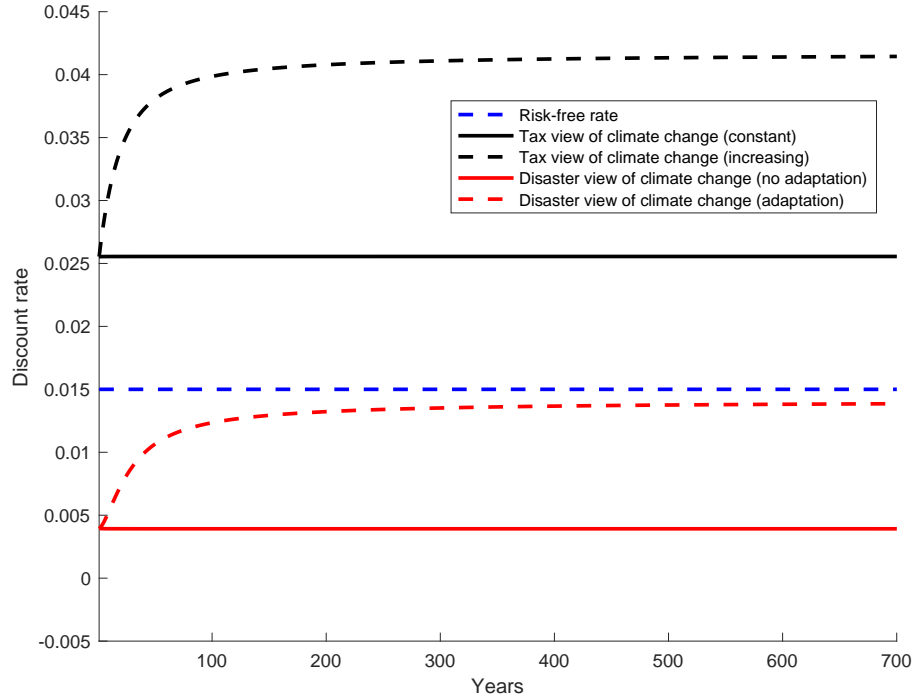
Note: The figure shows the per-period discount rate corresponding to different assets for maturities 1 to 1000 years, in our baseline model calibration. The top line represents the term structure of discount rates for the risky housing asset. The dashed line below it represents the real risk-free asset, that is, the real yield curve. The three black lines in the bottom represent different calibrations of the damage process with $\theta \in \{1, 0.5, 0.1\}$.

Figure XI: Leasehold Discounts



Note: The figure shows the discount rates for housing assets predicted by the model (left bars) and in the data (right bars), for the U.K. (left panel) and Singapore (right panel). The discounts are estimated from a hedonic regression and reported in log points.

Figure XII: Leading Models of Climate Change: Predictions for Discount Rates



Note: 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 $\pi_q = 1$.

Table I: Transaction Prices vs. Rent Prices: Hedonic Analysis

PANEL A. TRANSACTION PRICES DEPENDENT VARIABLE: LOG(TRANSACTION PRICES)						
	(1)	(2)	(3)	(4)	(5)	(6)
Flood Zone	0.004 (0.015)	0.014 (0.013)			0.085*** (0.006)	
log(Index by Zip-Year) × Flood Zone	-0.024*** (0.005)		-0.029** (0.010)			
log(Index by Zip-Quarter) × Flood Zone		-0.020*** (0.004)		-0.021** (0.007)		
Index by Zip-Quarter × Flood Zone					-0.210** (0.071)	-0.367*** (0.091)
Property Controls	✓	✓	.	.	✓	.
Zip × Quarter FE	✓	✓	✓	✓	✓	✓
Property FE	.	.	✓	✓	.	✓
R-squared	0.585	0.585	0.721	0.721	0.585	0.721
N	7,287,000	7,233,113	3,485,238	3,443,265	7,233,113	3,443,265

PANEL B. RENT PRICES DEPENDENT VARIABLE: LOG(RENT PRICES)						
	(1)	(2)	(3)	(4)	(5)	(6)
Flood Zone	0.041*** (0.012)	0.033** (0.011)			-0.034*** (0.006)	
log(Index by Zip-Year) × Flood Zone	0.018*** (0.004)		0.005 (0.005)			
log(Index by Zip-Quarter) × Flood Zone		0.015*** (0.004)		0.003 (0.003)		
Index by Zip-Quarter × Flood Zone					0.415*** (0.072)	0.016 (0.042)
Property Controls	✓	✓	.	.	✓	.
Zip × Quarter FE	✓	✓	✓	✓	✓	✓
Property FE	.	.	✓	✓	.	✓
R-squared	0.728	0.728	0.942	0.942	0.728	0.942
N	2,142,433	2,142,240	1,191,657	1,191,642	2,142,240	1,191,642

Note: 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 column 1, 2, and 5, we control for various property characteristics such as the property size, property age, and the number of bedrooms. In column 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.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table II: Real House Price Growth and Real Consumption Growth

	Period	Real HP Growth		Real Cons. Growth		Correlation	Cons. beta
		Mean	Std. Dev.	Mean	Std. Dev.		
Australia	1901-2009	2.51%	12.1%	1.51%	4.99%	0.102	0.248
Belgium	1975-2009	2.92%	6.06%	1.59%	1.51%	0.439	1.761
Canada	1975-2009	2.38%	7.69%	1.61%	1.73%	0.433	1.929
Denmark	1975-2009	1.99%	9.24%	0.98%	2.71%	0.538	1.838
Finland	1975-2009	2.17%	8.70%	2.07%	2.79%	0.710	2.214
France	1840-2009	2.06%	11.8%	1.49%	6.32%	-0.054	-0.101
Germany	1975-2009	-0.45%	2.33%	1.64%	1.52%	0.494	0.755
Italy	1975-2009	1.28%	8.10%	1.75%	2.18%	0.165	0.614
Japan	1975-2009	0.02%	4.45%	1.97%	1.60%	0.503	1.394
Netherlands	1814-2009	2.79%	20.8%	1.57%	7.49%	0.078	0.215
New Zealand	1975-2009	2.46%	8.09%	1.00%	2.30%	0.580	2.044
Norway	1830-2009	1.77%	11.6%	1.78%	3.83%	0.243	0.737
Singapore	1975-2009	7.18%	19.5%	3.43%	4.03%	0.348	1.685
South Africa	1975-2009	1.13%	10.1%	0.92%	3.02%	0.707	2.365
South Korea	1975-2009	0.58%	7.93%	4.62%	4.49%	0.370	0.652
Spain	1975-2009	3.14%	8.07%	1.56%	2.60%	0.593	1.837
Sweden	1952-2009	1.55%	6.04%	1.63%	1.99%	0.536	1.627
Switzerland	1937-2009	0.47%	7.17%	1.48%	3.82%	0.187	0.350
U.K.	1952-2009	2.89%	9.55%	2.26%	2.11%	0.700	3.169
U.S.	1890-2009	0.49%	7.36%	1.84%	3.41%	0.148	0.320

Note: The table shows time series properties of annual growth rates of real house prices (as described in 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.

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

	Real HP Growth		Real PDI Growth		Correlation	PDI beta
	Mean	Std. Dev.	Mean	Std. Dev.		
Australia	3.54%	6.67%	1.37%	2.10%	0.156	0.495
Belgium	2.53%	5.50%	0.92%	2.30%	0.431	1.031
Canada	2.91%	7.49%	1.35%	2.18%	0.466	1.604
Switzerland	1.12%	4.58%	1.17%	1.53%	0.425	1.275
Germany	0.07%	2.52%	1.28%	1.64%	0.237	0.365
Denmark	1.73%	8.58%	1.13%	2.29%	0.224	0.839
Spain	-0.09%	10.6%	0.81%	2.27%	0.409	1.909
Finland	1.90%	7.71%	1.92%	2.97%	0.470	1.219
France	2.28%	5.13%	1.12%	1.61%	0.332	1.056
U.K.	3.47%	8.49%	2.05%	2.26%	0.420	1.575
Ireland	3.36%	9.44%	1.89%	3.33%	0.574	1.627
Italy	0.33%	8.15%	0.89%	2.48%	0.363	1.195
Japan	-0.39%	4.24%	1.49%	1.44%	0.622	1.835
South Korea	0.64%	7.36%	3.97%	4.38%	0.245	0.412
Luxembourg	4.16%	6.40%	2.76%	3.63%	0.067	0.117
Netherlands	2.31%	9.12%	0.74%	3.01%	0.467	1.414
Norway	2.65%	6.92%	2.22%	2.05%	0.037	0.126
New Zealand	2.90%	7.73%	1.13%	3.41%	0.486	1.103
Sweden	2.00%	7.01%	1.40%	2.39%	0.467	1.371
U.S.	1.36%	3.88%	1.59%	1.54%	0.322	0.812
South Africa	0.49%	9.13%	0.34%	2.37%	0.474	1.824
Croatia	1.16%	12.3%	8.79%	27.0%	-0.345	-0.158
Israel	3.05%	8.83%	2.74%	7.37%	0.129	0.155

Note: This table shows 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 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.

Table IV: 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	10.3%	9.7%	8.9%	9.4%	10.0%	9.7%	9.3%	9.1%	8.9%	8.6%	9.5%	9.9%	10.3%	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%	7.2%	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%	6.2%	6.8%	6.3%	5.9%	5.7%	5.5%	5.2%	7.0%	7.4%	7.9%	7.5%	7.3%	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	

Note: This table shows our estimates for net real returns to housing and real rent growth in the U.S., the U.K., and Singapore 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 U.S., 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 U.K., LR stands for Land Registry, and CPIH is 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, changes in the total housing stock are controlled for 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 Appendix A.4 for further details on the estimation procedures and the underlying data sources used. Numbers may not add up due to rounding.