House Prices Under Rising Seas: Measuring Deep Uncertainty in Theory and Evidence*

Michael Barnett† Jacob Dice‡ Toàn Phan§ David Rodziewicz¶
Constantine Yannelis‖

30th April 2025
(preliminary and incomplete)

Abstract

We quantify ambiguity aversion by studying how real estate markets respond to long-run sea-level rise (SLR) risk. Using a novel high-resolution dataset covering over two million coastal properties, we show that housing prices reflect not only expected SLR risk but also deep uncertainty across climate scenarios. Leveraging a canonical asset-pricing framework, we recover one of the first market-based estimates of ambiguity aversion, finding substantial tilting of probability weights toward worst-case outcomes. We also estimate a relatively low long-run discount rate, suggesting strong valuation of distant climate risks. As an application, we quantify how ambiguity aversion amplifies the willingness to invest in SLR adaptation.

Keywords: ambiguity aversion, sea-level rise, climate risk, asset pricing, real estate

*The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Banks of Richmond, Kansas City, or the Federal Reserve System. We are grateful to Paul Ho, Cosmin Ilut and Ricardo Reis for their extensive suggestions. We also thank Chris Amante, David Anthoff, Job Boerma, Jarda Borovicka, Valentina Bosetti, Simon Dietz, Lars Hansen, Aleksei Oskolkov, Evi Pappa, Tom Sargent, Christian Traeger, and participants at the Blue Collar Working Group (University of Chicago), Richmond Fed Brownbag, Universt of South Carolina Darla Moore School of Business Finance Seminar, and the Mohammed VI Polytechnic University Climate Challenges Conference for helpful comments and suggestions. We thank Rosemary Coskrey, Claire Conzelmann, Brooke Hansbrough, Adin Hammond, and Nick Weitzel for excellent research assistance. All errors are our own.

†Arizona State University (michael.d.barnett@asu.edu).
‡Federal Reserve Bank of Kansas City (jacob.dice@kc.frb.org).
§Federal Reserve Bank of Richmond (toan.phan@rich.frb.org).
¶Federal Reserve Bank of Kansas City (david.rodziewicz@kc.frb.org).
‖University of Cambridge (cny21@cam.ac.uk)

1. Introduction

Economists have long recognized that risk—situations where probabilities are known—and ambiguity or deep uncertainty—situations where probabilities are hard to quantify—are fundamentally distinct, a point emphasized in the classic works of Knight (1921) and Keynes (1921). The distinction matters: while risk can be priced using expected utility, ambiguity often induces far more cautious, worst-case behavior, reshaping asset prices and macroeconomic dynamics (Caballero and Krishnamurthy, 2008; Hansen, 2014; Ilut and Schneider, 2023). However, because existing empirical settings have been unable to provide a clear way to separately measure risk–the variability within a given model–and ambiguity–the disagreement across competing models–empirically disentangling the effects of ambiguity from risk remains a major challenge.

The economic consequences of climate change provide an especially promising setting for characterizing the implications of risk and ambiguity because ambiguity is a particularly critical feature in such contexts. Climate projections typically involve multiple scenarios, such as the Representative Concentration Pathways (RCPs), which reflect divergent assumptions about future socio-economic behavior, including the pace of decarbonization, technological change, and economic growth. These scenarios are internally coherent but are not assigned objective probabilities, creating uncertainty about which trajectory the world will follow (Intergovernmental Panel on Climate Change, 2014). Future sea-level rise (SLR) projections vary widely across RCPs, reflecting uncertainty not only in emissions paths but also in geophysical feedbacks, such as ice sheet dynamics and the carbon cycle. For example, median SLR projections under RCP2.6 and RCP8.5 can diverge by several feet, with first-year inundation estimates for coastal properties spanning a century or more (Kopp et al., 2014; DeConto and Pollard, 2016; Rodziewicz et al., 2022). Because scenario developers do not provide formal probability distributions across RCPs—despite labels such as "business as usual," "most likely," or "least likely"–economic agents must evaluate both risks and ambiguity when making location and investment decisions. A more ambiguity-averse agent will react more strongly to worst-case projections. How markets capitalize these two distinct forms of uncertainty remains an open empirical question with important implications for climate adaptation and policy design.

Our paper confronts this challenge by focusing on the economically significant and deeply uncertain context of coastal housing markets exposed to sea-level rise (SLR). We construct a novel dataset that links over two million coastal property transactions from the proprietary CoreLogic database to detailed, property-specific projections of SLR inundation timing under RCP 2.6, 4.5, and 8.5, derived from state-of-the-art, high-resolution sea-level

rise projection maps produced by NOAA. Accurately assessing a property's inundation risk requires integrating localized elevation data, hydrologic connectivity, and regional sea-level projections, rather than relying on broad global estimates, to capture substantial local heterogeneity driven by topography, land subsidence, tidal variation, and ocean dynamics (Rodziewicz et al., 2022). For each property, we calculate the expected first year of inundation under each climate scenario, allowing a sharp empirical distinction between:

- *Risk*[1] as the variance of inundation timing *within* a given scenario, and
- *Ambiguity* as the divergence in expected inundation timing *across* scenarios.

This mapping closely mirrors the conceptual distinction articulated by Knight and Keynes, now operationalized in a modern, data-driven context. Because much of the projected SLR damage occurs far into the future—often beyond a century—the way markets discount uncertain outcomes becomes critical. Standard discounting approaches, which treat uncertainty purely as risk, may underestimate the economic costs of long-horizon climate damages. Allowing for ambiguity aversion enriches our understanding of how uncertainty affects asset prices and decision-making, including in contexts such as adaptation investment to long-run SLR risk, which we explore through an application of our model.

To interpret how markets price risk and ambiguity, we first analyze a canonical asset-pricing framework featuring both risk aversion and ambiguity aversion, building on the robust control approach of Hansen and Sargent (2001) and Hansen and Miao (2018). We adapt this framework to a parsimonious and transparent setting designed to fix ideas sharply in our context. In our model, housing prices reflect three components: the expected damages from SLR, a penalization for within-scenario variability (risk), and a penalization for disagreement across scenarios (ambiguity). Crucially, the model delivers a sharp, testable prediction: the capitalization of risk and ambiguity into housing prices is governed by distinct structural parameters—the risk aversion and ambiguity aversion coefficients, respectively. Under the null hypothesis of ambiguity neutrality, agents would treat uncertainty across scenarios as a compound lottery, and only the risk component would affect prices. Thus, the model not only separates the two effects conceptually but also provides a clear empirical test for the presence of ambiguity aversion in market behavior.

We then bring this structure to the data. Using our property-level measures of exposure, risk, and ambiguity, we estimate how coastal real estate prices respond to each component, following an empirical strategy similar in spirit to Baldauf et al. (2020). Specifically, we examine how housing prices vary with the expected first year of inundation (capturing

[1] In Proposition 3 of our main analysis we show that for our specific setting risk aversion will be isomorphic to misspecification aversion à-la Hansen and Sargent (2008). See Section 2 for complete details.

risk) and with the dispersion of inundation timing across scenarios (capturing ambiguity). Our approach is related in spirit to Giglio et al. (2021), who use real estate transactions to infer the discount rates relevant for evaluating long-run climate change mitigation, although we center our analysis on the effects of deep uncertainty (ambiguity) rather than on the discounting of known risks. Just as observed discount rates provide guidance on the appropriate social discount rate for climate policy, we argue that the degree of ambiguity aversion reflected in real estate prices offers a way to empirically discipline ambiguity preferences that should be incorporated into climate-economy models. If markets capitalize not only expected SLR risk but also the uncertainty surrounding it, we can use these pricing patterns to recover an implied ambiguity aversion parameter, offering new revealed-preference evidence on how deep uncertainty shapes economic behavior.

Our main results are as follows. First, by linking our empirical estimates to the theoretical model of decision-making under ambiguity, we recover an implied ambiguity aversion parameter for market participants. Consistent with the model's predictions, we find that greater SLR uncertainty leads to an additional discount in house prices, above and beyond the discount associated with expected inundation timing. Importantly, we reject the null hypothesis that agents treat uncertainty across scenarios as a mere compound lottery (i.e., $\alpha = 0$); instead, we estimate a positive ambiguity aversion parameter. Quantitatively, we find a market-implied ambiguity aversion coefficient of approximately 22. While hard to independent of the model, we show that this parameter value leads the agent to substantially tilt their subjective probabilities toward the worst-case SLR scenario. Specifically, starting from a uniform prior assigning one-third probability to each RCP scenario (2.6, 4.5, 8.5), ambiguity aversion effectively increases the weight on the worst-case RCP 8.5 to nearly 50%, while reducing the weight on the best-case RCP 2.6, when pricing a house under the ambiguity-aversion-tilted probability distribution. This provides one of the first revealed-preference, market-based estimates of ambiguity aversion and underscores that deep uncertainty is a first-order determinant of long-run asset values under climate risk.

Second, our empirical strategy delivers two additional contributions to the literature as byproducts of our data construction. One contribution is to add a time dimension to the growing literature on the capitalization of SLR risk in real estate markets (see Related Literature). Existing work typically relies on static measures of exposure, such as indicators for whether a property would be inundated under a fixed threshold of sea-level rise (e.g., six feet; Baldauf et al. 2020; Bakkensen et al. 2025), which convey little information about the *timing* of damages. By contrast, our construction of *property-level distributions of inundation timing* allows us to study not only whether but also when properties are expected to be exposed. This dynamic perspective, in turn, enables a second contribution: it allows us

to revisit a foundational question in climate economics, namely, what is the appropriate discount rate for valuing long-horizon climate damages (Weitzman, 2001, 2013; Stern, 2007; Dasgupta, 2008; Gollier, 2024). Because much of the projected SLR damage in our dataset occurs more than a century into the future, discounting plays a crucial role. We find that properties expected to be inundated 100 years earlier sell at approximately an 8.8% to 10.5% discount. Mapping this discount through our theoretical framework implies a long-run annual discount rate of about 2.34%—substantially lower than the 4% benchmark commonly used in integrated assessment models like DICE (Nordhaus, 2013), and slightly below the 2.6% rate estimated from UK and Singapore housing markets by Giglio et al. (2015, 2021). This suggests that future climate damages are valued more heavily—and thus weigh more on present economic decisions—than standard models typically assume.

Finally, to translate our empirical estimates into real-world policy-relevant valuations, we use our model to quantify the effects of ambiguity aversion on the willingness to invest in adaptation. For example, for a house currently located just 1 foot above the inundation threshold, the baseline willingness-to-adapt (i.e., the value of raising the house's elevation by 1 foot) is about $126,917. Accounting for ambiguity aversion raises this figure to approximately $130,574—an increase of about $3,657, or roughly 2.9% relative to the baseline. As elevation increases, the incremental willingness to adapt declines, but ambiguity aversion consistently amplifies adaptation incentives. This quantification is, to our knowledge, the first empirical link between market-based discount rates, ambiguity aversion, and adaptation costs, highlighting how deep uncertainty fundamentally reshapes the economic calculus of climate resilience investments.

1.1. Related Literature

Our paper is related to several strands of the economic and finance literature. To the best of our knowledge, our paper is the first to quantify ambiguity aversion based on market transaction data and projections from scientific models.

We contribute to the broader discussion on the economic valuation of long-run climate impacts, offering an asset-pricing-based complement to theoretical approaches in the valuation of the social costs of carbon (Weitzman, 2012; Golosov et al., 2014; Nordhaus, 2017; Cai and Lontzek, 2019; Giglio et al., 2021; Barrage and Nordhaus, 2024) and empirical estimates of the economic costs and consequences of climate change (Colacito et al., 2018; Hsiang et al., 2017; Burke et al., 2018). Our paper contributes to the growing literature on real estate and climate risk and the consequences of sea-level rise for various assets (Bernstein et al., 2019; Baldauf et al., 2020; Murfin and Spiegel, 2020; Painter, 2020; Goldsmith-Pinkham et al., 2023; Bakkensen and Barrage, 2022; Bakkensen et al., 2025) by quantifying how

markets price not just climate-related sea-level rise risk but also uncertainty. Furthermore, an extensive climate finance literature has highlighted the implications of climate change for various assets along diverse dimensions (Hong et al., 2019; Ilhan et al., 2021; Kruttli et al., 2021; Seltzer et al., 2020; Alok et al., 2020; Krueger et al., 2020; Bansal et al., 2019; Barnett, 2019; Choi et al., 2020; Engle et al., 2019; Giglio et al., 2021; Bolton and Kacperczyk, 2021b,a) to which we contribute a novel market-based analysis of ambiguity aversion.

Our work also contributes to the important work on uncertainty analysis and the application of dynamic decision theory. As noted previously, this literature includes work on multiple priors (Gilboa and Schmeidler, 1989), model misspecification (Anderson et al., 2003; Maccheroni et al., 2006; Hansen and Sargent, 2007; Cerreia-Vioglio et al., 2021), and smooth ambiguity (Klibanoff et al., 2009; Hansen and Miao, 2018; Ilut and Schneider, 2023). There is also a growing literature exploring uncertainty analysis as it pertains to climate change (Kelly and Kolstad, 1999; Crost and Traeger, 2010; Lemoine and Traeger, 2012; Brock and Hansen, 2018; Barnett et al., 2020; Rudik, 2020; Barnett et al., 2021; Barnett, 2023; Barnett et al., 2023). Prior work has attempted to quantify ambiguity aversion using survey-based methods, such as eliciting imprecise probabilities or presenting hypothetical choice experiments (Bhandari et al., 2024; Ilut and Schneider, 2023). While informative, these approaches are limited by hypothetical bias and framing effects. Work by Brenner and Izhakian (2018) provides one of the few market-based measures related to uncertainty, estimating ambiguity (not ambiguity aversion) based on high frequency return volatility, with a recent application of the methodology to climate ambiguity (Rocciolo et al., 2024). Instead, we take a market-based approach integrated with scientific model projections to directly estimate uncertainty aversion. By examining how housing markets capitalize both expected SLR risk and uncertainty in inundation timing, we provide the first revealed-preference estimate of ambiguity aversion related to climate change. Hence, our paper helps address a key gap in the literature on robust decision-making under uncertainty, the pricing of climate risk in real estate markets, and the empirical discipline of ambiguity aversion—specifically but not limited to the context of climate-related risks.

2. A (Stylized) Robust Control Problem

We formulate a robust control framework to characterize an uncertainty-averse agent evaluating the long-run impact of sea-level rise (SLR) on asset value. The agent faces uncertainty both across and within future SLR scenarios and is therefore concerned about potential model ambiguity and misspecification. This section outlines the primitives of the environment, including functional forms for preferences and the optimally robust

probability distortions for model ambiguity and misspecification concerns.

2.1. Model Primitives

Our robust control framework captures two distinct layers of uncertainty. First, the agent faces ambiguity regarding which sea-level rise (SLR) scenario will materialize. Second, within each scenario, there is concern over the potential misspecification of the underlying risk model. We describe these two components separately before combining them.

Ambiguity Across Scenarios

Let $\mathcal{N} \equiv \{0, 1, \ldots, N\}$ denote the set of possible future SLR scenarios. For each scenario $i \in \mathcal{N}$, the agent has an initial evaluation of the scenario's utility, denoted by v_i. However, uncertainty over which scenario will occur motivates the agent to consider alternative probability distributions. Specifically, the agent starts with a benchmark prior $p = (p_0, \ldots, p_N)$ over $\mathcal{N}$ but then seeks a distorted distribution q on the simplex

$$\Delta(\mathcal{N}) \equiv \{q \in [0,1]^{N+1} : \sum_i q_i = 1\},$$

which is chosen to minimize the expected utility subject to a penalization term, which together reflects ambiguity aversion. This ambiguity aversion penalty is captured by a divergence function $D_\alpha(q\|p)$. Formally, the agent's problem at the scenario level is

$$q^* \equiv \arg \min_{q \in \Delta(\mathcal{N})} \left\{ \sum_i q_i v_i + D_\alpha(q\|p) \right\}.$$

Misspecification Within Each Scenario

Within any given scenario i, the agent models the SLR risk index S using a distribution f_i on the support $[0, T_0]$. The preliminary evaluation of utility in scenario i is given by

$$u_i \equiv \mathbb{E}[u(S) \mid f_i],$$

where $u(S)$ is the payoff function (specified later). Recognizing that the model f_i might be misspecified, the agent entertains alternative distributions g_i (absolutely continuous with respect to f_i) and penalizes deviations with a second divergence $D_\gamma(g_i\|f_i)$. The worst-case (or distorted) measure g_i^* is thus determined by

$$g_i^* \equiv \arg \min_{g_i \ll f_i} \{\mathbb{E}[u(S) \mid f_i] + D_\gamma(g_i\|f_i)\}.$$

Under this worst-case measure, the valuation in scenario i is updated to

$$v_i \equiv \mathbb{E}[u(S) \mid g_i^*].$$

Combining the Two Layers

Incorporating the within-scenario distortions into the evaluation of each scenario, the overall robust control problem addressing ambiguity and misspecification concerns becomes

$$q^* = \arg\min_{q \in \Delta(\mathcal{N})} \left\{ \sum_i q_i \left[\min_{g_i \ll f_i} \{ \mathbb{E}[u(S) \mid f_i] + D_\gamma(g_i \| f_i) \} \right] + D_\alpha(q \| p) \right\}.$$

The robust valuation of the asset, accounting for both uncertainty channels, is therefore

$$v^* \equiv \sum_i q_i^* \mathbb{E}[u(S) \mid g_i^*]. \tag{1}$$

2.2. Functional Forms and Assumptions

Housing Utility Utility is linear in the SLR risk index:

$$u(S) = \bar{v} - \beta S.$$

This functional form is motivated both by empirical hedonic regressions (see empirical section) and by microfoundations discussed below. The risk index S captures expected losses from SLR and enters utility negatively.

Divergence Penalties Both ambiguity and model misspecification are penalized using χ^2 divergence:

$$D_\alpha(q \| p) = \frac{1}{2\alpha} \sum_i \frac{(q_i - p_i)^2}{p_i}, \quad D_\gamma(g \| f) = \frac{1}{2\gamma} \int_0^{T_0} \frac{(g(S) - f(S))^2}{f(S)} dS.$$

Alternative specifications using Kullback–Leibler divergence will be considered in a robustness section.

Remark 1. The linearity of v in S simplifies our analysis, and can be partially justified by the empirical observation that the log of housing prices are approximately linear in the projected first year of inundation, once amenities are controlled for (see Figure 4). The linearity can be microfounded as follows. Consider a house that, in the absence of SLR risk, yields housing utility flow between period $t = 0$ and a terminal period T_0. However, with SLR risk, the housing utility flow will stop as soon as the house is permanently inundated by the rising seas. Denote the first period of inundation under SLR scenario i by T_i. Suppose the homeowner discounts the future at rate ρ. Then the present discounted value of the housing unit in scenario i is given by $v_i = \frac{1}{\rho}(1 - \exp(-\rho T_i \wedge T_0))$. Let $S_i \equiv (T_0 - T_i)^+$. Then the housing valuation can be approximated linearly by $v_i \approx \bar{v} - \beta S_i$, where $\bar{v} \equiv \frac{1}{\rho}(1 - \exp(-\rho T_0)$ and $\beta \equiv \exp(-\rho T_0)$.[2]

[2]The approximation is reasonable in our context. For example, with $T_0 = 2301 - 2001$ (our SLR projection

2.3. Decomposition of Uncertainty

The agent's valuation of future housing utility depends not only on expected outcomes but also on how uncertain or volatile those outcomes are—both across and within scenarios.

Let the average SLR risk index under the benchmark prior p be defined as:

$$\bar{S} \equiv \sum_i p_i S_i,$$

where $S_i = \mathbb{E}[S \mid f_i]$ is the expected value of the SLR risk index under scenario i.

We decompose overall uncertainty about SLR into two distinct components:

- **Uncertainty across scenarios (ambiguity)**, denoted U_{across}, captures the dispersion in average outcomes S_i across different scenarios. It reflects the ambiguity about which scenario will actually occur:

  $$U_{\text{across}} \equiv \sum_i p_i (S_i - \bar{S})^2.$$

  An equivalent way to express this weighted variance is by noting the standard identity for weighted variances:

  $$U_{\text{across}} = \frac{1}{2} \sum_{i,j} p_i p_j (S_i - S_j)^2 = \sum_{i<j} p_i p_j (S_i - S_j)^2.$$

  In words, when probabilities are equal this expression is (up to a multiplicative constant) the sum of the squared differences between all pairs of scenario means. In our setting, U_{across} therefore captures how different the average SLR risk is across the various scenarios. Larger values of U_{across} indicate greater uncertainty (or ambiguity) about which state of the world will prevail. This form of uncertainty is handled via ambiguity aversion (with parameter α).

  To see U_{across} better, consider the case where the agent has an equal benchmark prior probability $p_1 = p_2 = p_3 = 1/3$ over three scenarios $1, 2, 3$.[3] Furthermore suppose that $S_1 < S_2 < S_3$, and S_2 is the average of S_1 and S_3 (i.e., $S_2 = \frac{S_1+S_3}{2}$). Then U_{across} simplifies to $\frac{1}{6}(S_1 - S_3)^2$. Hence, U_{across} is proportional to the squared divergence between the best-case (lowest SLR) and worst-case (highest SLR) scenarios. The term thus intuitively encapsulates how ambiguous the future is: the larger the gap between the extreme scenarios, the greater the ambiguity.

ends in 2300 and the first year in our data is 2001), $T = 2273 - 2001$ (the median projected first year of inundation due to SLR under RCP 4.5 is 2273), $\rho = 2\%$, the value of v_i is 49.783, which is well approximated by 49.807.

[3] We discuss the validity of this assumption, and the impact of relaxing it, in Section 4.2.

- **Uncertainty within scenarios (risk)**, denoted U_{within}, captures the average variance of the SLR risk index within each scenario:

$$U_{\text{within}} \equiv \sum_i p_i \sigma_i^2,$$

where $\sigma_i^2 = \mathbb{V}[S \mid f_i]$. This reflects the agent's concern that even if the agent is perfectly confidently about the underlying scenario, the specific outcome S may still be unpredictable or volatile. This form of uncertainty is handled via risk/misspecification aversion (with parameter γ).

These two forms of uncertainty are penalized differently by the agent: within-scenario variability is handled via misspecification or risk aversion (parameter γ), while across-scenario ambiguity is penalized via ambiguity aversion (parameter α).

2.4. Characterizing Distorted Probabilities

Distortion of Within-Scenario Probabilities The following lemma characterizes how the agent distorts the distribution f_i within each scenario to reflect concerns about model misspecification.

Lemma 1. The optimal distortion of the within-scenario distribution f_i is:

$$\frac{g_i^*(S)}{f_i(S)} = [1 + \gamma\beta(S - S_i)]^+.$$

This expression implies that the agent gives more weight to outcomes S that are worse than the scenario's average S_i, and less weight to more optimistic ones. The parameter γ governs how aggressively the agent tilts the distribution toward downside risks. Intuitively, when γ is large, even small deviations from the average are heavily penalized, reflecting strong aversion to within-scenario uncertainty.

Distortion of Scenario Probabilities The agent also distorts the prior p_i over scenarios to account for ambiguity aversion. The next lemma characterizes this distortion:

Lemma 2. The optimal distorted probability over scenarios q_i^* is:

$$q_i^* = p_i \frac{[1 - \alpha(v_i - \bar{v})]^+}{\sum_j p_j [1 - \alpha(v_j - \bar{v})]^+}.$$

Assuming an interior solution, this simplifies to:

$$\frac{q_i^*}{p_i} = 1 + \alpha\beta(S_i - \bar{S}) + \alpha\gamma\beta^2(\sigma_i^2 - U_{\text{within}}). \tag{2}$$

This result has an intuitive decomposition:

- The term $\alpha\beta(S_i - \bar{S})$ captures the leading-order distortion: scenarios with worse-than-average expected SLR risk (i.e., higher S_i) received augmented probability weight under the agent's distorted distribution q_i^*. This reflects the agent's ambiguity aversion toward scenario-level variation.
- The second-order term $\alpha\gamma\beta^2(\sigma_i^2 - U_{\text{within}})$ represents the influence of within-scenario variability on belief distortion. While conceptually important, it plays a smaller role for small values of β, and is dominated by the first-order across-scenario term.

In what follows, we will focus primarily on the dominant first-order effect—the deviation $S_i - \bar{S}$—which governs how the agent reshapes their probability distribution over scenarios in response to uncertainty. This term drives the main wedge between the benchmark prior p_i and the distorted probabilities q_i^*, and serves as a key link between ambiguity aversion and scenario-level tail risk.

2.5. Robust Valuation

Recall that by combining the distorted probabilities across scenarios q^* and the within-scenario probability distortions g_i^*, the agent's overall robust valuation of the asset is defined in (1) as:

$$v^* \equiv \sum_i q_i^* \mathbb{E}_{g_i^*}[u(S)].$$

This expression integrates both dimensions of concern: ambiguity over which SLR scenario will prevail and possible model misspecification within each scenario.

The following proposition provides a tractable second-order characterization that will serve as the backbone for our empirical analysis:

Proposition 3. Assuming interior solutions to the minimization problems, the robust valuation is given by:

$$v^* = \bar{v} - \beta\bar{S} - \gamma\beta^2 U_{\text{within}} - \alpha\beta^2 U_{\text{across}} + \mathcal{O}(\beta^3). \tag{3}$$

This decomposition offers intuitive insight into how the agent's valuation responds to risk and uncertainty:

- The term $\bar{v} - \beta\bar{S}$ reflects the benchmark valuation based on average SLR risk. Here, β captures the marginal sensitivity of value to expected inundation or damage, and $\bar{S}$ is the mean expected loss under the benchmark prior.

- The term $-\gamma\beta^2 U_{\text{within}}$ captures the agent's aversion to risk *within* each scenario—namely, the concern that the true outcome might be worse than suggested by the baseline model f_i. This misspecification or risk leads to a conservative reweighting of possible outcomes within each scenario, effectively discounting scenarios with greater internal dispersion.[4]

- The term $-\alpha\beta^2 U_{\text{across}}$ represents ambiguity aversion *across* scenarios. When expected damages vary substantially across plausible futures, the agent treats this divergence as a form of model uncertainty and discounts the asset accordingly. The more disagreement across expert scenarios (i.e., higher U_{across}), the greater the downward pressure on value.

While these latter two terms appear at second order in β, they are central to our empirical strategy. They provide the key identification channels for distinguishing:

1. the role of physical exposure to risk ($\bar{S}$),

2. the agent's sensitivity to *model risk* (U_{within}), and

3. the agent's sensitivity to *model uncertainty (or ambiguity)* across future climate scenarios (U_{across}).

It is natural that these uncertainty-related terms depend on β. If exposure to sea-level rise S has no impact on value—i.e., if $\beta = 0$—then neither risk (variance within scenarios) nor ambiguity (disagreement across scenarios) would matter for valuation. In other words, uncertainty only matters insofar as the agent is sensitive to the underlying risk. The quadratic dependence on β arises naturally because the valuation penalties are driven by variation in S—that is, second moments—and utility is linear in S with slope $-\beta$, so uncertainty enters the valuation through terms involving β^2.

In short, while these terms appear at second order in β, their structure is essential: they encode distinct responses to within-model (U_{within}) and across-model (U_{across}) uncertainty. As such, they play a central role in our empirical analysis, guiding how we interpret variation in observed prices and estimate the parameters governing robust preferences. For completeness, the full expression beyond the second-order expansion is provided in the appendix and is used in our robustness checks.

[4]This decomposition shows the isomorphism between misspecification aversion in our specific model context and risk aversion in a constant absolute risk aversion setting. As the variation within scenarios can be attributed to both internal variability resulting from the chaotic dynamical climate system and modeling differences, both labels are appropriate. We therefore use these labels interchangeably in our framework.

2.6. Corner Solutions and Maxmin Preferences

The closed-form expressions for g_i^* and q_i^* rely on interior solutions—i.e., the distortions must remain strictly positive across the relevant support. For within-scenario distortions g_i^*, this requires that

$$\frac{g_i^*(S)}{f_i(S)} = 1 + \gamma\beta(S - S_i) > 0 \quad \text{for all } S \in \text{supp}(f_i).$$

Since $S \geq 0$, a sufficient condition to guarantee positivity is:

$$\gamma < \frac{1}{\beta S_i}.$$

This bound ensures that the agent's pessimistic tilt toward higher S values does not result in negative (or zero) distorted probabilities. In practice, this condition is easily satisfied: in our data, β is small (approximately 10 basis points), and for typical values of $S_i > 100$, the bound becomes $\gamma < 10$. Thus, for moderate levels of misspecification aversion, interior solutions are well-behaved and analytically tractable.

Limiting Case: Maxmin Preferences More interestingly, the solution for the distorted prior q_i^* can approach a *corner solution* when the agent's ambiguity aversion is sufficiently high. From Lemma 2, the expression for q_i^* is:

$$q_i^* = p_i \frac{[1 - \alpha(v_i - \bar{v})]^+}{\sum_j p_j [1 - \alpha(v_j - \bar{v})]^+}.$$

If the ambiguity aversion parameter α becomes large enough, the agent effectively discards all scenarios except the one with the lowest utility v_i. Formally:

Corollary 4 (Maxmin preferences)**.** If $\alpha > (\bar{v} - \min_i\{v_i\})^{-1}$, then:

$$v^* = \min_{i \in \mathcal{N}} v_i.$$

This limiting case corresponds precisely to the Gilboa and Schmeidler (1989) maxmin expected utility model, in which the decision-maker evaluates alternatives by their worst possible outcome. Remarkably, this classical model emerges as a special case of our smooth robust control framework. While this connection may seem straightforward in hindsight, it has often been overlooked in applied work—even though the maxmin model plays a prominent role in macroeconomics models with ambiguity aversion (Ilut and Schneider, 2023).

2.7. Risk Aversion but No Ambiguity Aversion

Consider another special case where ambiguity aversion is absent, i.e. $\alpha = 0$. In this case, the agent does not distort the baseline probabilities across scenarios and treats the uncertainty over which scenario will occur as a compound lottery. Consequently, the overall distribution of the SLR risk index S is simply given by the mixture $\sum_{i=0}^{N} p_i f_i(S)$, where $f_i(S)$ is the density of S in scenario i and p_i is its benchmark probability.

Then, the robust valuation in (3) collapses to

$$v^* = \sum_{i=0}^{N} p_i v_i = \bar{v} - \beta \bar{S} - \gamma \beta^2 U_{\text{within}}.$$

Thus, without ambiguity aversion, the only source of uncertainty that matters for valuation is the within-scenario variation U_{within}. In climate model projections, this variation comes from both internal variability from the chaotic nature of dynamical climate system models, as well as differences in the underlying biophysical model used to generate the projection. The empirical distinction between risk and misspecification in this sense is not obvious, and our stylized model treats them in a somewhat interchangeable way.

This corner case implies a testable hypothesis: *If agents are not ambiguity averse, then housing valuations should only respond to risk (as measured by U_{within}), and not to ambiguity (captured by U_{across}).* In our data, for instance, if one finds that differences in scenario means (i.e. the gap between best and worst case outcomes) do not systematically affect valuations, this would be consistent with $\alpha = 0$. Conversely, significant effects of U_{across} would suggest that ambiguity aversion plays an important role.

2.8. Mapping the Model to Data

Is it reasonable to expect all homebuyers to be highly sophisticated in their understanding of sea-level rise (SLR) projections *and* to worry about potential model uncertainty both within and across scenarios? Likely not. In fact, survey research shows that climate beliefs in the U.S. are highly polarized: a significant portion of the population either does not believe that global warming is happening or does not believe that it will harm them personally (Howe et al., 2015). This polarization manifests in market behavior. For example, Baldauf et al. (2020) and Bakkensen et al. (2025) document that housing prices only capitalize future SLR risk in counties with sufficient belief or worry about global warming according to surveys.

To reflect this heterogeneity in beliefs, and to connect our model to observed housing market data, we introduce a concrete scenario structure. We define:

- Scenario 0: No SLR risk, i.e., $S = 0$ with certainty. This captures the belief of buyers who are unconcerned about climate change or skeptical of scientific projections.

- Scenarios 1, 2, and 3: Correspond to Representative Concentration Pathways (RCP) 2.6, 4.5, and 8.5, respectively. These are standard emission scenarios used by the Intergovernmental Panel on Climate Change (IPCC) to model climate futures.[5]

We assume that homebuyers fall into two belief types:

- **Optimists**: Assign full probability to scenario 0. That is, $p_0 = 1$, and $p_1 = p_2 = p_3 = 0$.
- **Pessimists**: Dismiss the no-SLR scenario, i.e., $p_0 = 0$, and assign probability only over the scientific RCPs.

Under this setup, the valuation of optimists is:

$$v^{\text{opt}} = \bar{v},$$

since according to optimists, $S = 0$ with certainty. However, the valuation of pessimists will be given by the robust value v^* in (3).

This heterogeneity in beliefs motivates an empirical specification of the following form:

$$v^* = \bar{v} + \beta_0 \bar{S} + \delta_0^a U_{\text{across}} + \delta_0^w U_{\text{within}} + \text{Pessimist}_i \times (\beta_1 S + \delta_1^a U_{\text{across}} + \delta_1^w U_{\text{within}}), \quad (4)$$

where Pessimist_i is an indicator for whether buyer i is classified as a pessimist, and $S, U_{\text{across}}, U_{\text{within}}$ are the average SLR exposure and the two uncertainty measures, using probability weights that will be discussed below.

Hypotheses: Optimists assign no weight to SLR risk or uncertainty. Hence, we expect:

$$\beta_0 = \delta_0^a = \delta_0^w = 0.$$

Pessimists, in contrast, discount housing value based on exposure, ambiguity, and risk:

$$\beta_1 < 0, \quad \delta_1^a < 0, \quad \delta_1^w < 0.$$

Furthermore, zooming in on ambiguity aversion, our discussion in Section 2.6 implies a testable null:

$$\text{H}_0 : \quad \delta_1^a = 0 \quad \text{(no ambiguity aversion)},$$

$$\text{H}_A : \quad \delta_1^a < 0 \quad \text{(ambiguity-averse pessimists dislike divergence across scenarios)}.$$

[5]RCPs were developed by the IPCC as part of the Coupled Model Intercomparison Project (CMIP). They represent different greenhouse gas concentration trajectories: RCP 2.6 assumes aggressive mitigation; RCP 4.5 represents a stabilization path; and RCP 8.5 is a high-emission, "business-as-usual" scenario. These scenarios are widely used in climate modeling, including in sea-level rise projections such as those by DeConto and Pollard (2016) which account for potential ice sheet destabilization, that we employ in our analysis.

Rejecting H_0 would provide empirical support for models of deep uncertainty and robust preferences. If buyers respond only to U_{within}, it would suggest either misspecification or risk aversion, but not ambiguity aversion. If they also respond to U_{across}, this supports the presence of deep uncertainty and belief heterogeneity in market pricing. This specification provides a clear empirical test for model uncertainty aversion in housing markets.

3. Empirical analysis

3.1. Data construction

Our analysis is based on high-resolution property transaction data from CoreLogic, which provides extensive mortgage and real estate transaction records spanning 2001 to 2016. By integrating CoreLogic's comprehensive property records with state-of-the-art SLR modeling, we construct a uniquely detailed dataset that allows for an in-depth analysis of how climate change risk is transmitted to housing markets. This dataset includes detailed property characteristics such as sale prices, transaction dates, number of bedrooms and bathrooms, building size, lot size, elevation, and year built. The availability of precise geolocation coordinates for each property enables us to merge these records with sea-level rise (SLR) projections and construct property-specific inundation risk estimates.

To estimate the first year of inundation for each property, we follow the methodology detailed in Rodziewicz et al. (2022). This process involves multiple steps. First, we determine property-specific elevations by matching each property's coordinates with NOAA's Digital Elevation Models (DEMs). These elevation values are adjusted to a consistent vertical datum and account for differences in land height relative to sea level. Unlike simple elevation-based assessments, we incorporate NOAA's Sea-Level Rise Viewer shapefiles, which delineate areas projected to be hydrologically connected to the ocean. This ensures that properties behind levees, dunes, or other protective barriers are not misclassified as vulnerable to near-term SLR.

These property-specific elevations are then integrated with well known SLR projections developed by DeConto and Pollard (2016). These projections estimate the timing of inundation under three Representative Concentration Pathways (RCPs): RCP2.6 (low emissions, slow SLR), RCP4.5 (moderate emissions, moderate SLR), and RCP8.5 (high emissions, accelerated SLR). The projections incorporate global and local uncertainties in future sea-level rise, accounting for ice sheet loss dynamics, thermal expansion, and land subsidence. The resulting property-level SLR inundation timing estimates range from 2050 to 2300, with uncertainty bands capturing projections from the 10th to the 90th percentile.

Figure 2 illustrates how we construct the projections of the first year of inundation for

each property. The solid lines represent the median SLR projections for RCP2.6, RCP4.5, and RCP8.5 (while the dotted lines plot the corresponding 10th and 90th percentile projections under each scenario). For each property, its elevation (represented for a single house by the horizontal solid line) determines the first year of inundation. Taking the RCP4.5 scenario as an example, the intersection of the horizontal line with the median SLR curve for RCP4.5 determines the first year ($T_{RCP4.5}$) that the property will be permanently inundated under the median projection. The same method is applied for the RCP2.6 and RCP8.5 scenarios, yielding the corresponding first years of inundation $T_{RCP2.6}$ and $T_{RCP8.5}$.

Figure 3 presents the spatial distribution of projected inundation risk and uncertainty for properties in the Miami subsample. The left panel maps the median projected first year of permanent inundation under the moderate emissions scenario (RCP4.5) for over 180,000 properties. The center panel visualizes cross-scenario uncertainty by showing the absolute difference in inundation timing between the low (RCP2.6) and high (RCP8.5) emissions scenarios. The right panel shows within-scenario uncertainty under RCP4.5, measured as the absolute difference between the 90th and 10th percentile projected inundation years. A key observation is that uncertainty across scenarios (middle panel) appears much more pronounced than uncertainty within the RCP4.5 scenario (right panel), with many areas showing differences of over a century depending on emission trajectories. Additionally, spatial patterns suggest that properties farther inland tend to experience later inundation dates and greater variability in projections, while coastal properties generally face earlier and more consistent risk.

Beyond physical exposure to SLR, we analyze how buyers' perceptions of climate risk influence housing transactions. We incorporate county-level climate opinion data from the 2014 Yale Climate Opinion Survey (Howe et al., 2015), which provides estimates of the percentage of residents who believe climate change is occurring. This dataset, based on over 13,000 survey responses, serves as a proxy for buyers' climate awareness at the time of purchase. By linking this measure to property transactions, we evaluate how climate beliefs impact pricing, demand, and market behavior.

To control for broader economic and demographic trends, we merge our dataset with county-level statistics from the U.S. Census Bureau, including population density, median income levels, and local economic indicators. These controls help disentangle the effects of SLR risk from other economic factors affecting property values.

We restrict our analysis to single-family properties from 17 major coastal metropolitan statistical areas (MSAs) (Figure 1). To reduce the influence of extreme outliers, we exclude transactions with sale prices outside the \$50,000–\$10,000,000 range and winsorize key housing variables, including house price, building size, elevation, and age, at the 1st and

99th percentiles.

Table 1 provides summary statistics for the properties in our dataset. The average home sale price is $464,496 (std: $539,648), with a median price of $325,500. Houses typically have 3.25 bedrooms and 2.31 bathrooms, with an average building size of 2,049.8 square feet. The average property elevation is 82.11 feet above sea level, but 10 percent of homes sit below 6.8 feet, highlighting the range of exposure to coastal flooding. SLR projections suggest substantial variation in the first year of inundation across properties and climate scenarios. Under RCP2.6, the average projected inundation year is 287 years from the base year of 2010, while under RCP8.5, it drops to 234. Uncertainty in these estimates increases under more extreme climate scenarios, with projected standard deviations ranging from 13.6 years (RCP2.6) to 86.5 years (RCP8.5).

3.2. Specification

To empirically test the implications of our model and map its structure to housing market data, we estimate the following regression specification. This is a direct empirical analogue of the robust valuation equation (4) developed in Section 2.8, where market participants are allowed to differ in their beliefs about climate change.

$$\begin{aligned}\log V_{h,c,t} = {} & \beta_0 S_h + \delta_0^w U_{\text{within},h} + \delta_0^a U_{\text{across},h} \\ & + (\beta_1 S_h + \delta_1^w U_{\text{within},h} + \delta_1^a U_{\text{across},h}) \times \text{Pessimist}_c \\ & + X_{h,c,t}\Gamma + \theta_{Z\times D} + \tau_t + \text{const} + \varepsilon_{h,c,t} \end{aligned} \tag{5}$$

Here, $\log V_{h,c,t}$ is the log sale price of property h in county c and year t. The key explanatory variables correspond directly to components of the theoretical model:

SLR risk measure The term S_h denotes sea-level rise risk under the middle (RCP 4.5) scenario for property h, calculated (consistently with Remark 1) as:

$$S_h \equiv (2301 - T^{\text{p50}}_{\text{RCP }4.5,h})^+,$$

where $T^{\text{p50}}_{\text{RCP }4.5,h}$ is the median projected year of inundation for house h. Properties not expected to be inundated by 2300 are assigned $S_h = 0$, reflecting long-horizon discounting and practical irrelevance for pricing.

Uncertainty measures To test for both risk and ambiguity aversion, we include two distinct measures of uncertainty, both of which are calculated at the property level (thus the subscript h):

- **Within-scenario uncertainty**:

$$U_{\text{within},h} \equiv \left(T^{\text{p90}}_{\text{RCP 4.5},h} - T^{\text{p10}}_{\text{RCP 4.5},h}\right)^2,$$

  which captures the variance in inundation year projections under a fixed emissions trajectory.

- **Across-scenario uncertainty**:

$$U_{\text{across},h} \equiv \left(T^{\text{p50}}_{\text{RCP 2.6},h} - T^{\text{p50}}_{\text{RCP 8.5},h}\right)^2,$$

  which measures how much the projected year of inundation shifts depending on whether aggressive mitigation or high emissions prevail.

Belief heterogeneity and interpretation The variable Pessimist_c equals one if the county's climate concern score (from the Yale Climate Opinion Survey) exceeds the national median. This allows us to distinguish between *optimists* (who effectively ignore climate risks) and *pessimists* (who take them seriously).

This structure maps cleanly to the model: under the null hypothesis, optimists set $p_0 = 1$ and are insensitive to SLR exposure or uncertainty, so we expect

$$\beta_0 = \delta_0^w = \delta_0^a = 0.$$

In contrast, pessimists assign beliefs across scenarios 1–3 (the RCPs) and value robustness to model error. Therefore, the coefficients on the interaction terms

$$\beta_1,\ \delta_1^w,\ \delta_1^a$$

represent the sensitivity of climate-concerned markets to projected inundation risk, model variance, and deep scenario divergence, respectively.

Testable implications Recall that the model yields several testable predictions:

- $\beta_1 < 0$: Pessimists discount homes based on augmented SLR inundation risk.
- $\delta_1^w < 0$: Pessimists value robustness to within-scenario variation.
- $\delta_1^a < 0$: Pessimists discount homes when there is variation across scenarios.

A particularly sharp test comes from evaluating the null hypothesis of the third condition of there being no ambiguity aversion:

$$\text{H}_0 :\ \delta_1^a = 0 \qquad \text{vs.} \qquad \text{H}_A :\ \delta_1^a < 0.$$

Rejecting H_0 would offer direct evidence that market participants act as ambiguity-averse agents responding to deep uncertainty in climate projections.

Remarks on selection and sorting. A relevant consideration in evaluating market-based transaction prices is the potential effects of selection and sorting by the households deciding to buy houses in our data. The literature has provided evidence that attenuates such concerns, as well as emphasized various considerations to account for in constructing the empirical model specification to address such concerns.

In order to identify the implications of sea-level rise on house prices, a common choice for the empirical estimation in the existing literature has been to split the sample of house-price transactions based on climate beliefs. This choice has been motivated by the theoretical analysis done by Baldauf et al. (2020) and Bakkensen and Barrage (2022). Specifically, Baldauf et al. (2020) highlight in a model of housing choice that includes homophily considerations that house prices have differential responses to climate change and SLR risk. Bakkensen and Barrage (2022) show in their theoretical framework characterizing the housing market when there is heterogeneity in beliefs about climate change and SLR risk and amenity values of coastal living that house price differences consistent with empirical analyses in the literature can arise. It is this theoretical work that motivates our choice to split our estimates between climate pessimists and climate optimists as has been done in those frameworks, as well as others.[6]

A concern related to splitting the sample based on climate beliefs is the possibility of introducing selection bias. Specifically, the standard belief measure used in the literature, which we use as our baseline measure, is a country-level variable that splits houses based on above- or below-median responses to questions about climate change beliefs in the Yale Climate Opinion Survey. One potential issue with this measure is that it could be correlated with other unobserved variables impacting house prices and could therefore be a biased proxy for individual-level beliefs. Numerous papers in the burgeoning literature on the implications of sea-level rise for house prices have tested and taken into account issues associated with selection and sorting. Bernstein et al. (2019) provide empirical evidence to show that this belief measure does not proxy for prior flood damage to a house, a lack of investment in home upkeep to a house, house characteristics like age or size, or missed controls for "views" amenity. Bakkensen and Barrage (2022) show that their results of SLR impacts on houses split by beliefs do not come from differential expectations about damages, government assistance, or insurance reimbursements. Furthermore they provide evidence showing that past flood experience does impact beliefs about SLR risk, and that households with high concern about flooding are more likely to plan on selling their home in the near future than other households. Finally, Bakkensen et al. (2025) show that

[6] In future work we plan to develop our own theoretical framework generalizing these settings to emphasize the relevant mechanisms related to model uncertainty implications for house price impacts of SLR.

individual- and county-level measures of climate beliefs are highly correlated, that counties where a higher fraction of buyers are from actually have higher climate change concerns, and that concerns about selection should actually attenuate results related to the impact of SLR on house prices as optimists should pay more for high risk homes.

To further address concerns about selection bias, there are a number of steps proposed by the literature which we have or plan to account for in our estimation. First, Bernstein et al. (2019), Baldauf et al. (2020) and Bakkensen et al. (2025) have emphasized the importance of including a rich set of controls and fixed effects to capture as much of the individual- and country-level unobservables that could confound our results. For this reason we include a rich set of controls related to houses purchased in our sample, as well as their interactions with the SLR risk measure, and various combinations zip-code, coastal distance, time, and other fixed effects. In addition, we include county-level measures of income, population, and political affiliation to capture the more aggregated unobservable economic trends that could be correlated with SLR beliefs and SLR risk. Bakkensen et al. (2025) also recommends using various alternative questions about climate change beliefs from the Yale Climate Opinion survey, which we do. Lastly, we plan to incorporate the following additional tests based on (Oster, 2019; Bernstein et al., 2019; Bakkensen et al., 2025) to further validate our estimation results: i) testing alternative measures of climate beliefs, including individual-level measures, ii) estimating our results separately for owner-occupied versus non-owner occupied houses, iii) extending our sample to exploit time-series variation related to climate-related disaster exposure, updates in SLR projections, and along other dimensions, and iv) implementing bounding tests to evaluate the robustness of our estimates to the potential impacts of omitted variable bias.

3.3. Results

This section presents our empirical findings on how housing markets price sea-level rise (SLR) risk and uncertainty. The specification corresponds to equation (5), and the results directly test the hypotheses laid out in Section 2.8, particularly regarding belief heterogeneity, risk or misspecification aversion, and ambiguity aversion.

3.3.1 SLR risk capitalization

We begin by estimating the average effect of SLR risk on house prices. Table 2 presents results for the full sample and for subsamples split by climate belief. In the full sample (column 1), the coefficient on SLR risk (S_h) is small and statistically insignificant, suggesting that, on average, projected inundation dates do not meaningfully affect prices.

However, in climate-conscious counties (column 2), the effect is significantly negative: homes expected to be inundated sooner sell for less. Figure 4 provides a visual counterpart to this result. The bin scatter plot shows the relationship between SLR risk and log house price, after controlling for the full set of property and county-level characteristics. The downward slope of the fitted line confirms the negative price response observed in column (2), reinforcing the finding that home values decline as inundation dates become nearer in climate-aware markets.[7]

Column (3) interacts SLR risk with the Pessimist indicator. The estimated coefficient on the interaction term is $\hat{\beta}_1 = -0.0010489$, or about -10 basis points. This implies that in climate-concerned counties, a house expected to be inundated 100 years earlier sells for approximately 10% less, relative to an otherwise comparable home. The estimate of β_1 is central to deriving the long-run discount rate, which we discuss in detail in Section 4.1.

3.3.2 Which scenarios matter most?

To examine how buyers process scenario-specific projections, Table 3 estimates specifications using SLR risk computed under RCP 2.6 (S1), RCP 4.5 (S2), and RCP 8.5 (S3), focusing exclusively on climate-concerned counties. Each column includes a different subset of the scenario-specific risk variables, allowing us to assess which scenario receives the most weight in buyers' pricing decisions.

In column (1), we estimate the effect of SLR risk under RCP 2.6 alone. The coefficient is negative and significant, indicating that buyers penalize properties expected to be inundated even under optimistic scenarios. However, when both RCP 2.6 and RCP 4.5 are included in column (2), the coefficient on RCP 4.5 becomes the dominant loading, while the coefficient on RCP 2.6 shrinks and loses significance. This shift suggests that when buyers are presented with a moderate and a more optimistic path, their pricing behavior puts more weight on the worse of the two—in this case, RCP 4.5.

The worst-case loading becomes even clearer in column (3), where all three RCP scenarios are included simultaneously. The effect of RCP 8.5 emerges as the largest in magnitude and is strongly statistically significant, while the coefficients on RCP 2.6 and 4.5 are smaller and insignificant. This indicates that once the true worst-case scenario is available, buyers assign the most importance to it—discounting home values based on the most extreme plausible pathway for sea-level rise.

This pattern is consistent with our model of ambiguity-averse buyers. As described in

[7]Notably, this bin scatter specification employs more stringent fixed effects—ZIP $\times$ distance-to-coast bin $\times$ year fixed effects—ensuring that the variation exploited is within highly localized markets and further reducing concerns about omitted variable bias.

Section 2.6, when agents face uncertainty over which scenario is correct, the endogenous worst-case belief distribution q^* shifts weight toward the scenario associated with the most severe outcome. In the limiting max–min case, buyers behave as if only the worst-case scenario matters. c These findings reinforce the idea that housing markets in climate-aware regions do not simply price expected risk. Instead, they reflect precautionary behavior consistent with worst-case thinking under ambiguity, with valuation anchored to the most adverse plausible future in the information set.

3.3.3 Capitalization of across- vs. within-scenario uncertainty

Table 4 introduces our two measures of uncertainty: across-scenario ambiguity (U_{across}) and within-scenario risk (U_{within}). The results strongly support the second and third hypotheses outlined in Section 2.8:

$$\delta_1^a < 0 \quad \text{(ambiguity aversion)}, \qquad \delta_1^w \lesssim 0 \quad \text{(risk aversion)}.$$

Column (1) includes only the interaction of U_{across} with the Pessimist indicator. The coefficient is negative and significant, indicating that homes in climate-concerned counties are discounted more heavily when there is greater disagreement across emissions pathways. This is consistent with ambiguity-averse buyers placing extra weight on model divergence—interpreting it as a signal of aversion towards deep uncertainty.

Column (2) includes only the within-scenario uncertainty interaction term. The coefficient on Pessimist $\times U_{\text{within}}$ is also negative and statistically significant at the 10% level, suggesting some aversion to risk or misspecification within a fixed scenario (RCP 4.5). However, the effect is smaller in magnitude than for U_{across}.

Column (3) includes both uncertainty measures simultaneously. Here, the interaction term for U_{across} remains negative and statistically significant, while the coefficient on U_{within} becomes small and statistically insignificant. This confirms that across-scenario uncertainty has a more robust and salient effect on prices than within-scenario variance.

The pattern of results allows us to formally reject the null hypothesis of no ambiguity aversion:

$$\mathrm{H}_0 : \delta_1^a = 0 \qquad \text{vs.} \qquad \mathrm{H}_A : \delta_1^a < 0,$$

providing strong support for the view that market participants are ambiguity averse, not just risk averse. Buyers in climate-aware markets behave as though they discount futures in which climate models disagree most—consistent with robust control preferences or the max–min utility framework described in Section 2.6.

In sum, these results reinforce a central insight of our model: SLR exposure is not just about expected damages—it is about probability distortions coming from deep uncertainty.

We dig deeper into this point in Section 4.2, where we use these estimates to recover an implied parameter of ambiguity aversion from observed market behavior.

4. Market-based estimates

4.1. Market-based estimate of the long-run discount rate

Our framework provides a unique opportunity to revisit the debate on what are appropriate long-run discount rates in climate economics, as we can estimate the discount rate against long-run SLR risk using our data. In our model, the parameter β captures the sensitivity of house prices to shifts in the expected first year of inundation due to sea-level rise. Our empirical estimates of $\hat{\beta}_1$ of about 10 basis points, implying that, all else equal, a property with an expected inundation date 100 years earlier sells at a 10% lower price.

To interpret the magnitude of $\hat{\beta}_1 = -10$ basis points (column 3 Table 2) in terms of forward-looking behavior, we map the estimate to the model's pricing kernel as described in Remark 1, where $|\beta| = \exp(-\rho T_0)$. Solving for ρ with $|\beta| = 0.9$ and $T_0 = 2301 - 2001 = 300$, we obtain:

$$\rho = -\frac{1}{T_0}\log(|\beta|) \approx -\frac{1}{300}\log(0.9) \approx 0.0023,$$

or a long-run annual discount rate of approximately 2.3%.

This market-implied discount rate of 2.3% is notably lower than the 4% benchmark commonly used in integrated assessment models such as DICE (Nordhaus, 2013), and slightly below the 2.6% rate inferred from long-term housing markets in the UK and Singapore by Giglio et al. (2015, 2021). These differences are economically meaningful. Because the discount rate has a role in determining how heavily future climate damages are weighted in present-value terms, even small reductions can translate into substantially higher estimates of the social cost of carbon. In this sense, our approach offers a new, market-based calibration of long-run climate discounting—one grounded in actual asset prices rather than theoretical assumptions or global averages. It highlights the potential of disaggregated housing market data to inform climate economics, particularly in evaluating long-horizon risks that are otherwise difficult to price.

4.2. Market-based Estimates of Ambiguity and Misspecification Aversion

Now we derive the values for our structural parameters related to ambiguity and misspecification aversion, α and γ. To do this, we need a number of plausible simplifying assumptions. First, we assume a uniform baseline prior across SLR scenarios (i.e., $p_1 = p_2 = p_3 = \frac{1}{3}$).

[8] In addition, as guided by our data, we assume that the middle scenario approximates the average of the two extreme scenarios:

$$T_2 \approx \frac{T_1 + T_3}{2} \quad \text{and hence} \quad S_2 \approx \frac{S_1 + S_3}{2},$$

since $S_i \equiv 2301 - T_i$.[9] As a result, we obtain

$$\bar{S} \approx S_2,$$

and the expression for the variance of outcomes across scenarios further simplifies to

$$U_{\text{across}} \approx \frac{1}{6}(S_1 - S_3)^2,$$

as the deviations of S_1 and S_3 from S_2 are symmetric. Similarly, we assume

$$\sigma_2^2 \approx \frac{\sigma_1^2 + \sigma_3^2}{2},$$

another empirically guided choice. Then, by assuming normality for the SLR risk index outcomes within a given scenario, which allows us to exploit the percentiles reported for our within scenario projections as in Lochstoer and Muir (2022) and Nagel and Xu (2023), we can approximate the expression for the variance of outcomes within scenarios by

$$U_{\text{within}} \approx \frac{1}{2.56^2}(S_2^{p10} - S_2^{p90})^2.$$

This reduction means that instead of including nine separate terms—S_1, S_2, S_3, $(S_1 - S_2)^2$, $(S_2 - S_3)^2$, $(S_3 - S_1)^2$, σ_1^2, σ_2^2, and σ_3^2—our specification needs only three key terms: the median inundation time S_2, the squared difference across extreme SLR projection scenarios $(S_1 - S_3)^2$, and the squared difference across extreme percentiles for the median SLR projection scenario $(S_2^{p10} - S_2^{p10})^2$.

In mapping our theoretical model to the econometric specification in (5), we interpret the deviation in valuations for a pessimistic agent as

$$V^{\text{pes}} - V^{\text{opt}} = -\beta\, S_h - \alpha\, \beta^2\, U_{\text{across},h} - \gamma\, \beta^2\, U_{\text{within},h},$$

where:

[8] There is a stylized fact in the forecasting literature, explored and tested by Genre et al. (2013); Diebold and Shin (2019); Diebold et al. (2025), that equal-weighted priors do quite well relative to alternative priors when optimally combining forecasts. We view this result as suggestive that an equal-weighted prior in this setting where weights across scenarios is a reasonable baseline assumption. Moreover, in Appendix B, we relax this assumption and instead only assume a symmetric prior where $p_1 = p_3 = \frac{1-p_2}{2}$ and find very similar results.

[9] This approximation is reasonable. Figure 7 plots the histogram of $S_2 - \frac{S_1+S_3}{2}$, which has a substantial concentration of mass at 0.

- β denotes the sensitivity to the central SLR risk S_h,
- α measures ambiguity aversion (i.e., aversion to uncertainty across scenarios), and
- γ captures misspecification aversion (i.e., aversion to uncertainty within a scenario).

Thus, the interaction coefficients in equation (5) can be approximated as

$$|\beta_1| \approx \beta, \quad |\delta_1^a| \approx \alpha \frac{\beta^2}{6}, \quad \text{and} \quad |\delta_1^w| \approx \gamma \frac{\beta^2}{2.56^2}.$$

Using Column 1 Estimates: Based on the estimated results in Table 4, we take Column 1, which includes only across–scenario uncertainty, as our baseline quantitative case. To derive the implied model parameters and probability distortions, we begin with the estimates

$$|\beta_1| \approx 0.0008869, \quad |\delta_1^a| \approx 0.0000029.$$

Then, setting

$$\alpha \frac{(0.0008869)^2}{6} = 2.9 \times 10^{-6},$$

and noting that $(0.0008869)^2 \approx 7.87 \times 10^{-7}$, we solve for α:

$$\alpha \approx \frac{2.9 \times 10^{-6} \times 6}{7.87 \times 10^{-7}} \approx 22.1.$$

Under these estimates the implied probability tilt is such that the subjective probability of the high-end scenario adjusts to

$$q_3^* - p_3 \approx \frac{|\delta_1^a|}{|\beta_1|}(S_3 - S_1) = 0.17 \quad \text{and} \quad q_1^* - p_1 \approx -0.17, \tag{6}$$

reflecting a substantial shift in beliefs relative to the baseline.

Figure 6 visualizes the distortion from prior p to q^*. The estimated distortion implies that market participants assign disproportionately higher probability weight to the worst-case climate scenario than an unbiased prior would suggest. In economic terms, this shift means that homebuyers in coastal markets act as if extreme sea-level rise outcomes are significantly more likely than indicated by an equal-weighted prior across scientific projections. This finding provides the first known market-based estimate of probability distortions due to ambiguity aversion, quantifying how uncertainty reshapes expectations and influences asset prices in an economically significant way for climate risk valuation and policy design.

Using Column 3 Estimates: Though U_{within} is statistically insignificant in Column 3 of Table 4, which includes both across– and within–scenario uncertainty, we can examine the implied model parameters and probability distortions in this setting for comparison to our baseline quantitative outcomes. The estimates here are

$$|\beta_1| \approx 0.0009334, \quad |\delta_1^a| \approx 0.0000025, \quad |\delta_1^w| \approx 0.0000010.$$

For ambiguity aversion we have

$$\alpha \frac{(0.0009334)^2}{6} = 2.5 \times 10^{-6} \quad \Longrightarrow \quad \alpha \approx \frac{2.5 \times 10^{-6} \times 6}{(0.0009334)^2} \approx 17.22.$$

For misspecification aversion, the mapping is given by

$$\gamma \frac{(0.0009334)^2}{2.56^2} = 1.0 \times 10^{-6} \quad \Longrightarrow \quad \gamma \approx \frac{1.0 \times 10^{-6} \times 2.56^2}{(0.0009334)^2} \approx 7.5.$$

Thus, under the Column 3 estimates we obtain $\alpha \approx 17.22$ and $\gamma \approx 7.52$.

It is important to note that the expression for the worst-case probability distortion with respect to ambiguity aversion,

$$q_3^* - p_3 \approx \frac{|\delta_1^a|}{|\beta_1|}(S_3 - S_1),$$

depends solely on the ambiguity channel—specifically, the ratio of the ambiguity coefficient to the SLR sensitivity and the difference $S_3 - S_1$. Although the updated specification now includes misspecification aversion via γ, this channel does not enter into the calculation of the probability tilt. As a result, even with the revised values of α and the inclusion of γ, the implied increase in the high–end probability remains robust at approximately 17 percentage points.

5. Application: Willingness to adapt to SLR

Our framework also allows us to quantify, for the first time to our knowledge, the role of ambiguity aversion in adaptation decisions by linking empirically estimated discount rates and ambiguity parameters to homeowners' willingness to raise their property's elevation. We have in mind the adaptation technology of raising a house on stilts (Fried, 2022). In the model, this technology simply maps to an increase in the house's elevation above the sea level.[10] Under our model, the housing utility flow is normalized to 1 until inundation

[10]We can also alternatively consider a version of the exercise where the decision maker is a local authority that is considering building a sea wall.

occurs at time T, so that the value of the house is given by

$$V(T) = \frac{1 - e^{-\rho T}}{\rho}.$$

Sea-level rise is modeled as

$$\mathcal{S}_t = \mathcal{S}_0 e^{\mu t}, \quad \mathcal{S}_0 = 1,$$

and for a house with current elevation h, the inundation time is

$$T(h) = \frac{1}{\mu} \log\left(1 + \frac{h}{\mathcal{S}_0}\right).$$

A homeowner can raise the elevation by increasing h, and the marginal willingness to adapt is given by $\frac{\partial V(h)}{\partial h} = \frac{(\mathcal{S}_0+h)^{-\frac{\rho}{\mu}-1}}{\mu}$. Similarly, the willingness to raise elevation by 1 unit is captured by the difference $V(h+1) - V(h)$.

As a baseline, we set $\rho = 2.4\%$ (as previously derived from our SLR capitalization estimates) and set the rat of SLR to be $\mu_1 = 1.60\%$, $\mu_2 = 1.75\%$, and $S_3 = 1.97\%$ (calibrated to average inundation timings under the RCP2.6, RCP4.5, and RCP8.5 scenarios given the average house elevation in the data). We obtain some quantitative predictions for the cost of adaptation. Table 5 summarizes our preliminary estimates: columns 2, 3, and 4 show the willingness to raise the current elevation by 1 ft as a percentage of the current house value, while columns 5, 6, and 7 report the corresponding dollar amounts based on an average house price of $464,000. For example, for a house with a current elevation of 1 ft, the baseline probability weights imply a willingness to raise elevation by 1 ft of approximately 27.2% of its value, while under distorted probability weights (reflecting ambiguity aversion) this increases to 28.08%—a difference of about 0.79 percentage points (or roughly $3,657). This is the first time, to the best of our knowledge, that the role of ambiguity aversion in adaptation decisions can be quantified using empirically estimated discount rates and ambiguity parameters.

6. Conclusion

Our findings provide direct empirical evidence that ambiguity aversion plays a significant role in how housing markets capitalize climate risks. By linking observed house price discounts to uncertainty in sea-level rise projections, we quantify the degree to which market participants exhibit aversion to ambiguous climate outcomes. The estimated ambiguity aversion parameter varies with assumed prior beliefs, but our benchmark estimates suggest that homebuyers in climate-conscious markets apply substantial discounts to properties with greater uncertainty in their inundation timing. This highlights the importance of

considering not just expected climate risks, but also the uncertainty surrounding those risks when assessing asset prices and policy interventions. Future research could explore how these effects evolve over time as more information about climate change becomes available and whether markets adjust their responses dynamically to new scientific projections and policy signals.

References

Alok, S., Kumar, N., and Wermers, R. (2020). Do fund managers misestimate climatic disaster risk. *The Review of Financial Studies*, 33(3):1146–1183.

Anderson, E. W., Hansen, L. P., and Sargent, T. J. (2003). A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, 1(1):68–123.

Bakkensen, L., Phan, T., and Wong, R. (2025). Leveraging the disagreement on climate change: Theory and evidence. *The Journal of Political Economy (forthcoming)*.

Bakkensen, L. A. and Barrage, L. (2022). Going underwater? flood risk belief heterogeneity and coastal home price dynamics. *The Review of Financial Studies*, 35(8):3666–3709.

Baldauf, M., Garlappi, L., and Yannelis, C. (2020). Does climate change affect real estate prices? only if you believe in it. *The Review of Financial Studies*, 33(3):1256–1295.

Bansal, R., Kiku, D., and Ochoa, M. (2019). Climate change risk. *preprint. URL https://www. frbsf. org/economic-research/events/2019/november/economics-of-climate-change/files/P*.

Barnett, M. (2023). Climate change and uncertainty: An asset pricing perspective. *Management Science*.

Barnett, M., Brock, W., Hansen, L. P., et al. (2021). Climate change uncertainty spillover in the macroeconomy. *Macoreconomics Annual*.

Barnett, M., Brock, W., Hansen, L. P., and Hong, Z. (2023). How should climate change uncertainty impact social valuation and policy? *Working Paper*.

Barnett, M., Brock, W. A., and Hansen, L. P. (2020). Pricing Uncertainty Induced by Climate Change. *Review of Financial Studies*, 33(3):1024–1066.

Barnett, M. D. (2019). *A run on oil: Climate policy, stranded assets, and asset prices*. PhD thesis, The University of Chicago.

Barrage, L. and Nordhaus, W. (2024). Policies, projections, and the social cost of carbon: Results from the dice-2023 model. *Proceedings of the National Academy of Sciences*, 121(13):e2312030121.

Bernstein, A., Gustafson, M. T., and Lewis, R. (2019). Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*.

Bhandari, A., Borovička, J., and Ho, P. (2024). Survey data and subjective beliefs in business cycle models. *The Review of Economic Studies*, page rdae054.

Bolton, P. and Kacperczyk, M. (2021a). Do investors care about carbon risk? *Journal of Financial Economics*.

Bolton, P. and Kacperczyk, M. (2021b). Global pricing of carbon-transition risk. Technical report, National Bureau of Economic Research.

Brenner, M. and Izhakian, Y. (2018). Asset pricing and ambiguity: Empirical evidence. *Journal of Financial Economics*, 130(3):503–531.

Brock, W. A. and Hansen, L. P. (2018). Wrestling with Uncertainty in Climate Economic Models. *University of Chicago Becker Friedman Institute Working Paper*, (2019-71).

Burke, M., Davis, W. M., and Diffenbaugh, N. S. (2018). Large Potential Reduction in Economic Damages Under UN Mitigation Targets. *Nature*, 557(7706):549.

Caballero, R. J. and Krishnamurthy, A. (2008). Collective risk management in a flight to quality episode. *The Journal of Finance*, 63(5):2195–2230.

Cai, Y. and Lontzek, T. S. (2019). The social cost of carbon with economic and climate risks. *Journal of Political Economy*, 127(6):2684–2734.

Cerreia-Vioglio, S., Hansen, L. P., Maccheroni, F., and Marinacci, M. (2021). Making Decisions under Model Misspecification.

Choi, D., Gao, Z., and Jiang, W. (2020). Attention to global warming. *The Review of Financial Studies*, 33(3):1112–1145.

Colacito, R., Hoffmann, B., and Phan, T. (2018). Temperatures and growth: A panel analysis of the united states. *Journal of Money, Credit, and Banking*, 51(2-3):2019.

Crost, B. and Traeger, C. P. (2010). Risk and Aversion in the Integrated Assessment of Climate Change. Technical report, UC Berkeley, Working Paper Series 1354288.

Dasgupta, P. (2008). Discounting climate change. *Journal of risk and uncertainty*, 37:141–169.

DeConto, R. M. and Pollard, D. (2016). Contribution of antarctica to past and future sea-level rise. *Nature*, 531(7596):591–597.

Diebold, F. X., Mora, A., and Shin, M. (2025). On the wisdom of crowds (of economists). *arXiv preprint arXiv:2503.09287*.

Diebold, F. X. and Shin, M. (2019). Machine learning for regularized survey forecast combination: Partially-egalitarian lasso and its derivatives. *International Journal of Forecasting*, 35(4):1679–1691.

Engle, Robert F, I., Giglio, S., Kelly, B. T., Lee, H., and Stroebel, J. (2019). Hedging Climate Change News. (25734).

Fried, S. (2022). Seawalls and stilts: A quantitative macro study of climate adaptation. *The Review of Economic Studies*, 89(6):3303–3344.

Genre, V., Kenny, G., Meyler, A., and Timmermann, A. (2013). Combining expert forecasts: Can anything beat the simple average? *International Journal of Forecasting*, 29(1):108–121.

Giglio, S., Maggiori, M., Rao, K., Stroebel, J., and Weber, A. (2021). Climate change and long-run discount rates: Evidence from real estate. *The Review of Financial Studies*, 34(8):3527–3571.

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

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, 18(2):141–153.

Goldsmith-Pinkham, P., Gustafson, M. T., Lewis, R. C., and Schwert, M. (2023). Sea-level rise exposure and municipal bond yields. *The Review of Financial Studies*, 36(11):4588–4635.

Gollier, C. (2024). Discounting. In *Handbook of the Economics of Climate Change*, volume 1, pages 113–142. Elsevier.

Golosov, M., Hassler, J., Krusell, P., and Tsyvinski, A. (2014). Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica*, 82(1):41–88.

Hansen, L. P. (2014). Nobel lecture: Uncertainty outside and inside economic models. *Journal of Political Economy*, 122:945–987.

Hansen, L. P. and Miao, J. (2018). Aversion to Ambiguity and Model Misspecification in Dynamic Stochastic Environments. *Proceedings of the National Academy of Sciences*, 115(37):9163–9168.

Hansen, L. P. and Sargent, T. J. (2001). Robust Control and Model Uncertainty. *The American Economic Review*, 91(2):60–66.

Hansen, L. P. and Sargent, T. J. (2007). Recursive Robust Estimation and Control without Commitment. *Journal of Economic Theory*, 136(1):1–27.

Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press, Princeton, NJ.

Hong, H., Li, F. W., and Xu, J. (2019). Climate Risks and Market Efficiency. *Journal of Econometrics*, 208(1):265–281.

Howe, P. D., Mildenberger, M., Marlon, J. R., and Leiserowitz, A. (2015). Geographic variation in opinions on climate change at state and local scales in the USA. *Nature Climate Change*, 5(6):596–603.

Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D., Muir-Wood, R., Wilson, P., Oppenheimer, M., et al. (2017). Estimating Economic Damage from Climate Change in the United States. *Science*, 356(6345):1362–1369.

Ilhan, E., Sautner, Z., and Vilkov, G. (2021). Carbon tail risk. *The Review of Financial Studies*, 34(3):1540–1571.

Ilut, C. and Schneider, M. (2023). Chapter 24 - ambiguity. In Bachmann, R., Topa, G., and van der Klaauw, W., editors, *Handbook of Economic Expectations*, pages 749–777. Academic Press.

Intergovernmental Panel on Climate Change (2014). *Climate Change 2014: Synthesis Report*. IPCC, Geneva, Switzerland. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.

Kelly, D. L. and Kolstad, C. D. (1999). Bayesian Learning, Growth, and Pollution. *Journal of Economic Dynamics and Control*, 23(4):491–518.

Keynes, J. M. (1921). *A treatise on probability*. London: MacMillan and Company.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2009). Recursive {Smooth} {Ambiguity} {Preferences}. *Journal of Economic Theory*, 144:930–976.

Knight, F. H. (1921). *Risk, uncertainty and profit*, volume 31. Cambridge: Houghton Mifflin Company.

Kopp, R. E., Horton, R. M., Little, C. M., Mitrovica, J. X., Oppenheimer, M., Rasmussen, D., Strauss, B. H., and Tebaldi, C. (2014). Probabilistic 21st and 22nd century sea-level projections at a global network of tide-gauge sites. *Earth's future*, 2(8):383–406.

Krueger, P., Sautner, Z., and Starks, L. T. (2020). The importance of climate risks for institutional investors. *The Review of Financial Studies*, 33(3):1067–1111.

Kruttli, M., Roth Tran, B., and Watugala, S. W. (2021). Pricing Poseidon: Extreme Weather Uncertainty and Firm Return Dynamics. *Working Paper. Available at SSRN 3284517*.

Lemoine, D. M. and Traeger, C. P. (2012). Tipping Points and Ambiguity in the Economics of Climate Change. Working Paper 18230, National Bureau of Economic Research.

Lochstoer, L. A. and Muir, T. (2022). Volatility expectations and returns. *The Journal of Finance*, 77(2):1055–1096.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Dynamic Variational Preferences. *Journal of Economic Theory*, 128(1):4–44.

Murfin, J. and Spiegel, M. (2020). Is the risk of sea level rise capitalized in residential real estate? *The Review of Financial Studies*, 33(3):1217–1255.

Nagel, S. and Xu, Z. (2023). Dynamics of subjective risk premia. *Journal of Financial Economics*, 150(2):103713.

Nordhaus, W. (2013). *The climate casino: Risk, uncertainty, and economics for a warming world*. Yale University Press.

Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204.

Painter, M. (2020). An inconvenient cost: The effects of climate change on municipal bonds. *Journal of Financial Economics*, 135(2):468–482.

Rocciolo, F., Billio, M., Guidolin, M., and Izhakian, Y. Y. (2024). Pricing climate ambiguity.

Rodziewicz, D., Amante, C. J., Dice, J., and Wahl, E. (2022). Housing market impairment from future sea-level rise inundation. *Environment Systems and Decisions*, 42(4):637–656.

Rudik, I. (2020). Optimal climate policy when damages are unknown. *American Economic Journal: Economic Policy*, 12(2):340–73.

Seltzer, L., Starks, L. T., and Zhu, Q. (2020). Climate regulatory risks and corporate bonds. *Nanyang Business School Research Paper*, (20-05).

Stern, N. (2007). *The Stern Review: The Economics of Climate Change*. Cambridge University Press.

Weitzman, M. L. (2001). Gamma discounting. *American Economic Review*, 91(1):260–271.

Weitzman, M. L. (2012). GHG Targets as Insurance Against Catastrophic Climate Damages. *Journal of Public Economic Theory*, 14(2):221–244.

Weitzman, M. L. (2013). Tail-hedge discounting and the social cost of carbon. *Journal of Economic Literature*, 51(3):873–882.

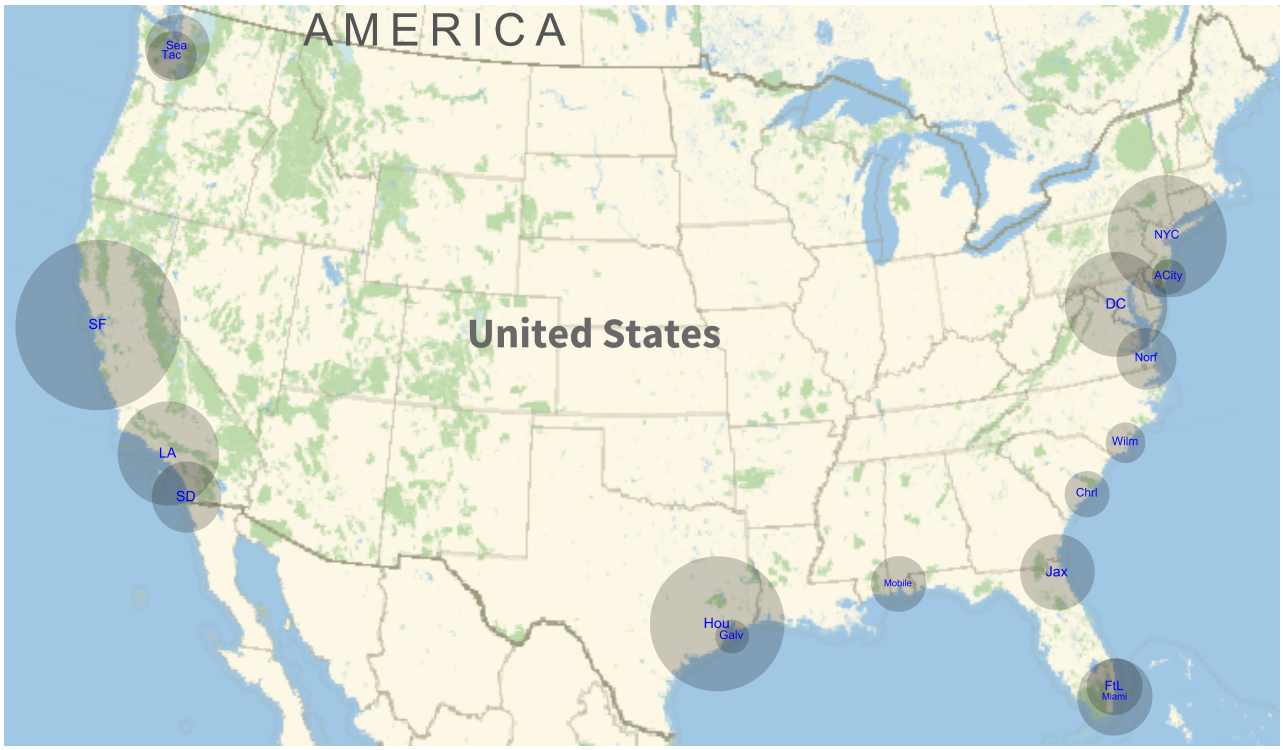


Figure 1: Geographic coverage: Our sample (of more than 2 million housing transactions) covers 17 major coastal metropolitan statistical areas (MSAs). Each MSA's circle radius is proportional to the sample size in that MSA.

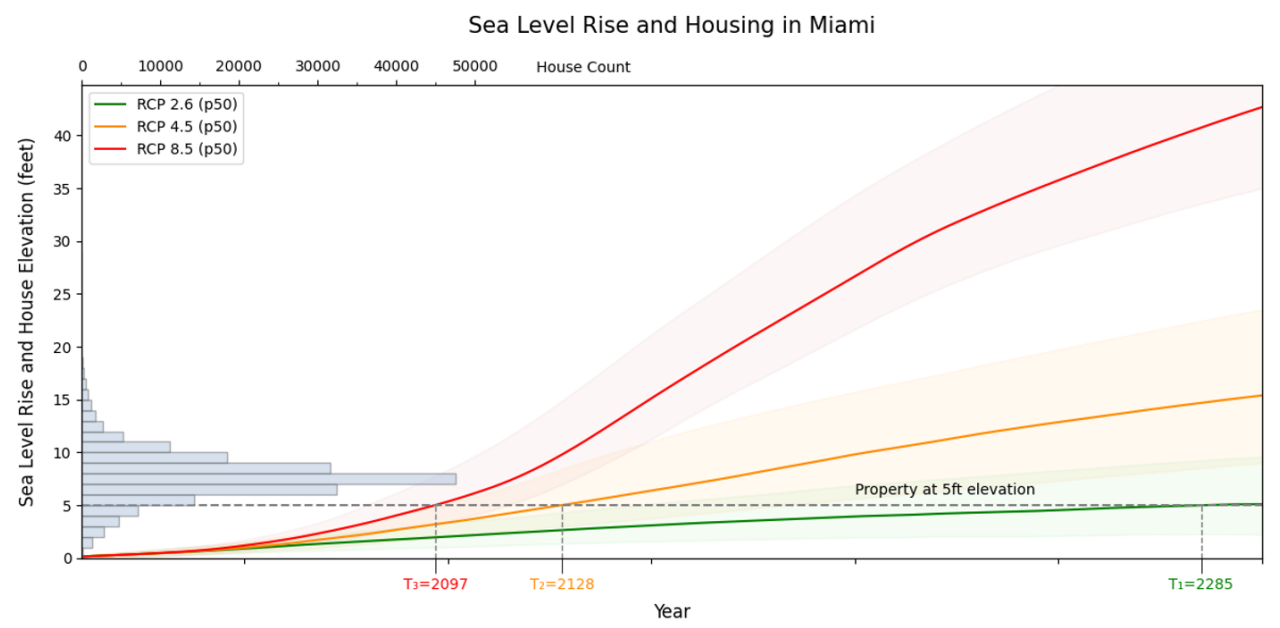


Figure 2: Example of how we construct each property's first year of inundation under different climate scenarios (RCP 2.6, 4.5 and 8.5). The histogram shows the distribution of the property elevation in the Miami subsample.

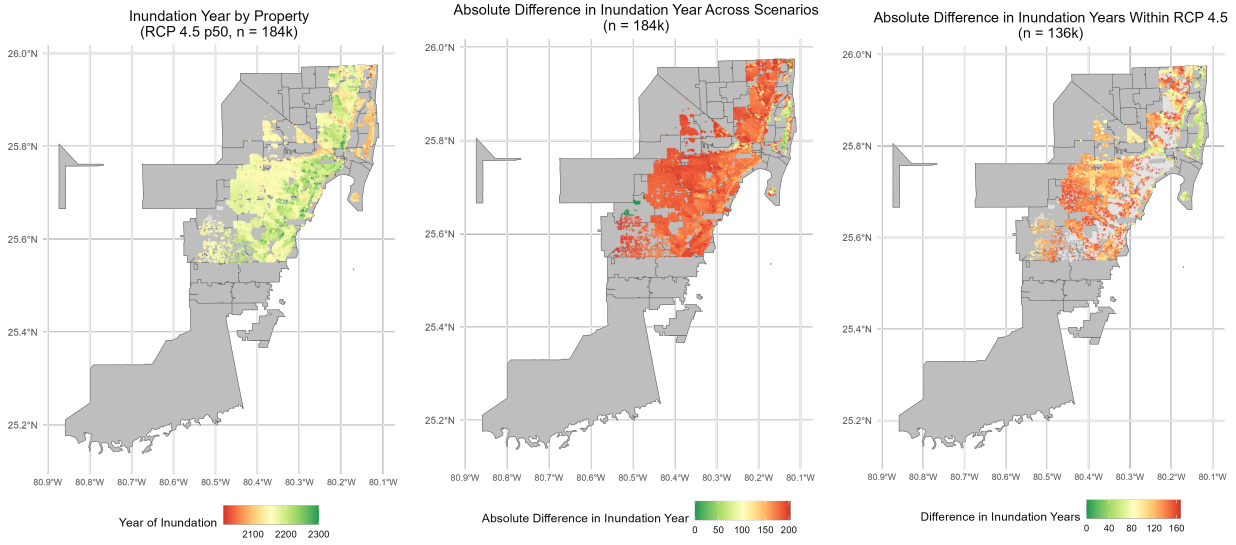


Figure 3: Spatial variation of our SLR measure, uncertainty, and risk in Miami subsample. The left map shows the median projected first year of permanent inundation under the middle scenario (RCP 4.5). The middle map shows a measure of uncertainty across scenarios: the absolute difference between the projected first year of inundation under the best (RCP 2.6) and worst (RCP 8.5) scenarios. The right map shows a measure of uncertainty within a scenario: the absolute difference between the projected first year of inundation at the 10th and 90th percentiles under the middle scenario (RCP 4.5).

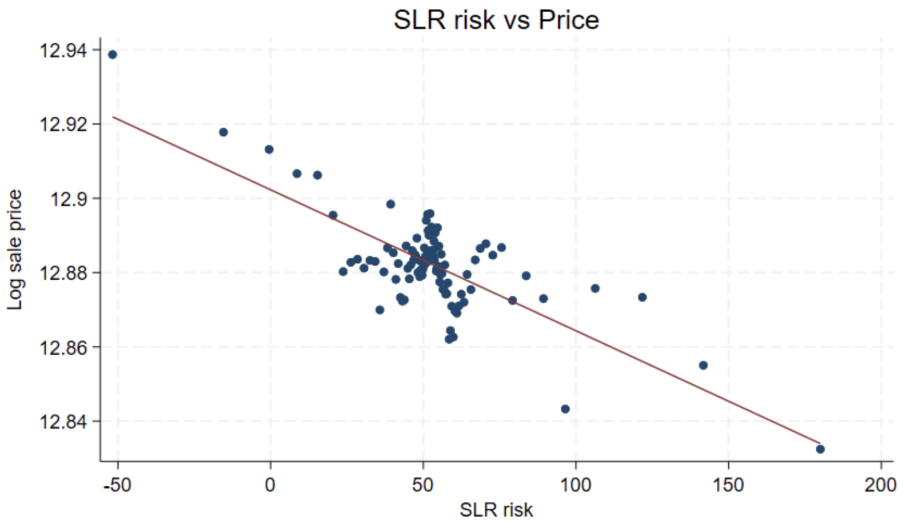


Figure 4: Bin scatter of SLR Risk and Log of House Sale Price, after controlling for property controls (building age, elevation, number of bedrooms, number of bathrooms, building square footage), the house's county controls (average county income and county population for the year of transaction, and GOP voter share in the last presidential election), and ZIP by distance to coast bin by sale year fixed effects.

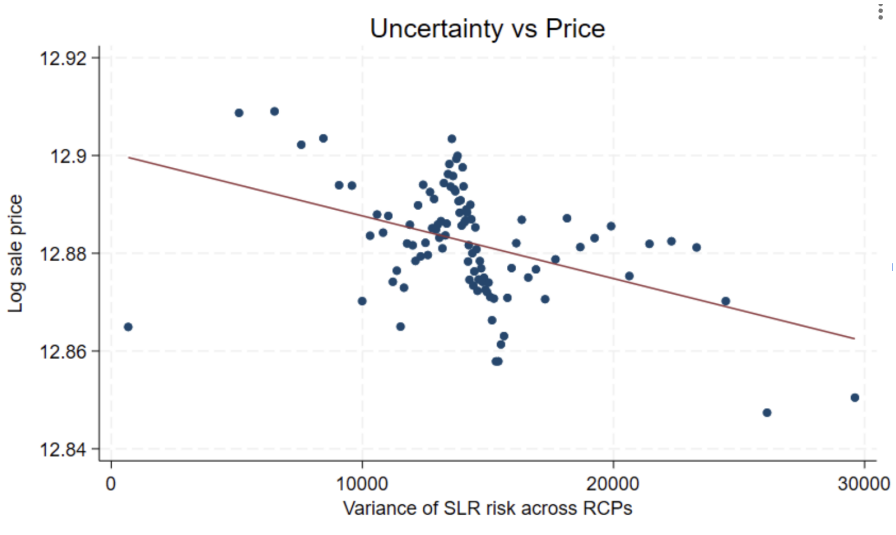


Figure 5: Bin scatter of SLR Uncertainty ($:= (T^{\text{median}}_{\text{RCP 2.6}} - T^{\text{median}}_{\text{RCP 8.5}})^2$) and Log of House Sale Price. The controls and fixed effects include those described in the caption of Figure 4 and additionally the interactions of the controls with *SLR Risk*.

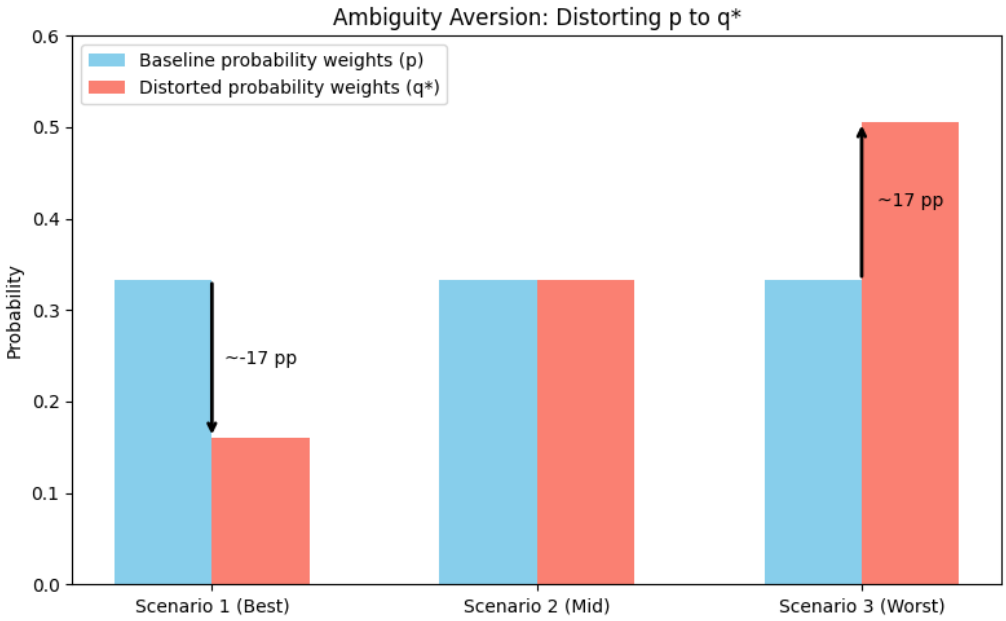


Figure 6: Model-implied belief distortion due to estimated ambiguity distortion parameter $1/\xi \approx 22$.

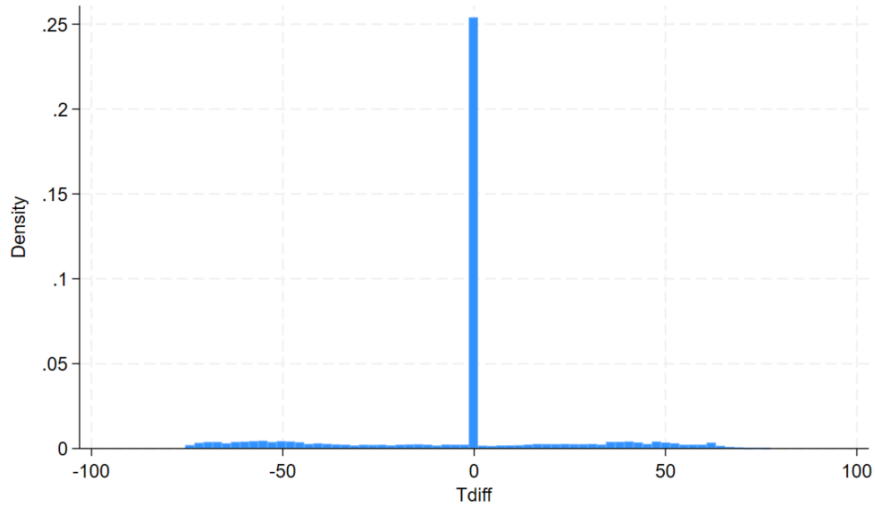


Figure 7: Histogram of $S_2 - \frac{S_1+S_3}{2}$ (i.e., $S^{p50}_{RCP4.5} - \frac{S^{p50}_{RCP2.6}+S^{p50}_{RCP8.5}}{2}$).

Housing characteristics						
	Mean	Std	p10	p50	p90	N
House sale price ($)	464,496.00	539,648.10	117,000.00	325,500.10	879,000.40	2,278,863
Number of bedrooms	3.25	0.88	2.00	3.00	4.00	2,449,749
Number of bathrooms	2.31	1.04	1.00	2.00	4.00	2,680,114
Building square feet	2,049.80	975.59	1,082.00	1,824.00	3,284.00	2,628,989
Property elevation (feet)	82.11	106.91	6.80	32.53	260.61	2,100,303
House age (years)	39.29	28.36	1.00	39.00	78.00	2,172,837
First year of SLR inundation, median projection undereach RCP scenarios, up to 2300						
	Mean	Std	Corr	$T_{rcp2.6}$	$T_{rcp4.5}$	$T_{rcp8.5}$
$T_{rcp2.6}$	2,297.27	23.52	$T_{rcp2.6}$	1.0000		
$T_{rcp4.5}$	2,273.39	56.30	$T_{rcp4.5}$	0.5113	1.0000	
$T_{rcp8.5}$	2,244.11	78.03	$T_{rcp8.5}$	0.3438	0.8422	1.0000
$\sigma_T = T_{rcp2.6} - T_{rcp8.5}$	53.16	73.34				
σ_T^2	8205.05	12735.54				

Table 1: Summary statistics of key selected variables.

	log House Sale Price		
S2	0.0001455	-0.0003327**	0.0000290
	(0.0001305)	(0.0001464)	(0.0007480)
Pessimist # S2			-0.0010489***
			(0.0003752)
Sample	Full	Pessimist	Full
Controls (house & county)	✓	✓	✓
Controls x S			✓
FE	ZxD Y	ZxD Y	ZxD Y
N	1058703	488638	1058703
R2	0.807	0.810	0.808

Table 2: SLR Risk Capitalization

	log House Sale Price		
S1	-0.0003869***	-0.0001918	-0.0002203
	(0.0001462)	(0.0001725)	(0.0001748)
S2		-0.0002891*	0.0001549
		(0.0001664)	(0.0002030)
S3			-0.0006452***
			(0.0001565)
Sample	Pessimist	Pessimist	Pessimist
Controls	✓	✓	✓
FE	ZxD Y	ZxD Y	ZxD Y
N	488582	488582	488582
R2	0.810	0.810	0.811

Table 3: Worst case (RCP 8.5) matters most.

	log House Sale Price		
S2	0.0002067	0.0003681	0.0003412
	(0.0007598)	(0.0007779)	(0.0007785)
Uncertainty across (RCP 2.6-8.5)	0.0000017		0.0000015
	(0.0000012)		(0.0000012)
Uncertainty within (p10-p90 RCP4.5)		0.0000016	0.0000010
		(0.0000015)	(0.0000013)
Pessimist # S2	-0.0008869**	-0.0010780***	-0.0009334**
	(0.0003575)	(0.0003785)	(0.0003663)
Pessimist # Uncertainty across	-0.0000029**		-0.0000025**
	(0.0000012)		(0.0000013)
Pessimist # Uncertainty within		-0.0000024*	-0.0000010
		(0.0000013)	(0.0000013)
Sample	Full	Full	Full
Controls	✓	✓	✓
Controls x S	✓	✓	✓
FE	ZxD Y	ZxD Y	ZxD Y
N	1,058,703	1,058,703	1,058,703
R2	0.808	0.808	0.808

Table 4: Capitalization of Across-Scenario vs. Within-Scenario Uncertainty

Current house elevation (ft)	Willingness to raise elevation by 1ft (% of current house value) Baseline prob weights	Distorted prob weight	Difference (pp)	Willingness to raise elevation by 1ft ($, based on average house price) Baseline prob weights	Distorted prob weight	Difference
1	27.29%	28.08%	0.79%	$126,917	$130,574	$3,657
2	9.51%	9.89%	0.38%	$44,217	$45,991	$1,774
3	4.78%	5.01%	0.24%	$22,221	$23,317	$1,097
4	2.84%	3.00%	0.16%	$13,213	$13,968	$755
5	1.87%	1.99%	0.12%	$8,681	$9,236	$554
6	1.31%	1.40%	0.09%	$6,097	$6,523	$425
7	0.97%	1.04%	0.07%	$4,493	$4,830	$337
8	0.74%	0.80%	0.06%	$3,434	$3,708	$274
9	0.58%	0.63%	0.05%	$2,700	$2,927	$227
10	0.47%	0.51%	0.04%	$2,173	$2,364	$191

Table 5: Model-implied willingness to raise elevation by 1 ft, expressed both as a percentage of current house value (columns 2–4) and in dollars (columns 5–7), based on an average house price of $464,000 in our sample.

A. Appendix: Theory

A.1. Proof of Lemmas 1

Proof. For each scenario i, the agent minimizes

$$\min_{g_i \ll f_i} \mathbb{E}_{g_i}[u(S)] + D_\gamma(g_i \| f_i),$$

with

$$u(S) = \bar{v} - \beta S, \quad D_\gamma(g_i \| f_i) = \frac{1}{2\gamma} \int \frac{\big(g_i(S) - f_i(S)\big)^2}{f_i(S)}\, dS.$$

Introducing a Lagrange multiplier λ_i for the constraint $\int g_i(S)dS = 1$, the Lagrangian is

$$\mathcal{L}_i(g_i, \lambda_i) = \int \left\{ (\bar{v} - \beta S) g_i(S) + \frac{1}{2\gamma} \frac{(g_i(S) - f_i(S))^2}{f_i(S)} \right\} dS + \lambda_i \left(\int g_i(S) dS - 1 \right).$$

Taking the first-order condition (FOC) with respect to $g_i(S)$ yields

$$\bar{v} - \beta S + \frac{1}{\gamma} \frac{g_i(S) - f_i(S)}{f_i(S)} + \lambda_i = 0,$$

so that

$$g_i^*(S) = f_i(S) \Big[1 - \gamma \big(\bar{v} - \beta S + \lambda_i \big) \Big].$$

Enforcing $\int g_i^*(S) dS = 1$ and using

$$\int f_i(S) dS = 1, \quad \int S\, f_i(S) dS = S_i,$$

we obtain

$$\lambda_i = \beta S_i - \bar{v}.$$

Thus, the distorted density becomes

$$g_i^*(S) = f_i(S)\Big[1 + \gamma\beta(S - S_i)\Big].$$

The distorted expected utility in scenario i is then

$$v_i = \mathbb{E}_{g_i^*}[u(S)] = \int (\bar{v} - \beta S) g_i^*(S) dS = \bar{v} - \beta S_i - \gamma\beta^2\sigma_i^2,$$

where $\sigma_i^2 = \mathrm{Var}_{f_i}(S)$.

Intuitively, the distortion g_i^* tilts the baseline density f_i so that outcomes with higher losses (i.e. larger S) receive more weight. This leads to a penalty of $\gamma\beta^2\sigma_i^2$ in the expected payoff. □

A.2. Proof of Lemma 2

Proof. The agent now distorts the scenario probabilities by choosing $q \in \Delta\mathcal{N}$ to minimize

$$\min_{q \in \Delta\mathcal{N}} \left\{ \sum_i q_i v_i + D_\alpha(q\|p) \right\},$$

with the χ^2 divergence

$$D_\alpha(q\|p) = \frac{1}{2\alpha} \sum_i \frac{(q_i - p_i)^2}{p_i}.$$

Setting up the Lagrangian (with multiplier μ for $\sum_i q_i = 1$) and differentiating, we find the FOC implies

$$q_i^* = p_i \left[1 + \alpha\Big(v_i - \sum_j p_j v_j \Big) \right].$$

Substitute

$$v_i = \bar{v} - \beta S_i - \gamma\beta^2\sigma_i^2,$$

and define the benchmark quantities

$$\bar{S} = \sum_i p_i S_i, \quad U_{\text{within}} = \sum_i p_i \sigma_i^2,$$

so that

$$\sum_j p_j v_j = \bar{v} - \beta\bar{S} - \gamma\beta^2 U_{\text{within}}.$$

It follows that

$$q_i^* = p_i\Big[1 + \alpha\beta(S_i - \bar{S}) + \alpha\gamma\beta^2(\sigma_i^2 - U_{\text{within}})\Big].$$

Intuitively, the distortion q_i^* shifts the scenario weights from the benchmark p_i by penalizing scenarios with higher-than-average S_i (more loss) and larger risk σ_i^2. The terms $\alpha\beta(S_i - \bar{S})$ and $\alpha\gamma\beta^2(\sigma_i^2 - U_{\text{within}})$ capture the sensitivity to uncertainty across scenarios. □

A.3. Proof of Proposition 3

Proof. The robust valuation is given by

$$v^* = \sum_i q_i^*\, v_i = \sum_i p_i\Big[1 + \alpha\beta(S_i - \bar{S}) + \alpha\gamma\beta^2(\sigma_i^2 - U_{\text{within}})\Big]\,\Big[\bar{v} - \beta S_i - \gamma\beta^2\sigma_i^2\Big].$$

Expanding and using the properties

$$\sum_i p_i(S_i - \bar{S}) = 0, \quad \sum_i p_i(\sigma_i^2 - U_{\text{within}}) = 0,$$

one finds that the leading terms yield

$$v^* = \bar{v} - \beta\bar{S} - \gamma\beta^2 U_{\text{within}} - \alpha\beta^2 U_{\text{across}},$$

with

$$U_{\text{across}} = \sum_i p_i(S_i - \bar{S})^2.$$

The higher-order corrections come out as

$$-2\alpha\gamma\beta^3 \sum_i p_i\, S_i(\sigma_i^2 - U_{\text{within}}) - \alpha\gamma^2\beta^4 \sum_i p_i\, \sigma_i^2(\sigma_i^2 - U_{\text{within}}).$$

We can expressing higher-order terms in terms of covariance and variance:
Define

$$\text{Cov}_p(S_i, \sigma_i^2) = \sum_i p_i(S_i - \bar{S})(\sigma_i^2 - U_{\text{within}}),$$

and

$$\text{Var}_p(\sigma_i^2) = \sum_i p_i(\sigma_i^2 - U_{\text{within}})^2.$$

Since

$$\sum_i p_i\, S_i(\sigma_i^2 - U_{\text{within}}) = \text{Cov}_p(S_i, \sigma_i^2)$$

and

$$\sum_i p_i\, \sigma_i^2(\sigma_i^2 - U_{\text{within}}) = \text{Var}_p(\sigma_i^2),$$

the combined higher-order corrections become

$$-2\alpha\gamma\beta^3 \operatorname{Cov}_p(S_i, \sigma_i^2) - \alpha\gamma^2\beta^4 \operatorname{Var}_p(\sigma_i^2).$$

Thus, the full robust valuation is

$$v^* = \bar{v} - \beta\bar{S} - \gamma\beta^2 U_{\text{within}} - \alpha\beta^2 U_{\text{across}} \underbrace{-\alpha\gamma\beta^3 \left[2 \operatorname{Cov}_p(S_i, \sigma_i^2) + \gamma\beta \operatorname{Var}_p(\sigma_i^2)\right]}_{\mathcal{O}(\beta^3)}.$$

□

B. Appendix: Robustness Check with Alternative Priors

In this appendix we show that the implied probability distortion remains robust even if we relax the uniform prior assumption. Instead, assume that the baseline priors are symmetric, so that

$$p_1 = p_3 = \frac{1 - p_2}{2},$$

and that the middle scenario is approximately the average of the two extremes, i.e.,

$$\frac{S_1 + S_3}{2} \approx S_2.$$

Under these assumptions the uncertainty across scenarios is given by

$$U = \sum_{i=1}^{3} p_i \, (S_i - \bar{S})^2 \approx p_1 \, (S_1 - S_2)^2 + p_3 \, (S_3 - S_2)^2.$$

Because $S_3 - S_2 \approx S_2 - S_1 \approx \frac{S_3 - S_1}{2}$ and $p_1 = p_3 = \frac{1-p_2}{2}$, it follows that

$$U \approx (1 - p_2) \frac{(S_3 - S_1)^2}{4}.$$

In the theoretical model the ambiguity penalty is given by $-\beta^2 \, \alpha \, U$, where β is the sensitivity to SLR risk and α denotes ambiguity aversion. In the empirical specification the corresponding term is captured by the coefficient on $(S_3 - S_1)^2$ (denoted $\hat{\delta}_1^a$). Equating the two we have

$$\beta^2 \, \alpha \, \frac{(1 - p_2)(S_3 - S_1)^2}{4} = |\hat{\delta}_1^a|.$$

Since we identify β with $|\hat{\beta}_1|$, this mapping implies

$$\alpha = \frac{4 \, |\hat{\delta}_1^a|}{|\hat{\beta}_1|^2 \, (1 - p_2)(S_3 - S_1)^2}. \tag{7}$$

The model further implies that the distorted probability for scenario i is

$$q_i^* \approx p_i + p_i \, \beta \, \alpha \, (S_i - \bar{S}),$$

so that for the high-end scenario we have

$$q_3^* \approx p_3 + p_3 \, \beta \, \alpha \, (S_3 - S_2).$$

Using the symmetry assumption $S_3 - S_2 \approx \frac{S_3 - S_1}{2}$ and substituting $p_3 = \frac{1-p_2}{2}$, we obtain

$$q_3^* \approx \frac{1-p_2}{2} + \frac{1-p_2}{2} \, \beta \, \alpha \, \frac{S_3 - S_1}{2}.$$

Thus, the distortion in the high-end probability is

$$q_3^* - p_3 \approx \frac{(1-p_2) \, \beta \, \alpha \, (S_3 - S_1)}{4}.$$

Substituting the expression for α from (7) into the above, the $(1-p_2)$ terms cancel, yielding

$$q_3^* - p_3 \approx \frac{|\hat{\delta}_1^q|(S_3 - S_1)}{|\hat{\beta}_1|}.$$

This final expression shows that the probability tilt depends only on the ratio of the empirical coefficients and the difference $S_3 - S_1$; notably, the value of p_2 cancels out. In our application, plugging in the empirical estimates leads to an adjustment of approximately 17 percentage points (i.e., $q_3^* \approx p_3 + 0.17$ and, by symmetry, $q_1^* \approx p_1 - 0.17$).

This derivation confirms that the key result regarding the magnitude of the probability distortion is robust to alternative assumptions about the prior, as the dependence on p_2 cancels out under the symmetry condition.