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

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

26th February 2026
(preliminary and incomplete)

Abstract

We develop a tractable asset-pricing model in which investors price long-run disaster risk while facing ambiguity about the underlying probability distribution. Using variational calculus, the model delivers sharp, testable implications—especially a clear null for ambiguity neutrality. We test these implications in a large coastal real-estate market using more than 13.6 million U.S. single-family transactions (2001–2024) matched to high-resolution SLR inundation projections. Our key empirical innovation is a probability-based term structure of SLR disaster risk—the probability of permanent inundation as a function of horizon (e.g., 100 years). We reject ambiguity neutrality in price levels and find strong evidence that ambiguity amplifies risk capitalization. The horizon profile of window-based estimates also yields a novel estimate of the long-run discount rate for disaster risk at century horizons.

*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 Jarda Borovicka, Paul Ho, Cosmin Ilut and Ricardo Reis for their extensive suggestions. We also thank Chris Amante, David Anthoff, Job Boerma, Valentina Bosetti, Simon Dietz, Stefano Giglio, Lars Hansen, Aleksei Oskolkov, Evi Pappa, Monika Piazzesi, Esteban Rossi-Hansberg, Tom Sargent, Immo Schott, Christian Traeger, and participants at the Blue Collar Working Group (University of Chicago) 2025, Richmond Fed Brownbag, UofSC Darla Moore School of Business, Mohammed VI Polytechnic University Climate Challenges Conference, Stanford Institute of Theoretical Economics (SITE) Asset Pricing: Models, Solution Methods, and Applications session, SITE Climate Finance and Banking session, and System Committee Meeting on Macroeconomics for helpful comments and suggestions. We thank Christopher Amante for his guidance on sea-level rise projections and coastal elevation data. 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)

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

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 characterized and priced using expected utility, ambiguity typically requires a distinct toolkit from decision theory to address and 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, existing empirical settings have yet to provide a clear way to separately measure risk–variability within a known stochastic model–and ambiguity–dubiety or disagreement across competing potential models. As a result, 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. An agent that is averse to ambiguity is likely to react more robustly 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 studying coastal housing markets exposed to SLR using a novel, large-scale dataset and a probability-based measurement of long-run climate risk. We use the universe of all arm's-length transactions of single-family homes located within five miles of the U.S. coastline, totaling more than 13.6 million transactions, and link each sale to high-resolution, parcel-level projections of permanent inundation under RCP 2.6, 4.5, and 8.5, constructed from NOAA inundation layers and state-of-the-art local SLR simulations. Rather than characterizing exposure in physical terms (e.g., feet of SLR), we measure risk in economically meaningful space–time units as the probability that a property, given its precise geography and local SLR projections, is permanently inundated within a given horizon (such as 100 years). This construction allows us to sharply distinguish between *risk*—not knowing if a house will be inundated by a given horizon for a known probability model or scenario—and *ambiguity*— not knowing which probability model or scenario to use across a set of possibilities, captured by the divergence between best- and worst-case projections. By operationalizing risk and ambiguity directly in probability space at long horizons, we provide a clean empirical mapping of the Knight–Keynes distinction and a natural setting to test how markets price long-run climate risk under ambiguity.

To interpret how markets price inundation risk and ambiguity, we develop a stylized asset-pricing framework with 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 setting in which housing values depend on the probability of permanent inundation at long horizons. Rather than modeling damages directly, the object priced by the market is the probability distribution over inundation timing. Using Gâteaux perturbations in probability space, we characterize how marginal changes in (i) the mean probability across scenarios of inundation and (ii) the dispersion of probabilities across scenarios affect asset values. This formulation delivers three sharp, testable implications: higher mean inundation risk lowers prices; greater ambiguity independently depresses prices; and ambiguity amplifies the price response to risk. Under ambiguity neutrality, scenario disagreement collapses into a compound lottery and only mean probabilities matters for valuation. The model therefore provides a transparent mapping from probability-based SLR exposure to prices and a clear empirical test for ambiguity aversion in market behavior.

We test the model's three implications using hedonic regressions that map horizon-specific inundation *probabilities* and cross-scenario *disagreement* (the spread in the probabilities across scenarios) into transaction prices. Identification comes from comparing homes that are observationally similar along finely gridded spatial, temporal, and physical dimensions, but differ in their long-run probability of permanent inundation and the extent to which that probability varies across emissions scenarios. The empirical design

relies on highly saturated location-by-time fixed effects and exploits residual variation in probability-based SLR exposure to isolate the distinct price effects of inundation probability, scenario disagreement, and their interaction.

First, we find robust evidence that long-run SLR inundation probabilities (our novel term-structure measure of SLR risk) are negatively capitalized into housing prices. After absorbing stringent fixed effects that control for coastal amenities and micro-topography, higher probabilities of permanent inundation by a given horizon are associated with significantly lower prices. For example, in the price-level specification under RCP 4.5, a 10 percentage-point increase in the probability of permanent inundation within 100 years is associated with a price decline of roughly $5,000. Using window probabilities (e.g., the increment in inundation probability between years 100 and 110), we also find a sizable additional discount, consistent with the idea that markets price not only whether inundation occurs by a given horizon but also the timing of that risk.

Second, the *direct* effect of scenario disagreement holding the mean inundation probability fixed is not theoretically signed and is empirically less robust. Conceptually, disagreement can raise or lower prices depending on whether buyers place relatively more weight on optimistic versus pessimistic scenarios. In the data, the estimated coefficient on disagreement varies across specifications and horizons, and we therefore treat it as suggestive rather than a core test.

Third—and most importantly—we provide direct evidence for the amplification mechanism and reject ambiguity neutrality. The model delivers a sharp null: under ambiguity neutrality, scenario disagreement should *not* change the marginal price impact of inundation probability, so the interaction between inundation probability and disagreement in the *price-level* regression should be zero. We reject this null decisively. In the level regressions, the interaction between mean inundation probability and scenario disagreement is negative and highly statistically significant across stringent fixed-effect specifications and across alternative constructions of inundation probability. Economically, the estimates imply that the marginal dollar price of inundation risk becomes substantially more negative in more ambiguous locations. A margins figure visualizes this mechanism by plotting the implied marginal effect of inundation probability on price over the support of disagreement: the relationship tilts downward sharply as disagreement rises, highlighting how ambiguity rotates the risk–price relationship rather than merely shifting price levels.

We treat the log-price interaction results as corroborating robustness checks rather than the primary test. The main reason is conceptual: even if the interaction in price levels were zero, taking logs can mechanically induce interaction patterns due to the curvature of the log transformation. Accordingly, our central test of ambiguity neutrality is conducted in

price levels.

Finally, our probability-based construction delivers two byproducts that broaden the contribution. First, it adds a time dimension to the SLR-capitalization literature, which typically relies on static exposure thresholds (e.g., whether a property is inundated under six feet of SLR; Baldauf et al. 2020; Bakkensen et al. 2025). By building property-level distributions of inundation timing, we can study not only whether but also *when* exposure occurs. Second, the horizon profile of the estimated window coefficients provides an empirical handle on long-run discounting: the magnitude of capitalization declines with horizon in a manner consistent with discounting of distant housing-service flows, and a log-linear decay provides a useful approximation over the horizons we study, implying a long-run annual discount rate on the order of a few percent.

Taken together, the results provide unified empirical support for a robust-control asset-pricing view of climate risk in a large real-asset market: markets price the probability of long-run SLR inundation, and—crucially—scenario disagreement amplifies how strongly inundation risk is capitalized into today's housing prices.

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., 2014a; 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 SLR 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 SLR 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., 2020; 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.

Finally, our empirical approach relates to recent theoretical work on pricing changes in probabilities. Dávila et al. (2025) develop an asset-pricing framework for valuing marginal changes in probabilities by mapping probability perturbations into equivalent cash flows priced with standard discounting tools. We provide an empirical counterpart by constructing horizon-specific inundation probabilities and testing how markets price both expected risk and ambiguity in those probabilities using transaction data. More broadly, our analysis contributes to an emerging literature that uses tools from the calculus of variations to study how economic objects respond to changes in probability distributions. We rely on directional (Gateaux) derivatives to characterize how asset values respond to infinitesimal perturbations of inundation probabilities at specific horizons, following the formal treatment in Clarke (1990). This variational perspective not only clarifies the structure of the asset-pricing model but also directly guides our empirical design by identifying the economically relevant objects—horizon-specific probabilities and their cross-scenario spreads—to be taken to the data. Recent work has used functional derivatives to study policy and pricing problems in macroeconomics, finance, and public economics (Borovička

et al., 2011; Golosov et al., 2014b; Bilal, 2023; Bilal and Rossi-Hansberg, 2023; Dávila and Walther, 2023; Dávila and Schaab, 2025; Hansen and Souganidis, 2025). To the best of our knowledge, we are the first to apply these tools to map asset pricing theory to an empirical investigation of the market pricing of (climate) risk.

2. Stylized Model

This section develops a stylized framework to clarify how long-run SLR risk and ambiguity are capitalized into coastal housing values. The model is deliberately parsimonious. Its purpose is not to fully characterize housing markets, but to generate sharp, testable comparative statics and to guide interpretation of the empirical estimates.

The starting point is a standard present-value representation in the spirit of Gordon (1962), augmented to allow for stochastic disaster timing and permanent damage. We then extend the framework to multiple SLR scenarios and ambiguity aversion.

A defining feature of SLR risk is the presence of a *tipping point*: once a critical threshold is crossed, damages are persistent rather than transitory. In the context of coastal real estate, this tipping point corresponds to permanent inundation. The single-scenario model below isolates how beliefs about disaster timing enter asset prices in the simplest possible setting.

2.1. Single-scenario asset pricing model

Time is continuous, $t \geq 0$. The single-scenario model has three primitives:

(i) an asset (a house) that pays a dividend flow (rent or housing utility) $\{\text{div}(t)\}_{t\geq 0}$,

(ii) a stochastic discount factor $\{\text{SDF}(t)\}_{t\geq 0}$,

(iii) a one-time disaster process summarized by a random arrival time $T \in [0, \infty]$.

2.1.1 Primitives: dividend, SDF, and a tipping-point disaster

Key assumption 1. A disaster arrives once at random time T and causes permanent damage to the dividend process.

Specifically, let $\{h(t)\}_{t\geq 0}$ denote the *no-disaster* dividend flow, and let $\{\hat{h}(t)\}_{t\geq 0}$ denote the post-disaster flow. The realized dividend process is

$$\text{div}(t) = \begin{cases} h(t), & t < T, \\ \hat{h}(t), & t \geq T, \end{cases} \qquad t \geq 0.$$

We also allow discounting to change after the disaster:

$$\mathrm{SDF}(t) = \begin{cases} m(t), & t < T, \\ \widehat{m}(t), & t \geq T, \end{cases} \qquad t \geq 0.$$

The time-0 value of the house is therefore

$$v := \mathbb{E}_0\left[\int_0^\infty \mathrm{SDF}(t)\,\mathrm{div}(t)\,\mathrm{d}t\right].$$

Key assumption 2. In the absence of the disaster, the processes $\{h(t)\}_{t\geq 0}$ and $\{m(t)\}_{t\geq 0}$ are independent of the disaster timing T.

Economically, $h(t)$ represents the rent or housing-utility flow of an otherwise comparable *safe* house—one not exposed to inundation—and

$$v^{\text{safe}} := \int_0^\infty \mathbb{E}_0[m(t)h(t)]\,\mathrm{d}t$$

is the corresponding benchmark value. In the empirical analysis, this benchmark is absorbed by hedonic controls and granular location-by-time fixed effects, so that remaining price variation can be interpreted as reflecting beliefs about T and disaster damages.

To isolate the geometry of disaster timing and facilitate interpretation, we impose two simplifying assumptions that will be relaxed later.

Simplifying assumption 1. Post-disaster housing services are a constant fraction of pre-disaster services:

$$\hat{h}(t) \equiv (1-\ell)\,h(t), \qquad \ell \in (0,1) \text{ constant}.$$

To further strip down the benchmark, we set housing services to a constant:

$$h(t) \equiv 1 \qquad \Longrightarrow \qquad \hat{h}(t) \equiv 1-\ell.$$

Thus, the disaster causes a permanent proportional loss ℓ once it occurs.

Simplifying assumption 2. Post-disaster discounting coincides with pre-disaster discounting:

$$\widehat{m}(t) \equiv m(t),$$

so that $m(t) \perp T$ for all t.

Under these assumptions, the only channel through which SLR risk affects valuation is the distribution of T, which we summarize by the cumulative disaster probability function $p(\cdot)$.

2.1.2 Two pricing kernels: cash flows vs. probabilities

Before deriving valuations, it is useful to distinguish two distinct pricing kernels that arise in this framework.

In standard asset pricing, the stochastic discount factor $m(t)$ is the kernel that prices dividend strips. In our disaster model, the *probability pricing kernel* $\ell\mathcal{K}(t)$ prices the *probability term structure* $p(\cdot)$, where $\mathcal{K}$ is defined as the dividend strip price:

$$\mathcal{K}(t) := \mathbb{E}_0[m(t)h(t)], \tag{1}$$

which is the time-0 price of a claim to the undamaged dividend $h(t)$ at horizon t.[1] Intuitively:

- $\mathcal{K}(t)$: the value of dividends at risk at horizon t,
- ℓ: the fraction of those dividends lost upon disaster.

Disaster risk at horizon t is more costly when $\mathcal{K}(t)$ is high (dividends are valuable) and when ℓ is large (more is destroyed). As we show below, the disaster-induced discount to house value takes the form

$$\text{Disaster discount} = \langle \ell\mathcal{K}, p \rangle = \ell \int_0^\infty \mathcal{K}(t)\, p(t)\, \mathrm{d}t,$$

so that $\ell\mathcal{K}(t)$ prices the probability of disaster by horizon t. This parallels the role of $\mathcal{K}(t)$ in pricing dividend strips, but with probabilities replacing cash flows.

2.1.3 Geometric representation and calculus of variations

Denote the cumulative disaster probability and the survival function at horizon t by:

$$p(t) := \mathbb{P}(T \le t) \qquad \text{and} \qquad s(t) := \mathbb{P}(T > t) = 1 - p(t).$$

Under the benchmark simplifications, asset valuation becomes

$$\begin{aligned} v(p) &= \mathbb{E}_0\left[\int_0^\infty m(t)\Big(h(t)\mathbf{1}\{t < T\} + (1-\ell)h(t)\mathbf{1}\{t \ge T\}\Big)\mathrm{d}t\right] \\ &= \int_0^\infty \mathcal{K}(t)\Big((1-p(t)) + (1-\ell)p(t)\Big)\mathrm{d}t \\ &= \int_0^\infty \mathcal{K}(t)\,\mathrm{d}t - \ell\int_0^\infty \mathcal{K}(t)\,p(t)\,\mathrm{d}t \\ &= v^{\text{safe}} - \langle \ell\mathcal{K}, p\rangle. \end{aligned} \tag{2}$$

[1] Under the simplifying assumption $h(t) \equiv 1$, this reduces to

$$\mathcal{K}(t) = \mathbb{E}_0[m(t)] = \frac{1}{R_t^f},$$

the price of a zero-coupon risk-free bond maturing at t.

Equation (2) expresses house value as the safe benchmark v^{safe} minus the disaster discount $\langle \ell\mathcal{K}, p\rangle$. The disaster discount is the inner product of the probability pricing kernel $\ell\mathcal{K}$ with the disaster probability term structure p.

It is convenient to view asset valuation through the lens of the calculus of variations: $v(\cdot)$ is a *functional* (a function of a function) defined on probability curves $p : [0,\infty) \to [0,1]$.[2]

Remark 1 (Connection to probability pricing and risk-price dynamics). Our framework augments the probability pricing approach of Dávila et al. (2025) to continuous time, enabling analysis of the *term structure* of probability pricing. In their two-period setting, Dávila et al. show that the willingness-to-pay for a change in probabilities can be expressed as the price of a hypothetical asset with consumption-equivalent cash flows. Our continuous-time extension reveals that:

- The "state" in their framework corresponds to "time" t in ours,
- Their survival function change $\frac{d(1-F)}{d\theta}$ corresponds to perturbations $\delta(t)$ of the probability curve,
- Their consumption sensitivity $\frac{dc_1}{ds}$ corresponds to the damage-weighted dividend $\ell \cdot h(t)$.

The continuous-time formulation connects naturally to the term structure in risk-price dynamics (Borovička et al., 2011), where pricing kernels are decomposed by horizon. Here, the probability pricing kernel $\ell\mathcal{K}(t)$ plays an analogous role for disaster risk: it determines how disaster probability at each horizon t is capitalized into current prices.

2.1.4 Directional derivatives

Because asset valuation (2) is affine in p, directional derivatives take a simple form.[3] For any admissible[4] direction $\delta(\cdot)$, the derivative is:

$$\frac{d}{d\varepsilon} v(p + \varepsilon\delta)\Big|_{\varepsilon=0} = -\ell \int_0^\infty \mathcal{K}(t)\,\delta(t)\,\mathrm{d}t = -\langle \ell\mathcal{K}, \delta\rangle. \tag{3}$$

The negative sign reflects that increasing disaster probability reduces house value.

[2] Formally, the pairing $\langle \ell\mathcal{K}, p\rangle = \ell \int_0^\infty \mathcal{K}(t)p(t)\,\mathrm{d}t$ is well-defined whenever $\mathcal{K} \in L^1([0,\infty))$ and p is bounded (which holds since $p(t) \in [0,1]$). We do not require $p \in L^2([0,\infty))$.

[3] Recall from the calculus of variations (see, e.g., Clarke, 1990) that the directional (Gâteaux) derivative of a functional v at p in direction δ is $\frac{d}{d\varepsilon} v(p + \varepsilon\delta)\big|_{\varepsilon=0}$. When v is affine in p, this derivative is a continuous linear functional of δ.

[4] *Admissible* means that the perturbation direction δ is integrable against $\mathcal{K}$ and that, for sufficiently small $|\varepsilon|$, the perturbed curve $p + \varepsilon\delta$ remains a valid probability function (bounded in $[0,1]$, nondecreasing, right-continuous with $p(0) \geq 0$).

For convenience, for each horizon $\tau \geq 0$, we define the *annuity term* or continuation value

$$\mathcal{A}(\tau) := \int_\tau^\infty \mathcal{K}(t)\,\mathrm{d}t = \int_\tau^\infty E_0[m(t)h(t)], \qquad \tau \geq 0, \tag{4}$$

which is the time-0 price of a claim to the *stream of undamaged dividends from τ onward*.

Horizon-τ perturbation. Fix a horizon $\tau \geq 0$. We consider a perturbation that increases $p(\tau)$ by a small amount $\varepsilon > 0$. Since cumulative probabilities must be nondecreasing, increasing $p(\tau)$ also increases $p(t)$ for all $t \geq \tau$. The perturbation direction is thus:

$$\delta(t) = \mathbf{1}\{t \geq \tau\}, \qquad t \geq 0,$$

so that $p(t) \mapsto p(t) + \varepsilon\,\mathbf{1}\{t \geq \tau\}$.[5]

Figure 1 illustrates the key geometric insight: because v is affine in p, perturbing the probability curve in any direction produces a proportional change in value.

[Figure 1 about here.]

Lemma 1 (Sensitivity to horizon-τ disaster risk)**.** Fix $\tau \geq 0$. The directional derivative of v with respect to a perturbation that increases $p(\tau) = \mathbb{P}(T \leq \tau)$ is

$$\frac{\partial v}{\partial p(\tau)} := \frac{d}{d\varepsilon} v(p + \varepsilon \mathbf{1}\{t \geq \tau\})\Big|_{\varepsilon=0} = -\ell \underbrace{\int_\tau^\infty \mathcal{K}(t)\,\mathrm{d}t}_{=:\,\mathcal{A}(\tau)} = -\ell\,\mathcal{A}(\tau). \tag{5}$$

Proof. Appendix A.1. □

Lemma 1 has a simple economic interpretation. Increasing $p(\tau)$ shifts probability mass from "survive beyond τ" to "disaster hits by τ," which reduces the fraction of the post-τ cash-flow tail that remains undamaged. The annuity $\mathcal{A}(\tau)$ measures the present value of that tail, and ℓ scales how much of the tail is destroyed when the disaster arrives: only the wedge ℓ between $h(t)$ and $(1-\ell)h(t)$ is at stake. An immediate corollary is:

Corollary 2 (Sensitivity to disaster risk in a finite window)**.** Fix $0 \leq \tau_1 < \tau_2 < \infty$. The directional derivative of v with respect to a perturbation that increases disaster probability only in the window $[\tau_1, \tau_2]$ is

$$\frac{\partial v}{\partial p([\tau_1, \tau_2])} = -\ell \underbrace{\int_{\tau_1}^{\tau_2} \mathcal{K}(t)\,\mathrm{d}t}_{=:\,\mathcal{A}(\tau_1,\tau_2)} = -\ell\,\mathcal{A}(\tau_1, \tau_2). \tag{6}$$

[5] Formally, probability curves must remain valid (bounded in $[0,1]$, nondecreasing, right-continuous). To implement an admissible horizon-τ change, we use a compensated step direction. Fix any $L > \tau$ and define $\delta^{(\tau,L)}(t) := \mathbf{1}\{t \geq \tau\} - \mathbf{1}\{t \geq L\}$. For sufficiently small $|\varepsilon|$, $p + \varepsilon\delta^{(\tau,L)}$ remains a valid probability function and satisfies $p(\tau) \mapsto p(\tau) + \varepsilon$. Letting $L \to \infty$ delivers the pure tail direction $\mathbf{1}\{t \geq \tau\}$ in the valuation pairing (which is well-defined whenever $\mathcal{K} \in L^1$).

Term structure of probability pricing. An important property of equation (5) is that the sensitivity $|\partial v/\partial p(\tau)| = \ell\,\mathcal{A}(\tau)$ is *declining in the horizon* τ. This follows immediately from the fact that $\mathcal{A}(\tau) = \int_\tau^\infty \mathcal{K}(t)\,\mathrm{d}t$ is a decreasing function of τ (since $\mathcal{K}(t) > 0$ and we are shrinking the domain of integration).

Economically, disaster risk at distant horizons has a smaller price impact than disaster risk at near horizons because:

1. Distant cash flows are more heavily discounted (smaller $\mathcal{K}(t)$ for large t),
2. Fewer future dividends are at stake (smaller integration domain $[\tau, \infty)$).

This declining term structure of probability price sensitivities is a testable prediction of the model.

2.2. Multiple scenarios and robust control

We now allow for N mutually exclusive climate scenarios indexed by $i \in \{1, \ldots, N\}$. Scenario 1 is the most optimistic and scenario N is the most pessimistic. In our setting, "more pessimistic" means both (i) earlier inundation risk (a larger cumulative probability of permanent inundation by any given horizon) and (ii) a larger permanent proportional damage once inundation occurs.

Scenario ordering. Let $p_i(t) := \mathbb{P}_i(T_i \leq t)$ denote the cumulative probability of inundation by horizon t in scenario i. We assume a pointwise ordering of inundation risk across scenarios:

$$p_1(t) \;<\; p_2(t) \;<\; \cdots \;<\; p_N(t), \qquad \forall t \geq 0. \tag{7}$$

In addition, each scenario features a (constant) proportional loss parameter $\ell_i \in (0, 1)$, and we assume damages are larger in more pessimistic scenarios:

$$\ell_1 \;<\; \ell_2 \;<\; \cdots \;<\; \ell_N. \tag{8}$$

Thus, scenario N both reaches inundation sooner *and* delivers the largest permanent wedge between pre- and post-inundation housing services.

Let $\pi = (\pi_1, \ldots, \pi_N)$ denote the agent's prior over scenarios,

$$\pi_i > 0, \qquad \sum_{i=1}^{N} \pi_i = 1.$$

Simplifying assumption 3. To focus attention on differences in inundation timing and damage severity, we impose that the dividend strip price $\mathcal{K}(t) = \mathbb{E}_0[m(t)h(t)]$ is common across scenarios. Consequently, the annuity (continuation-value) function is also common:

$$\mathcal{A}(\tau) := \int_\tau^\infty \mathcal{K}(t)\,\mathrm{d}t, \qquad \tau \geq 0.$$

Scenario-specific valuation. Conditional on scenario i, primitives are as in the single-scenario model, with probability curve $p_i(\cdot)$ and constant proportional loss ℓ_i. Under the benchmark assumptions, the scenario-i value takes the affine form

$$v_i = v^{\text{safe}} - \langle \ell_i \mathcal{K}, p_i \rangle = \int_0^\infty \mathcal{K}(t)\,\mathrm{d}t \;-\; \ell_i \int_0^\infty \mathcal{K}(t)\,p_i(t)\,\mathrm{d}t. \tag{9}$$

Relative to the common "safe" benchmark v^{safe}, the only scenario-specific inputs are the probability curve $p_i(\cdot)$ and the damage parameter ℓ_i.

2.2.1 Ambiguity aversion

Equation (9) provides a valuation v_i for the house under each SLR scenario $i \in \{1, \ldots, N\}$. If the agent were confident that the prior $\pi = (\pi_1, \ldots, \pi_N)$ correctly described the likelihood of each scenario, she would simply compute the expected value under the prior distribution $\bar{v} := \sum_i \pi_i v_i$. The challenge, however, is that long-run greenhouse gas emissions, their implications, and therefore climate and SLR dynamics, are deeply uncertain. Thus, the agent may not fully trust her prior probabilities.

To reflect this concern, we model the agent as ambiguity-averse over the scenario index. Rather than committing to the benchmark prior π, the agent allows for alternative probability distributions $q = (q_1, \ldots, q_N)$ over scenarios, constrained to lie on the simplex

$$\Delta^N := \Big\{ q \in \mathbb{R}_+^N : \; \sum_{i=1}^N q_i = 1 \Big\}.$$

Each such q represents a plausible distortion of the prior, capturing the idea that the true scenario probabilities may differ from π in ways that matter for valuation.

The agent prices the house under these distorted probabilities by forming the q-weighted value $\sum_i q_i v_i$. However, she does not consider all distortions equally plausible given her prior. Deviations from the benchmark prior incur a penalty measured by a divergence function

$$D_f(q\|\pi) \;:=\; \sum_{i=1}^N \pi_i\, f\!\left(\frac{q_i}{\pi_i}\right), \tag{10}$$

where $f : \mathbb{R}_+ \to \mathbb{R}$ is convex and normalized by $f(1) = 0$. This penalty plays the role of a *trust region* around the prior: small deviations are inexpensive, while large distortions are

increasingly costly. The class of f-divergences nests familiar cases such as Kullback–Leibler (KL) and χ^2, both of which we use in our analysis.

Robust-control problem. The agent's robust valuation problem trades off pessimism against plausibility. Formally, she chooses a distorted distribution q to minimize

$$v^* := \min_{q \in \Delta^N} \left\{ \sum_{i=1}^{N} q_i v_i + \frac{1}{\alpha} D_f(q \| \pi) \right\}, \qquad \alpha > 0. \tag{11}$$

The first term favors probability mass on low-v_i (high-risk) scenarios, while the second term discourages excessive departures from the benchmark prior. The parameter α governs the degree of ambiguity aversion: a larger α weakens the penalty and allows the agent to consider more pessimistic distortions.

Let q^* denote an optimizer. The robust value v^* is therefore the value of the house under the most adverse probability distortion that remains sufficiently close to the prior,

$$v^* = \sum_{i=1}^{N} q_i^* v_i + \frac{1}{\alpha} D_f(q^* \| \pi). \tag{12}$$

Geometrically, the expected value $\sum_i q_i v_i$ is a linear functional on the simplex, while the divergence penalty defines a curved neighborhood centered at π. The optimal distortion q^* lies where the incentive to tilt probability mass toward worse SLR scenarios is exactly balanced by the cost of departing from the prior. In the next subsection, we characterize this tilt explicitly and show how ambiguity aversion systematically amplifies the weight placed on pessimistic SLR scenarios.

2.2.2 Optimal tilts

Lemma 3. In the N-scenario robust-control problem, the optimal distortion q^* tilts the prior π toward low-v_i (low-value, high-risk) scenarios. In particular,

$$\frac{q_i^*}{\pi_i} = \begin{cases} \dfrac{e^{-\alpha v_i}}{\sum_{k=1}^{N} \pi_k e^{-\alpha v_k}} \propto e^{-\alpha v_i} & \text{(KL)}, \\ 1 + \alpha(\bar{v} - v_i) & (\chi^2 \text{ interior}). \end{cases} \tag{13}$$

where the prior mean value is

$$\bar{v} := \sum_{i=1}^{N} \pi_i v_i. \tag{14}$$

Proof. Appendix A.2. □

Intuitively, Lemma 3 shows that ambiguity aversion operates by tilting probability mass toward scenarios with lower valuation v_i. The agent thus cautiously prices the asset as if more pessimistic SLR scenarios are more likely than suggested by the benchmark prior π.

The ambiguity-aversion parameter α governs the strength of this pessimistic tilt. When α is small, deviations from the prior are costly and q^* remains close to π; when α is large, the agent is willing to place substantially more weight on low-v_i scenarios.

In the KL case, the distortion takes a soft-min form, $q_i^*/\pi_i \propto e^{-\alpha v_i}$, so lower-valuation scenarios receive exponentially more weight. As $\alpha \to 0$, $q^* \to \pi$; as $\alpha \to \infty$, probability mass concentrates on the worst scenario, approaching max–min behavior. In the χ^2 case (under an interior solution), the distortion is linear, $q_i^*/\pi_i = 1 + \alpha(\bar{v} - v_i)$, so scenarios worse than average are up-weighted in direct proportion to how far their valuation falls below $\bar{v}$.

Remark 2 (Interior vs. corner solutions and the maxmin limit). With KL divergence, the optimal tilt is always interior: the Gibbs form in Lemma 3 implies $q_i^* > 0$ for all i whenever $\pi_i > 0$ and $v_i < \infty$. In contrast, under the χ^2 divergence, the affine tilt $q_i^* = \pi_i\big[1 + \alpha(\bar{v} - v_i)\big]$ can hit the boundary of the simplex when ambiguity aversion is sufficiently strong, so some scenarios receive zero distorted weight.[6] As α grows and the distortion concentrates on the most pessimistic scenario:

$$v^* \;\to\; \min_{i \in \{1,\ldots,N\}} v_i.$$

This limiting case corresponds precisely to the Gilboa and Schmeidler (1989) maxmin expected utility model, where the decision-maker evaluates alternatives by their worst-case scenario. This classical model, which plays a prominent role in macroeconomic models with ambiguity aversion (Ilut and Schneider, 2023), emerges here as a special case of our ambiguity-aversion framework.

2.2.3 Directional derivatives

Let $p = (p_1, \ldots, p_N)$ denote the vector of scenario-specific probability curves, where $p_i(t) = \mathbb{P}_i(T_i \le t)$ for each scenario i. Under the benchmark assumptions of a common dividend strip price $\mathcal{K}(t)$ and constant proportional loss ℓ_i, the scenario-i value (9) can be

[6] Formally, one can equivalently write the χ^2 solution with a positive-part operator as

$$q_i^* \;=\; \pi_i \frac{\big[1 + \alpha(\bar{v} - v_i)\big]^+}{\sum_{j=1}^N \pi_j \big[1 + \alpha(\bar{v} - v_j)\big]^+},$$

which makes clear that scenarios with v_i far above $\bar{v}$ may be discarded.

written as the affine functional

$$v_i(p_i) = v^{\text{safe}} - \langle \ell_i \mathcal{K}, p_i \rangle = \int_0^\infty \mathcal{K}(t)\,\mathrm{d}t \;-\; \ell_i \int_0^\infty \mathcal{K}(t)\, p_i(t)\,\mathrm{d}t. \tag{15}$$

Only the second term depends on p_i, so $v_i(\cdot)$ is linear in p_i up to an irrelevant constant. Hence, for any admissible direction δ_i, the directional derivative is:

$$\frac{d}{d\varepsilon} v_i(p_i + \varepsilon\delta_i)\Big|_{\varepsilon=0} = -\ell_i \int_0^\infty \mathcal{K}(t)\,\delta_i(t)\,\mathrm{d}t = -\langle \ell_i \mathcal{K}, \delta_i \rangle, \tag{16}$$

where $\langle a, b\rangle := \int_0^\infty a(t)b(t)\,\mathrm{d}t$ denotes the usual pairing.

Recall from the robust-control problem (11) that the robust value v^* depends on $(v_1, \ldots, v_N)$ through the distorted probabilities $q^* = (q_1^*, \ldots, q_N^*)$. By the envelope theorem,

$$\frac{\partial v^*}{\partial v_i} = q_i^*. \tag{17}$$

Combining (17) with (16) yields the directional derivative of the robust value v^* as a functional of the *entire vector* p:

$$\frac{d}{d\varepsilon} v^*(p + \varepsilon\delta)\Big|_{\varepsilon=0} = -\sum_{i=1}^N q_i^*\, \ell_i\, \langle \mathcal{K}, \delta_i \rangle, \tag{18}$$

for any admissible perturbation direction $\delta = (\delta_1, \ldots, \delta_N)$.

Lemma 4 (Gradient representation of directional derivatives)**.** Let $\mathcal{N} := (L^2([0,\infty)))^N$ be the Hilbert space with inner product

$$\langle a, b\rangle_{\mathcal{N}} := \sum_{i=1}^N \int_0^\infty a_i(t) b_i(t)\,\mathrm{d}t.$$

Then the directional derivative of v^* admits the Riesz representer

$$\nabla_p v^*(p)(t) = -\big(q_1^*\,\ell_1\,\mathcal{K}(t), \ldots, q_N^*\,\ell_N\,\mathcal{K}(t)\big) \in \mathcal{N},$$

so that, for every admissible $\delta \in \mathcal{N}$,

$$\frac{d}{d\varepsilon} v^*(p + \varepsilon\delta)\Big|_{\varepsilon=0} = \langle \nabla_p v^*(p),\, \delta\rangle_{\mathcal{N}} = -\sum_{i=1}^N q_i^*\, \ell_i\, \langle \mathcal{K}, \delta_i \rangle. \tag{19}$$

Proof. Appendix A.3. □

The gradient has a natural interpretation: the i-th component is the probability pricing kernel for scenario i, weighted by the robust probability q_i^*:

$$[\nabla_p v^*]_i(t) = -q_i^* \cdot \underbrace{\ell_i\,\mathcal{K}(t)}_{\text{Scenario-}i\text{ prob. pricing kernel}}.$$

Scenarios with higher robust weights q_i^* and larger damage ℓ_i have more price-relevant probability curves.

Lemma 4 makes comparative statics transparent: any economically meaningful change in beliefs corresponds to a direction $\delta = (\delta_1, \ldots, \delta_N)$, and the induced marginal value change is the single inner product $\langle \nabla_p v^*(p), \delta \rangle_{\mathcal{N}}$.

2.2.4 Sensitivity to disaster probability

This subsection applies Lemma 4 to obtain horizon-τ risk sensitivities with minimal algebra. The compensated-step constructions and $L \to \infty$ limits that ensure admissibility are collected in Appendix A.5.

Shift in a single scenario. Fix a horizon $\tau \geq 0$ and a scenario $i \in \{1, \ldots, N\}$. Consider a perturbation that increases the cumulative inundation probability $p_i(\tau) = \mathbb{P}_i(T_i \leq \tau)$ by a small amount, leaving the other scenarios unchanged. In terms of probability curves, this is a (compensated) upward shift of $p_i(t)$ for all $t \geq \tau$, i.e., a direction proportional to $\mathbf{1}\{t \geq \tau\}$ in coordinate i.

Applying (19) along this tail direction yields

$$\frac{\partial v^*}{\partial p_i(\tau)} = -\, q_i^*\, \ell_i \underbrace{\int_\tau^\infty \mathcal{K}(t)\, \mathrm{d}t}_{\mathcal{A}(\tau)} = -\, q_i^*\, \ell_i\, \mathcal{A}(\tau). \tag{20}$$

Dividing by v^* gives the semi-elasticity,

$$\frac{\partial \ln v^*}{\partial p_i(\tau)} = -\frac{q_i^*\, \ell_i}{v^*}\, \mathcal{A}(\tau). \tag{21}$$

Common shifts across scenarios. Next consider a perturbation that raises horizon-τ inundation probabilities uniformly across scenarios, i.e., increases

$$\bar{p}(\tau) := \frac{1}{N} \sum_{i=1}^N p_i(\tau)$$

by a small amount. In probability space, this corresponds to moving each p_i in the same tail direction $\mathbf{1}\{t \geq \tau\}$. Applying (19) yields

$$\frac{\partial v^*}{\partial \bar{p}(\tau)} = -\left(\sum_{i=1}^N q_i^*\, \ell_i \right) \int_\tau^\infty \mathcal{K}(t)\, \mathrm{d}t = -\left(\sum_{i=1}^N q_i^*\, \ell_i \right) \mathcal{A}(\tau). \tag{22}$$

Hence,

$$\frac{\partial \ln v^*}{\partial \bar{p}(\tau)} = -\frac{\left(\sum_{i=1}^N q_i^*\, \ell_i \right) \mathcal{A}(\tau)}{v^*}. \tag{23}$$

When $\ell_i \equiv 1$ (pure termination at inundation), the bracketed term equals $\sum_i q_i^* = 1$ and (22) collapses to the single-scenario formula. More generally, with heterogeneous disaster damage ℓ_i, a common increase in inundation risk is weighted by the robust severity index $\sum_i q_i^* \ell_i$.

2.2.5 Sensitivity to extreme spread

We now isolate the effect of an increase in ambiguity. In our empirical setting, ambiguity naturally manifests as greater dispersion across a small number of plausible SLR scenarios (e.g., optimistic, intermediate, and pessimistic emissions paths). A simple way to capture this is to increase the gap between the most optimistic and most pessimistic scenarios at a fixed horizon.

Fix $\tau \geq 0$ and define the *extreme spread* as

$$p_d(\tau) := p_N(\tau) - p_1(\tau).$$

In this subsection, we define $\partial/\partial p_d(\tau)$ as a perturbation that increases $p_N(\tau)$ and decreases $p_1(\tau)$ by the same amount, leaving the intermediate scenarios unchanged.[7]

At the probability-curve level, the perturbation decreases p_1 beyond τ and increases p_N beyond τ:

$$\delta(t) = \big(-\mathbf{1}\{t \geq \tau\}, 0, \ldots, 0, \mathbf{1}\{t \geq \tau\} \big).$$

Applying Lemma 4 immediately gives

$$\frac{d}{d\varepsilon} v^*(p + \varepsilon\delta)\Big|_{\varepsilon=0} = -\sum_{i=1}^{N} q_i^* \, \ell_i \, \langle \mathcal{K}, \delta_i \rangle = \big(q_1^* \, \ell_1 - q_N^* \, \ell_N \big) \, \mathcal{A}(\tau).$$

Since increasing $p_d(\tau) = p_N(\tau) - p_1(\tau)$ corresponds to moving in direction δ, we obtain:

$$\frac{\partial v^*}{\partial p_d(\tau)} = \big(q_1^* \, \ell_1 - q_N^* \, \ell_N \big) \, \mathcal{A}(\tau). \tag{24}$$

Dividing by v^* gives the semi-elasticity,

$$\frac{\partial \ln v^*}{\partial p_d(\tau)} = \frac{\big(q_1^* \, \ell_1 - q_N^* \, \ell_N \big) \, \mathcal{A}(\tau)}{v^*}. \tag{25}$$

The sign of (24) is *a priori* ambiguous: it depends on whether the robustly weighted disaster severity is larger in the pessimistic scenario than in the optimistic one, i.e., on the sign of $q_1^* \, \ell_1 - q_N^* \, \ell_N$.

[7] Formally, we take the directional derivative along the (compensated) tail direction that shifts probability mass beyond τ in opposite directions for scenarios 1 and N. A convenient representative direction is $\delta(t) = \big(-\mathbf{1}\{t \geq \tau\}, 0, \ldots, 0, \mathbf{1}\{t \geq \tau\} \big)$ (with the same compensated-step/$L \to \infty$ construction as in Appendix A.5 to ensure admissibility). This perturbation increases $p_d(\tau)$ by construction at first order.

Economically, the perturbation makes the best scenario (1) safer beyond τ and the worst scenario (N) riskier beyond τ. If the agent remains sufficiently optimistic in the robust sense—so that the robust weight on scenario 1 is large relative to the robust weight on scenario N once scaled by severities—then $q_1^* \ell_1 - q_N^* \ell_N > 0$ and increasing the spread raises value. Conversely, if the agent is sufficiently pessimistic (large robust weight on scenario N and/or large ℓ_N), then $q_1^* \ell_1 - q_N^* \ell_N < 0$ and increasing the spread lowers value because it makes the robustly salient worst-case tail worse.

2.2.6 Interaction: cross partials

A central implication of ambiguity aversion is that the robust weights q^* are *endogenous*: they move when scenario valuations move. This generates curvature in v^* that is absent under ambiguity neutrality. In particular, when $\alpha = 0$ we have $q_i^* = \pi_i$ fixed, so any cross derivative that operates only through movements in q^* is mechanically zero.

We focus on the interaction between (i) a common increase in horizon-τ disaster risk, summarized by $\bar{p}(\tau)$, and (ii) an increase in dispersion across the two extremes, summarized by $p_d(\tau) = p_N(\tau) - p_1(\tau)$.

Recall from (22) that

$$\frac{\partial v^*}{\partial \bar{p}(\tau)} = -\Big(\sum_{i=1}^{N} q_i^* \, \ell_i\Big) \, \mathcal{A}(\tau).$$

Because the dividend strip price $\mathcal{K}$ (and hence $\mathcal{A}$) does not depend on disaster probabilities, the cross partial $\partial^2 v^* / (\partial \bar{p}(\tau) \, \partial p_d(\tau))$ is driven entirely by how the endogenous weights $\{q_i^*\}_{i=1}^N$ respond to the spread perturbation.

KL robust control and a covariance formula. From Lemma 3, under KL divergence the optimal distortion q^* tilts the prior π exponentially toward low-v_i (high-risk) scenarios. Differentiating those soft-min weights with respect to any scalar perturbation z gives

$$\frac{\partial q_i^*}{\partial z} = -\alpha \, q_i^* \Big(\frac{\partial v_i}{\partial z} - \sum_{k=1}^{N} q_k^* \, \frac{\partial v_k}{\partial z}\Big). \tag{26}$$

Intuitively, this says that when scenario i becomes relatively less attractive (its $\partial v_i / \partial z$ exceeds the q^*-average across scenarios), the robust decision maker reallocates probability mass *toward* that scenario, proportionally to the degree of ambiguity aversion α.

Combining (26) with the envelope condition $\partial v^* / \partial v_i = q_i^*$ yields the general cross-partial identity for any two variables z and w:

$$\frac{\partial^2 v^*}{\partial z \, \partial w} = \sum_{i=1}^{N} q_i^* \, \frac{\partial^2 v_i}{\partial z \, \partial w} \; - \; \alpha \, \mathrm{Cov}_{q^*}\Big(\frac{\partial v_i}{\partial z}, \, \frac{\partial v_i}{\partial w}\Big), \tag{27}$$

where $\mathrm{Cov}_{q^*}(\cdot,\cdot)$ denotes the q^*-weighted covariance across scenarios.[8]

Specializing the partial derivatives. Under the common-kernel and constant-loss assumptions,

$$\frac{\partial v_i}{\partial \bar{p}(\tau)} = -\ell_i \, \mathcal{A}(\tau) \qquad (i = 1, \ldots, N),$$

because a common increase in horizon-τ risk reduces the undamaged continuation value beyond τ in every scenario, scaled by ℓ_i.

For the extreme-spread perturbation at horizon τ, we use the direction that raises $p_N(\tau)$ and lowers $p_1(\tau)$ (holding the intermediate scenarios fixed). This implies

$$\frac{\partial v_1}{\partial p_d(\tau)} = +\ell_1 \, \mathcal{A}(\tau), \qquad \frac{\partial v_N}{\partial p_d(\tau)} = -\ell_N \, \mathcal{A}(\tau), \qquad \frac{\partial v_i}{\partial p_d(\tau)} = 0 \ (2 \le i \le N-1).$$

Plugging these expressions into (27) yields a closed-form covariance representation of the interaction:

$$\begin{aligned}
\frac{\partial^2 v^*}{\partial \bar{p}(\tau)\, \partial p_d(\tau)} &= -\alpha \, \mathrm{Cov}_{q^*}\Big(\frac{\partial v_i}{\partial \bar{p}(\tau)}, \, \frac{\partial v_i}{\partial p_d(\tau)} \Big) \\
&= -\alpha \, \mathrm{Cov}_{q^*}\Big(-\ell_i \mathcal{A}(\tau), \ \ell_1 \mathbf{1}\{i=1\} \mathcal{A}(\tau) - \ell_N \mathbf{1}\{i=N\} \mathcal{A}(\tau) \Big) \\
&= \alpha \, \mathcal{A}(\tau)^2 \, \mathrm{Cov}_{q^*}\Big(\ell_i, \ \ell_1 \mathbf{1}\{i=1\} - \ell_N \mathbf{1}\{i=N\} \Big). \qquad (28)
\end{aligned}$$

For the χ^2 model, we have a similar expression, except that the covariance now uses the prior π:

$$\frac{\partial^2 v^*}{\partial \bar{p}(\tau)\, \partial p_d(\tau)} = \alpha \, \mathcal{A}(\tau)^2 \, \mathrm{Cov}_{\pi}\Big(\ell_i, \ \ell_1 \mathbf{1}\{i=1\} - \ell_N \mathbf{1}\{i=N\} \Big). \qquad (29)$$

Lemma 5. Assume

$$\ell_1 < \ell_2 < \cdots < \ell_N,$$

then

$$\mathrm{Cov}_{q^*}\Big(\ell_i, \ \ell_1 \mathbf{1}\{i=1\} - \ell_N \mathbf{1}\{i=N\} \Big) \ < \ 0. \qquad (30)$$

As a corollary, the cross derivative is *negative*:

$$\frac{\partial^2 v^*}{\partial \bar{p}(\tau)\, \partial p_d(\tau)} < 0.$$

Proof. Appendix A.4. □

[8] For any scenario-indexed vectors $\{a_i\}$ and $\{b_i\}$, define the covariance under discrete measure q^* as

$$\mathrm{Cov}_{q^*}(a_i, b_i) := \mathbb{E}_{q^*}[a_i b_i] - \mathbb{E}_{q^*}[a_i]\mathbb{E}_{q^*}[b_i]$$

Intuitively, the random variable $\ell_1 \mathbf{1}\{i = 1\} - \ell_N \mathbf{1}\{i = N\}$ is *positive* only in the best scenario and *negative* only in the worst scenario. Since ℓ_i is increasing in i, high-ℓ states coincide with negative realizations of the selector and low-ℓ states coincide with positive realizations, generating a negative comovement under any strictly positive weighting measure q^*.

Cross partial of $\ln v^*$. Applying the chain rule yields

$$\frac{\partial^2 \ln v^*}{\partial \bar{p}(\tau)\, \partial p_d(\tau)} = \frac{1}{v^*} \frac{\partial^2 v^*}{\partial \bar{p}(\tau)\, \partial p_d(\tau)} - \frac{1}{(v^*)^2} \Big(\frac{\partial v^*}{\partial \bar{p}(\tau)} \Big) \Big(\frac{\partial v^*}{\partial p_d(\tau)} \Big). \tag{31}$$

The first term is the ambiguity-driven interaction; in our setting it is strictly negative under KL with $\alpha > 0$ and $\ell_1 < \cdots < \ell_N$ because $\partial^2 v^* / (\partial \bar{p}(\tau)\, \partial p_d(\tau)) < 0$.

The second term is the purely mechanical curvature from the log transform. Since $\partial v^* / \partial \bar{p}(\tau) < 0$ always, this term is negative whenever the spread lowers value, i.e., whenever $\partial v^* / \partial p_d(\tau) < 0$. Hence, in the empirically relevant case where $p_d(\tau)$ reduces valuation, the log transformation adds an additional negative component on top of the already-negative level interaction, making

$$\frac{\partial^2 \ln v^*}{\partial \bar{p}(\tau)\, \partial p_d(\tau)} < 0.$$

2.3. Summary of testable implications

We have proved the following proposition, which summarizes the main testable implications of the model.

Proposition 6 (Testable implications for probability pricing of disaster risk)**.** Fix a horizon $\tau \geq 0$. Under the common-kernel and constant-loss assumptions, the robust value v^* satisfies:

(i) **Negative risk sensitivity:**

$$\frac{\partial v^*}{\partial p_i(\tau)} < 0, \qquad \frac{\partial v^*}{\partial \bar{p}(\tau)} < 0.$$

Higherdisaster probability at any horizon τ reduces house value. Furthermore, the sensitivity features a declining term structure: its absolute value declines in τ (near-term disaster risk has a larger price impact than distant risk).

(ii) **Negative ambiguity interaction:** Under robust control with $\alpha > 0$ and ordered damages $\ell_1 < \cdots < \ell_N$,

$$\frac{\partial^2 v^*}{\partial \bar{p}(\tau)\, \partial p_d(\tau)} < 0, \qquad \frac{\partial^2 \ln v^*}{\partial \bar{p}(\tau)\, \partial p_d(\tau)} < 0.$$

Greater scenario dispersion amplifies the negative effect of disaster risk on value. Under ambiguity neutrality ($\alpha = 0$), this cross-partial equals zero.

In the subsequent empirical analysis, we will test these predictions by:

- Mapping v^* to transaction prices of coastal real estate,
- Mapping $\bar{p}(\tau)$ to the average inundation probability across RCP scenarios 2.6, 4.5, and 8.5 at, say, 100 years from the transaction date,
- Mapping $p_d(\tau) = p_{8.5}(\tau) - p_{2.6}(\tau)$ to the spread between pessimistic and optimistic scenario probabilities,

3. Data construction

We assemble a property-level dataset by combining proprietary deed transactions from CoreLogic with high-resolution coastal inundation layers from the National Oceanic and Atmospheric Administration (NOAA). Building on Bakkensen et al. (2025), our analysis substantially expands both the scope and the precision of the data. We extend geographic coverage to all single-family homes located within five miles of the U.S. ocean coastline, lengthen the sample period to 2001–2024, and exploit newly released NOAA Sea Level Rise Viewer layers that report inundation thresholds at half-foot increments. Crucially, these data improvements enable a new probability-based construction of SLR exposure defined at arbitrary horizons (e.g., the probability of permanent inundation within 100 years), rather than the elevation- or threshold-based measures commonly used in the prior literature (such as indicators for inundation at a given number of feet of SLR). This shift from physical thresholds to horizon-specific probabilities allows us to measure long-run climate risk in economically meaningful time units and plays a central role in the empirical analysis that follows.

Transaction data Our primary data source is CoreLogic, which is a large propriety database that links the universe of deed transactions to property tax assessor ("tax roll") records in the United States. For each transaction, we observe the sale date and price, along with property characteristics including property size (square feet), number of stories, building age, and ground elevation. CoreLogic also provides precise parcel-level geographic coordinates (latitude and longitude). These coordinates allow us to compute each property's distance to the nearest ocean coastline and, critically, to assign property-specific exposure to SLR inundation.

We restrict attention to the universe of arm's-length transactions of single-family properties located within five miles of the ocean coast. We exclude transactions with sale prices below $100,000. We winsorize the data at the 1 and 99 percentiles. The sample spans the period 2001–2024.

NOAA inundation maps To measure each property's exposure to future SLR, we use state-of-the-art inundation layers from NOAA's Sea Level Rise Viewer. This database provide high-resolution maps indicating areas that would be permanently inundated under alternative hypothetical SLR amounts.[9] For each property i, we overlay its exact coordinates on NOAA's inundation polygons and record whether the parcel would be inundated under each SLR increment

$$x \in \{0, 0.5, 1.0, \ldots, 9.5, 10.0\} \text{ feet.}$$

The resulting parcel-level inundation profile forms the basis for the SLR risk and ambiguity measures used in the empirical analysis.

Figure 2 illustrates the spatial structure of these inundation layers for three major coastal housing markets—New York City, Miami, and Charleston—at SLR levels of 0, 5, and 10 feet. The figure highlights the sharp spatial heterogeneity in exposure that arises from micro-topography and hydrologic connectivity. Even within narrow geographic areas, neighboring parcels can differ substantially in their inundation thresholds. This fine-grained variation underlies our identification strategy: the econometric specifications exploit comparisons among properties that are otherwise highly similar in location and physical characteristics, but differ in predicted long-run exposure to permanent coastal inundation.

[Figure 2 about here.]

Constructing horizon-specific inundation probabilities A central contribution of our analysis is a shift from elevation-based SLR exposure measures to horizon-specific probabilities of permanent inundation that are directly interpretable in economic time. Rather than asking whether a property is exposed at a particular SLR threshold, we ask how likely it is to be permanently inundated within a given time horizon (for example, 50 or 100 years) under alternative climate scenarios. This probability-based approach aligns long-run climate risk with the time horizons relevant for asset pricing and allows us to separately quantify risk within scenarios and ambiguity across scenarios.

We convert NOAA's parcel-level inundation thresholds (in feet) into horizon-specific probabilities of inundation under alternative SLR scenarios.

[9] Retrieved from https://coast.noaa.gov/slrdata/Sea_Level_Rise_Vectors/index.html in February 2025.

Step 1: Parcel-level inundation threshold. Using the NOAA Sea Level Rise Viewer vector inundation layers (available at half-foot increments), we assign to each property i a minimum inundation threshold,

$$minSLR_i \in \{0, 0.5, 1.0, \ldots, 9.5, 10.0\},$$

defined as the smallest SLR amount (in feet) that would permanently inundate the parcel given local elevation and hydrologic connectivity. Intuitively, $minSLR_i$ summarizes each parcel's physical exposure to permanent coastal inundation on a vertical scale.

Step 2: Local SLR projections by scenario. To translate vertical exposure into timing and probability, we combine $minSLR_i$ with local SLR projections based on the well-known SLR simulation framework of DeConto and Pollard (2016). These projections provide distributions of future relative SLR under three Representative Concentration Pathways (RCPs): RCP2.6 (low emissions), RCP4.5 (moderate emissions), and RCP8.5 (high emissions).[10] For each U.S. tide gauge location, the projections deliver an uncertain SLR path over long horizons (extending to year 2300), which we summarize using percentile bands (e.g., 10th–90th percentiles).[11]

Figure 3 illustrates the local projections for Miami under the three RCPs, based on our implementation of the DeConto and Pollard (2016) simulation outputs. This figure also clarifies the conceptual distinction between *risk* and *ambiguity* that motivates our empirical design. Holding a scenario fixed (e.g., RCP4.5), the projection delivers a well-defined distribution of future local SLR at each horizon—and thus a *known* inundation probability for any parcel-specific threshold $minSLR_i$; this within-scenario uncertainty corresponds to *risk*. By contrast, the appropriate scenario itself is not known ex ante. We therefore treat the scenario dimension as *ambiguous* and quantify ambiguity at each horizon by the dispersion in implied inundation probabilities across scenarios, with RCP2.6 representing a slow-SLR benchmark and RCP8.5 representing a fast-SLR benchmark. In the empirical analysis, this risk–ambiguity decomposition maps naturally into (i) horizon-specific inundation probabilities computed within a given scenario and (ii) a horizon-specific ambiguity measure constructed from cross-scenario divergence.

[10] Each pathway is defined by a target level of additional radiative forcing (in W/m^2) by year 2100 relative to preindustrial conditions; for example, RCP2.6, RCP4.5, and RCP8.5 correspond to approximately 2.6, 4.5, and 8.5 W/m^2 in 2100, respectively. RCPs are widely adopted by the IPCC and the climate-science community as standardized scenarios for producing long-run climate and SLR projections.

[11] Relative SLR is location-specific: the same global forcing can translate into different local sea levels due to factors such as vertical land motion (subsidence/uplift), ocean circulation and dynamic sea level, gravitational and rotational "fingerprints" from ice mass loss, and local tidal/hydrologic conditions. Using a dense set of all tidal gauges across the US coastline, we match each transaction to the local SLR projection at the tide gauge geographically closest to the transacted property.

[Figure 3 about here.]

Step 3: Probability of inundation by horizon. Consider a transaction for property i in sale year t. Fix an RCP scenario $r \in \{2.6, 4.5, 8.5\}$ and a horizon H (in years since sale). Let $g(i)$ denote the nearest tide gauge to property i. Let $SLR_{g,r,t+H}$ be the (random) projected relative sea level at gauge g in calendar year $t + H$ under scenario r, with cumulative distribution function $F_{g,r,t+H}(\cdot)$. We define the probability that property i is permanently inundated by horizon H as

$$p_{r,H}(i \mid t) \;:=\; \Pr\Big(SLR_{g(i),r,t+H} \geq minSLR_i\Big) \;=\; 1 - F_{g(i),r,t+H}(minSLR_i)\,.$$

In practice, we compute $p_{r,H}(i \mid t)$ using the simulated distribution of $SLR_{g(i),r,t+H}$ for the nearest tide gauge, evaluating the fraction of simulation draws (or equivalently, the implied quantile) that exceed $minSLR_i$. The half-foot resolution of $minSLR_i$ reduces discretization error in this mapping and will later allow us to construct probability-based risk measures at arbitrary horizons.

Figure 4 provides a concrete illustration for RCP4.5. For a property with $minSLR_i = 5$ feet sold in year $t = 2020$, the horizon $H = 100$ corresponds to calendar year 2120; the figure shows how we read off the implied probability of exceeding 5 feet by 2120 from the local (Miami) RCP4.5 projection distribution.

[Figure 4 about here.]

Summary statistics Table 1 reports summary statistics for the transaction-level analysis dataset, which comprises more than 13.6 million arm's-length transactions of single-family homes located within five miles of the U.S. coastline. To our knowledge, this is among the largest datasets used to study the capitalization of long-run climate risk into housing prices. The scale and breadth of the sample span a wide range of housing characteristics and SLR exposures—from properties with essentially zero inundation risk within 100 years to those facing near-certain inundation under high-emissions scenarios—allowing us to exploit fine-grained spatial variation while maintaining substantial statistical power in highly saturated fixed-effects specifications.

The homes in the sample are broadly representative of single-family housing in U.S. coastal markets, with prices that are higher than inland averages but not unusually extreme for these locations. Median properties are approximately 1,700 square feet, predominantly one-story, and about 40 years old. Consistent with the long-run nature of SLR, inundation risk is highly right-skewed: most transactions have near-zero projected inundation probability at horizons up to 120 years, but the upper tail becomes economically meaningful and grows quickly with the horizon.

Figure 5 summarizes the term structure of projected inundation probabilities. Scenario-specific probabilities rise monotonically with the horizon, with the increase steepest under RCP8.5. The gap between high- and low-emissions scenarios widens with H, which is reflected in an increasing disagreement measure $p_d(H)$ over this range. In parallel, the average risk measure $\bar{p}(H)$ also increases with H, capturing the growing long-run exposure of coastal properties to SLR in the sample.

[Table 1 about here.]

[Figure 5 about here.]

[Figure 6 about here.]

Remark 3 (Concentration of real estate sales near the coast). Figure 6 reveals an extremely right-skewed spatial distribution of transactions: activity is heavily concentrated at the shoreline and then declines smoothly with distance, with a long tail out to five miles. In the first 0.1-mile bin alone, there are about 1.0 million transactions; summing the first ten bins implies roughly 6 million transactions within one mile of the coast (Panel (a)). This concentration is also where exposure is most acute: within the first 0.1 mile, about 0.43 million transactions involve properties that would be permanently inundated at 6 ft of SLR (a common benchmark in the literature). Panel (b) shows that the same narrow coastal band also concentrates housing wealth: in the more recent 2010–2024 subsample, total nominal sale volume exceeds $500 billion in the first 0.1-mile bin alone. Taken together, the figure highlights that both coastal market activity and housing wealth are disproportionately located in the narrow band where long-run SLR risk is highest.

Remark 4 (RCP 8.5 truncation). NOAA's high-resolution inundation layers are available only up to 10 feet of sea-level rise. Consequently, for many locations the implied cumulative inundation probability under RCP8.5, $p_{\text{rcp8.5}}(H)$, increases with the horizon H but mechanically plateaus once projected SLR exceeds 10 feet—on average around $H \approx 130$ years in our sample. This censoring can attenuate or distort long-horizon estimates, particularly for disagreement measures that rely on RCP8.5 (e.g., $p_d(H) = p_{\text{rcp8.5}}(H) - p_{\text{rcp2.6}}(H)$) and for the simple average $\bar{p}(H)$. We therefore avoid using RCP8.5-based spread measures for horizons beyond 130 years. For very long-run horizons, scenario measures that are less affected by truncation—such as $p_{\text{rcp4.5}}(H)$—provide a more reliable proxy for long-run SLR risk than $\bar{p}(H)$.

4. Empirical evidence

4.1. Pricing of SLR risk

Table 2 evaluates the first implication of the model, that house prices declining in the probability of SLR inundation. Throughout, we measure SLR risk using horizon-specific probabilities of permanent inundation. Panel A uses sale price levels (so coefficients are in dollars), whereas Panel B uses log sale price (so coefficients are semi-elasticities and can be read as percent changes for a unit change in the probability).

Column (1): Naive cross section. Column (1) relates sale price to $p_{\text{rcp4.5}}(100)$ with baseline controls but without granular coastal fixed effects. The coefficient is positive and precisely estimated, reflecting spatial sorting: high-amenity coastal locations tend to have both higher prices and higher modeled long-run inundation probabilities, so raw cross-sectional correlations confound risk with coastal desirability.

Columns (2)–(3): Stringent fixed effects and the sign flip. Columns (2) and (3) add progressively stringent fixed effects. Column (2) introduces $ZDEY$ (ZIP5 $\times$ distance-to-coast bin $\times$ elevation bin $\times$ year-of-sale), and Column (3) further tightens time controls by replacing year-of-sale with year$\times$month-of-sale ($ZDEM$). Once these high-dimensional fixed effects absorb local amenity and time variation, the coefficient on $p_{\text{rcp4.5}}(100)$ flips sign and becomes economically and statistically negative in Panel A and Panel B. A literal unit-change interpretation is mechanically large (moving from $p(100) = 0$ to $p(100) = 1$), so a more informative scaling is a marginal change in risk: using Column (2), a 10 percentage-point increase in $p_{\text{rcp4.5}}(100)$ (i.e., $\Delta p = 0.1$) is associated with a price decline of about \$5,100 in levels (Panel A) and about 0.35% (35 basis points) in logs (Panel B).[12]

It is tempting to interpret the log estimate in Column (2) as the "tail" contribution of housing services beyond 100 years—i.e., as the effect of truncating the continuation value after 100 years. However, that reading relies on extrapolating a linear regression far beyond typical empirical variation in $p(100)$; we therefore treat it as a suggestive model-mapping heuristic rather than a literal truncation experiment. The marginal-change interpretation above is the safer economic summary.

Column (4): Window risk and the Gateaux-derivative mapping. Column (4) includes both $p_{\text{rcp4.5}}(100)$ and the window probability $p_{\text{rcp4.5}}(100, 110) := p_{\text{rcp4.5}}(110) - p_{\text{rcp4.5}}(100)$.

[12]The sample size declines once we add high-dimensional fixed effects because the estimator drops singleton observations within the absorption structure.

This specification more closely matches the Gateaux derivative emphasized by the model: the coefficient on $p_{\text{rcp4.5}}(100, 110)$ captures the incremental price effect of increasing the probability of inundation specifically in the $(100, 110]$ window, holding fixed the probability of inundation by year 100 (and hence holding fixed earlier-horizon risk). Empirically, the window coefficient is negative in both levels and logs, consistent with markets discounting risk that arrives in the far tail even after accounting for baseline 100-year risk.

Columns (5)–(6): Cross-scenario average risk. Columns (5) and (6) repeat the Column (2) and Column (4) exercises using the cross-scenario average risk measure $\bar{p}$ in place of $p_{\text{rcp4.5}}$. The resulting coefficients are again negative in both panels, indicating that the capitalization of long-run inundation risk is not driven by reliance on a single emissions scenario. Column (6) additionally includes the window probability $\bar{p}(100, 110)$, delivering the analogous "window" mapping under the averaged risk measure. Overall, across specifications and risk definitions, the results support the model prediction that higher long-run inundation risk is capitalized into lower house prices once local amenity and time variation are absorbed.

[Table 2 about here.]

4.1.1 Term structure

Figure 7 tests the model's term-structure prediction that the price impact of SLR risk should decline as the relevant horizon moves farther into the future. In Panel (a), we estimate a sequence of hedonic regressions using window risk $p(H, H{+}10) := p(H{+}10) - p(H)$. The coefficient $\hat{\beta}_{H,H+10}$ maps to the Gateaux derivative associated with shifting probability mass into the window $(H, H{+}10]$ (holding earlier-horizon risk fixed). In the model, this derivative is proportional to the negative of the "annuity weight"

$$-\mathcal{A}(H, H{+}10) \;:=\; -\int_{H}^{H+10} \mathcal{K}(t)\, dt,$$

which can be interpreted as the $t{=}0$ price of a claim that pays a safe unit flow over $t \in [H, H{+}10]$. Because discounting makes distant cash flows cheaper today, $\mathcal{A}(H, H{+}10)$ should decay with H, implying that $\hat{\beta}_{H,H+10}$ should become smaller in magnitude as H increases.

Panel (b) makes this decay pattern transparent by plotting $\ln|\hat{\beta}_{H,H+10}|$ against H. Let $r(t)$ be the effective discount rate defined by $\mathbb{E}_0[m(t)] = \exp\big(-r(t)\, t\big)$. If $r(t)$ is approximately constant over long horizons, then annuity prices (and hence $|\hat{\beta}_{H,H+10}|$) decline approximately exponentially in H, so $\ln|\hat{\beta}_{H,H+10}|$ should be close to linear in H;

systematic departures from linearity would instead suggest a non-flat long-run discount-rate term structure. Empirically, the coefficients decay strongly with horizon, and the log-linear fit in Panel (b) is tight, implying that an approximately flat long-run discount rate is a good approximation in this range. The estimated slope is about -0.028, i.e., an exponential decay rate of roughly 2.8% per year, consistent with discounting in the pricing kernel process $m(t)$ being well summarized by a constant long-run rate near 2.8%.

[Figure 7 about here.]

4.2. Ambiguity amplifies the price effect of risk

This section tests the model's *amplification* mechanism: ambiguity (cross-scenario disagreement) makes a given amount of SLR risk more negatively capitalized into house prices. The clean empirical object to map to the theory is the *level* of the value function v^* (dollars), not $\ln v^*$. The reason is conceptual: even if the cross-partial of v^* were zero, the cross-partial of $\ln v^*$ need not be, purely due to the curvature of the log transformation. Accordingly, Table 3 (price levels) is our main test, while Table 4 (log price) is a robustness check.

Null: Ambiguity neutrality. Proposition 6 delivers a crisp null hypothesis. Under ambiguity neutrality ($\alpha = 0$), the optimal tilted beliefs collapse to the baseline beliefs and ambiguity does not distort pricing, implying

$$\frac{\partial^2 v^*}{\partial \text{risk}\, \partial \text{ambiguity}} = 0.$$

In our empirical implementation, we operationalize risk by $p(H)$ (either $\bar{p}(H)$ or $p_{\text{rcp4.5}}(H)$) and ambiguity by $\text{spread}(H) := p_{\text{rcp8.5}}(H) - p_{\text{rcp2.6}}(H)$. Consider the level regression

$$\text{Price} = \beta_p\, p(H) \;+\; \beta_s\, \text{spread}(H) \;+\; \beta_{ps}\, p(H) \times \text{spread}(H) \;+\; \text{controls} + \text{FE}.$$

The null of ambiguity neutrality corresponds to $\beta_{ps} = 0$ in levels (up to the regression's scaling of p and spread), because β_{ps} is precisely the empirical analog of the cross-partial of v^* with respect to risk and ambiguity.

We robustly reject ambiguity neutrality. Table 3 shows that the interaction coefficient is negative and highly statistically significant across the main specifications. Using $\bar{p}(100)$ as the risk measure, the interaction estimate is $\widehat{\beta}_{ps} = -133{,}198$ (SE 24,356) in Column (1), $\widehat{\beta}_{ps} = -151{,}999$ (SE 32,679) in Column (2) with the more stringent $ZDEM$ fixed effects, and $\widehat{\beta}_{ps} = -136{,}584$ (SE 20,880) in the centered specification of Column (3). All three reject $\beta_{ps} = 0$ with t-statistics well in excess of conventional thresholds. This is direct evidence

against $\alpha = 0$: empirically, ambiguity *systematically* steepens the (negative) marginal price of risk.

In the interaction regression, the marginal price of risk is

$$\frac{\partial \,\text{Price}}{\partial p(H)} = \beta_p + \beta_{ps}\,\text{spread}(H).$$

Because $\widehat{\beta}_{ps} < 0$, the capitalization of risk becomes more negative as ambiguity increases. Figure 8 makes this rotation concrete by plotting the fitted marginal effect over the support of spread(100), along with confidence bands. The figure provides a numerical illustration of the amplification channel: moving to more ambiguous locations shifts the implied dollar price of risk sharply downward, even when β_p by itself is modest.

Column (6) implements the window design using $p(H, H{+}10)$ and spread$(H, H{+}10)$. This specification is useful because it maps most tightly to the Gateaux derivative associated with moving probability mass into the window $(H, H{+}10]$, holding earlier-horizon risk fixed. Consistent with the amplification mechanism, the window interaction term remains negative and statistically significant: $\widehat{\beta}_{ps} = -111{,}908$ (SE 41,784).

Robustness: log-price specifications. Table 4 repeats the same designs with $\ln(\text{Price})$ as the dependent variable. The interaction term remains negative and statistically significant in the main columns, consistent with amplification. However, we interpret these log results as corroborating evidence rather than the primary test, because the model's sharp null concerns the cross-partial of v^* and because $\ln(\cdot)$ can mechanically induce nonzero cross-effects even when the level cross-partial is zero.

Remark 5 (On the direct effect of ambiguity.). The model does *not* generally sign $\partial v^* / \partial \text{spread}$: its sign depends on whether buyers are sufficiently "optimistic" in the sense that the tilted weight on the good scenario dominates the bad one (e.g., whether $q_1^* \ell_1$ is large relative to $q_3^* \ell_3$). Empirically, the coefficient on spread(H) is mixed across specifications (and less robust than the interaction term), with some columns suggesting a positive direct association (a pattern consistent with local optimism) and others suggesting the opposite once risk is summarized differently or horizons shift. The robust and central fact is instead the interaction: regardless of the level effect of ambiguity, higher ambiguity reliably amplifies the negative capitalization of risk, strongly rejecting ambiguity neutrality.

[Table 3 about here.]

[Table 4 about here.]

[Figure 8 about here.]

5. Conclusion

This paper develops and empirically evaluates a stylized robust-control framework to study how long-run climate risk and ambiguity are capitalized into housing prices. The model distinguishes between *risk*—unknown realizations for known probabilities conditional on a given climate scenario—and *ambiguity*—not knowing which potential model or scenario from a set of possible climate scenarios will ultimately prevail. Within this framework, risk and ambiguity jointly shape asset prices, with ambiguity amplifying discounting of long-horizon climate risk.

We test the model's core implications using transaction-level data from U.S. coastal housing markets. Exploiting high-resolution spatial variation in exposure to permanent coastal inundation, we construct horizon-specific probabilities of inundation under alternative Representative Concentration Pathways and an ambiguity measure based on cross-scenario divergence. The empirical analysis documents three robust patterns consistent with the model: housing prices decline with long-run inundation risk; prices decline further with greater ambiguity holding expected risk fixed; and ambiguity amplifies the capitalization of risk through a negative interaction effect.

To our knowledge, this paper provides one of the first concrete mappings from a robust-control model of risk and ambiguity to observed market prices in a large asset market. A key contribution is the disciplined translation of physical climate projections into economically meaningful, horizon-specific probabilities that can be taken directly to the data. This mapping allows us to evaluate the theory of risk and ambiguity using observed prices rather than stylized experiments or calibration alone.

Empirically, the analysis leverages a uniquely large and carefully constructed dataset, combining more than 13.6 million coastal housing transactions with parcel-level inundation thresholds and local SLR projections. The scale of the data allows us to estimate highly saturated specifications that exploit fine-grained spatial variation while maintaining substantial statistical power.

This paper is a work in progress. Our ongoing and future work will explore broader implications of the estimated risk and ambiguity premia, including how they relate to households' willingness to invest in SLR adaptation. More generally, the framework developed here provides a foundation for studying how markets price long-run climate risks when both probabilities and models themselves are uncertain.

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, T.-N. (2025). Leveraging the disagreement on climate change: Theory and evidence. *Journal of Political Economy*, 133(10):3132–3166.

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.

Bilal, A. (2023). Solving heterogeneous agent models with the master equation. Technical report, National Bureau of Economic Research.

Bilal, A. and Rossi-Hansberg, E. (2023). Anticipating climate change across the united states. Technical report.

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.

Borovička, J., Hansen, L. P., Hendricks, M., and Scheinkman, J. A. (2011). Risk-price dynamics. *Journal of Financial Econometrics*, 9(1):3–65.

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. Technical report.

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.

Clarke, F. H. (1990). *Optimization and nonsmooth analysis*. SIAM.

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.

Dávila, E., Parlatore, C., and Walther, A. (2025). Probability pricing. Technical report, National Bureau of Economic Research.

Dávila, E. and Schaab, A. (2025). Welfare assessments with heterogeneous individuals. *Journal of Political Economy*, 133(9):2918–2961.

Dávila, E. and Walther, A. (2023). Prudential policy with distorted beliefs. *American Economic Review*, 113(7):1967–2006.

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

Engle, R. F., Giglio, S., Kelly, B., Lee, H., and Stroebel, J. (2020). Hedging climate change news. *The Review of Financial Studies*, 33(3):1184–1216.

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.

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.

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

Golosov, M., Tsyvinski, A., and Werquin, N. (2014b). A variational approach to the analysis of tax systems. Technical report, National Bureau of Economic Research.

Gordon, M. J. (1962). The investment, financing, and valuation of the corporation. *(No Title)*.

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 Souganidis, P. (2025). Stochastic responses and marginal valuation. *Proceedings of the National Academy of Sciences*, 122(48):e2520857122.

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

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.

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.

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

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. Technical report.

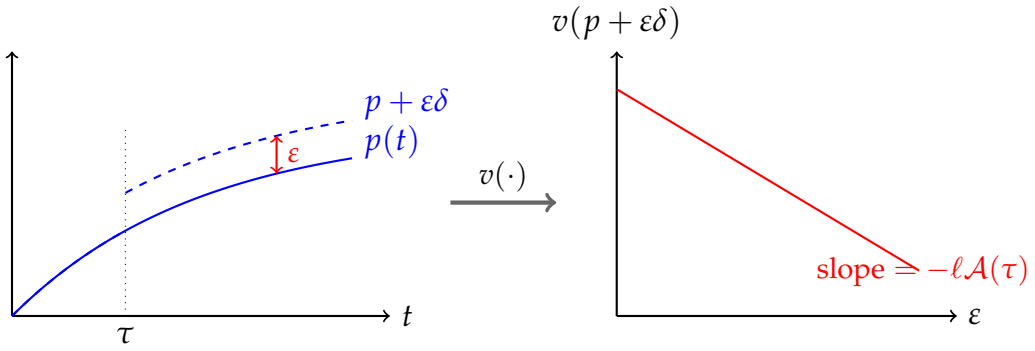
Rodziewicz, D., Amante, C. J., Dice, J., and Wahl, E. (2022). Housing market impairment from future sea-level rise inundation. *Environment Systems and Decisions*, 42(4):637–656.

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. Technical report.

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

Figures



Perturbation: $p \mapsto p + \varepsilon\delta$
where $\delta(t) = \mathbf{1}\{t \geq \tau\}$

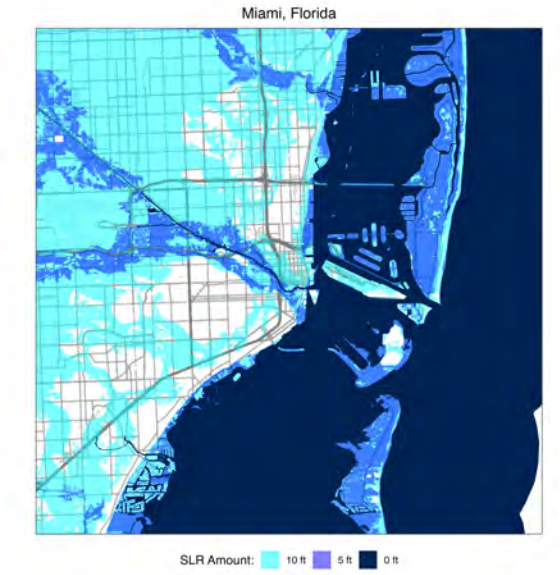
Value response along direction δ
is linear in perturbation size ε

Figure 1: Linearity of the valuation functional. *Left panel:* A perturbation $\delta(t) = \mathbf{1}\{t \geq \tau\}$ shifts the probability curve upward for all $t \geq \tau$, increasing the cumulative disaster probability $p(\tau)$ by ε. *Right panel:* Because $v(p) = v^{\text{safe}} - \langle \ell\mathcal{K}, p \rangle$ is affine in p, the value $v(p + \varepsilon\delta)$ responds linearly in ε, with slope $-\ell\mathcal{A}(\tau)$ where $\mathcal{A}(\tau) = \int_\tau^\infty \mathcal{K}(t)\,\mathrm{d}t$ is the continuation value.

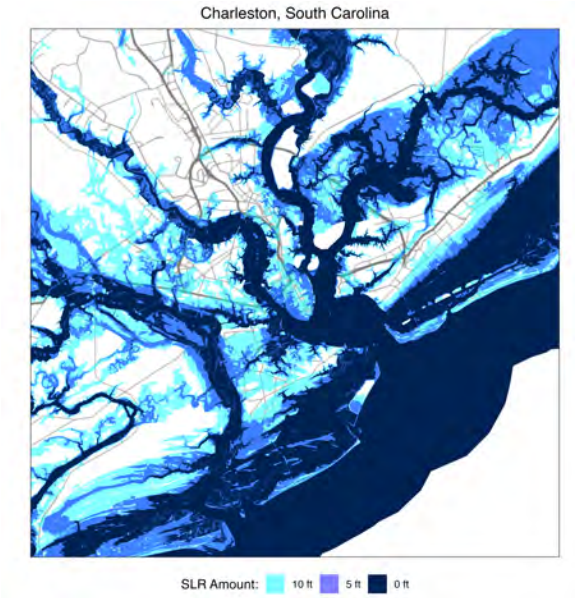
Figure 2: **High-resolution SLR inundation maps.** This figure illustrates NOAA Sea Level Rise Viewer inundation layers for three major coastal housing markets: New York City (panel A), Miami (panel B), and Charleston, South Carolina (panel C). Shaded areas indicate permanent inundation under hypothetical SLR amounts of 0, 5, and 10 feet. The maps highlight substantial within-neighborhood spatial heterogeneity in exposure driven by local elevation and hydrologic connectivity.



(a) New York City, NY



(b) Miami, FL



(c) Charleston, SC

Figure 3: **Local sea-level-rise projections for Miami (author's calculations).** This figure plots projected relative SLR for the Miami tide gauge under three Representative Concentration Pathways (RCP2.6, RCP4.5, and RCP8.5). Solid lines show median (50th percentile) projections and shaded regions depict uncertainty bands (10th–90th percentiles). The series are constructed by the authors from simulation outputs in the replication materials of DeConto and Pollard (2016).

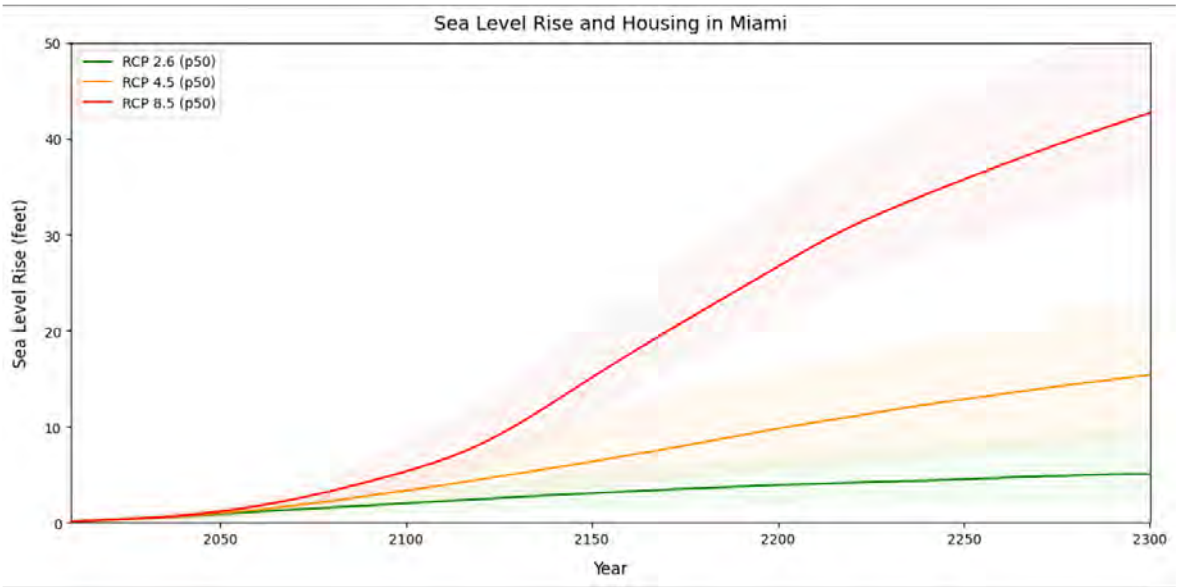
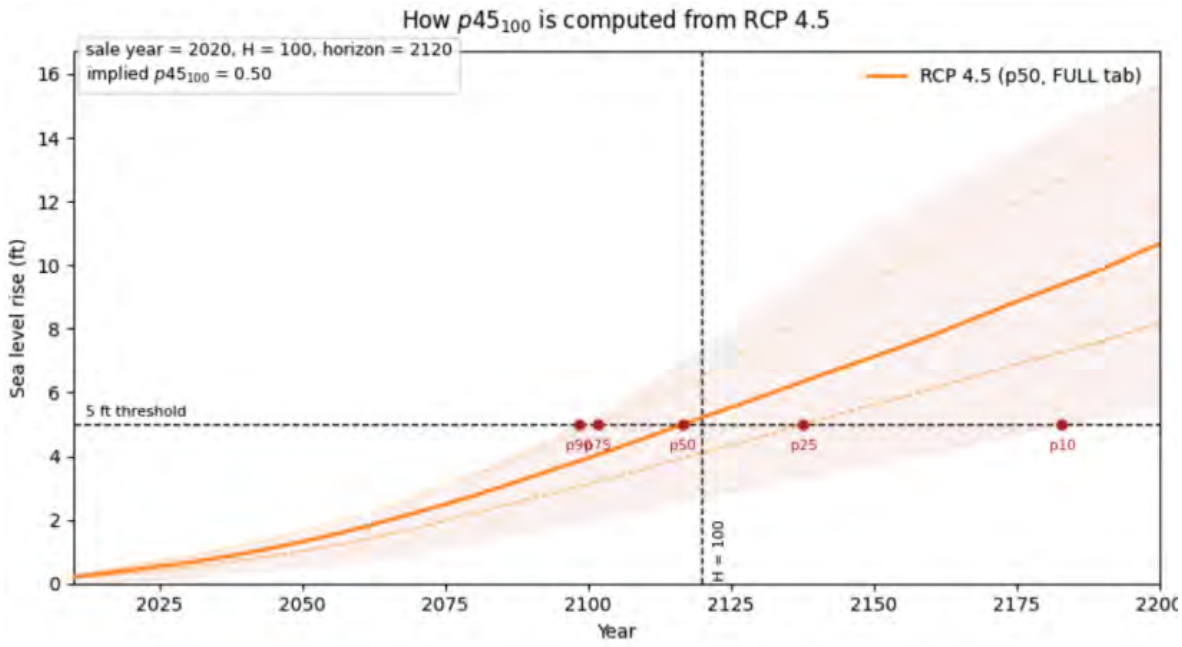


Figure 4: **Illustration: computing $p45_{100}$ from local RCP4.5 projections (author's calculations).** This figure illustrates the construction of the 100-year inundation probability under RCP4.5 for a representative transaction. The example considers a sale in year 2020 ($t = 2020$) and horizon $H = 100$, corresponding to calendar year $t + H = 2120$. The property's NOAA-based inundation threshold is $minSLR_i = 5$ feet (horizontal dashed line). Using the local RCP4.5 projection distribution for the nearest tide gauge (Miami), we compute $p45_{100} = \Pr(SLR_{g,4.5,2120} \geq 5)$ as the share of projected outcomes exceeding 5 feet at year 2120 (equivalently, one minus the CDF evaluated at 5 feet). The plotted percentiles and uncertainty band are constructed by the authors from simulation outputs in the replication materials of DeConto and Pollard (2016).



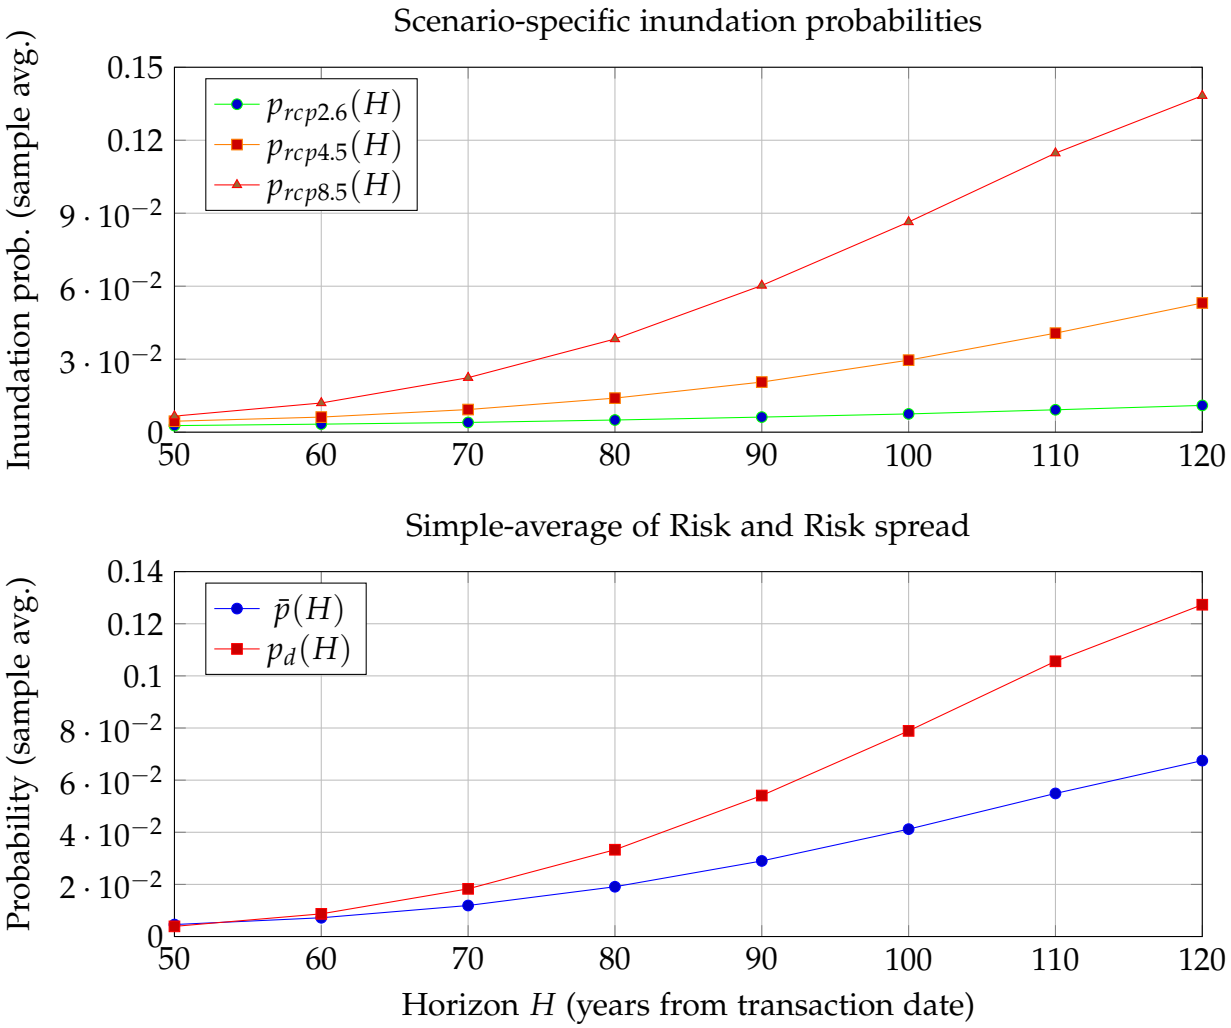
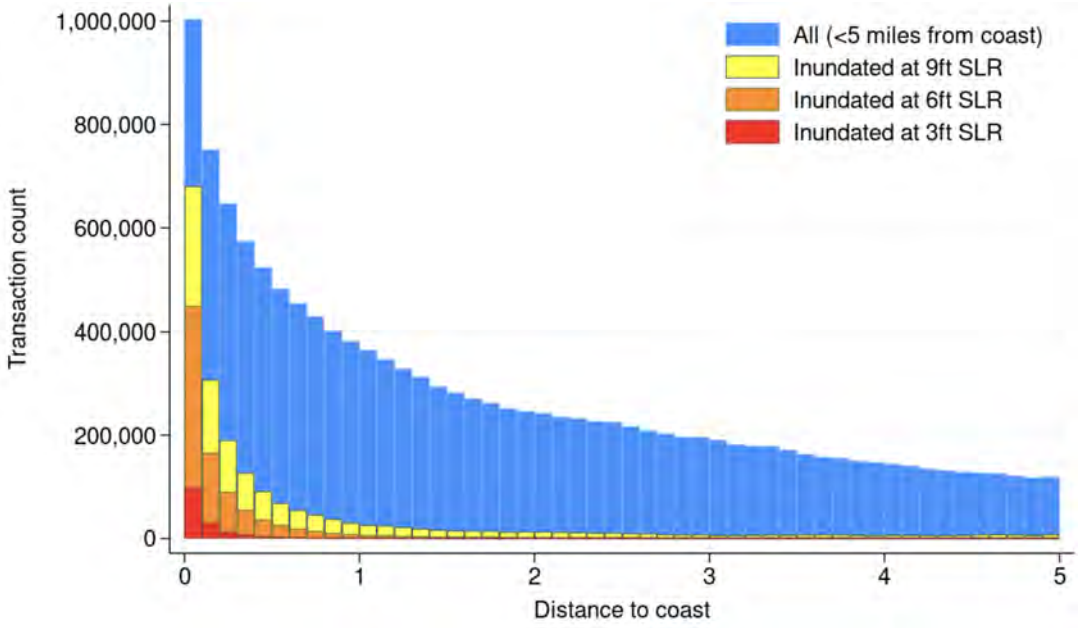
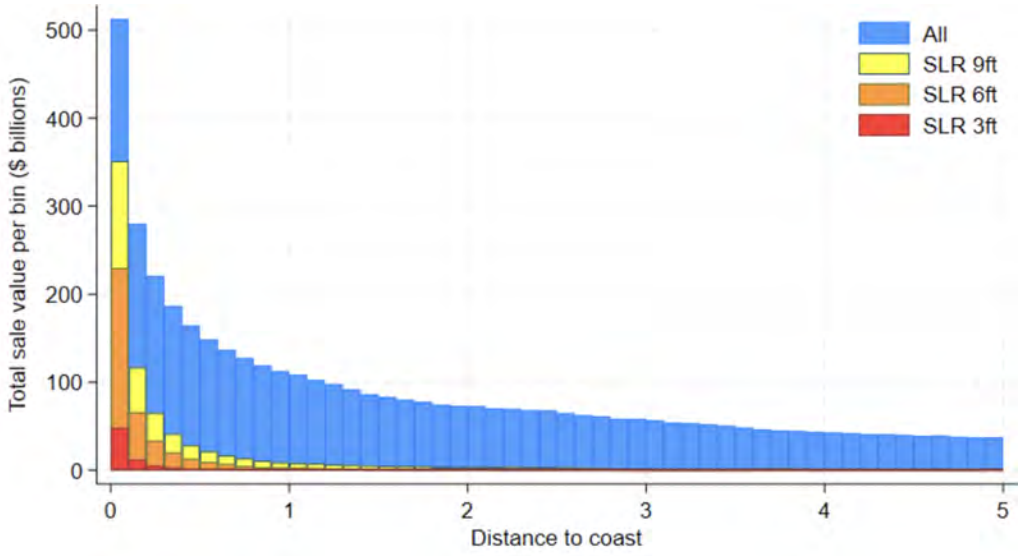


Figure 5: Projected SLR inundation probabilities in the transaction-level sample. The top panel plots cross-sectional mean inundation probabilities under RCP2.6, RCP4.5, and RCP8.5. The bottom panel plots the corresponding average risk measure $\bar{p}(H) := \frac{1}{3}(p_{rcp.26}(H) + p_{rcp4.5}(H) + p_{rcp8.5}(H))$ and disagreement measure $p_d(H) := p_{rcp8.5}(H) - p_{rcp2.6}(H)$. Horizons H are measured in years from the transaction date.

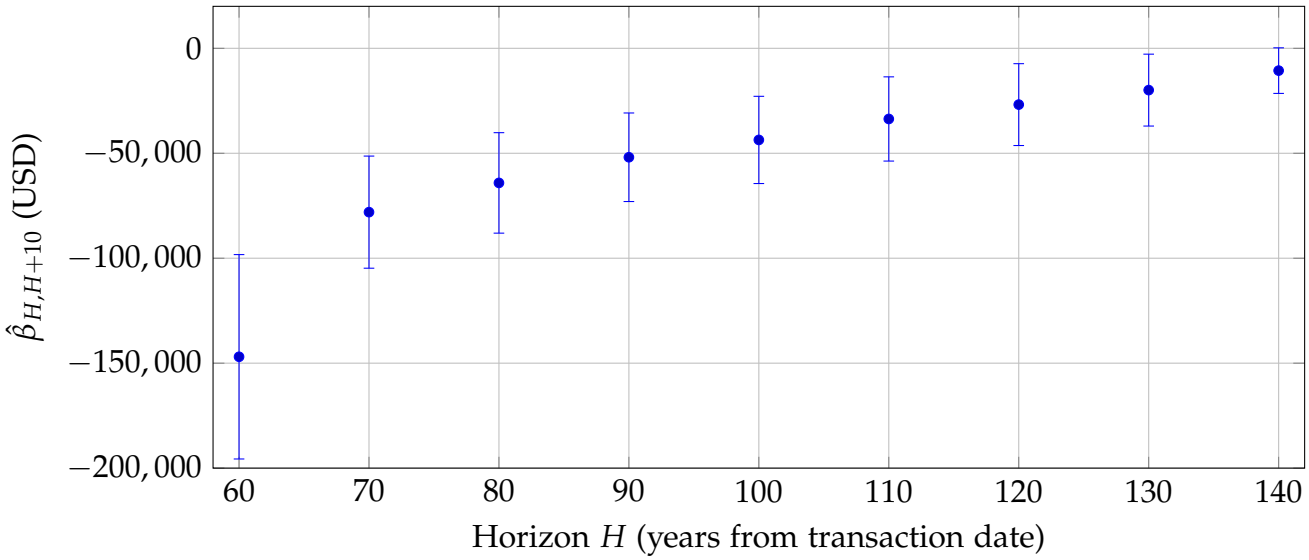
Figure 6: **Sales activity and housing wealth by distance to the coast (bin width 0.1 mile).** Panel (a) plots transaction counts for all single-family home sales from January 2001 to December 2024 within 5 miles of the U.S. coastline (after dropping outliers with sale price $<$ \$100,000). Panel (b) plots the corresponding dollar volume per bin, computed as the sum of transaction prices within each distance bin (nominal, not inflation-adjusted) for the 2010–2024 subsample. In both panels, colored overlays indicate transactions for properties that would be permanently inundated at x feet of sea-level rise (SLR), for $x \in \{3, 6, 9\}$.



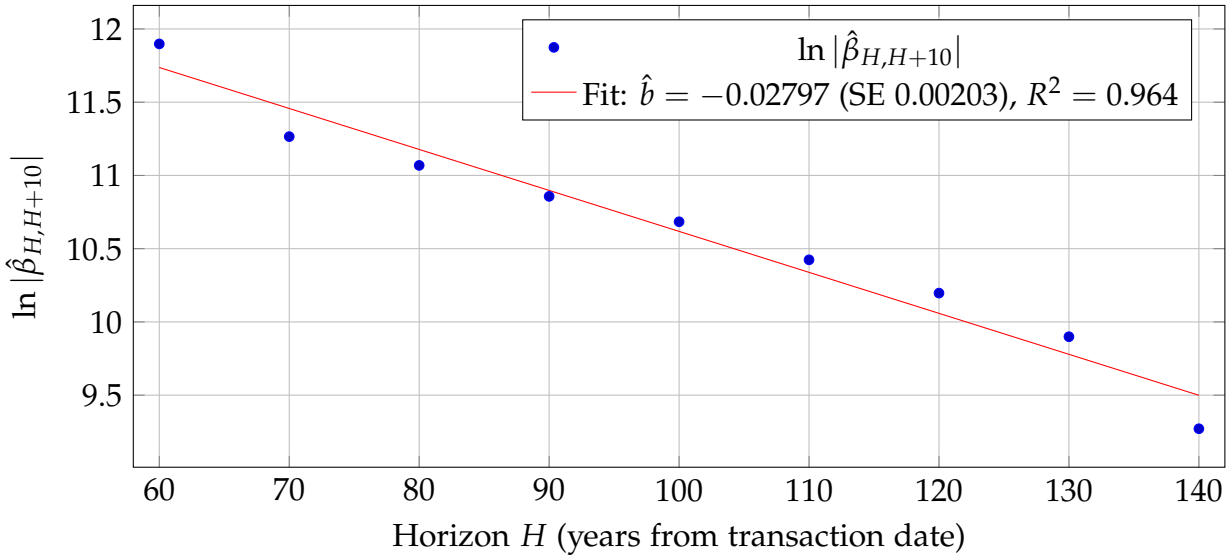
(a) Count of sales in full 2001-2024 sample.



(b) Total volume of sales (in nominal dollars) in 2010-2024 subsample.

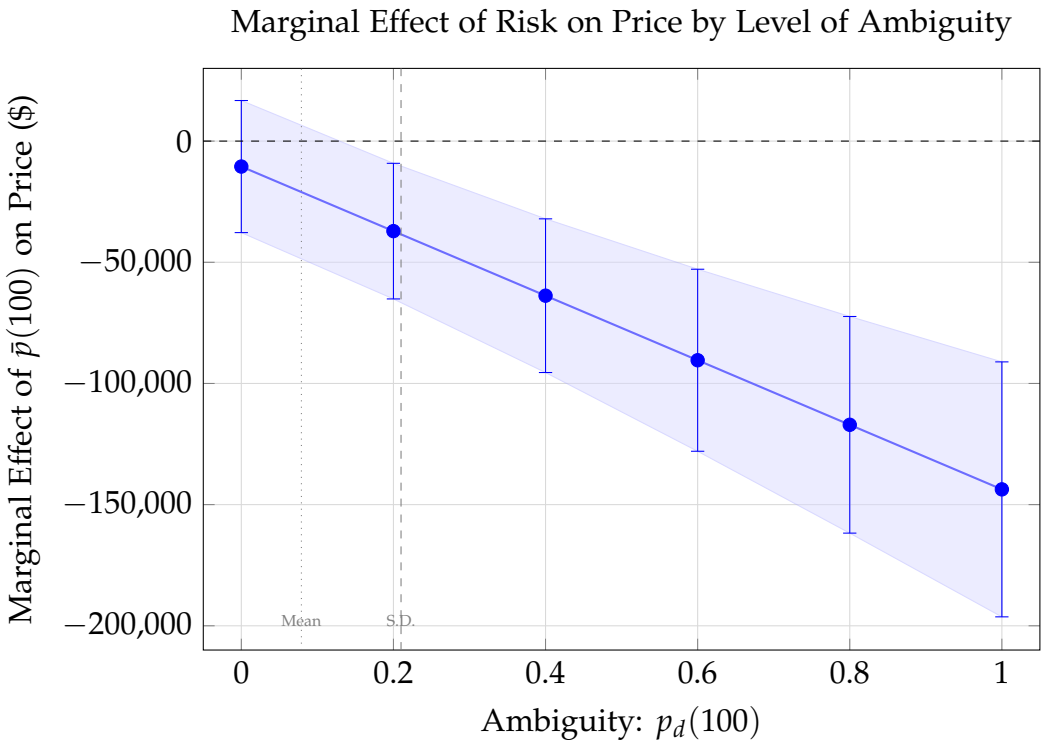


(a) Window-by-window coefficients in levels with 95% confidence intervals.



(b) Log absolute coefficients and linear fit across horizons.

Figure 7: **Window-by-window SLR capitalization and decay with horizon.** Panel (a) plots coefficient estimates $\hat{\beta}_{H,H+10}$ from regressions of residualized sale price on window risk $p(H, H+10)$ for $H \in \{60, 70, \ldots, 140\}$, with 95% confidence intervals based on ZIP5-clustered standard errors. Panel (b) plots $\ln|\hat{\beta}_{H,H+10}|$ and a fitted line $\ln|\hat{\beta}_{H,H+10}| = \hat{a} + \hat{b}H$ with $\hat{b} = -0.02797$ (SE 0.00203).



$p_d(100)$	Effect ($)	t-stat
$= 0$	$-10{,}504$	-0.76
$= 0.21$ (s.d.)	$-38{,}475^{***}$	-2.66

Figure 8: **Marginal Effect of Inundation Risk by Ambiguity Level.** This figure plots the marginal effect of mean inundation probability $\bar{p}(100)$ on sale price as a function of cross-scenario disagreement $p_d(100)$. The regression specification corresponds to Column (1) of Table **??**, Panel A: Price $= \beta_1 \bar{p}(100) + \beta_2 p_d(100) + \beta_3 \bar{p}(100) \times p_d(100) + \text{controls} + \text{FE}$. The marginal effect is $\partial \text{Price} / \partial \bar{p}(100) = \hat{\beta}_1 + \hat{\beta}_3 \cdot p_d(100) = -10{,}504 - 133{,}198 \times p_d(100)$. Shaded region shows 95% confidence interval.

Tables

Table 1: **Summary statistics.** This table reports summary statistics for key variables in the transaction-level analysis dataset, constructed by combining proprietary deed transactions from CoreLogic with high-resolution SLR inundation maps from NOAA and local SLR projections based on DeConto and Pollard (2016). The sample consists of the universe of arm's-length transactions of single-family homes located within five miles of the U.S. ocean coastline over the period 2001–2024. `sale_amount` is the transaction price in dollars; `sqft` is building square footage; `stories` is the number of stories; and `prop_age` is property age in years at the time of sale. For each property and horizon H, $p_{rcp2.6}(H)$, $p_{rcp4.5}(H)$, and $p_{rcp8.5}(H)$ denote the probability that the property is permanently inundated due to SLR within H years of the sale date under RCP2.6, RCP4.5, and RCP8.5, respectively. We define the average risk measure $\bar{p}(H) := \frac{1}{3}\left(p_{rcp2.6}(H) + p_{rcp4.5}(H) + p_{rcp8.5}(H)\right)$ and the disagreement (ambiguity) measure $p_d(H) := p_{rcp8.5}(H) - p_{rcp2.6}(H)$. Panel A reports prices and property characteristics. Panel B reports SLR risk measures at horizons $H \in \{80, 100, 120\}$.

Variable	Mean	SD	Median	P90	Max	N
Panel A. Prices and property characteristics						
`sale_amount`	455,898.1	482,334.2	305,000	867,000	3,025,000	13,658,396
`sqft`	1,900.2	823.0	1,703	2,977	5,018	13,166,384
`stories`	1.4	0.5	1.0	2.0	3.0	11,983,632
`prop_age`	41.6	29.9	38.0	84.0	124.0	12,527,485
Panel B. SLR risk measures at selected horizons						
$p_{rcp2.6}(80)$	0.005,0	0.053,5	0.000,0	0.000,0	0.900,0	13,758,322
$p_{rcp4.5}(80)$	0.014,0	0.090,4	0.000,0	0.000,0	0.900,0	13,758,322
$p_{rcp8.5}(80)$	0.038,3	0.146,7	0.000,0	0.100,0	0.900,0	13,758,322
$\bar{p}(80)$	0.019,1	0.089,0	0.000,0	0.033,3	0.900,0	13,758,322
$p_d(80)$	0.033,3	0.125,7	0.000,0	0.000,0	0.900,0	13,758,322
$p_{rcp2.6}(100)$	0.007,5	0.063,3	0.000,0	0.000,0	0.900,0	13,758,322
$p_{rcp4.5}(100)$	0.029,6	0.126,4	0.000,0	0.000,0	0.900,0	13,758,322
$p_{rcp8.5}(100)$	0.086,4	0.226,9	0.000,0	0.250,0	0.900,0	13,758,322
$\bar{p}(100)$	0.041,2	0.125,3	0.000,0	0.083,3	0.900,0	13,758,322
$p_d(100)$	0.078,9	0.208,0	0.000,0	0.250,0	0.900,0	13,758,322
$p_{rcp2.6}(120)$	0.011,0	0.075,3	0.000,0	0.000,0	0.900,0	13,758,322
$p_{rcp4.5}(120)$	0.053,1	0.165,1	0.000,0	0.250,0	0.900,0	13,758,322
$p_{rcp8.5}(120)$	0.138,3	0.311,3	0.000,0	0.900,0	0.900,0	13,758,322
$\bar{p}(120)$	0.067,5	0.164,6	0.000,0	0.383,3	0.900,0	13,758,322
$p_d(120)$	0.127,3	0.292,9	0.000,0	0.800,0	0.900,0	13,758,322

Table 2: **Capitalization of Inundation Risk.** Estimates from hedonic regressions of sale price (Panel A) and log sale price (Panel B) on SLR risk: $p_{\text{rcp4.5}}(H)$ denotes the probability that a property is permanently inundated by horizon H years from the transaction date under RCP4.5; $\bar{p}(H) := \frac{1}{3}\big(p_{\text{rcp2.6}}(H) + p_{\text{rcp4.5}}(H) + p_{\text{rcp8.5}}(H)\big)$; and $p(H_1, H_2) := p(H_2) - p(H_1)$. ZDEY fixed effects correspond to ZIP5×distance-to-coast bin (p5)×elevation bin (p5)×year of sale; ZDEM replaces year with month-year. All specifications include standard property controls (square footage, number of stories, and property age). Standard errors clustered by ZIP5 reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
Fixed effects		ZDEY	ZDEM	ZDEY	ZDEY	ZDEY
Panel A: Dependent variable is Sale Price ($)						
$p_{\text{rcp4.5}}(100)$	$330{,}940^{***}$	$-50{,}738^{***}$	$-73{,}894^{***}$	$-54{,}723^{***}$		
	(28,911)	(10,836)	(15,599)	(11,320)		
$p_{\text{rcp4.5}}(100, 110)$				$-42{,}667^{***}$		
				(7,725)		
$\bar{p}(100)$					$-59{,}511^{***}$	$-61{,}405^{***}$
					(13,168)	(13,343)
$\bar{p}(100, 110)$						$-25{,}386^{**}$
						(12,174)
Observations	11,418,736	10,504,699	6,103,086	10,504,699	10,504,699	10,504,699
Adj. R^2	0.26	0.78	0.80	0.78	0.78	0.78
Prop. Controls	Y	Y	Y	Y	Y	Y
Panel B: Dependent variable is $\ln(Price)$						
$p_{\text{rcp4.5}}(100)$		-0.035^{**}		-0.038^{**}		
		(0.016)		(0.017)		
$p_{\text{rcp4.5}}(100, 110)$				-0.030^{**}		
				(0.012)		
$\bar{p}(100)$					-0.040^{**}	-0.042^{**}
					(0.020)	(0.020)
$\bar{p}(100, 110)$						-0.017
						(0.016)
Observations		10,504,699		10,504,699	10,504,699	10,504,699
Adj. R^2		0.83		0.83	0.83	0.83
Prop. Controls	Y	Y	Y	Y	Y	Y

Table 3: **Ambiguity Amplification of Risk Capitalization (Price Levels).** Estimates from hedonic regressions of sale price on SLR risk and ambiguity. The risk measure $p(H)$ denotes either the cross-scenario average inundation probability $\bar{p}(H) := \frac{1}{3}\big(p_{\text{rcp2.6}}(H) + p_{\text{rcp4.5}}(H) + p_{\text{rcp8.5}}(H)\big)$ or the scenario-specific probability $p_{\text{rcp4.5}}(H)$, as indicated in the table. $\text{spread}(H) := p_{\text{rcp8.5}}(H) - p_{\text{rcp2.6}}(H)$ measures cross-scenario disagreement (ambiguity), and $p(H_1, H_2) := p(H_2) - p(H_1)$ denotes the probability of inundation in window $(H_1, H_2]$. ZDEY fixed effects correspond to ZIP5×distance-to-coast bin (p5)×elevation bin (p5)×year of sale; ZDEM replaces year of sale with year×month of sale. Specifications marked *centered* use mean-centered regressors. All specifications include standard property controls. Standard errors clustered by ZIP5 are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
Risk measure	$\bar{p}$	$\bar{p}$	$\bar{p}$	$p_{\text{rcp4.5}}$	$\bar{p}$	$\bar{p}$
Centering			Y			
Horizon	H=100	H=100	H=100	H=100	H=80	H=100
Fixed effects	ZDEY	ZDEM	ZDEY	ZDEY	ZDEY	ZDEY
Dependent variable: Sale Price ($)						
$p(H)$	−10,504	−38,138*	−29,475**	−11,172	21,459	−43,083***
	(13,903)	(20,574)	(13,733)	(13,913)	(14,155)	(14,723)
$\text{spread}(H)$	21,811**	32,302**	48,443***	−10,062	−42,515***	−18,143***
	(10,306)	(13,354)	(8,844)	(6,347)	(11,020)	(6,532)
$p(H) \times \text{spread}(H)$	−133,198***	−151,999***	−136,584***	−73,566***	−92,734***	
	(24,356)	(32,679)	(20,880)	(18,417)	(26,186)	
$p(H, H{+}10)$						−53,818***
						(17,151)
$\text{spread}(H, H{+}10)$						38,825***
						(9,573)
$p(H, H{+}10) \times \text{spread}(H, H{+}10)$						−111,908***
						(41,784)
Observations	10,504,699	6,103,086	10,504,699	10,504,699	10,504,699	10,504,699
Adj. R^2	0.78	0.80	0.78	0.78	0.78	0.78
Prop. Controls	Y	Y	Y	Y	Y	Y

Table 4: **Ambiguity Amplification of Risk Capitalization (Log Price).** This table reports log price specifications corresponding to Table **??**. All variable definitions, fixed effects, and controls are identical.

	(1)	(2)	(3)	(4)	(5)	(6)
Risk measure	$\bar{p}$	$\bar{p}$	$\bar{p}$	$p_{\text{rcp4.5}}$	$\bar{p}$	$\bar{p}$
Centering			Y			
Horizon	H=100	H=100	H=100	H=100	H=80	H=100
Fixed effects	ZDEY	ZDEM	ZDEY	ZDEY	ZDEY	ZDEY
Dependent variable: $\ln(Price)$						
$p(H)$	0.004	-0.027	-0.006	0.003	0.032^{*}	-0.025
	(0.018)	(0.027)	(0.018)	(0.018)	(0.019)	(0.020)
spread(H)	0.019	0.025	0.014	-0.006	-0.035^{*}	-0.017^{**}
	(0.013)	(0.018)	(0.012)	(0.009)	(0.018)	(0.009)
$p(H) \times$ spread(H)	-0.119^{***}	-0.112^{**}	-0.119^{***}	-0.073^{***}	-0.084^{**}	
	(0.036)	(0.048)	(0.036)	(0.025)	(0.033)	
$p(H, H{+}10)$						-0.038
						(0.024)
spread$(H, H{+}10)$						0.032^{**}
						(0.015)
$p(H, H{+}10) \times$ spread$(H, H{+}10)$						-0.099
						(0.064)
Observations	10,504,699	6,103,086	10,504,699	10,504,699	10,504,699	10,504,699
Adj. R^2	0.83	0.85	0.83	0.83	0.83	0.83
Prop. Controls	Y	Y	Y	Y	Y	Y

A. Proofs

A.1. Proof of Lemma 1

Proof. Fix a horizon $\tau \geq 0$. An increase in the cumulative disaster probability $p(\tau)$ by $\varepsilon > 0$ corresponds to an increase in probability mass for all $t \geq \tau$. To define an admissible perturbation within the space of probability functions, fix $L > \tau$ and define the perturbed probability curve

$$p^{(\varepsilon,L)}(t) := p(t) + \varepsilon\,\mathbf{1}\{t \geq \tau\} - \varepsilon\,\mathbf{1}\{t \geq L\}.$$

For sufficiently small ε, this perturbation preserves monotonicity and ensures $p^{(\varepsilon,L)}(t) \in [0,1]$ for all t. Moreover, $p^{(\varepsilon,L)}(\tau) = p(\tau) + \varepsilon$, so $p(\tau)$ increases by ε.

By linearity of $v(\cdot)$ in p,

$$\begin{aligned}\frac{v\left(p^{(\varepsilon,L)}\right) - v(p)}{\varepsilon} &= -\ell \int_0^\infty \mathbb{E}_0[m(t)]\left(\mathbf{1}\{t \geq \tau\} - \mathbf{1}\{t \geq L\}\right) \mathrm{d}t \\ &= -\ell \int_\tau^\infty \mathbb{E}_0[m(t)]\,\mathrm{d}t + \ell \int_L^\infty \mathbb{E}_0[m(t)]\,\mathrm{d}t \\ &= -\ell\,\mathcal{A}(\tau) + \ell\,\mathcal{A}(L).\end{aligned}$$

Taking the limit as $\varepsilon \to 0$ and then letting $L \to \infty$, using $\mathcal{A}(L) \to 0$, yields

$$\frac{\partial v}{\partial p(\tau)} = -\ell\,\mathcal{A}(\tau),$$

establishing (5). □

A.2. Proof of Lemma 3

Proof. The Lagrangian for (11) is

$$\mathcal{L}(q,\lambda) = \sum_{i=1}^N q_i v_i + \frac{1}{\alpha}\sum_{i=1}^N \pi_i f\left(\frac{q_i}{\pi_i}\right) + \lambda\left(\sum_{i=1}^N q_i - 1\right).$$

For an interior optimizer ($q_i^* > 0$), the first-order condition for each i is

$$v_i + \frac{1}{\alpha} f'\left(\frac{q_i^*}{\pi_i}\right) + \lambda = 0, \qquad i = 1,\ldots,N. \tag{32}$$

Equivalently,

$$f'\left(\frac{q_i^*}{\pi_i}\right) = -\alpha\,(v_i + \lambda), \qquad i = 1,\ldots,N. \tag{33}$$

Since f is convex, f' is invertible on $\mathbb{R}_+$, and

$$q_i^* = \pi_i\,(f')^{-1}\big(-\alpha\,(v_i + \lambda)\big), \qquad i = 1,\ldots,N, \tag{34}$$

where λ is chosen so that $\sum_{i=1}^N q_i^* = 1$. (Interior feasibility requires parameters such that $q_i^* > 0$ for all i.)

KL divergence. For KL, take

$$f_{\mathrm{KL}}(x) = x \log x, \quad \Rightarrow \quad D_{f_{\mathrm{KL}}}(q\|\pi) = \sum_{i=1}^N q_i \log \frac{q_i}{\pi_i}.$$

Since $f'_{\mathrm{KL}}(x) = \log x + 1$, (33) gives

$$\log\left(\frac{q_i^*}{\pi_i}\right) + 1 = -\alpha(v_i + \lambda),$$

hence the exponential tilt

$$q_i^* = \frac{\pi_i e^{-\alpha v_i}}{\sum_{k=1}^N \pi_k e^{-\alpha v_k}}, \qquad i = 1, \ldots, N. \tag{35}$$

The minimized objective has the standard log-sum-exp form:

$$v^* = -\frac{1}{\alpha} \log\left(\sum_{k=1}^N \pi_k e^{-\alpha v_k}\right), \tag{36}$$

which is the familiar *entropic (or certainty-equivalent) aggregator*: a smooth, ambiguity-averse analogue of the minimum operator that interpolates between the Bayesian expected value and the worst-case valuation as α increases.

The distorted expected value is $v^\circ := \sum_{i=1}^N q_i^* v_i$.

χ^2 divergence (interior solution). For χ^2, take

$$f_{\chi^2}(x) = \frac{1}{2}(x-1)^2, \quad \Rightarrow \quad D_{f_{\chi^2}}(q\|\pi) = \frac{1}{2}\sum_{i=1}^N \frac{(q_i - \pi_i)^2}{\pi_i}.$$

Since $f'_{\chi^2}(x) = x - 1$, (33) yields

$$\frac{q_i^*}{\pi_i} - 1 = -\alpha(v_i + \lambda), \qquad i = 1, \ldots, N.$$

Imposing $\sum_{i=1}^N q_i^* = 1$ gives $\lambda = -\bar{v}$ with $\bar{v}$ from (14), and therefore

$$q_i^* = \pi_i\Big[1 - \alpha\,(v_i - \bar{v})\Big], \qquad i = 1, \ldots, N. \tag{37}$$

Define the prior variance $\mathrm{Var}_\pi(v) := \sum_{i=1}^N \pi_i (v_i - \bar{v})^2$. Then the distorted expected value simplifies to

$$v^\circ = \sum_{i=1}^N q_i^* v_i = \bar{v} - \alpha\,\mathrm{Var}_\pi(v), \tag{38}$$

and the minimized objective value is

$$v^* = v^\circ + \frac{1}{\alpha} D_{f_{\chi^2}}(q^* \| \pi) = \bar{v} - \frac{\alpha}{2} \operatorname{Var}_\pi(v). \tag{39}$$

□

A.3. Proof of Lemma 4

Proof. Let $\mathcal{N} := (L^2([0,\infty)))^N$ with inner product

$$\langle a, b \rangle_{\mathcal{N}} = \sum_{i=1}^{N} \int_0^\infty a_i(t) b_i(t)\, \mathrm{d}t.$$

From (18), for any admissible direction $\delta = (\delta_1, \ldots, \delta_N) \in \mathcal{N}$,

$$\mathrm{D}v^*(p)[\delta] = -\sum_{i=1}^{N} q_i^*\, \ell_i \int_0^\infty \mathbb{E}_0[m(t)]\, \delta_i(t)\, \mathrm{d}t. \tag{40}$$

Step 1: Bounded linear functional on $\mathcal{N}$. Linearity in δ is immediate from (40). To verify boundedness, apply Cauchy–Schwarz in each coordinate:

$$\left| \int_0^\infty \mathbb{E}_0[m(t)]\, \delta_i(t)\, \mathrm{d}t \right| \le \|\mathbb{E}_0[m]\|_{L^2}\, \|\delta_i\|_{L^2}.$$

Therefore,

$$\begin{aligned} \left|\mathrm{D}v^*(p)[\delta]\right| &\le \sum_{i=1}^{N} q_i^*\, \ell_i\, \|\mathbb{E}_0[m]\|_{L^2}\, \|\delta_i\|_{L^2} \\ &\le \|\mathbb{E}_0[m]\|_{L^2} \Big(\sum_{i=1}^{N} q_i^*\, \ell_i \Big) \Big(\sum_{i=1}^{N} \|\delta_i\|_{L^2}^2 \Big)^{1/2}, \end{aligned}$$

where the last inequality uses Cauchy–Schwarz in $\mathbb{R}^N$ together with $\sum_{i=1}^N (q_i^* \ell_i)^2 \le \big(\sum_{i=1}^N q_i^* \ell_i\big)^2$. Since $\sum_{i=1}^N q_i^* = 1$ and $\ell_i \in (0,1)$, we have $\sum_{i=1}^N q_i^* \ell_i \le 1$, hence

$$\left|\mathrm{D}v^*(p)[\delta]\right| \le \|\mathbb{E}_0[m]\|_{L^2}\, \|\delta\|_{\mathcal{N}}, \qquad \|\delta\|_{\mathcal{N}} := \langle \delta, \delta \rangle_{\mathcal{N}}^{1/2}.$$

Thus $\mathrm{D}v^*(p)$ is a bounded linear functional on $\mathcal{N}$.

Step 2: Riesz representation. Because $\mathcal{N}$ is a Hilbert space, the Riesz representation theorem implies that there exists a unique element $\nabla_p v^*(p) \in \mathcal{N}$ such that

$$\mathrm{D}v^*(p)[\delta] = \langle \nabla_p v^*(p), \delta \rangle_{\mathcal{N}} \qquad \text{for all } \delta \in \mathcal{N}.$$

Define

$$\nabla_p v^*(p)(t) := -\big(q_1^*\,\ell_1\,\mathbb{E}_0[m(t)], \ldots, q_N^*\,\ell_N\,\mathbb{E}_0[m(t)]\big).$$

Then, for any $\delta \in \mathcal{N}$,

$$\begin{aligned}\langle \nabla_p v^*(p), \delta\rangle_{\mathcal{N}} &= -\sum_{i=1}^{N} \int_0^\infty \big(q_i^*\,\ell_i\,\mathbb{E}_0[m(t)]\big)\,\delta_i(t)\,\mathrm{d}t \\ &= -\sum_{i=1}^{N} q_i^*\,\ell_i \int_0^\infty \mathbb{E}_0[m(t)]\,\delta_i(t)\,\mathrm{d}t \\ &= \mathrm{D}v^*(p)[\delta],\end{aligned}$$

where the last equality follows from (40). This verifies that the stated $\nabla_p v^*(p)$ is indeed the Riesz representer, and uniqueness follows directly from the Riesz theorem. □

A.4. Proof of Lemma 5

Proof. Let $X_i := \ell_i$ and $Y_i := \ell_1 \mathbf{1}\{i = 1\} - \ell_N \mathbf{1}\{i = N\}$ be scalar random variables on $\{1, \ldots, N\}$ under probability mass function q^*. By definition,

$$\mathrm{Cov}_{q^*}(X_i, Y_i) = \mathbb{E}_{q^*}[X_i Y_i] - \mathbb{E}_{q^*}[X_i]\mathbb{E}_{q^*}[Y_i].$$

We have

$$\mathbb{E}_{q^*}[X_i] = \sum_{i=1}^{N} q_i^*\,\ell_i =: \bar{\ell}_{q^*}, \qquad \mathbb{E}_{q^*}[Y_i] = q_1^*\ell_1 - q_N^*\ell_N,$$

and

$$\mathbb{E}_{q^*}[X_i Y_i] = \sum_{i=1}^{N} q_i^*\,\ell_i\big(\ell_1 \mathbf{1}\{i=1\} - \ell_N \mathbf{1}\{i=N\}\big) = q_1^*\ell_1^2 - q_N^*\ell_N^2.$$

Therefore,

$$\begin{aligned}\mathrm{Cov}_{q^*}(X_i, Y_i) &= \big(q_1^*\ell_1^2 - q_N^*\ell_N^2\big) - \bar{\ell}_{q^*}\big(q_1^*\ell_1 - q_N^*\ell_N\big) \\ &= q_1^*\ell_1(\ell_1 - \bar{\ell}_{q^*}) + q_N^*\ell_N(\bar{\ell}_{q^*} - \ell_N).\end{aligned}$$

Because $q_i^* > 0$ for all i and $\ell_1 < \cdots < \ell_N$, the weighted mean satisfies

$$\ell_1 < \bar{\ell}_{q^*} < \ell_N.$$

Hence $\ell_1 - \bar{\ell}_{q^*} < 0$ and $\bar{\ell}_{q^*} - \ell_N < 0$. Since $q_1^*\ell_1 > 0$ and $q_N^*\ell_N > 0$, both terms in the final expression are strictly negative, so the covariance is strictly negative:

$$\mathrm{Cov}_{q^*}\Big(\ell_i,\ \ell_1 \mathbf{1}\{i=1\} - \ell_N \mathbf{1}\{i=N\}\Big) < 0.$$

This proves (30). □

A.5. Derivations for Section 2.2.4

A convenient compensated-step direction. Fix $\tau \geq 0$ and $L > \tau$. Define the compensated tail step

$$\varphi^{(\tau,L)}(t) \;:=\; \mathbf{1}\{t \geq \tau\} - \mathbf{1}\{t \geq L\}.$$

Then $\varphi^{(\tau,L)} \in L^2([0,\infty))$ and

$$\langle \mathbb{E}_0[m], \varphi^{(\tau,L)} \rangle = \int_\tau^\infty \mathbb{E}_0[m(t)]\, dt - \int_L^\infty \mathbb{E}_0[m(t)]\, dt = \mathcal{A}(\tau) - \mathcal{A}(L).$$

1) Derivation of $\partial v^*/\partial p_i(\tau)$. Fix $i \in \{1,\dots,N\}$. Consider the admissible direction $\delta^{(\tau,L;i)} \in \mathcal{N}$ defined by

$$\delta_i^{(\tau,L;i)}(t) \;:=\; \varphi^{(\tau,L)}(t), \qquad \delta_j^{(\tau,L;i)}(t) \equiv 0 \;\; (j \neq i).$$

For small $\varepsilon > 0$, the perturbed probability curve $p_i^\varepsilon := p_i + \varepsilon\, \delta_i^{(\tau,L;i)}$ raises probability by ε on $[\tau, L)$ and restores it by ε for $t \geq L$. Thus it increases the horizon-τ disaster probability $p_i(\tau)$ at first order, while remaining an admissible (square-integrable) perturbation.

Applying Lemma 4 gives

$$\mathrm{D}v^*(p)\big[\delta^{(\tau,L;i)}\big] = -\sum_{k=1}^N q_k^*\, \ell_k\, \langle \mathbb{E}_0[m], \delta_k^{(\tau,L;i)} \rangle = -q_i^*\, \ell_i\, \langle \mathbb{E}_0[m], \varphi^{(\tau,L)} \rangle = -\, q_i^*\, \ell_i \big(\mathcal{A}(\tau) - \mathcal{A}(L)\big).$$

Taking $L \to \infty$ and using $\mathcal{A}(L) \to 0$ yields

$$\frac{\partial v^*}{\partial p_i(\tau)} = \lim_{L\to\infty} \mathrm{D}v^*(p)\big[\delta^{(\tau,L;i)}\big] = -\, q_i^*\, \ell_i\, \mathcal{A}(\tau).$$

2) Derivation of $\partial v^*/\partial \bar{p}(\tau)$. Define the common-shift direction $\delta^{(\tau,L;\mathrm{com})} \in \mathcal{N}$ by

$$\delta_i^{(\tau,L;\mathrm{com})}(t) \;:=\; \varphi^{(\tau,L)}(t), \qquad i = 1,\dots,N.$$

This raises each probability curve by ε on $[\tau, L)$ and restores the same mass at $t \geq L$, corresponding to a uniform first-order increase in horizon-τ risk across scenarios.

By Lemma 4,

$$\begin{aligned} \mathrm{D}v^*(p)\big[\delta^{(\tau,L;\mathrm{com})}\big] &= -\sum_{i=1}^N q_i^*\, \ell_i\, \langle \mathbb{E}_0[m], \varphi^{(\tau,L)} \rangle \\ &= -\Big(\sum_{i=1}^N q_i^*\, \ell_i\Big)\big(\mathcal{A}(\tau) - \mathcal{A}(L)\big). \end{aligned}$$

Letting $L \to \infty$ yields

$$\frac{\partial v^*}{\partial \bar{p}(\tau)} = -\Big(\sum_{i=1}^N q_i^*\, \ell_i\Big)\mathcal{A}(\tau).$$

3) A useful identity for KL: cross partials and a covariance term. Under KL, q^* is interior and satisfies the standard derivative identity (a direct differentiation of the soft-min weights from Lemma 3): for any scalar perturbation z,

$$\frac{\partial q_i^*}{\partial z} = -\alpha\, q_i^* \Big(\frac{\partial v_i}{\partial z} - \sum_{k=1}^{N} q_k^* \frac{\partial v_k}{\partial z} \Big).$$

Combining this with the envelope condition $\partial v^*/\partial v_i = q_i^*$ implies that for any pair (z, w),

$$\frac{\partial^2 v^*}{\partial z\, \partial w} = \sum_{i=1}^{N} q_i^* \frac{\partial^2 v_i}{\partial z\, \partial w} - \alpha\, \mathrm{Cov}_{q^*}\Big(\frac{\partial v_i}{\partial z}, \frac{\partial v_i}{\partial w} \Big).$$

In our application, v_i is affine in (p_i, ℓ_i), so the mixed second derivatives $\partial^2 v_i/(\partial z\, \partial w)$ are zero for the step perturbations used below. Hence, the cross partial of v^* is entirely the (negative) covariance term.

4) Cross partial $\partial^2 v^*/(\partial \bar{p}(\tau)\, \partial p_d(\tau))$. Use the (compensated) extreme-spread direction that increases $p_N(\tau)$ and decreases $p_1(\tau)$ (holding $2, \dots, N-1$ fixed):

$$\delta^{(\tau, L; d)}(t) := \big(-\varphi^{(\tau, L)}(t),\, 0, \dots, 0,\, \varphi^{(\tau, L)}(t) \big).$$

Under this direction, the scenario-level derivatives are

$$\frac{\partial v_i}{\partial \bar{p}(\tau)} = -\ell_i\, \mathcal{A}(\tau), \qquad \frac{\partial v_i}{\partial p_d(\tau)} = \begin{cases} +\ell_1\, \mathcal{A}(\tau), & i = 1, \\ 0, & 2 \le i \le N-1, \\ -\ell_N\, \mathcal{A}(\tau), & i = N. \end{cases}$$

Therefore, under KL,

$$\begin{aligned} \frac{\partial^2 v^*}{\partial \bar{p}(\tau)\, \partial p_d(\tau)} &= -\alpha\, \mathrm{Cov}_{q^*}\Big(\frac{\partial v_i}{\partial \bar{p}(\tau)}, \frac{\partial v_i}{\partial p_d(\tau)} \Big) \\ &= -\alpha\, \mathrm{Cov}_{q^*}\Big(-\ell_i\, \mathcal{A}(\tau),\ \mathcal{A}(\tau)\big(\ell_1 \mathbf{1}\{i=1\} - \ell_N \mathbf{1}\{i=N\}\big) \Big) \\ &= \alpha\, \mathcal{A}(\tau)^2\, \mathrm{Cov}_{q^*}\Big(\ell_i,\ \ell_1 \mathbf{1}\{i=1\} - \ell_N \mathbf{1}\{i=N\} \Big). \end{aligned}$$

Finally, assume $q_i^* > 0$ for all i and $\ell_1 < \cdots < \ell_N$. Let $\bar{\ell}_{q^*} := \sum_{i=1}^{N} q_i^* \ell_i$, so $\ell_1 < \bar{\ell}_{q^*} < \ell_N$. Because the selector $\ell_1 \mathbf{1}\{i=1\} - \ell_N \mathbf{1}\{i=N\}$ is nonzero only in states $i = 1$ and $i = N$, the covariance reduces to

$$\mathrm{Cov}_{q^*}\Big(\ell_i,\ \ell_1 \mathbf{1}\{i=1\} - \ell_N \mathbf{1}\{i=N\} \Big) = q_1^*(\ell_1 - \bar{\ell}_{q^*})\ell_1 \ +\ q_N^*(\ell_N - \bar{\ell}_{q^*})(-\ell_N),$$

which is strictly negative since the first term is $(+)\cdot(-)\cdot(+)<0$ and the second term is $(+)\cdot(+)\cdot(-)<0$. Hence,

$$\frac{\partial^2 v^*}{\partial \bar{p}(\tau)\,\partial p_d(\tau)} \;<\; 0 \qquad \text{for KL with } \alpha > 0.$$

Under ambiguity neutrality ($\alpha = 0$), the robust weights collapse to the fixed prior ($q^* = \pi$) and do not respond to $(\bar{p}, p_d)$, so the covariance channel is shut down and the cross partial equals 0.